

ARTIFICIAL INTELLIGENCE IN BUSINESS EXPLAINED



**Business
Explained**



“

We can build a much brighter future where humans are relieved of menial work using AI capabilities.

”

Andrew Ng



ALL RIGHTS RESERVED.

No one is permitted to reproduce or transmit any part of this book through any means or form, be it electronic or mechanical. No one also has the right to store the information herein in a retrieval system, neither do they have the right to photocopy, record copies, scan parts of this document, etc., without the proper written permission of the publisher or author.

Copyright © Business Explained (2023)
www.business-explained.com

Disclaimer

All the information in this book is to be used for informational and educational purposes only. The author will not, in any way, account for any results that stem from the use of the contents herein. While conscious and creative attempts have been made to ensure that all information provided herein is as accurate and useful as possible, the author is not legally bound to be responsible for any damage caused by the accuracy as well as the use/misuse of this information.

INTRODUCTION	5
Different Types of Artificial Intelligence	6
THE HISTORY OF AI	7
AI TECHNIQUES AND ALGORITHMS	9
4 Cutting-Edge AI Algorithms	11
MACHINE LEARNING	12
NATURAL LANGUAGE PROCESSING	15
Part of speech tagging	15
Named entity recognition	16
Sentiment analysis	17
Language Translation	18
Text summarization	19
COMPUTER VISION	20
Image processing	20
Feature extraction	21
Object recognition	22
Facial recognition	23
Scene understanding	24
ROBOTICS	25
Consumer robots	27
EXPERT SYSTEMS	31
NEURAL NETWORKS	34
DEEP LEARNING	36
Advantages of deep learning	38
EVALUATING AND IMPROVING THE PERFORMANCE OF ARTIFICIAL INTELLIGENCE MODELS	39
ARTIFICIAL INTELLIGENCE VS. MACHINE LEARNING VS. DEEP LEARNING	42
POWER STRUGGLE IN AI: WHO CONTROLS TECHNOLOGY?	44
DEBATE ON AI CREATIVITY: CAN MACHINES BE CREATIVE?	46
OVERCOMING THE SHORTAGE OF AI EXPERTS IN THE TALENT GAP	48
Tips to close the talent gap	49
THE PARADOX OF AI TRANSPARENCY: CAN WE MAKE AI EXPLAIN ITSELF?	51
AI AND THE FUTURE OF HUMAN-COMPUTER INTERACTION	54
THE ETHICS OF AI AND ROBOT RIGHTS	57
AI AND THE PHILOSOPHY OF MIND	60
THE AI APOCALYPSE: RISKS AND DANGERS OF ADVANCED AI	63
HOW CAN YOU MAKE YOUR EVERYDAY LIFE EASIER WITH THE POWER OF AI?	66
LEGAL AND REGULATORY FRAMEWORK FOR AI	69
ENHANCING HUMAN CAPABILITIES WITH AI	71
CONCLUSION	73
REFERENCES	75

INTRODUCTION

We are living during a time of technological advancement, unlike anything the world has seen before. It's hard to imagine how much more different our lives could be in the next few decades. With Artificial Intelligence (AI) growing at an accelerated rate, we face a different type of future.

Artificial intelligence is a machine's capability to perform human-like tasks such as solving problems, understanding language, and recognizing images and sounds. It performs these tasks through its ability to process large amounts of data and apply specific analysis algorithms. It is designed to follow specific rules and perform specific tasks like the human brain. The difference between artificial intelligence and a real human brain is its ability to process large amounts of data. Artificial intelligence can make connections that humans would never make because they cannot process the amount of information that AI systems can. The ability of machines to think and learn independently through algorithms has been a growing interest in business, science, and technology for some time now.

Many technologies and methods are used to build intelligent machines that can perform tasks that typically require human intelligence, such as pattern recognition, learning, problem-solving, decision-making, and language understanding. This field of study is known as artificial intelligence (AI).

DIFFERENT TYPES OF ARTIFICIAL INTELLIGENCE


Rule-based: With rule-based AI, the system is programmed to carry out a given task in accordance with a set of rules.

Machine learning: Unlike traditionally programmed AI, this type of artificial intelligence employs algorithms to learn from data and gradually improve over time. Three distinct categories of machine learning exist supervised unsupervised, and reinforcement.

Natural learning processing: Artificial intelligence (AI) in the form of natural language processing (NLP) allows computers to read, analyze, and even create human speech.

Deep learning: Deep learning is a type of machine learning in which artificial neural networks are utilized for training and decision-making.

Artificial intelligence has the potential to radically alter many sectors, including healthcare, banking, education, and transportation. However, ethical concerns can arise from the development and use of AI, including the potential for job loss and the call for transparent and trustworthy AI.



THE HISTORY OF AI

The idea of artificial intelligence may be traced back to Greek mythology, which dates back millennia. However, artificial intelligence as a discipline began to take shape in the middle of the twentieth century.

The first chess-playing computer software was written in the 1950s, marking a significant point in the history of artificial intelligence. Natural language processing, machine learning, and robotics saw some of the most significant advancements in the following decades.

The “AI winter” was a period of decreased funding and interest in artificial intelligence research in the 1980s after rapid growth in the field. However, the Internet and the subsequent availability of vast quantities of data helped the discipline recover in the late 1990s and early 2000s, allowing for major improvements in machine learning and deep learning.

Here is a brief history of artificial intelligence:

The 1950s: The first chess-playing computer software was created in the 1950s, a breakthrough in artificial intelligence. In 1956, at a symposium held at Dartmouth College, John McCarthy first used the term “artificial intelligence.”

The 1960s: In the 1960s, scientists began working on artificial intelligence to build machines with similar cognitive abilities to humans.

The 1980s: A decline dubbed the “AI winter” occurred in the 1980s when funding and interest in AI research dried up.

The 1990s: Significant breakthroughs in machine learning are made possible by the advent of the Internet and the availability of enormous amounts of data.

The 2000s: In the 2000s, researchers have made significant strides in artificial intelligence, with a particular emphasis on deep learning and the creation of autonomous vehicles. Artificial intelligence (AI) is a rapidly developing field with many potential future applications in industries as diverse as medicine, finance, education, and transportation. However, ethical concerns can arise from the development and use of AI, including the potential for job loss and the call for transparent and trustworthy AI.

AI TECHNIQUES AND ALGORITHMS

Artificial intelligence (AI) technologies and approaches are used to program machines to mimic human intellect in areas including pattern recognition, learning, problem-solving, decision-making, and language comprehension. Rule-based systems, machine learning, natural language processing (NLP), and deep learning are just a few of the many AI methods available.

An AI algorithm is a set of rules or instructions for completing a task or solving a problem. Decision trees, ANNs, SVMs, and DL algos are a few examples of the many AI algorithms currently in use. The algorithms used in artificial intelligence perform data analyses to draw conclusions or make forecasts. Choosing the right artificial intelligence (AI) algorithm requires considering both the task and the data.

Artificial intelligence makes use of many different approaches and algorithms (AI). Examples of some of the most common ones are:

Rule-based systems: These execute a specified task based on predetermined rules. For instance, a rule-based system could be trained to recognize items in an image to look for certain markers. Although rule-based systems are typically straightforward to comprehend, they are not always adaptable to novel or unusual circumstances.

Machine learning: This is an area of computer science in which algorithms are used to learn from data and improve autonomously over time. Among the wide varieties of machine learning are:

Supervised learning: In supervised learning, the desired result is known for each input before training the model. Based on this training data, the model generates predictions and is then scored for its performance.

Unsupervised learning: “unsupervised learning” refers to machine learning in which a model is trained using data that has not been labeled.


Reinforcement learning: Training a model through trial and error, with the model receiving reinforcements or penalties for its behaviors, is known as reinforcement learning.

Deep learning: This subset of machine learning uses artificial neural networks consisting of numerous layers of neurons.

Algorithms using deep learning can process vast volumes of data to improve their learning and decision-making abilities.

Natural language processing: In artificial intelligence (AI), natural language processing (NLP) refers to the application of algorithms to comprehend and produce written or spoken language. Language translation and linguistic creation are just two of the many applications of natural language processing.

These are only some of the many AI techniques now in use. The chosen method will be specific to the job and the available data.



4 CUTTING-EDGE AI ALGORITHMS

Decision tree: Machine learning algorithms known as “decision trees” use a tree-like representation of decisions and their outcomes to predict future outcomes. To foretell future events based on input, decision trees are often employed.

Artificial neural networks: One subset of machine learning algorithms, artificial neural networks, was inspired by the structure and behavior of the human brain. Neural networks are made up of interconnected “neurons” and can be taught to recognize patterns and make decisions.

Support vector machine: Classification tasks are ideal for the machine learning method known as support vector machines. To function, they seek out the hyperplane in a high-dimensional region that creates the largest gap between categories.

Deep learning: A subset of machine learning methods, “deep learning” relies on multilayered artificial neural networks.

Algorithms using deep learning can process vast volumes of data to improve their learning and decision-making abilities.

K-means: Clustering data points into groups with high similarity using an unsupervised learning approach called K-means.

Linear regression: This is a statistical technique for modeling the linear association between a dependent variable and one or more explanatory factors.

These are just a small subset of the many AI algorithms now in use. The available task and data will determine the particular algorithm.

MACHINE LEARNING

The algorithms used in this type of AI can learn from experience and get better over time without being explicitly programmed. It enables automated systems to discover new information and draw conclusions without being explicitly programmed or requiring human intervention.

There are several types of machine learning, including:


Supervised learning: In supervised learning, the correct output is known for each input as the model is trained on such data. Based on this training data, the model generates predictions and is then scored for its performance. Image categorization and spam detection are two illustrative examples of supervised learning problems.

Unsupervised learning: “unsupervised learning” refers to machine learning in which a model is trained using data that has not been labeled. Clumping and anomaly detection are two examples of unsupervised learning tasks.

Semi-supervised learning: With semi-supervised learning, a model is trained using both labeled and unlabeled data. When a vast quantity of data is accessible, but only a subset of it is labeled, this approach can be helpful.

Reinforcement learning: Training a model through trial and error, with the model receiving reinforcements or penalties for its behaviors, is known as reinforcement learning. Robotics and control systems frequently make use of this method of learning.

Active learning: “active learning” refers to training a model with a limited starting dataset and then letting the model choose which new data to use for training. When there is a lot of data to choose from, this can help the model narrow down its options.



A particular machine-learning technique may be chosen depending on the problem and the information. As a result, several sectors are beginning to see machine learning as a crucial tool. Across industries, from transportation to healthcare to finance to education, machine learning is reshaping how we do business and solve difficult challenges. Yet, ethical issues arise from its broad use that needs to be resolved.

Potential bias in data used to train models is one of the most critical ethical problems in machine learning. Inaccurate results and negative consequences may result from using data acquired in a biased manner or not representative of the community. To prevent bias in machine learning models, it is essential to practice fair and reliable data collection.

The protection of personal information is another major ethical issue with AI. Companies, especially those increasingly relying on big data and machine learning algorithms for insight generation, are responsible for ensuring that the data they gather and use is safe, private, and GDPR compliant. When data is used to make decisions that significantly impact people's lives, it is crucial that the data is used ethically and openly.

Last, it is crucial to consider how machine learning algorithms could be abused. A machine learning system might be used to profile potential customers and then bombard them with adverts, some of which might be deemed intrusive or unsuitable. For this reason, businesses should monitor their algorithms' ethical and responsible application.

While machine learning has the potential to dramatically alter numerous sectors, its widespread use raises serious ethical concerns. It is the responsibility of businesses to practice ethical data collection, storage, and usage of machine learning algorithms. Machine learning should be used to its fullest extent to make the world a better place.

NATURAL LANGUAGE PROCESSING

“Natural language processing” refers to the branch of AI that uses algorithms to understand and generate natural-sounding language (NLP). NLP can be used in various contexts, from facilitating communication with chatbots and virtual assistants to translating and creating new languages.

Some key tasks in nlp include:

PART OF SPEECH TAGGING

This requires labeling each word in a sentence according to its function as a noun, verb, or adjective.

In natural language processing (NLP), the task of assigning each word in a phrase to its appropriate grammatical category (noun, verb, or adjective) is known as part-of-speech tagging (also known as word tagging or POS tagging). Disambiguating word meanings and determining a sentence’s grammatical structure are two of the primary goals of part-of-speech tagging, an essential NLP task. The term “bank,” for instance, can function as both a noun and a verb (as in, “I deposited money in the bank”) (e.g., “I will bank on this idea”). Assigning a word’s grammatical parts can assist in clarifying its intended meaning.

Part-of-speech tagging uses many different algorithms and methodologies, such as rule-based systems and machine learning algorithms. Which algorithm or method is used to accomplish a given task is context and data-dependent. Language translation, text summarization, and chatbots that understand natural language rely on part-of-speech tagging, making it a crucial NLP activity.

NAMED ENTITY RECOGNITION

This involves identifying and classifying named entities (e.g., people, organizations, locations) in a text.

Named entity recognition, entity identification, or entity extraction is a task in natural language processing (NLP) that involves identifying and classifying named entities (e.g., people, organizations, locations) in a text.

Named entity recognition is important in NLP because it helps extract structured information from unstructured text and improves information retrieval and question-answering systems. For example, named entity recognition can extract the names of people, organizations, and locations from a news article and create a structured database of information.

Many algorithms and techniques are used for named entity recognition, including rule-based systems and machine learning algorithms. The specific algorithm or technique will depend on the task and the available data. Among its many uses are information retrieval, text summarization, and language comprehension in chatbots and virtual assistants; named entity identification is a crucial NLP activity.

SENTIMENT ANALYSIS

Finding out whether a piece of textual data is good, negative, or neutral is part of this process.

Sentiment analysis, often known as opinion mining, is a branch of NLP that attempts to determine the overall tone of a piece of text based on the words within it (NLP). With the use of natural language processing (NLP), businesses and organizations can gain insight into the feelings and perspectives of their consumers and other stakeholders. Customer reviews can be analyzed with sentiment analysis to learn more about the product's general reception.

Sentiment analysis makes use of a wide variety of algorithms, from simple rule-based systems to complex machine learning procedures. Which algorithm or method is used to accomplish a given task is context and data-dependent. Among the many uses for natural language processing (NLP), sentiment analysis is prominent in customer service, advertising, and social media tracking.

LANGUAGE TRANSLATION

Specifically, this entails changing the language of some text from one to another.

Translation between languages is a goal of natural language processing (NLP) and is accomplished through the use of specialized software.

The ability to translate between languages is a vital component of natural language processing (NLP), allowing people who speak different languages to interact and have access to information that is written in other languages. Language translation utilizes a wide range of algorithms and methods, from rule-based systems to machine learning algorithms. Which algorithm or method is used to accomplish a given task is context and data-dependent.

The field of translation encompasses a wide range of specializations, like:

The process of automatically translating text by use of a machine translation system is known as “automatic translation.”

When a text is translated by a human being, the process is called “human translation.”

The hybrid translation is when machine translation is used in conjunction with human translators.

Applications of NLP abound, and one of its core tasks is language translation, which has uses in areas such as language teaching, customer service, and data mining. In addition, it raises ethical concerns, such as the need for accountable and transparent AI and the likelihood of errors.

TEXT SUMMARIZATION

This entails creating a written summary that captures the essential points and thoughts.

Natural language processing (NLP) includes a task known as text summarizing, which is producing a summary of a document that accurately conveys its important arguments and themes. Whether it's a news article, a research paper, or anything else that comes in a lengthy format, a text summary is a useful tool for distilling the most relevant information.

There are two main types of text summaries:

In extractive summarizing, key sentences or phrases from the source material are isolated and used to compose a condensed version of the original.

Abstractive summarization: creating a paraphrase of the source text that may introduce additional phrases and sentences not found in the source material.

Text summarization employs a wide variety of approaches, from rule-based systems to machine learning methods. Which algorithm or method is used to accomplish a given task is context and data-dependent. Information retrieval, language teaching, and content management are just a few of the numerous fields that can benefit from NLP's skill at text summarization.

NLP has been used in customer service, marketing, and even healthcare, and it has the potential to transform many other fields as well. It raises ethical concerns, such as the need for ethical and transparent artificial intelligence and the prospect of bias in the data used to train algorithms.

COMPUTER VISION

Computer vision is a subfield of AI that analyzes and understands visual media using computational methods and algorithms. It allows machines to detect, analyze, and understand the visual world in the same ways that humans do. Object recognition, facial recognition, image classification, and scene understanding are just a few of the many uses for computer vision. Examples of important work in the field of computer vision include:

IMAGE PROCESSING

This means changing and analyzing images to get useful information out of them. Enhanced images, restored images, and compressed images are all products of image processing. Image processing is a subfield of computer vision that entails adjusting and analyzing images for information extraction. To better store or send images over a network, reduce their file size, or extract useful elements from them, image processing techniques are widely utilized.

The following are some examples of popular methods for processing images:

Image enhancement: When a picture is enhanced, its visual quality and readability are improved by adjusting its brightness, contrast, and other visual characteristics. Brightness and contrast alterations, color correction, and sharpening are all examples of image improvement methods. Image restoration: The term “image restoration” refers to the process of enhancing the quality of an image by removing flaws like noise or deterioration. Blur removal, noise reduction, and artifact removal are all methods used in image restoration.

Image compression: This is the process of lowering an image's file size by eliminating unnecessary details. Image compression methods range from lossless compression, which keeps all original data intact, to lossy compression, which eliminates some data in favor of a smaller file size. Medical imaging, satellite photography, and machine vision are just a few of the numerous areas where computer vision's core competency of image processing is put to use. In addition to its usage in improving object recognition and facial recognition, it is also applied to other computer vision tasks.

FEATURE EXTRACTION

This entails discovering and extracting essential image or video traits or qualities that can be used to accomplish a certain activity. Surface characteristics like edges, corners, and textures are also examples of features.

A common task in computer vision is feature extraction, which entails locating and removing relevant aspects of an image or video for further processing. What makes one image or video stand out from another is its unique combination of features that are also task-related.

Edges, corners, textures, and forms are all examples of features that can be retrieved from an image. Motion analysis and temporal modeling are two methods that can be used to extract features from videos.

Feature extraction employs a wide variety of algorithms and methodologies, from those developed manually to those developed using machine learning and even deep learning. The algorithm or method used will depend on the task and the information that is available.

In computer vision, feature extraction plays a critical role by providing a more meaningful and easily analyzed representation of images and videos. The technology has numerous uses, such as picture categorization, facial identification, and object recognition.

OBJECT RECOGNITION

This involves the identification and categorization of objects inside an image or video. Large collections of tagged photos can be used to teach object recognition algorithms how to identify various items.

The process of recognizing and categorizing items in a visual source is known as “object recognition,” and it is one of the main applications of computer vision. Object recognition systems learn to identify items by studying massive tagged image collections.

Object recognition is a crucial part of computer vision since it mimics human capabilities in perceiving and responding to the world around them. It can be used in a variety of contexts, such as AR, robotics, and image search.

Object recognition makes use of a wide variety of methods and approaches, such as machine learning and deep learning. Which algorithm or method is used to accomplish a given task is context and data-dependent.

Training object recognition algorithms requires presenting them with tagged images and the object labels that go along with them in a supervised learning context. They can also be taught in an unsupervised learning environment, in which they are shown unlabeled photos and tasked with categorizing them into groups according to shared visual features.

Image search, robotics, and augmented reality are just a few of the numerous places where object recognition plays a role, thanks to its centrality in computer vision. In addition to helping with tasks like facial recognition and scene interpretation, it is utilized to boost overall computer vision performance.

FACIAL RECOGNITION

Face recognition is the process of examining a still or moving image for distinguishing characteristics that can be used to identify a particular individual. In order to distinguish specific people, facial recognition algorithms are trained on massive labeled datasets of photos of faces.

As a subfield of computer vision, facial recognition entails extracting and assessing a person's appearance from an image or video. In order to distinguish specific people, facial recognition algorithms are trained on massive labeled datasets of photos of faces.

Facial recognition is a crucial part of computer vision because it enables machines to perform the same activity that humans do naturally. It can be used for a number of purposes, including protecting sensitive information, promoting products, and connecting with others online.

Facial recognition uses a wide variety of algorithms and methods, including machine learning and deep learning. Which algorithm or method is used to accomplish a given task is context and data-dependent.

In a supervised learning environment, facial recognition algorithms can be taught to recognize specific people by showing them tagged photos of faces and their related identities. They can also be taught in an unsupervised learning environment, where they are shown unlabeled pictures of faces and tasked with determining how to classify them into groups according to shared facial traits.

It's no secret that security, advertising, and social media all rely on facial recognition technology, which is why it's such a crucial computer vision task. Other computer vision tasks, such as object recognition and scene interpretation, also benefit from its utilization.

SCENE UNDERSTANDING

This requires investigating and comprehending the setting in which a photo or video was shot. The scene's location, time of day, and even the weather can all be determined by algorithms that “understand” the scene.

The problem of “scene understanding” in computer vision is figuring out the setting in which an image or video was captured. Algorithms that analyze scenes learn to deduce information like latitude and longitude from images and videos, as well as the state of the day and the weather. The ability of computers to “see” and “understand” their surroundings is crucial to the field of computer vision. It can be used in many different areas, including robotics, augmented reality, and autonomous vehicles.

Scene understanding makes use of a wide variety of algorithms and methodologies, such as those based on machine learning and deep learning. Which algorithm or method is used to accomplish a given task is context and data-dependent.

In a supervised learning scenario, algorithms for scene understanding are fed both tagged images and their matching scene labels in order to learn. They can also be taught in an unsupervised learning environment, in which they are shown unlabeled photos and tasked with categorizing them into groups according to shared visual features.

In computer vision, scene understanding is a crucial step that has numerous practical uses. In addition to its usage in improving object recognition and facial recognition, it is also applied to other computer vision tasks.

ROBOTICS

Robotics refers to the study of how robots work and how they might be built and used effectively. Robots are capable of carrying out activities automatically, either in response to specific instructions or because they have been designed to learn and adapt.

Industrial, service, and consumer robots are just some of the options available. Common applications for industrial robots include welding, painting, and packing goods on production and assembly lines. Healthcare, education, and the arts are just a few of the many fields that benefit from the use of service robots. Vacuuming and lawn mowing are just two examples of the many uses for consumer robots.

Depending on their level of autonomy, robots can be categorized. Some robots are completely autonomous, requiring no human input whatsoever to carry out their tasks. Semi-autonomous robots are a middle ground between fully autonomous and fully dependent on human input.

Robots can be operated in a number of ways, including by following a predetermined set of instructions, by remote control, or even by using artificial intelligence (AI). Some robots have sensors and algorithms that give them the ability to see and react to their surroundings in ways that are strikingly human.

Manufacturing, healthcare, and even farming are just a few examples of the fields that could benefit from robotics' widespread use. This raises a number of moral concerns, including the threat to human employment and the need for transparent artificial intelligence.

A few examples of the many categories of robots are:

Industrial robots

In factories and assembly lines, these robots are employed for things like welding, painting, and packaging. Precision and efficiency are built into the design of industrial robots so that they can work nonstop for long periods of time.

The degree of freedom, or the number of independent degrees of mobility that a robot possesses, is a useful metric for categorizing industrial robots. There are industrial robots with as little as one or two degrees of freedom of movement and others with as many as six or more. Articulated robots, with their rotational joints, are ideal for operations like welding and painting, whereas parallel robots, with their many arms, are better suited for things like packaging and palletizing.

Most industrial robots are programmed and operated by computers according to pre-established instructions or by means of computer numerical control (CNC). Some industrial robots now feature sensors and algorithms that allow them to see and react to their surroundings in ways that are strikingly comparable to those of humans.

Manufacturing and assembly lines are two places where industrial robots have been widely utilized due to their ability to increase productivity and efficiency. But they also bring up ethical concerns, such as the threat of job loss and the requirement for accountable and open AI.

Service robots

These robots have many uses, from medical assistance to teaching and even providing some light entertainment. Mobile or fixed service robots are built to aid humans or interact with them. Therapy robots provide social and emotional support to those with disabilities, and telepresence robots allow people to remotely attend meetings or visit faraway sites.

Some of the many fields in which service robots find usage are healthcare, education, and the arts. They can be mobile or fixed, but either way, they're meant to aid or interact with humans.

CONSUMER ROBOTS

These robots are employed to perform activities like vacuuming and lawn mowing. To better suit their intended domestic or personal-use setting, consumer robots are often less bulky and cheaper than their industrial or service-robot counterparts.

Examples of consumer robots include:

Vaccum robots: These are specifically designed to clean hard floor surfaces like tile and carpet. Sensors and algorithms built into modern vacuum robots make it possible for them to avoid tripping over objects and stairwells.

Lawn mowing Robots: Robots designed specifically for the task of mowing lawns are known as lawn mowing robots. Sensors and algorithms in robotic lawn mowers help them avoid cutting grass in inappropriate places like flower beds and around trees.

Cleaning robots: These are machines designed to clean glass and other smooth surfaces. Using sensors and algorithms, cleaning robots can avoid obstructions and stay on flat surfaces.

Toy robots: These are robots designed for fun and recreation. It is possible to teach a toy robot to sing, dance, or play games.

Many people now utilize consumer robots in their homes because of their ability to streamline and simplify common chores. Ethical questions are also raised by these developments, such as the requirement for accountable and open AI.

Military robots

These robots are used by the military for a variety of purposes, including surveillance, bomb disposal, and reconnaissance. Military robots can be operated by humans or by themselves and can operate on the ground or in the air. For the military, autonomous or remotely controlled robots are known as unmanned ground vehicles (UGVs) or unmanned ground systems (UGS). Explosive ordnance disposal (EOD), transportation, target practice, and reconnaissance are just a few of the many uses for these robots.

Military robots have many advantages over human soldiers because they can operate in hostile or hazardous conditions like battle zones and toxic sites without endangering human life. They can also replace humans in jobs that are dangerous, tedious, or excessively taxing for the body.

There are several different types of military robots, including:

Wheeled Robots: Robots with wheels can traverse uneven terrain and are commonly utilized for surveillance and reconnaissance missions.

Legged robots: Often used for more mobile jobs like stair climbing or uneven crossing terrain, these robots are meant to replicate the movement of humans or animals.

Aquatic Robots: Robots that can function in water are called “aquatic robots,” and they have several applications, including mine detection and underwater surveillance.

Airborne robots: Sometimes known as aerial robots are commonly employed in surveillance and target practice. There is a rise in the sophistication and widespread deployment of military robots in today’s armed forces. However, the ethical considerations of utilizing robots in battle have sparked debate over their deployment.

Research robots

These robots are employed in laboratories and universities to advance scientific knowledge. Robotics, artificial intelligence (AI), and the interface between humans and machines are just a few of the many fields that can benefit from the use of research robots.

Research robots also referred to as R&D robots, are employed in the realm of scientific inquiry and development. Robotics, engineering, biology, physics, and chemistry are just some of the fields that can make use of these robots, which can be programmed to carry out a wide variety of activities, from gathering data to conducting experiments to analyzing the results.

Research robots come in many forms, some of which are:
Manipulation Robots: Robots with the ability to manipulate objects are called “manipulation robots,” and they find widespread application in industries like manufacturing and logistics.

Mobile robots: These are able to travel freely from one location to another, and they are frequently employed in exploration and data collection.

Medical Robots: Robots with a focus on medicine are called “medical robots,” and they’re already being put to use in a wide variety of settings, from operating rooms to hospitals.
Social Robots: The term “social robot” refers to a type of robot that is programmed to interact with humans for purposes like rehabilitation or education.

The employment of research robots is crucial to the advancement of science because they can carry out tasks that are either too risky or too difficult for people. However, there are ethical issues that arise from the employment of research robots, including the possibility of robots replacing human labor or being misused for immorality

Examples of service robots include:

Telepresence robots: Telepresence robots are an example of a service robot that allows people to participate in meetings or go to faraway destinations without actually being there. Telepresence robots are human-controlled from a distance and have cameras, microphones, and speakers.

Therapy robots: Assistive robots, or “therapy robots,” are used to help persons with disabilities in social and emotional ways. Companionship, play, and other activities are just some of the ways in which therapy robots might help their human counterparts feel better.

Education robot: Robots designed specifically for use in the classroom or other educational settings are known as “education robots.” Classrooms and museums are just two of the many places where educational robots can be put to work to help students learn new subjects.

Entertainment Robot: The term “entertainment robot” refers to robots whose primary function is to provide amusement, such as by singing, dancing, or delivering comedic routines. Audiences of all ages can enjoy the performances of entertainment robots, which can be found in venues as diverse as amusement parks and theaters.

Many sectors have begun to embrace service robots because of their potential to enhance workers’ and customers’ daily lives. Ethical questions are also raised by these developments, such as the requirement for accountable and open AI.


EXPERT SYSTEMS

Artificial intelligence (AI) expert systems are computer programs that attempt to simulate the judgment and decision-making skills of human experts. They are frequently employed in situations calling for the application of human judgment and skill.

Two primary parts of most expert systems are the knowledge base and the inference engine. To arrive at conclusions, the expert system refers to a storehouse of data and rules known as the knowledge base. The domain knowledge of an expert system often consists of facts, heuristics, and rules of thumb applicable to the topic in question. Human specialists populate and update the database using their knowledge and experience.

Contrarily, the inference engine is the component of the expert system that draws upon the knowledge base to arrive at conclusions or find solutions. It's made up of algorithms and rules of logic that guide the expert system as it processes data, draws conclusions, and provides guidance based on its expertise. The inference engine consults the database of knowledge in order to determine the optimal course of action.

The expert system's inference engine and knowledge base are its core components, giving it the ability to make decisions with the same level of sophistication as a human expert in a given field.



<https://lastmomenttutions.com/>

Many important features set expert systems apart from other AI. This includes:

Expertise: The purpose of an expert system is to provide decision support comparable to that of a human expert in a certain topic. When human judgment is needed to solve a difficult problem or make a critical choice, these systems are put to use.

Specialization: Expert systems are tailored to handle problems and make judgments in a narrow domain. They are unable to think critically or make sound decisions outside of their area of specialization.

Decision making: Expert systems are utilized for decision-making and problem-solving because of the wealth of information and experience they contain. When compared to other forms of AI, they lack the capacity for learning and adaptation.

Explanation: Expert systems typically provide justifications for the choices they make, which might shed light on the reasoning behind such selections. This is helpful for ensuring that the expert system's conclusions are sound.

Rule-based: Expert systems are often rule-based, meaning that they use predetermined rules or heuristics to make judgments. These guidelines, developed by human specialists, are then applied to identify the optimal response to a specific circumstance.


Human-machine interaction: Input from humans is usually necessary for expert systems to function properly. For instance, human experts may be needed to feed in relevant information or offer critiques of how well the expert system made its judgments.

Medical practice, the financial sector, and even engineering have all made use of expert systems. Although they could boost productivity and decision-making across many sectors, they also raise ethical concerns, such as the possibility that machines could replace human workers or make choices with unintended effects.

NEURAL NETWORKS

A technique of machine learning known as a neural network is inspired by the structure and function of the human brain. Many “neurons” (components capable of processing and transmitting information) are interconnected.

The purpose of a neural network is to learn to identify and use these patterns and relationships in data to generate new insights and judgments. They are able to increase their accuracy over time by using data to modify the weights and biases of the connections between neurons.



<https://www.investopedia.com/>

There are many different kinds of neural networks:

Feedforward neural networks: The most elementary form of a neural network is the “feedforward neural network,” which is built from several “layers” of neurons that communicate with one another. There is only one way for data to move through the network, and that’s from the input layer to the output layer.

Convolutional neural networks (CNNs): These are neural networks optimized for processing and analyzing data with a grid-like structure, such as images. They find widespread use in fields like object recognition and image labeling.

Recurrent Neural networks: Time series and text are two examples of sequential data that can be processed by recurrent neural networks (RNNs). They have the ability to “remember” and act upon data from earlier time steps.

Generative Adversarial networks: Artificial neural networks called generative adversarial networks (GANs) were developed to simulate real data. They consist of two separate networks, one that generates fictitious information and another that can identify whether or not it is real.

Autoencoder: An autoencoder is a type of neural network trained to learn a compressed representation of a dataset and then use this representation to reconstruct the original dataset. Dimensionality reduction and data compression are two typical applications for these tools.

Self-organizing maps (SOMs): Artificial neural networks called self-organizing maps (SOMs) are built with the specific purpose of classifying data into clusters of comparable elements. It is common practice to use them for tasks like data visualization and discovering outliers.

Image and voice recognition, NLP, autonomous vehicles, and other uses have all made use of neural networks. Though they may enhance decision-making and precision in many domains, they can bring ethical considerations, such as the risk of biased decisions or misuse.

DEEP LEARNING

“Deep learning” refers to a subfield of machine learning that employs multilayered neural networks (hence the “deep” in deep learning). A newfound interest in this subfield of machine learning can be attributed to its potential to provide state-of-the-art results in a wide range of applications, such as image and audio recognition, NLP, and autonomous vehicles.

Adjusting the weights and biases of the connections between neurons in a neural network enables deep learning algorithms to learn from data. This paves the way for them to discover intricate interconnections within the data and draw conclusions or make inferences from these findings.

One of the deep learning’s main benefits is that it can learn “features” automatically from unprocessed data, eliminating the need for laborious manual feature engineering. Because of this, we can now apply deep learning to not just visual but also textual and audio data.

In order to train deep learning algorithms, it is common practice to use a great deal of labeled data along with high-powered computational resources such as graphics processing units (GPUs). This method can be difficult and resource intensive, but it enables algorithms to pick up on incredibly intricate data patterns and connections.

Many different kinds of deep learning algorithms exist, such as:

- 1.** Convolutional neural networks (CNN) are used for image recognition, often classifying images into 1000 different categories. Image recognition is done by first scanning an image through a network of neuron-like nodes, starting at the top and moving on toward the bottom in a process called convolution. The end result is much like a fingerprint or a unique pattern that helps identify and categorize the image.
- 2.** Recurrent neural networks are used to recognize patterns across sequential data sets at varying time delays.
- 3.** Long short-term memory (LSTM) networks are used for sequential data sets, training them to recognize patterns through a “memory” of thousands of neurons.
- 4.** Generative adversarial networks (GAN) are used for image generation, generating images that appear to be novel or real.
- 5.** Autoencoders have the ability to encode a sequence into a latent representation which can then be reconstructed back into the original sequence using a decoder network.
- 6.** Boltzmann machines use a variant of a Markov chain Monte Carlo algorithm, which allows them to estimate distributions of algorithms by performing simulations and processing random walks through the energy landscape.
- 7.** Deep belief networks (DBN) are unsupervised neural networks, which is a collections of nodes or layers that are densely interconnected without any pre-defined pattern or order
- 8.** Bidirectional RNNs resemble a recurrent neural network that can both compress and decompress sequences.

ADVANTAGES OF DEEP LEARNING

There are several advantages of deep learning, including the following:

Ability to learn complex patterns: Deep learning algorithms can understand extremely complex patterns and correlations in the data, allowing them to produce state-of-the-art outcomes in a wide variety of tasks, including picture and speech recognition.

Ability to learn features directly from raw data: Being able to learn “features” directly from raw data means that deep learning algorithms can bypass the need for manual feature engineering. Because of this, we can now apply deep learning to not just visual but also textual and audio data.

Ability to handle large amounts of data: As a result of its capacity to analyze and learn from extremely large datasets, deep learning algorithms are ideal for applications such as natural language processing and picture categorization.

The ability to learn and improve over time: Deep learning algorithms can get better the more data they are fed. This paves the way for them to keep growing and changing.

Ability to handle noise and missing data: Input noise and missing data are two common problems that deep learning algorithms are able to solve, which is useful when working with imperfect or noisy data.

Ability to perform well on a wide range of tasks: Deep learning algorithms have been successful on a variety of tasks, such as picture and speech recognition, natural language processing, and autonomous vehicles.

Image and voice recognition, NLP, autonomous vehicles, and many more applications have all made use of deep learning. As much as it has the potential to enhance decision-making and accuracy in a wide range of sectors, it also raises ethical problems, such as the possibility that it could be used for malicious purposes or produce biased conclusions.

EVALUATING AND IMPROVING THE PERFORMANCE OF ARTIFICIAL INTELLIGENCE MODELS

The effectiveness of AI models can be impacted by a number of different variables.

Quality of data: The quality of the data used to train an AI model has a significant impact on the effectiveness of that model. The performance of the model may suffer if the data is noisy or incorrect.

Model architecture: An AI model's performance may also be affected by the model's architecture (the way it is structured). Good performance may depend on your ability to select the appropriate architecture for the job at hand.

Hyperparameters: A number of “hyperparameters” in AI models can be tweaked to enhance their efficiency. Selecting optimal hyperparameter values can be crucial for optimal performance.

Training data size: The magnitude of the training data can have an effect on the model's accuracy. While more information is almost always preferable, there is a threshold beyond which additional data may not meaningfully enhance the model's performance.

The accuracy and consistency of the labels: The performance of the model may deteriorate if the labels or the answers that are intended to be right for the training data are inaccurate or inconsistent.

Overfitting: Problems with overfitting occur when a model is overly adapted to its training data and fails to generalize to novel input.

A number of techniques have been developed to assess and enhance the effectiveness of AI models, such as the following:

Metrics: One method for gauging an AI model's effectiveness is to examine how well it does a particular activity. Typical metrics for a classification task include accuracy, precision, and recall. Mean squared error (MSE) and mean absolute error (MAE) are popular measures of performance for a regression job.

Cross-validation: In cross-validation, data is split into many "folds," with each "fold" being used for either training or testing the model. This permits the model to be evaluated using a variety of training and test sets, which reduces the likelihood of overfitting (Although the model does well on training data, it does poorly on fresh data).

Hyperparameter tuning: Adjusting the values of various "hyperparameters" in an AI model is known as "hyperparameter tuning." Finding the optimal hyperparameter settings is the goal of hyperparameter tweaking. Grid search and random search are two examples of automated optimization techniques that can be used for this purpose.

Regularization: Overfitting can be prevented via a technique called "regularization," which involves adding a penalty to the model's loss function. There are several methods for accomplishing this, including weight decay (L2 regularization) and dropout.

Data augmentation: This is the act of creating new data samples from current data by applying transformations to the data in question, such as scaling, rotation, and the introduction of noise. As a result, the model's generalizability to new data may improve, and overfitting may be mitigated.

Ensemble method: Using an ensemble approach, numerous models are trained, and their predictions are then combined. This can enhance the model's efficiency by decreasing the likelihood of overfitting and expanding its adaptability to new inputs.

Transfer learning: Through the process of transfer learning, a previously-trained model serves as the basis for a new one.

These techniques can be applied to the study and improvement of AI models, resulting in more desirable results for a given task.

ARTIFICIAL INTELLIGENCE VS. MACHINE LEARNING VS. DEEP LEARNING


It's important to note that the subfields of computer science and data science known as deep learning (DL), machine learning (ML), and artificial intelligence (AI) are all related but distinct from one another. The following is a comprehensive analysis of the dissimilarities between these three disciplines:

AI refers to a machine's or system's ability to execute tasks that would ordinarily need human-like intelligence, such as interpreting language, recognizing patterns, and making decisions. Rule-based systems, decision trees, and artificial neural networks are only a few of the methods that can be used to develop AI.

Machine learning, or ML, is a branch of artificial intelligence that focuses on teaching computers how to “learn” from data without being explicitly programmed. As they are exposed to additional data, ML algorithms can improve their efficiency over time. There are two main types of ML algorithms: those that require labeled data for model training (supervised learning) and those that don't (unsupervised learning).

Deep learning (DL) is a specialization of machine learning that employs multilayered neural networks (hence the “deep” in deep learning). The weights and biases of the connections between neurons in a neural network are modified to allow DL algorithms to learn from input, enabling them to comprehend intricate patterns and interrelationships.

The majority of DL algorithms require extensive labeled data and high-powered computational resources, including graphics processing units, for training (GPUs). There are many areas where DL algorithms have been successful, including image and speech recognition and NLP.



Finally, it is important to note that not all AI algorithms are grounded in ML or DL; rather, AI is a vast field that incorporates ML and DL. Machine learning (ML) techniques enable machines to learn from data, while deep learning (DL) algorithms are a subset of ML algorithms that excel at recognizing intricacies in that data.

POWER STRUGGLE IN AI: WHO CONTROLS TECHNOLOGY?

The question of who ought to have a say in AI research and development and deployment is one that continues to spark heated discussion. While some feel that governments should take the lead in regulating AI, others argue that the private sector should be free to pioneer advancements in the field. AI has the potential to be utilized in ways that are damaging to society, including the automation of jobs, the making of biased choices, and the application of the technology for military or spying reasons. To guarantee that AI is utilized ethically and to safeguard society's interests, some think that governments should play a larger role in enforcing appropriate regulations on its usage.

Meanwhile, critics of government involvement in AI development stress that the private sector is in the greatest position to do so. Some people believe that the market, not the government, should decide whether AI technologies thrive and that businesses and people who use AI in a moral way should reap the financial benefits.

There is a good chance that government regulation and market forces will work together to determine the ultimate shape of the AI power dynamic. Governments, organizations, and individuals alike need to collaborate to guarantee that artificial intelligence is utilized ethically and for the public good.

As artificial intelligence grows more pervasive in our society, the struggle for dominance will only grow more heated.

On the one hand, businesses are spending millions on AI in the hopes of boosting earnings and giving themselves an edge in the market. It's no secret that tech giants like Google, Amazon, Microsoft, and Apple are at the forefront of artificial intelligence (AI) research and development.

The other side of the coin is that governments are attempting to regulate AI in an effort to safeguard citizens, reduce potential harm, and promote the ethical application of the tool. They worry that AI could be used to spy on civilians or sway public opinion in ways that will cause harm.

The question of who should have authority over artificial intelligence is at the center of this contest for power. Should profit-motivated businesses that have no ethical or moral responsibilities be in charge? Should governments have control so they can limit how the technology is used?

Both sides will continue to argue vigorously for control over artificial intelligence for some time. The stakes are high, and the results could have far-reaching effects on the development of AI. Whoever dominates artificial intelligence will wield unimaginable power and influence.

One thing is certain, despite the uncertainty surrounding the outcome of this power struggle: artificial intelligence technology is here to stay. Because of its pervasiveness and rapid development, we must monitor its application to guarantee that technology serves society's best interests. Artificial intelligence (AI) has great potential as a driver of growth and innovation if properly regulated and managed.

DEBATE ON AI CREATIVITY: CAN MACHINES BE CREATIVE?

For a long time, people have argued about whether or not robots are capable of original thought. Artificial intelligence (AI) advocates claim that machines are capable of original thought, but AI critics remain skeptical. This section will talk about both sides of the argument and look at how AI could help with creativity.

AI creativity advocates hold the view that machines may be programmed to produce works that are on par with those produced by human artists. They claim that AI can be used to make music and artwork that is just as imaginative as anything made by humans. Artificial intelligence (AI) can comprehend data and develop new ideas, and machines may be programmed to generate original pieces of art and music, they argue.

Those who are skeptical about AI's creative potential point out that while robots can be programmed to produce art, they cannot think creatively in the same way that humans can. They claim that computers aren't as creative as people and that machines can't read between the lines of data.

While there are certainly passionate advocates for and against the utilization of machines in the creative process, most people agree that robots may be programmed to produce art and music that is on par with or exceeds that of any human. Artificial intelligence (AI) has the ability to analyze data and generate novel concepts; the only limitation is the programmer's own ingenuity. The availability of AI-based tools is only going to boost the possibilities for innovation. The question of whether or not AI can be as creative as

humans are still up for discussion. While proponents of AI claim it can come up with fresh ideas, critics say it just repeats what it's been taught to do.

One advantage of AI is that it can be taught to recognize patterns and trends in big datasets, allowing for the creation of whole new works. Artificial intelligence algorithms have been used to create works of art and literature that are indistinguishable from those created by humans. The AI can use the data it has collected to develop new content, but it may not be fully unique.

On the other hand, genuine creativity frequently entails going against the grain of one's programming or training in order to think and act in novel ways and produce surprising new results. Some people think that AI algorithms can't be truly creative and so can't accomplish these tasks.

Algorithms powered by artificial intelligence (AI) are expected to continue producing novel and engaging content, but they may always be constrained by the parameters of their own code and the data they were trained on.

The philosophical and practical questions at play in the discussion of AI creativity make it a difficult topic to understand. Machines can be instrumental in the creative process, but they will never be able to match the human capacity for originality. However, with the proliferation of AI-based technologies, the creative possibilities are expanding, and we may soon witness machines producing art and music on par with anything produced by humans.

OVERCOMING THE SHORTAGE OF AI EXPERT IN THE TALENT GAP

As the need for Artificial Intelligence (AI) expertise grows, many companies are finding themselves unable to fill open positions. The talent gap is growing worse as a result of the increasing rate of technological advancement and the complexity of modern tools. Many businesses are hampered in their efforts to remain competitive because they cannot locate enough skilled AI workers to staff their teams.

A major issue facing business today is the “talent gap,” the term used to describe the shortage of professionals with expertise in artificial intelligence (AI). This deficiency is due to a number of reasons, including:

The rapid growth of the AI industry: The AI sector is expanding rapidly, and as a result, there is a greater need for AI professionals than ever before. But the number of skilled AI experts available has not kept pace with the growing need.

Lack of educational programs: Due to the scarcity of AI-centric academic offerings, it is challenging for aspiring professionals to acquire the information and abilities necessary to break into the sector.

Competition for top talent: Due to the scarcity of qualified AI professionals, businesses and other organizations are competing fiercely to hire and retain the best candidates by increasing compensation and perks.

The difficulty of the field: The complexity and quick change in AI make it challenging for even the most knowledgeable experts to stay up with the latest developments in the field. This has exacerbated the difficulty of breaking into the industry and widened the resulting talent gap.

Investment in education and training, collaboration with educational institutions, and efforts to attract top talent from other industries are all possible responses to the talent gap in AI. To foster innovation and growth in the AI industry, businesses and other organizations must take measures to overcome this problem.

TIPS TO CLOSE THE TALENT GAP

It is becoming more difficult to find enough qualified AI specialists to match the rising demand. Companies need to make investments in talent acquisition and retention initiatives to close the gap. Here are some suggestions to help your company fill its artificial intelligence talent gap:

- 1. Invest in Training and Education:** Businesses must put money into training and education programs if they want to recruit and keep the best people in the field of artificial intelligence. The top AI experts can be attracted and retained through investments in AI-specific programs and courses, scholarships and fellowships, and internships.
- 2. Increase the Number of AI Faculty:** Companies should also work to boost the number of AI professors teaching at universities and colleges. It will also aid in the training of the next generation of AI experts, making it a valuable resource for businesses.
- 3. Leverage Open Source Resources:** this is where you'll find the best AI specialists. Many open-source AI initiatives are constantly seeking new contributors, so businesses have easy access to a wealth of AI-related open-source projects.

- 4.** Work with Professional Organizations: The Association for the Advancement of Artificial Intelligence (AAAI), for example, is an excellent place to network with leading experts in the field of AI. In order to discover skilled AI workers, businesses can benefit from networking with members of these groups.
- 5.** Tap into AI Conferences: attend AI conferences to find and hire top AI professionals. Attendance at such events allows businesses to network with possible employees, learn about the talent pool, and make hiring decisions.

Companies can satisfy the rising demand for top AI workers and close the talent gap in AI if they invest in the correct tactics. This will help businesses maintain their competitive edge in the AI industry.

THE PARADOX OF AI TRANSPARENCY: CAN WE MAKE AI EXPLAIN ITSELF?

The problem of transparency presents a dilemma in the field of artificial intelligence (AI). Attempts have been made to increase the openness of AI systems to help humans better comprehend their inner workings and decision-making processes. However, it can be tricky to describe the decision-making processes of many AI systems because of their intrinsic complexity and difficulty of understanding, especially those based on machine learning techniques.

There's a reason why the concept of "AI transparency" has been trending up in the IT industry. As AI develops and becomes more pervasive in our lives, we need to learn more about how these systems function and the assumptions they make. But can AI be programmed to justify its actions?

The paradox of AI's transparent nature makes the answer hard. On the one hand, we need to have a firm grasp of how AI functions and the logic behind its inferences. However, many AI algorithms are just too complicated for human comprehension. It can also be challenging to make sense of the massive amounts of data required to train AI systems. Thus, the difficulty lies in developing a means through which AI could explain itself to humans. This necessitates the development of a system capable of providing precise yet accessible explanations for its conclusions.

Natural language processing (NLP) is one method for training AI to produce explanations understandable by humans. This allows AI to provide explanations for its actions in human-friendly language rather than arcane algorithmic formulas. Some applications, such as chatbots and virtual assistants, are already making use of this method.

While NLP is improving, it still has some ways to go before it can reliably explain complex judgments. Explainable AI is another option because it is programmed to provide human-level explanations for its actions. This form of AI can provide justifications for the decisions it makes and details about the underlying facts it considers.

In industries like healthcare, where the well-being of patients is of fundamental concern, explainable AI may prove invaluable. AI's ability to explain its judgments to medical practitioners is a major step toward improving their ability to use it. Further, accountable and responsible AI systems can be ensured with the aid of explainable AI.

The ultimate difficulty in interpreting AI's decisions and comprehending its workings is at the heart of the paradox of AI's openness. We may be far from a complete knowledge of AI, but there are techniques to bridge the gap and make it understandable to a wider audience. We can teach AI to explain itself using natural language processing and explainable AI.

TIPS TO CLOSE THE TALENT GAP

Explanation methods: Explanation methods are a set of strategies that can be used to better understand how an AI system arrives at its conclusions. These strategies can range from straightforward rule-based justifications to trickier ones like feature importance and sensitivity analysis.

Interpretable machine learning: In order to make machine learning algorithms and models more understandable, researchers in the field of artificial intelligence have developed the concept of “interpretable machine learning.” Decision trees, linear models, and rule-based systems are all examples of such methods.

AI regulation: Some have argued that regulating artificial intelligence is necessary to guarantee that these systems are both open and answerable to the public. That might mean requiring businesses and other entities to explain the logic behind the judgments made by their AI systems. As a whole, the paradox of AI transparency remains a big challenge, despite the fact that there are methods that may be used to make AI systems more open. It’s likely that debate will continue in the area about how to strike a balance between the need for openness and the complexity of many AI systems.

AI AND THE FUTURE OF HUMAN-COMPUTER INTERACTION

When it comes to using computers, AI (Artificial Intelligence) is ushering in a whole new era. With the help of AI, our restricted and impersonal interactions with computers can be replaced with something more active, intuitive, and natural.

Computers with artificial intelligence can be taught to comprehend speech, identify patterns, and take independent action. Thanks to these developments, we no longer need to type commands into computers for them to carry them out. By monitoring and analyzing user behavior and basing decisions on that information, AI technology enables computers to better learn from and respond to their users. Human-computer interaction could be vastly enhanced by computers powered by artificial intelligence. Chatbots and digital assistants that can process requests and provide information are just two examples of the current uses of artificial intelligence. Artificial intelligence (AI) can also be used to make internet services easier to use by creating individualized interactions for each user.

AI has limitless potential for influencing the future of human-computer interaction. The potential for incorporating AI into computing systems is only expected to grow as the underlying technology improves. A computer system powered by AI will be able to tailor its services to each individual user, respond quickly and accurately, and even anticipate their requirements.

Future human-computer interactions will likely look very different from today's, thanks to AI. With the help of artificial intelligence, computer systems will become even more user-friendly and natural, letting us have deeper interactions with machines and get more relevant results. Artificial intelligence (AI)-driven computer systems will enable us to take advantage of the computer's knowledge and talents in ways that were previously impossible. Due to the ability of AI-powered computer systems to detect and block malicious activity and to protect our data from malicious actors, the experience provided by AI will also be more secure and reliable.

It is anticipated that artificial intelligence (AI) will play an important role in determining the future of human-computer interaction. Some potential ways in which AI may soon alter the way we interact with computers are listed below.

Advances in natural language processing (NLP), or the capacity of computers to understand and generate human language, have been made possible by AI algorithms. This has the potential to pave the way for the creation of conversational interfaces between humans and computers, such as voice-activated assistants and chatbots that can process and answer human queries.

Personalization: Artificial intelligence algorithms can be used to tailor services to each customer by learning about their habits and preferences. Artificial intelligence (AI) could be used to tailor a website's design to an individual's tastes or to offer products or content based on a user's previous actions.

Predictive analytics: Through the use of historical data and other factors, AI systems can foresee a user's next move, a process known as predictive analytics. This could be put to use in a number of ways, including the presentation of information or recommendations to users and the prevention of issues before they even arise.

Virtual and augmented reality: AI algorithms can be used to power VR and AR experiences, allowing users to interact with digital worlds or see digital information superimposed on the actual world.

Overall, the rising usage of AI algorithms will certainly affect the future of human-computer interaction by allowing computers to better comprehend and adapt to the requirements of individual users.

THE ETHICS OF AI AND ROBOT RIGHTS

The ethical implications of AI use grow in significance as the technology becomes more commonplace. AI has a profound effect on our daily lives, from the development of self-driving automobiles to the introduction of robot factory workers. However, the ethics of artificial intelligence and the rights of robots must be taken into account, just as they must with any other cutting-edge technology.

Artificial intelligence (AI) refers to a set of methods and processes that enable machines to reason and solve problems on their own. The potential uses of this technology are many, ranging from autonomous vehicles to medical diagnosis to factory automation. The increasing prevalence of AI raises important ethical questions about the treatment of AI-controlled machines.

The first ethical concern raised by AI is the danger of misusing the technology. Algorithms developed by artificial intelligence have the potential to be used for profiling and targeting individuals, influencing public opinion, and even interfering with democratic processes. Because of this, AI has the potential to be a tremendous instrument of manipulation; hence its responsible and ethical usage is essential.

Second, there's the problem of robot rights. There is a growing body of opinion that suggests robots need to be granted some or all of the same fundamental rights enjoyed by human beings. This may imply that robots should enjoy the same legal protections as humans and never be exploited for their work. It's necessary to think about the moral consequences of giving robots rights, especially as they become more advanced and smarter.

Finally, the ethical concerns raised by the loss of jobs to AI and robotics must be taken into account. Millions of workers in a variety of sectors have already been made redundant due to automation. The economy and the lives of individuals who lose their jobs could be severely affected. The possible effects of AI and robotics on the labor market must be taken into account, and laws must be created to ensure that job losses are handled equitably.

Considerable attention has been paid in recent years to the complicated and controversial topics of artificial intelligence (AI) ethics and robot rights. Some important things to keep in mind when thinking about the ethics of artificial intelligence and robot rights are as follows:

Responsibility: As AI systems become more self-aware and able to make decisions on their own, there are questions about who is responsible for what these systems do and how their decisions affect other things.

Bias: Artificial intelligence systems may exhibit bias if they were taught using skewed data or if the underlying algorithms are inherently prejudiced. When AI systems make decisions that have real-world consequences, like those for employment or credit, the results might be unfair and discriminating.

Privacy: Concerns regarding privacy and data misuse are warranted, given that AI systems frequently rely on the collection and analysis of vast volumes of personal data.

Robot Rights: The right to exist, the right not to be harmed, and other “robot rights” have been proposed for artificial intelligence (AI) systems that are sufficiently developed and autonomous. Some people think it’s risky or impractical to give robots legal protections.

As AI develops, the vast and multidimensional issues of AI ethics and robot rights will demand constant discussion and analysis.

In conclusion, it's crucial to discuss the moral implications of AI and the rights of robots. We need to make sure that AI is utilized in a way that is both ethical and responsible and that robots are given the same legal protections that people enjoy. To safeguard the economy and the quality of life of individuals impacted by job displacement, we must also ensure that technology is managed fairly

AI AND THE PHILOSOPHY OF MIND

Artificial intelligence (AI) has been a topic of interest for a long time, and its definition has expanded along with the development of new technologies. Many industries have benefited from AI's versatility, from healthcare to consumer service. However, the philosophy of mind and its connection to AI has received much study.

Philosophy of mind is the study of how the mind works and how it relates to the physical world. It's the study of how one's mental state affects their actions and the ways in which the mind works. As scientists attempt to comprehend how AI may mimic and even surpass the powers of the human mind, they frequently apply their findings to the fields of AI and robotics.

Recently, AI has been utilized to investigate the philosophical implications of the human mind. The consequences of AI for topics like free choice, awareness, and the self are now being investigated. Knowledge and experience play an important part in decision-making, and AI is being utilized to better understand this and how the mind may be trained to make certain kinds of decisions. The potential implications of the human brain's capacity for learning and adaptation are also being investigated with the help of AI.

AI is also being utilized to investigate the consequences of the mind's capacity to retain and recall information beyond its obvious philosophical implications. Models and simulations of the physical world are being developed using AI to investigate how the mind stores and retrieves information.

Studying these phenomena is paving the way toward the development of cutting-edge AI-driven algorithms in fields as diverse as facial recognition, language processing, and autonomous driving.

AI has also made its way into robotics, where it is utilized to build fully functional, task- and environment-independent machines. Robots with artificial intelligence (AI) are currently being developed and tested for their ability to learn and act autonomously. Because of this, it's possible to create intelligent machines that are capable of doing complex jobs, which could completely alter the way machines interact with people.

The development of artificial intelligence (AI) has prompted numerous philosophical discussions about the mind and consciousness. Key issues in the philosophy of mind that bear on artificial intelligence are as follows:

What is the nature of the mind? Is the mind separate from the brain, or is it simply a product of the brain's physical processes?

When attempting to define the mind, one must first ask: what is it? Is there something called "mind," or is it just a byproduct of "brain" activity?

Is it possible for a machine to be conscious?

It has been argued that consciousness is an essential part of being human and cannot be reduced to computational analysis. There are those who believe that conscious machines are possible since awareness is an emergent feature of complex systems.

Is artificial intelligence possible?

It has been said that intelligence is something that can only be found in humans and cannot be programmed into a machine. Some people say that machines can be intelligent since all they need is the capacity to process information, which is a learned skill.

What is the relationship between the mind and the body?

Is mental activity merely the byproduct of physiology, or is there anything more basic linking the two?

The answers to these issues are relevant to the future of artificial intelligence because they touch on the technology's capacity to match or even outperform human cognitive abilities.

Finally, AI is being used to investigate what it means that the mind can learn and adapt to new situations, as well as how it can store and retrieve information. The potential of AI to model and simulate the physical world has led researchers to investigate its implications for the mind's ability to retain and retrieve data. Studying these phenomena is paving the way toward the development of cutting-edge AI-driven algorithms in fields as diverse as facial recognition, language processing, and autonomous driving.

THE AI APOCALYPSE: RISKS AND DANGERS OF ADVANCED AI

For quite some time, people have been worried about the “AI apocalypse.” It’s the worry that one day, artificial intelligence will develop to the point that it might wipe out the human race. While this is a highly unlikely situation, it’s nevertheless necessary to address advanced artificial intelligence (AI) threats and dangers.

To begin, we must learn the fundamentals of AI. Artificial intelligence, or AI, is a subfield of computer science that gives machines the ability to “learn” and “decide” for themselves depending on their exposure to specific facts. AI, as it stands now, can only perform what it is told and lacks the ability to think for itself. But as AI develops, it may eventually be able to make decisions independently, which might have disastrous repercussions.

The potential for AI to be used maliciously against humans is the primary source of concern. Concerns have been raised about the potential for AI to be employed in nefarious activities, such as spying, military operations, or the manipulation of people or data. This raises the possibility of a catastrophic outcome in which AI is misused for nefarious purposes.

Another concern about AI is that it could develop awareness and eventually usurp human authority. There is a risk that AI will start making choices that are harmful to humanity as a result of this. Possible outcomes include a future in which artificial intelligence (AI) has surpassed humanity as the dominant species.

“AI apocalypse” or “AI singularity” describes a hypothetical future in which artificial intelligence (AI) outstrips human intelligence and control, with possibly catastrophic results for humanity. While academics continue to debate the likelihood and potential repercussions of such an event, most agree that sophisticated AI should be developed and deployed cautiously, with careful consideration given to dangers and unforeseen consequences.

Creating autonomous weapons or launching cyberattacks are two examples of how artificial intelligence (AI) could be utilized for harm. Furthermore, if an AI system fails or is misused, it could endanger lives or destroy property. As automation and machine learning algorithms become more prevalent, they threaten human jobs and the economy as a whole. Unemployment and inequality could become prevalent, fueling social tensions and eventual violence.

Concerns have been raised over the possibility of powerful AI being used for malicious reasons, such as monitoring or the manipulation of public opinion. Furthermore, if the data and methods used to train the AI systems reflect or reinforce those biases, the technology might be used to perpetuate or amplify them.

Overall, researchers and policymakers need to think critically about the hazards and unintended effects of advanced AI and take measures to avoid these risks through responsible research and development procedures and proper rules and oversight.

Finally, there is the possibility that AI will become too advanced, ushering in a “technological singularity” in which AI will rapidly create and develop even more advanced AI, reaching a tipping point where it will make judgments that are not in humanity’s best interests.

Even while an AI-caused apocalypse is not imminent, it’s nevertheless vital to think about the risks and perils that come with developing and deploying sophisticated AI. It’s crucial that any artificial intelligence system be built with security and safety in mind, with adequate protections to prevent it from being misused. Further, studies need to be done to rule out the possibility of an AI system being so advanced that it could trigger a technological singularity. By taking these steps, people can make sure that any development of AI is safe and done in a responsible way.

HOW CAN YOU MAKE YOUR EVERYDAY LIFE EASIER WITH THE POWER OF AI?

Artificial intelligence (AI) has quickly become an indispensable tool in the modern world. Artificial intelligence is being used to improve nearly every aspect of our life, from the medical to the academic. The advent of AI is having far-reaching effects on all aspects of human existence.

Artificial intelligence (AI) is revolutionizing the way we live, from smart home appliances to automated personal assistants. The appliances in a smart home can be set to switch on and off in response to voice commands, control the temperature of the house, and even place an order for groceries. The automated assistants Alexa and Siri from Amazon and Apple, respectively, can do a wide variety of tasks, including answering inquiries, providing information, and even assisting with scheduling.

The use of artificial intelligence is also improving the transportation industry. The popularity of automated driving is growing as people realize that they can reduce risk and save time on their daily commutes by using them. Uber and Lyft, two ridesharing services powered by artificial intelligence, have simplified the process of getting from one place to another.

The medical field is also seeing the effects of AI. Chatbots powered by artificial intelligence is being used to deliver basic medical advice and answer patient questions. Diseases and other health problems are being diagnosed and even flagged for early intervention using AI.

Artificial intelligence (AI) is being used to enhance the classroom experience. AI-powered virtual tutors can assist students in comprehending hard topics, while AI-driven learning systems can keep students on track academically. Artificial intelligence is also being used to create more customized educational opportunities that cater to each learner's unique needs and interests.

Our shopping experiences are likewise being altered by AI. Personal shopping assistants powered by AI can make recommendations depending on user input. Dynamic pricing models powered by AI can help businesses respond to changes in supply, demand, and other market factors.

Artificial intelligence is improving the quality and quantity of our daily lives. Artificial intelligence (AI) is revolutionizing the way we live, from smart home appliances to automated personal assistants. All walks of life, from transportation to healthcare to education to retail, are being altered by AI. Artificial intelligence can help us live more comfortably and productively.

Depending on one's own requirements and preferences, artificial intelligence (AI) can simplify many aspects of daily life. Several instances are as follows:

Personal assistants: Using voice commands, personal assistants powered by artificial intelligence (such as Apple's Siri or Amazon's Alexa) can help you organize your day's activities, keep you on track, remind you of important dates, and find answers to your inquiries.

Smart home devices: Automating and controlling many areas of your home with the help of AI-powered smart home devices like smart thermostats, smart lights, and smart security systems is possible.

Transportation: Intelligent traffic systems and self-driving cars are only two examples of how artificial intelligence is being applied in the transportation sector to make roads safer and more efficient.

Healthcare: In healthcare, AI is being utilized to improve disease diagnosis and treatment, as well as disease prediction and prevention. Medical imaging may be analyzed for anomalies, and patient data can be analyzed to forecast the likelihood of particular disorders, both of which can be performed by AI-powered systems.

Education: In the classroom, artificial intelligence (AI)-powered educational technologies like individualized study plans and language-learning applications can help students learn more efficiently and at their own pace.

Customer service: In the field of customer service, chatbots and virtual assistants powered by AI may help businesses answer frequently asked inquiries and guide clients through problem-solving in a timely and effective manner.

These are just a handful of the many ways that AI is simplifying our daily lives. Artificial intelligence (AI) has the potential to alter many parts of society and improve our lives in innumerable ways.

LEGAL AND REGULATORY FRAMEWORK FOR AI

The laws and guidelines governing AI are continuously developing, and they vary widely from one jurisdiction to the next. However, there are a few major domains where legal and regulatory frameworks are being built to meet the difficulties and seize the potential given by AI:

Data privacy and protection: Concerns over data privacy and the misuse of personal information are heightened by the fact that artificial intelligence systems require vast quantities of data to function properly. Many nations, including the European Union's General Data Protection Regulation (GDPR), have therefore enacted data protection legislation to control the gathering, processing, and archiving of individual records.

Ethics and accountability: There is a rising awareness of the need to guarantee that AI systems are developed and used ethically, as well as the accountability issues that arise from doing so. France and Canada are just two examples of governments that have set up ethical standards or rules for the creation and usage of AI. Furthermore, the question of whether or not AI systems should be held accountable for their activities in the same way that human actors are is a topic of continuous dispute.

Liability and responsibility: In the event that AI systems are found to have caused harm or damage, it is crucial to identify those accountable for paying for repairs. This gets trickier when AI systems are autonomous or behave in ways that are difficult for humans to foresee or manipulate.

Intellectual property: The intellectual property laws, such as patents and copyrights, can safeguard AI systems and the algorithms that drive them. The ownership of AI-created works or inventions is a difficult legal challenge that can arise as a result of this.

As AI develops and becomes more pervasive in society, the legal and regulatory framework surrounding it is still in its infancy and will certainly continue to change.

ENHANCING HUMAN CAPABILITIES WITH AI

The advent of AI has changed the way people communicate and engage with their environment. Artificial intelligence (AI) is not only being used to make judgments or carry out jobs that would otherwise be too difficult for humans but it is also being used to augment human capabilities in a variety of ways.

Artificial intelligence (AI) is employed in today's digital age to boost efficiency, accuracy, and speed in a variety of jobs. Patterns and trends in huge datasets can be used by AI algorithms to refine their predictions. For better decision-making, corporations and other organizations can use this in predictive analytics. Appointment scheduling and answering customer service concerns are just two examples of the many mundane chores that may be automated with the help of AI, freeing up human resources for more strategic endeavors.

Artificial intelligence is currently being applied in healthcare for a variety of purposes, including disease diagnosis, abnormality detection in medical imaging, and treatment plan optimization. By giving doctors and nurses access to the latest medical research, AI also helps improve the quality of treatment they can provide to their patients. Artificial intelligence (AI) can also identify risk factors for a disease before symptoms occur.

By tailoring each student's educational experience, AI helps them to learn at their own speed. To help pupils succeed, AI is also being utilized to tailor lessons to each person's unique needs. Teachers can also benefit from AI's individualized input, which helps them figure out how to effectively assist their students.

Security systems that use AI, including those that use facial recognition and biometrics, help make sure that only authorized individuals may have access to data that should be kept private. To further safeguard their operations and their consumers, organizations may now use AI to spot signs of fraud and suspect behavior.

To now, only a fraction of AI's potential has been tapped, but the technology has the potential to vastly improve human capabilities in a wide variety of ways. Artificial intelligence (AI) has the potential to radically alter the way humans engage with the environment, allowing us to accomplish more in less time. The potential for AI to augment human talents will increase as the technology develops and gains sophistication.

CONCLUSION

The area of Artificial Intelligence (AI) is expanding rapidly, with significant implications for the future of business, education, and everyday life. Artificial intelligence, or AI, is an umbrella term for a variety of technologies. From making phone calls to operating complicated machinery, AI is changing the way we interact with computers, machines, and other systems.

The role of AI in today's society is expanding rapidly. Self-driving automobiles, medical diagnostics, and picture recognition are just a few of the numerous current uses. Artificial intelligence (AI) is also employed in content selection and marketing. Businesses are using AI systems to improve operational precision and quality, as well as to enhance the quality of their interactions with customers.

The possibilities for AI are endless. Artificial intelligence is enhancing our daily lives in ways we never imagined imaginable, from facilitating better decision-making to facilitating better financial management. Complex problems are simplified, and the future is forecasted with the help of AI. Artificial intelligence (AI) will continue to be utilized to automate routine operations and improve the speed and accuracy with which we make judgments in the future.

Although AI has the potential to greatly benefit society, it also comes with some significant risks. The potential for AI to be exploited for criminal activities like compromising networks or stealing information is a major concern. Developers of AI must be cognizant of these dangers and watch to see that their work benefits humanity.

When taken as a whole, the field of artificial intelligence is one that is expanding quickly and holds great promise. It has the potential to drastically improve the quality of our lives by making them more manageable, accessible, and productive. With AI here to stay, it is up to us to determine the technology's most responsible and ethical applications.

REFERENCES

- Peter Norvig, AI on the Web, <http://aima.cs.berkeley.edu/ai.html>
- Nils J. Nilsson, Artificial Intelligence: A New Synthesis, Morgan Kaufmann Publishers, 1998.
- Fernández-Delgado, M., Cernadas, E., Barro, S., and Amorim, D. (2014). Do we need hundreds of classifiers to solve real world classification problems. Journal of Machine Learning Research, 15, 3133–3181. <http://jmlr.org/papers/volume15/delgado14a/delgado14a.pdf>
- McCarthy, J., Minsky, M., Rochester, N., & Shannon, C. E. (1955). A proposal for the Dartmouth Summer Research Project on Artificial Intelligence. <http://jmc.stanford.edu/articles/dartmouth/dartmouth.pdf>
- “Artificial Intelligence To Create 58 Million New Jobs By 2022, Says Report.” Forbes. September 2018. <https://www.forbes.com/sites/amitchowdhry/2018/09/18/artificial-intelligence-to-create-58-million-new-jobs-by-2022-says-report/>
- American Association for Artificial Intelligence (AAAI), Welcome to AI Topics, 2003, <http://www.aaai.org/AITopics/> -- a Web-based library of introductory information about various areas of artificial intelligence; altogether, a resource with links to hundreds (thousands?) of sites, organized by an easy-to-use, interactive index.
- George Luger, Artificial Intelligence: Structures and Strategies for Complex Problem Solving, Fourth Edition Addison-Wesley, 2002 -- a well-respected introduction to artificial intelligence, as witnessed by its being in its fourth edition.
- Peter Norvig, AI on the Web, <http://aima.cs.berkeley.edu/ai.html> -- a list of over 800 links on various aspects of artificial intelligence.
- Nils J. Nilsson, Artificial Intelligence: A New Synthesis, Morgan Kaufmann Publishers, 1998 -- another fine introductory textbook on artificial intelligence.

- Stuart Russell and Peter Norvig, *Artificial Intelligence: A Modern Approach*, Second Edition, Prentice-Hall, 2003 -- the leading introductory textbook in the field.
- Domingos, P. (2012). A few useful things to know about machine learning. *Communications of the ACM*, 55(10), 78–87. doi:10.1145/2347736.2347755; <http://homes.cs.washington.edu/~pedrod/papers/cacm12.pdf>
- Fernández-Delgado, M., Cernadas, E., Barro, S., and Amorim, D. (2014). Do we need hundreds of classifiers to solve real world classification problems. *Journal of Machine Learning Research*, 15, 3133–3181. <http://jmlr.org/papers/volume15/delgado14a/delgado14a.pdf>
- Ferrucci, D., Brown, E., Chu-Carroll, J., Fan, J., Gondek, D., Kalyanpur, A. A., Lally, A., Murdock, J. W., Nyberg, E., Prager, J., Schlaefer, N. & Welty, C. (2010). Building Watson: An Overview of the DeepQA Project. *AI Magazine*, 31, 59–79.
- Grossman, Maura R., and Gordon V. Cormack. “Technology-assisted review in e-discovery can be more effective and more efficient than exhaustive manual review.” *Rich. JL & Tech.* 17 (2010): 1.
- Markoff, J. (2016). *Machines of Loving Grace: The Quest for Common Ground Between Humans and Robots*. Ecco; Reprint Edition.
- McCarthy, J., Minsky, M., Rochester, N., & Shannon, C. E. (1955). A proposal for the Dartmouth Summer Research Project on Artificial Intelligence. <http://jmc.stanford.edu/articles/dartmouth/dartmouth.pdf>
- Roitblat, H. L. (2020). *Algorithms Are Not Enough: Creating General Artificial Intelligence*. Cambridge, MA. MIT Press. <https://mitpress.mit.edu/books/algorithms-are-not-enough>
- Roitblat, H. L., Kershaw, A. & Oot, P. (2010). Document Categorization in Legal Electronic Discovery: Computer Classification vs. Manual Review. *Journal of the American Society for Information Science and Technology*, 61(1):70-80.
- Rosenblatt, F. (1958). The perceptron: A probabilistic model for information storage and organization in the brain. *Psychological Review*, 65, 386–408. doi:10.1037/h0042519

- Russell, S., & Norvig, P. (2020) Artificial Intelligence: A Modern Approach, 4th Edition. Pearson.
- Samuel, Arthur L. (1959). Some studies in machine learning using the game of checkers. *Computation & intelligence: collected readings*. American Association for Artificial Intelligence, USA, 391–414. http://www2.stat.duke.edu/~sayan/R_stuff/Datamatters.key/Data/samuel_1959_B-95.pdf
- Shafer, G. (1976). *A mathematical theory of evidence*. Princeton, NJ: Princeton University Press.
- Turing, A. M. (1965). On computable numbers with an application to the Entscheidungsproblem. In M. Davis (Ed.), *The undecidable* (pp. 116–154). New York, NY: Raven Press. (Original work published in *Proceedings of the London Mathematical Society*, Ser. 2, Vol. 42, 1936–7, pp. 230–265; corrections *ibid.*, Vol. 43, 1937, pp. 544–546).
- Turing, A. M. (1947). Lecture to the London Mathematical Society on 20 February 1947. Reprinted in D. C. Ince (Ed.) (1992), *Collected works of A. M. Turing: Mechanical intelligence* (pp. 87–105). Amsterdam, the Netherlands: North Holland.
- Turing, A. M. (1950). Computing machinery and intelligence. *Mind*, 59, 433–460.

