Agents

Accelerator

AI Governance

AI Detector/Detection

AI Safety

AI Plugin

Alignment

Algorithm

Application Programming Interface (API)

Anthropomorphism

A class of microprocessor designed to accelerate AI applications.

Software that can perform certain tasks independently and proactively without the need for human intervention, often utilizing a suite of tools like calculators or web browsing.

This is a tool designed to identify whether something was created using generative AI tools. While these tools are not completely reliable, they can provide an indication of how likely it is that a given text is AI-generated.

Refers to the frameworks, policies, and practices established to oversee the development, deployment, and use of artificial intelligence systems. Its purpose is to ensure that AI is used responsibly, ethically, and safely, minimising risks and promoting transparency, accountability, and fairness in AI applications.

A specialised software component that integrates artificial intelligence and machine learning capabilities into external applications and services. These enable the automation of tasks, personalisation of user experiences, and optimisation of workflows by leveraging AI technologies.

An interdisciplinary field that focuses on understanding the long-term effects of artificial intelligence and the potential risks of its rapid advancement. It seeks to find ways to ensure that AI systems are safe and beneficial for society.

A set of step-by-step instructions that a computer follows to solve a specific problem or perform a task. In the context of artificial intelligence and machine learning, algorithms are used to process data, make decisions, and generate predictions.

The task of ensuring that the goals of an AI system are in line with human values.

This refers to when humans attribute human-like characteristics to non-human objects. In AI, this means perceiving a chatbot or virtual assistant as more aware, emotional, or even sentient than it truly is - such as believing it feels happy, sad, or conscious.

A set of rules and tools that allows different software applications to communicate with each other.

Artificial intelligence systems

Artificial General Intelligence (AGI)

Attention

ASI (Artificial Super Intelligence)

Autoencoders

Audit

Automated Speech Recognition

Autonomous Systems

Bias

Back Propagation

A hypothetical form of AI that could understand, learn, and apply knowledge across diverse tasks at a human-like level.

Systems capable of perceiving an environment through data acquisition, and interpreting that information to take an action or to imitate intelligent behavior given a specific goal; and learning and adapting behavior by analyzing how the environment is affected by prior actions.

Though subject to debate, this is commonly defined as artificial intelligence that surpasses the capabilities of the human mind.

In the context of neural networks, mechanisms that help the model focus on relevant parts of the input when producing an output.

Legislation that discusses an audit or evaluation of how the artificial intelligence use is functioning.

A neural network that compresses data into a smaller representation and then reconstructs it.

Indepdently uses data to make decisions and carry out steps to achieve a goal, free of or with limited human intervention.

A type of natural language processing that focuses on converting spoken language into text, enabling technologies like voice assistants to understand and respond to human speech.

An algorithm often used in training neural networks, referring to the method for computing the gradient of the loss function with respect to the weights in the network.

Refers to the tendency of artificial intelligence systems to produce unfair or prejudiced results due to the data they are trained on or the way they are designed.

Chain of Thought

Big Data

ChatGPT

Chatbot

CLIP (Contrastive Language–Image Pretraining)

Claude

Compute

Cognitive computing

Computerised Learning

Computer vision

Extremely large and complex datasets that cannot be easily processed using traditional methods. It comes from various sources, such as social media and sensors.

In AI, this term is often used to describe the sequence of reasoning steps an AI model uses to arrive at a decision.

A software application that uses artificial intelligence to simulate human-like conversations. It interacts with users via text or voice, providing information, answering questions, or assisting with tasks in real-time.

An AI language model developed by OpenAI that can understand and generate human-like text. It's designed to assist with various tasks like answering questions, writing content, and engaging in conversations, making it useful for everything from casual chats to more complex problem- solving.

An AI chatbot developed by Anthropic, designed to assist with a range of tasks such as answering questions, providing recommendations, and generating content.

An AI model developed by OpenAI that connects images and text, allowing it to understand and generate descriptions of images.

A technology that mimics human thought processes by using artificial intelligence to analyse data, learn from patterns, and make decisions, often applied to tasks like language understanding, decision-making, and problem-solving.B

The computational resources (like CPU or GPU time) used in training or running AI models.

A field of artificial intelligence that enables computers to interpret and understand visual information from the world, such as images and videos, to perform tasks like object recognition, image analysis, and facial detection.

Refers to the use of artificial intelligence technologies to enhance or automate the learning process. It involves AI-driven tools like adaptive learning platforms, intelligent tutoring systems, and personalised educational programs that adjust content based on a learner's progress and needs.

Convolutional Neural Network (CNN)

Conversational AI

Data Mining

Data Augmentation

Deep Learning

Deepfakes

Diffusion Models

Diffusion

Distortions

Disclaimer

A sub-field of AI that refers to technologies, such as chatbots or virtual assistants or agents, that users can talk to. For example, a chatbot can understand and respond to customer inquiries in a natural and human-like manner.

A type of deep learning model that processes data with a grid-like topology (e.g., an image) by applying a series of filters. Such models are often used for image recognition tasks.

Uses existing data to generate new, diverse samples, which helps enhance a model's optimisation and generalisability. By increasing the dataset size, it allows models to perform better on new, unseen data.

The process of discovering patterns, trends, and useful information from large sets of data. It involves analysing and extracting valuable insights from data, often using techniques from statistics and machine learning.

Videos, photos, or audio recordings that seem real but have been manipulated with AI. The underlying technology can replace faces, manipulate facial expressions, synthesize faces, and synthesize speech. Deepfakes can depict someone appearing to say or do something that they in fact never said or did.

A type of AI that teaches computers to learn from a large amount of data. It uses layers of artificial "neurons" to recognize patterns in things like pictures, sounds, or text. The more data it has, the better it becomes at tasks such as identifying objects in images or understanding speech.

In AI and machine learning, a technique used for generating new data by starting with a piece of real data and adding random noise. A diffusion model is a type of generative model in which a neural network is trained to predict the reverse process when random noise is added to data. Diffusion models are used to generate new samples of data that are similar to the training data.

A type of AI model that creates data like images by starting with random noise (like static on a TV) and slowly removing the noise to form a clear picture. It's used to generate new and realistic data, such as artwork or photos.

A writing that dictates a legal note on political material that discloses whether artificial intelligence had any input into the political material at hand.

Refers to errors, biases, or inaccuracies in AI systems that lead to skewed or unfair outcomes. These distortions can arise from biassed data, flawed algorithms, or improper system design, causing the AI to make incorrect or prejudiced decisions.

Embedding

Double Descent

End-to-End Learning

Emergence/Emergent Behaviour

Ethics

Environmental Impact

Explainable AI (XAI)

Expert Systems

Forward Propagation

Fine Tuning

A phenomenon in machine learning in which model performance improves with increased complexity, then worsens, then improves again.

The representation of data in a new form, often a vector space. Similar data points have more similar embeddings.

This refers to advanced artificial intelligence systems that exhibit capabilities or behaviours that were not explicitly programmed or anticipated during their development. These behaviours emerge from complex interactions within the system, often as a result of training on large datasets.

A type of machine learning model that does not require hand-engineered features. The model is simply fed raw data and expected to learn from these inputs.

Refers to the effects that developing, deploying, and maintaining artificial intelligence systems have on the environment. This includes energy consumption, carbon emissions from data centres, and the resource-intensive processes required for training large AI models, highlighting the need for sustainable practices in AI.

Refers to the principles and guidelines that govern the development and use of artificial intelligence to ensure that it is fair, transparent, and does not cause harm. It addresses issues like bias, privacy, accountability, and the societal impact of AI.

An application of artificial intelligence technologies that provides solutions to complex problems within a specific domain.

A subfield of AI focused on creating transparent models that provide clear and understandable explanations of their decisions.

The process of taking a pre-trained AI model and making small adjustments so it works better for a specific task.

In a neural network, forward propagation is the process where input data is fed into the network and passed through each layer (from the input layer to the hidden layers and finally to the output layer) to produce the output. The network applies weights and biases to the inputs and uses activation functions to generate the final output.

General Adversarial Network (GAN)

Foundation Models

Generative AI (GenAI)

Generative Adversarial Network (GAN)

GPT (Generative Pretrained Transformer)

GitHub Copilot

Gradient Descent

GPU (Graphics Processing Unit)

Hidden Layer

Hallucinations

Large AI models trained on a wide range of data, making them capable of understanding and performing many tasks. They serve as a base for other, more specific AI applications, like chatbots or language translators, by fine-tuning them for particular tasks.

A type of machine learning model used to generate new data similar to some existing data. It pits two neural networks against each other: a “generator,” which creates new data, and a “discriminator” which tries to distinguish that data from real data.

A type of AI model with two parts: one creates new data (like images), and the other checks if it looks real or fake. They work against each other, improving over time until the generated data looks real. GANs are often used to create realistic images, videos, or music.

Artificial intelligence systems designed to create new content, such as text, images, or music, based on patterns learned from existing data. Unlike traditional AI that only analyses or classifies data, it generates original outputs.

An AI-powered coding assistant embedded within development environments, designed to help developers with tasks like writing, completing, and refactoring code, enhancing productivity through intelligent code suggestions and support.

A large-scale AI language model developed by OpenAI that generates human-like text.

A specialized type of microprocessor primarily designed to quickly render images for output to a display. GPUs are also highly efficient at performing the calculations needed to train and run neural networks.

In machine learning, gradient descent is an optimization method that gradually adjusts a model’s parameters based on the direction of largest improvement in its loss function. In linear regression, for example, gradient descent helps find the best-fit line by repeatedly refining the line’s slope and intercept to minimize prediction errors.

Occur when an AI generates incorrect or made-up information that sounds believable. This happens because the AI predicts patterns without truly understanding or verifying facts, leading to confident but false responses.

Layers of artificial neurons in a neural network that are not directly connected to the input or output.

Image Recognition

Hyperparameter Tuning

Instruction Tuning

Inference

Latent Space

Large Language Model (LLM)

Machine Learning

Loss Function (or Cost Function)

Microsoft Copilot

Microsoft Azure

The process of selecting the appropriate values for the hyperparameters (parameters that are not learned from the data) of a machine learning model.

Image recognition in AI is the ability of machines to identify and classify objects or features in images using deep learning algorithms. It's used in applications like facial recognition and object detection.

The process of making predictions with a trained machine learning model.

A technique in machine learning where models are fine-tuned based on specific instructions given in the dataset.

A type of AI model that can generate human-like text and is trained on a broad dataset.

In machine learning, this term refers to the compressed representation of data that a model (like a neural network) creates.

A function that a machine learning model seeks to minimize during training. It quantifies how far the model's predictions are from the true values.

A branch of artificial intelligence that enables computers to learn from data and improve their performance on tasks without being explicitly programmed.

A cloud computing platform offering services like application hosting, data analytics, AI tools, and storage, enabling organisations to build, deploy, and manage solutions across cloud and hybrid environments.

This is an AI-powered assistant embedded within Microsoft applications, designed to help users with tasks such as drafting, summarising, and analysing content, enhancing productivity through intelligent automation and support.

Multimodal

Mixture of Experts

NeRF (Neural Radiance Fields)

Natural Language Processing (NLP)

Objective Function

Neural Networks

Optical Character Recognition (OCR)

OpenAI

Overfitting

Optimisation

A machine learning technique where several specialized submodels (the “experts”) are trained, and their predictions are combined in a way that depends on the input.

This is a subfield of machine learning focused on interpreting and integrating multiple types of data, such as text, images, and audio, to build models that can process and relate information from various sources.

A field of artificial intelligence that focuses on the interaction between computers and human language. It enables machines to understand, interpret, and respond to written or spoken language.

A method for creating a 3D scene from 2D images using a neural network. It can be used for photorealistic rendering, view synthesis, and more.

A type of artificial intelligence model inspired by the structure of the human brain. They consist of layers of interconnected nodes, or “neurons,” that process information.

A function that a machine learning model seeks to maximize or minimize during training.

A leading AI research organisation that creates and promotes AI technologies, including tools like ChatGPT.

This is the process of converting images of printed, handwritten, or typed text into a machine-readable text format. This technology allows computers to recognise and interpret text from images, making it easier to edit, search, and store information digitally.

Involves improving the performance and efficiency of AI algorithms and models to achieve better results.

A modeling error that occurs when a function is too closely fit to a limited set of data points, resulting in poor predictive performance when applied to unseen data.

Pre-training

Parameters

Prompt

Predictive analytics

Recursive Prompting

Prompt Engineering

Reinforcement Learning

Regularization

RLHF (Reinforcement Learning from Human Feedback)

Responsible AI

In machine learning, parameters are the internal variables that the model uses to make predictions. They are learned from the training data during the training process. For example, in a neural network, the weights and biases are parameters.

Refers to the initial phase of training a model on a large dataset before it is fine-tuned for specific tasks. During this, the model learns general patterns, language structures, and contextual relationships from the data, helping it to understand and generate text more effectively.

The use of AI to analyse current and past data to make predictions about future events. It helps businesses and organisations anticipate trends or outcomes, such as predicting customer behaviour or sales trends.

The initial context or instruction that sets the task or query for the model.

The skill of creating questions or instructions for AI to get the answers or responses you want. By carefully choosing the right words and structure, you help the AI understand what you're asking and produce better results.

A strategy for guiding AI models to produce better quality output. It works by giving the model a sequence of prompts or questions that build on its previous responses, gradually refining the context and the AI's understanding to achieve the desired outcome.

In machine learning, regularization is a technique used to prevent overfitting by adding a penalty term to the model's loss function. This penalty discourages the model from excessively relying on complex patterns in the training data, promoting more generalizable and less prone-to-overfitting models.

A type of machine learning where an AI learns to make decisions by interacting with an environment and receiving feedback in the form of rewards or penalties.

The development and use of artificial intelligence systems in an ethical, transparent, and accountable way. The aims of this are to ensure fairness, prevent biases, and safeguard privacy.

A method to train an AI model by learning from feedback given by humans on model outputs.

Sentiment analysis

Robotics (Autonomous Robot)

Singularity

Single-shot prompting

Supervised Learning

Stochastic Parrot

Synthetic Media

Symbolic Artificial Intelligence

Token and Tokenisation

TensorFlow

A branch of technology that involves designing and creating robots that can perform tasks on their own without human control.S

A method in AI that examines text (like reviews or social media posts) to determine the writer’s feelings or opinions. It helps identify whether the sentiment is positive, negative, or neutral, allowing businesses to understand customer feedback and public opinion.

This involves providing a model with a single example or prompt to understand a specific task and generate a response or prediction. This approach allows the model to learn from just one instance, enabling it to perform the desired task effectively without needing extensive training data. Also known as one-shot prompting.

In the context of AI, the singularity (also known as the technological singularity) refers to a hypothetical future point in time when technological growth becomes uncontrollable and irreversible, leading to unforeseeable changes to human civilization.

AI systems that use statistical relationships from massive datasets to generate human-like text, but they lack true semantic understanding behind the word patterns.

A type of machine learning where an AI model is trained on labelled data, meaning each training example includes the correct answer. The model learns to make predictions based on this data.

A type of AI that utilizes symbolic reasoning to solve problems and represent knowledge.

Any human being's voice, photograph, image, video or other human likeness created, reproduced, or modified by computer, using artificial intelligence or software algorithm, to be indistinguishable to a reasonable viewer from a natural person.

An open-source machine learning platform developed by Google that is used to build and train machine learning models.

A fundamental process in Natural Language Processing (NLP) and machine learning that involves breaking down text into smaller units called tokens. These tokens can be words, subwords, or even characters. The primary goal of tokenisation is to convert text into a format that computational models can more easily analyse and understand.

Training Data

TPU (Tensor Processing Unit)

Transformer

Transfer Learning

Underfitting

Turing Test

Validation Data

Unsupervised Learning

Voice-cloning

Voice Processing

A type of microprocessor developed by Google specifically for accelerating machine learning workloads.

The dataset used to train an AI or machine learning model. The quality, size, and representativeness of the training data are crucial factors in the model's performance.

A method in machine learning where a pre-trained model is used on a new problem.

A specific type of neural network architecture used primarily for processing sequential data such as natural language. Known for their ability to handle long-range dependencies in data, thanks to a mechanism called "attention," which allows the model to weigh the importance of different inputs when producing an output.

An evaluation of a machine's ability to exhibit human-like intelligence. An AI passes the test if it can engage in a conversation with a human without the human realising they are talking to a machine.S (continued) Strong AI (also known as artificial general intelligence – AGI) Refers to AI systems that possess generalised intelligence and capabilities on par with human cognition. Unlike narrow or weak AI, which is designed for specific tasks, strong AI aims to replicate human-like reasoning, learning, and

A modeling error in statistics and machine learning when a statistical model or machine learning algorithm cannot adequately capture the underlying structure of the data.

A type of machine learning where the model is not provided with labeled training data, and instead must identify patterns in the data on its own.

A subset of the dataset used in machine learning that is separate from the training and test datasets. It's used to tune the hyperparameters (i.e., architecture, not weights) of a model.

Involves converting spoken language into text (speech-to-text) and then converting text back into spoken language (text-to-speech). This enables AI systems to comprehend and produce human speech, facilitating natural and efficient interactions between humans and machines.

A technology that uses AI to create a digital copy of a person's voice. It can replicate the tone, pitch, and speech patterns, allowing the AI to generate new speech that sounds like the original person.

Weak AI

Watermarking

XAI (Explainable AI)

Whisper

Zero-shot Prompting

Zero-shot Learning

The act of embedding information, which is typically difficult to remove, into outputs created by AI – including into outputs such as photos, videos, audio clips, or text – for the purposes of verifying the authenticity of the output or the identity or characteristics of its provenance, modifications, or conveyance.

Also known as narrow AI, this refers to artificial intelligence systems designed to perform specific tasks without general understanding or awareness. Examples include chatbots or recommendation systems.

An AI system developed by OpenAI to perform automatic speech recognition (ASR), the task of transcribing spoken language into text. Whisper is trained on a large dataset of diverse audio, allowing it to effectively manage various accents, background noise and other audio variations.

A subfield of AI focused on creating transparent models that provide clear and understandable explanations of their decisions.

A type of machine learning where the model makes predictions for conditions not seen during training, without any fine-tuning.

This is a technique that leverages the generalisation capabilities of LLMs to perform new tasks without any prior specific training or examples.