Glossary of Common AI Terms

AI Agent

An AI agent is a software program that interacts with its environment, perceives information, and takes actions based on its observations to achieve specific goals. AI agents can range from simple systems like chatbots to more complex systems such as autonomous drones or virtual assistants like Google Assistant.

2. Algorithm

An algorithm is a set of rules or instructions followed by a computer to solve a problem or perform a task. In AI, algorithms process input data to make decisions or predictions. For instance, recommendation systems on streaming platforms use algorithms to suggest content based on user behavior.

3. Artificial Intelligence (AI)

AI refers to the simulation of human intelligence in machines, enabling them to perform tasks like learning, reasoning, and problem-solving. AI includes subfields like machine learning and natural language processing. Examples include autonomous vehicles, voice assistants, and image recognition systems.

4. Artificial General Intelligence (AGI)

AGI is a theoretical form of AI that could understand, learn, and apply knowledge across a wide range of tasks, similar to a human. Unlike today's AI, which excels at narrow, specific tasks, AGI would be capable of generalizing knowledge and adapting to new, unfamiliar tasks without retraining.

5. Bias

Bias in AI occurs when a model produces skewed results due to imbalances in its training data or design. Bias can lead to unfair outcomes, such as facial recognition systems misidentifying people from certain ethnic backgrounds. Mitigating bias involves using diverse, representative datasets and refining algorithms.

6. Bot

Origin: The term "bot" originates from "robot" and initially referred to automated programs performing repetitive tasks.

A bot is an automated software application designed to perform tasks over the internet, such as answering customer queries or crawling websites for search engines. Bots can be useful (e.g., virtual assistants) or malicious (e.g., spam bots).

7. Chatbot

A chatbot is a type of bot specifically designed to simulate conversation with users, usually through text or voice. Powered by AI and natural language processing, chatbots are commonly used in customer service to automate responses to routine queries or provide support around the clock.

8. Deep Learning

Deep learning is a subset of machine learning that uses artificial neural networks with multiple layers (hence "deep") to model complex patterns in data. It excels at tasks like image classification, speech recognition, and language translation. An example is how deep learning enables self-driving cars to recognize pedestrians and road signs.

9. Generative AI

Generative AI refers to systems that create new data (text, images, music) based on the patterns they've learned from existing data. For instance, GPT models can generate coherent text or even entire articles based on a user prompt, while tools like DALL·E create images from textual descriptions.

10. Hallucination

In AI, hallucination refers to when a model, particularly a language model, generates information that is factually incorrect or nonsensical. This can happen when the model produces responses that seem plausible but are not grounded in the training data or real-world facts. Hallucinations are common in generative AI models when they "invent" details.

11. Machine Learning (ML)

Machine learning is a subset of AI where algorithms learn from data to make decisions or predictions without being explicitly programmed for every outcome. ML models improve their accuracy over time by finding patterns in the data. Common uses include recommendation systems, fraud detection, and personalized marketing.

12. Model

A model in AI is a mathematical representation trained on data to perform specific tasks, such as classification or prediction. Once trained, the model can analyze new inputs and produce outputs based on the patterns it has learned. For example, image recognition models classify images by analyzing pixel patterns.

13. Natural Language Processing (NLP)

NLP is a branch of AI that focuses on enabling machines to understand, interpret, and respond to human language. It combines linguistics with machine learning to analyze text and speech, powering applications like voice assistants, machine translation, and text summarization. An example of NLP in action is how chatbots understand and respond to user queries.

14. Neural Network

A neural network is a type of AI model inspired by the structure of the human brain, consisting of layers of interconnected nodes (neurons). These networks are used to recognize patterns in data. Neural networks are fundamental to deep learning and are employed in tasks such as image recognition, speech processing, and language translation.

15. Overfitting

Overfitting occurs when an AI model learns not only the general patterns in its training data but also the noise or irrelevant details, leading to poor performance on new, unseen data. In other words, the model becomes too specialized to the training data, limiting its ability to generalize to real-world scenarios.

16. Predictive Analytics

Predictive analytics involves using AI and statistical techniques to analyze historical data and predict future outcomes. It is widely used in industries like finance (for stock market predictions), healthcare (to predict disease outbreaks), and marketing (to anticipate customer behavior). AI models analyze trends and patterns to make accurate forecasts.

17. Reinforcement Learning

Reinforcement learning (RL) is a type of machine learning where an AI agent learns by interacting with an environment and receiving feedback in the form of rewards or penalties. RL is often used in applications like game-playing AI (e.g., AlphaGo) and robotics, where the agent learns to perform tasks by trial and error.

18. Supervised Learning

Supervised learning is a machine learning technique in which a model is trained on labeled data. The model learns by comparing its predictions to the correct outputs provided during training. Supervised learning is commonly used for classification tasks (e.g., identifying spam emails) and regression tasks (e.g., predicting house prices).

19. Turing Test

The Turing Test, proposed by Alan Turing, is a measure of a machine's ability to exhibit behavior indistinguishable from a human. If a human evaluator cannot reliably tell whether they are interacting with a machine or a person, the machine is considered to have passed the test. The test remains a benchmark for evaluating the "intelligence" of AI systems.

20. Training Data

Training data refers to the dataset used to train an AI model. It consists of examples (labeled or unlabeled) that teach the model how to perform a task, such as recognizing images or predicting outcomes. The quality, size, and diversity of the training data significantly impact the model's performance and accuracy.

21. Virtual Assistant

A virtual assistant is an AI-powered software application designed to help users perform tasks such as scheduling appointments, sending messages, or controlling smart home devices. Examples include Apple's Siri, Google Assistant, and Amazon's Alexa, which can understand voice commands and provide personalized responses based on user preferences.