Glossary of Generative AI Terms

Artificial Intelligence (AI):
A branch of computer science focused on creating systems capable of performing tasks that typically require human intelligence, such as learning, reasoning, problem-solving, and understanding natural language.

Machine Learning (ML):
A subset of AI that enables computers to learn from data and improve their performance on tasks without being explicitly programmed. Machine learning algorithms use data to identify patterns and make decisions or predictions.

Deep Learning:
A subset of machine learning that uses neural networks with many layers (often referred to as "deep" neural networks). It's particularly effective for tasks like image recognition, natural language processing, and generating text.

Neural Networks:
A series of algorithms modeled after the human brain, designed to recognize patterns. Neural networks are composed of layers of interconnected "neurons," where data passes through and is processed to make predictions or decisions.

Generative AI (GenAI):
A type of AI that can create new content such as text, images, audio, or video, based on patterns it has learned from large datasets. ChatGPT, which generates human-like text responses, is an example of generative AI.

Large Language Model (LLM):
A type of AI model, like ChatGPT, trained on vast amounts of text data to understand and generate natural language. LLMs can answer questions, complete sentences, translate text, and assist with various language-based tasks.

Natural Language Processing (NLP):
A field of AI focused on the interaction between computers and humans through natural language. NLP enables AI models to understand, interpret, and generate human language. ChatGPT is an NLP-based tool.

Transformer Model:
A type of neural network architecture that has become the backbone of most state-of-the-

art AI models, including GPT (Generative Pre-trained Transformer). Transformers process input data in parallel and can handle complex language tasks more efficiently.

Pre-training:
The process where an AI model, such as GPT, is trained on vast amounts of text data before being fine-tuned for specific tasks. Pre-training allows the model to learn language patterns, grammar, facts, and general knowledge.

Fine-Tuning:
A process that takes a pre-trained AI model and adjusts it to perform specific tasks or specialize in certain areas, often using a smaller dataset that is relevant to the task.

Prompt:
The input or question that a user provides to a generative AI model to generate a response. The more specific and clear a prompt is, the better the output from the AI.

Token:
A token is a piece of text (word or sub-word) used by AI models to process input and generate responses. In languages, tokens can be as short as one letter or as long as a complete word.

Response:
The output generated by an AI model based on the prompt or input it received. For ChatGPT, the response would be the text it generates in reply to a user's question or request.

Bias in AI:
AI models can reflect the biases present in the data they are trained on. Bias in AI refers to the unintentional favoring of certain viewpoints, demographics, or data types over others, which can affect the fairness of AI outputs.

Overfitting:
A scenario where an AI model is too closely tuned to its training data, making it perform well on that data but poorly on new, unseen data. Overfitting reduces the model's ability to generalize its learning to broader contexts.

Turing Test:
A test proposed by Alan Turing to assess whether a machine can exhibit intelligent behavior indistinguishable from that of a human. If an AI passes the Turing Test, it means its responses are human-like to the point where people can't tell it's a machine.

Chatbot:
A software application designed to simulate conversation with human users, especially over the internet. ChatGPT is an advanced form of a chatbot that uses AI to generate sophisticated, natural-sounding responses.

Training Data:
The dataset used to train an AI model. In the case of generative AI, training data typically includes text, images, or other types of content that the model learns from to generate outputs.

Big Ideas in AI

Algorithms

- In computing, an **algorithm** is a precise list of operations that could be done by a machine.

Artificial Intelligence

Artificial Intelligence (AI) is the theory and development of computer systems able to perform tasks that normally require human intelligence, such as visual perception, speech recognition, decision-making, and translation between languages.[1]

Machine Learning

Machine learning (ML) is an application of artificial intelligence (AI) that provides systems the ability to automatically learn and improve from experience without being explicitly programmed. Machine learning focuses on the development of computer programs that can access data and use it.[2]

Types of Machine Learning

Deep Learning

Deep learning is a subset of machine learning and functions in a similar way. However, its capabilities are different. While basic machine learning models do become progressively better at whatever their function is, they still need some guidance. If an AI algorithm returns an inaccurate prediction, then an engineer has to step in and make adjustments. With a deep learning model, an algorithm can determine on its own if a prediction is accurate or not through its own neural network.[3]

Imitation Learning Model

In **imitation learning** models, a machine observes an expert demonstrate the desired behavior and then attempts to mimic it. Applications for this type of AI include the development of machines that can perform typically human tasks.

Language Models

Language models are a form of AI machine learning where algorithms are used to learn the probabilities of a sequence of words that occur in a commonly spoken language (say,

English) and predict the next possible word in that sequence. In practice, it gives the probability of a certain word sequence being "valid." Validity in this context does not refer to grammatical validity. Instead, it means that it resembles how people write, which is what the language model learns.[4,5]

Neural Network

A **neural network** is a method in artificial intelligence that teaches computers to process data in a way that is inspired by the human brain. It is a type of machine learning process, called deep learning, that uses interconnected nodes or neurons in a layered structure that resembles the human brain. It creates an adaptive system that computers use to learn from their mistakes and improve continuously. Thus, artificial neural networks attempt to solve complicated problems, like summarizing documents or recognizing faces, with greater accuracy.[6]

Reinforcement Learning

In **reinforcement learning**, AI agents are attempting to find the optimal way to accomplish a particular goal or improve performance on a specific task. As the agent takes action that goes toward the goal, it receives a reward. The overall aim is to predict the best next step to take to earn the biggest final reward.[7]

Supervised Learning Model

In a **supervised learning model** (as opposed to an unsupervised model), the algorithm learns on a labeled dataset, providing an answer key that the algorithm can use to evaluate its accuracy on training data.[8]

Unsupervised Learning Model

An **unsupervised model** (as opposed to a supervised model) provides unlabeled data that the algorithm tries to make sense of by extracting features and patterns on its own.[9]