

AI and generative AI cheat sheet

Artificial Intelligence

Artificial intelligence (AI):

The ability of machines to perform tasks that typically require human intelligence, such as learning, reasoning, problem-solving, and decision-making.

Artificial intelligence for IT operations (AIOps):

The application of AI, machine learning (ML), and analytics to improve the day-to-day operational work for IT operations teams.

Deep learning:

A subfield of neural networks that has many layers, allowing it to learn significantly more complex relationships than other machine learning algorithms.

Machine learning (ML):

A branch of AI that focuses on the use of data and algorithms to imitate the way humans learn, gradually improving accuracy over time. One way they do this is with neural networks that utilize interconnected nodes in a layered structure that resembles the human brain.

Natural language processing (NLP):

A subfield of artificial intelligence that focuses on enabling machines to understand, interpret, and generate human language.

Neural networks:

A type of machine learning algorithm that consists of interconnected layers of nodes that process and transmit information. It is inspired by the structure and function of the human brain.

Generative AI

Generative AI:

A branch of AI centered around computer models capable of generating original content that mimics human creativity. By leveraging the power of large language models, neural networks, and ML, generative AI models are trained to learn the underlying structures, relationships, and patterns to produce new and unique outputs like images, video, code, and more.

Large language model (LLM):

A deep learning algorithm that can perform a variety of natural language processing (NLP) tasks.

Prompting:

A prompt is an instruction given to an LLM. Few-shot prompting teaches the model to predict outputs through the use of examples.

Retrieval augmented generation (RAG):

A framework that enables users to “feed” an LLM private or proprietary, external data so it has the most up-to-date information.

Hallucinations:

When an LLM produces a false or nonsensical output or one that does not match the user’s intent. Because large language models are not search engines or databases — they only predict the next syntactically correct word or phrase — they can appear to produce results that are factually incorrect or contradictory, especially if the data set they are trained on contains contradictory information.

