

Assessing a vendor's specific AI features is tough, particularly when the terms may be unfamiliar. This primer provides relatively simple explanations of some common AI terms.

Autonomous Agents

Generative AI models that run "on a loop" to generate self-directed instructions and actions with each iteration (loop). They are highly scalable because they do not rely on humans to guide their conversations. Examples of autonomous agents include the "Westworld" simulation, Camel, BabyAGI and AutoGPT.

Deep Learning

A term associated with neural network learning: it's when a neural network that is more than three layers deep ingests and processes large datasets of unstructured data, thus reducing some of the human labor involved in training an AI model.

Deep learning is often used to do complex data processing that allows an AI model to analyze and learn from data to detect patterns, do predictive analytics, and rapidly assess large volumes of data.

Diffusion Model

This is how people give instructions and otherwise speak with an AI. This is typically done via natural language – writing an instruction in text, but it can also be done with computer code. Prompt engineering is the process of revising and/or tailoring prompts to get the AI to generate a (high quality) desired result.

Foundation model (also known as transformer model; FM)

An artificial intelligence model that is pretrained on a broad range of data and can then be adapted to other tasks. FMs are (pre)trained on vast datasets and parameters in an unlabeled / unsupervised fashion.

Neural networks and Deep Learning help make FMs possible. Foundation models use a deep learning model called the transformer model, in which sequentially input data, like a sentence, is processed all at once – much in the same way we read or hear sentences and rapidly process how the individual words all relate to one another. So sometimes, a foundational model is called a transformer model.

Generative AI	Generative AI works by identifying patterns in its training data, then creating output that matches the patterns it identifies. The most well-known implementation of generative AI is ChatGPT for text, but there's also DALL-E for images. Generative AI tools are also being incorporated into products like Microsoft's Copilot, Salesforce's Einstein, Slack, and many others. To date, most generative AI implementations are designed to assist humans (like a power tool) versus function without humans.
Hallucination	This happens when LLMs generate text or images that are factually incorrect or nonsensical.
Large language model (LLM)	A neural network trained on large quantities of unlabeled text. These next-step neural networks use a method of data training called unsupervised learning, where a model is given a dataset without explicit instructions on what to do with the data, and the model learns and infers from the data, setting up its own classification systems and performing analytic outcomes based on how it's processed the data.
Model	The technology and concept at the heart of artificial intelligence: a model is the program that analyzes data sets to find patterns and provide data-driven analytics in response.
Modeling	A model is the concept, modeling is the process of creating, training, and deploying a model that will analyze available datasets and provide data-driven analytics in response.

Neural network (NN)

First proposed in 1944 by Warren McCulloch and Walter Pitts, neural networks (NNs) are a computing model based on thousands of processing nodes that are both multilayered and densely connected.

NNs are comprised of multiple node layers (input, hidden and output); each layer is itself comprised of individual nodes (artificial neurons). These nodes connect to each other in multiple ways; an individual node can receive data from a node it's connected to in the layer below it, and relay data to a node it's connected to in the layer above it. Labeled training data (i.e., a set of photos of animals in which the animal is labeled) goes into the input layer, then is processed by passing through the (multiple) hidden layers until the output is achieved. Using our animal example, a NN might go through multiple rounds of training before it can correctly distinguish between a Persian cat, a llama, a polar bear and a Samoyed dog.

Each node has its own linear regression model comprised of input data, weights, a bias (or threshold) and an output. "Training" a neural net describes the process of adjusting the weights and threshold assigned to the nodes as the training data passes through each layer until the training data exercises consistent yield similar outputs.

Prompting

This is how people give instructions and otherwise speak with an AI. This is typically done via natural language – writing an instruction in text, but it can also be done with computer code. Prompt engineering is the process of revising and/or tailoring prompts to get the AI to generate a (high quality) desired result.

Predictive AI	A foundation model built to infer information about different data points so they can make decisions. These models have mostly been used to automate human processes to a greater or lesser extent. An example of the latter would be email filtering. An example of the former might be a medical application. These predictive AI models must exhibit (very) high levels of accuracy because they can operate with minimal human oversight.
Training Data	Also referred to as a dataset, models can be trained on text, images, sound, sensor data, video data. LLMs were trained on text-based material gathered from the Internet.
Unsupervised learning	This type of training neural networks relies on algorithms to analyze and cluster unlabeled data sets, and this method helps discover hidden patterns or groupings without the need for human intervention. There are three types of ML algorithms that pertain to UL: Clustering, Association and Dimensionality Reduction. An example of clustering algorithms being used includes customer segmentation (e.g., purchase history, social media, geography, etc.). Association uses different rules to find relationships between variables in a dataset (e.g., customers who bought this also bought XYZ). Dimensionality reduction involves techniques that reduce number of input variables in a dataset so redundant parameters do not overrepresent the impact on the outcome (e.g., auto encoders remove noise from visual data to improve picture quality). Somewhat confusingly, sometimes this practice of using unlabeled data to train a neural network is also called semi-supervised learning.