


A Glossary of Common Artificial Intelligence Terms for Business Leaders

By Shahriar Parvarandeh, Ned Calder, and Freddy Solis

November 2023



A foundational, nontechnical, and shared understanding of AI concepts is vital for empowering leadership teams to understand the potential and challenges of this technology. Because AI is pervasive in its reach as a general-purpose technology, executives from different functions—marketing, HR, and R&D—are exposed to distinct tools, use cases, and impacts. This makes them likely to interpret terms differently, at the expense of recognizing the true breadth and depth of AI’s implications for an organization. The definitions provided below are meant to serve as a reference guide and help corporate leaders and stakeholders as they explore common ground and develop alignment around their aspirations and goals for incorporating AI into their workflow and organization. (Readers are also encouraged to explore Innosight’s publication *Leading into the Age of AI: A Five-Part Blueprint for Empowering Corporate Transformation*, from which this glossary is derived, available [here](#).)

Levels of AI

Artificial intelligence (AI): A field of computer science dedicated to creating systems capable of performing tasks that usually require human intelligence, such as visual perception and decision-making.

Artificial narrow intelligence (ANI): AI systems that are designed and trained for a particular task, like voice assistants or image recognition systems, representing the majority of existing AI applications today.



Artificial capable intelligence (ACI): Also referred to as intelligent agents, these AI systems can understand, learn, and apply knowledge in different domains, making decisions and solving problems across various contexts and tasks, marking a transitional stage towards more generalized AI abilities.

Artificial general intelligence (AGI): Also known as broad AI, this refers to AI that can understand, learn, and apply knowledge across diverse domains, essentially possessing broad cognitive abilities similar to human intelligence. It does not yet exist. Opinions among leading AI experts vary widely: some believe its arrival is imminent, while others contend that it is impossible.

Artificial super intelligence (ASI): Hypothetical AI that surpasses human intelligence, possessing the ability to improve itself rapidly and potentially outperforming the best human brains in most economically valuable work, which is purely speculative and not present in our current technological landscape.

Turing Test: A test that evaluates a machine's ability to exhibit intelligent behavior equivalent to, or indistinguishable from, that of a human, assessing whether human interrogators can distinguish between responses from a machine and a human.

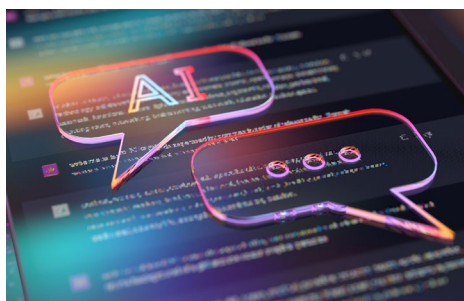
Fields and Types of AI

Machine learning: A subset of AI that provides systems with the ability to automatically learn and improve from experience; for example, predicting customer churn based on a variety of factors like purchase history and customer service interactions.

Natural language processing: Helping machines understand and interact with human language, allowing applications like chatbots to understand and respond to user requests.

Computer vision: Enables machines to interpret and make decisions based on visual data, like image recognition systems used in self-driving cars to identify objects and navigate roads.

Robotics: Integrating AI models to control robots, facilitating autonomous actions and adaptations to new environments and tasks.



Generative AI models: AI models that can generate creative content such as text, images, or music and are often used for applications like chatbots, content creation, and more.

Discriminative models: AI models that differentiate between different types of data, often used in classification tasks, like spam filtering.

AI Methodologies and Processes

Neural networks: A system of algorithms modeled after the human brain, neural networks discern patterns in data and form the foundation for most modern AI, enabling applications from image recognition to language translation by adjusting their structures during training to make accurate predictions and decisions.

Deep learning: Utilizes neural networks with many layers (deep neural networks) and has been vital in advancing fields like computer vision and natural language processing.


Reinforcement learning: A type of machine learning where an agent learns how to behave in an environment by performing actions and receiving rewards or penalties. For instance, AlphaGo, developed by DeepMind, used reinforcement learning to master the complex game of Go by playing millions of games against itself.

Unsupervised learning: Engaging with unlabeled data to discern hidden patterns and structures without predefined labels. For instance, unsupervised learning can be used to

identify different customer segments in e-commerce by analyzing shopping patterns, time spent on different product pages, and purchase history, even when the specific customer categories are not predefined.

Transfer learning: Applying knowledge learned in one domain to a different but related domain; for instance, using a model trained on general images to recognize specific types of objects by retraining it on a smaller dataset of those objects.

Training: The process where an AI model is taught to make decisions by feeding it data and allowing it to adjust its internal parameters to improve its performance; for example, training a spam filter model using a dataset of emails labeled as “spam” or “not spam.”



Deployment: Implementing the AI model into production, where it starts taking real-world data, making decisions, and producing results; for instance, integrating a trained recommendation model into an e-commerce website to suggest products to users.

Fine-tuning: Adjusting the parameters of an already trained model to improve its performance on a slightly different task; for example, modifying a pre-trained image recognition model to recognize a new category of objects.

Emergent capabilities: The abilities or features that arise during the development or utilization of an AI system that were not explicitly programmed or expected. This might include the system developing new strategies, understanding new types of data, or finding novel solutions to problems without being explicitly programmed to do so. These capabilities emerge from the system’s interactions with data and its environment.

Ethics and Trust

Black box: The term “black box” describes AI systems in which the internal mechanisms or decision-making processes are not transparent or comprehensible to humans. This can impede understanding and validation of how the system derives its results, presenting challenges in ensuring accountability and fairness in applications.

Explainability: The degree to which the functioning and decision-making processes of AI are clear and understandable to humans, ensuring that stakeholders can interpret AI outcomes and potentially question them.

Alignment: Ensuring AI models act in ways that are aligned with human values and can be controlled by human operators.

AI bias: AI bias occurs when algorithms produce unfair or skewed outcomes, often stemming from using prejudiced training data or from unintended consequences of the algorithm's decision-making rules, creating results that may unintentionally favor one group over others.

Hallucination: Hallucination in AI involves the system perceiving patterns or features in data that don't actually exist, leading it to make decisions based on these inaccurate perceptions. For instance, an AI interpreting medical images might "see" a condition that isn't present, potentially leading to misdiagnoses and emphasizing the need for careful oversight and validation of AI-generated insights.

Types of Data

Core proprietary data: Internal, unique data assets, like customer transactions, that are generated within and owned by the company.

External proprietary data: Data sourced from external entities through agreements or partnerships and is not publicly available.

External non-proprietary data: Publicly accessible data that any organization or individual can utilize.

Latent data: Available data that has not been leveraged or analyzed for certain purposes previously.

Synthetic data: Computer-generated data created to model specific conditions or scenarios, which can be used to augment real-world data or create data where none exists.

ABOUT THE AUTHORS

Shahriar Parvarandeh is a Senior Director at Innosight based in London.
sparvarandeh@innosight.com

Ned Calder is a Managing Director at Innosight based in Boston. ncalder@innosight.com

Freddy Solis is a Senior Director at Innosight based in Boston. fsolis@innosight.com

ABOUT INNOSIGHT

Innosight is a strategy and innovation consulting firm that helps organizations navigate disruptive change and manage strategic transformation. Now a member of the Huron Consulting Group, we work with leaders to create new growth strategies, accelerate critical innovation initiatives, and build innovation capabilities. Discover how we can help your organization navigate disruption at www.innosight.com.

Contact Us