

COURSE GLOSSARY

Generative AI for Business

AI literacy: The knowledge and skills required to understand, use, and critically evaluate AI tools and their outputs effectively and responsibly

Artificial Intelligence (AI): A field of computer science that designs systems capable of performing tasks that normally require human intelligence, such as perception, reasoning, learning, and decision-making

Augmentation: Use of AI to assist and extend human capabilities by performing parts of a task while a human remains responsible for judgment and final decisions

Cloud computing: The delivery of compute, storage, and other services over the internet, enabling scalable access to the hardware and infrastructure required for training and deploying AI systems

Co-creation: A collaborative workflow where humans and AI iteratively contribute to and refine a shared creative or problem-solving process

Dataset: A curated collection of data examples (e.g., images, text, videos) used to train, validate, or test machine learning models

Deepfake: Synthetic media (audio, images, or video) created by AI that convincingly imitates real people or events and can be used to mislead or defame

Hallucination: When a model produces confident but incorrect or fabricated information, often due to gaps or biases in its training data or reasoning process

Hallucination: When a model produces confident but incorrect or fabricated information, often due to gaps or biases in its training data or reasoning process

Foundational model: A large pre-trained model intended as a general-purpose starting point that can be adapted to many downstream tasks via fine-tuning or prompting

Generative Adversarial Network (GAN): A pair of neural networks—one generator and one discriminator—that are trained adversarially so the generator learns to produce realistic synthetic data

Generative AI: A subset of AI focused on creating new content (text, images, audio, or video) that resembles examples from its training data

Human-in-the-loop: A design principle where human oversight, feedback, or intervention is incorporated into AI workflows to ensure safety, quality, and alignment with human values

Intellectual property (copyright) in AI: The legal and ethical questions about who owns or controls content generated by AI—whether the user, developer, or other party holds rights to generated works

Large Language Model (LLM): A transformer-based model trained on vast amounts of text to predict and generate human-like language across many tasks

Mitigation techniques: Methods used to detect, reduce, or correct harmful biases or other negative impacts in AI systems, including data diversification, algorithmic adjustments, auditing, and human review

Model training: The process of teaching a machine learning model by exposing it to large amounts of labeled or unlabeled data so it can learn patterns and make predictions or generate content

Open-source project: A software or model initiative whose source code or weights are publicly available for anyone to use, modify, and redistribute, often accelerating collaboration and innovation

Prompting: The practice of giving a generative model a user-provided input (a prompt) that guides the content the model produces

Reinforcement Learning from Human Feedback (RLHF): A training approach that refines model behavior by using human evaluations or preferences as rewards to guide the model toward more desirable outputs

Replacement (automation): A scenario in which AI fully automates a task that previously required human effort, potentially eliminating the need for human intervention in that task

Social bias: Systematic and unfair treatment or representation of certain groups that can arise in AI systems due to biased training data, design choices, or deployment contexts

Transformer: A neural network architecture that uses self-attention to model relationships in sequential data, and which underpins many modern language and multimodal generative models

Watermarking (AI content watermarking): Techniques for embedding detectable signals into AI-generated content to indicate provenance and help distinguish synthetic from authentic media