

Implementing AI Solutions in Business

AI Ethics Board: A cross-functional group established by an organization to oversee, define, and enforce Responsible AI principles, policies, and governance at scale

Artificial Intelligence (AI): The field of computer science focused on creating systems that perform tasks that typically require human intelligence, such as perception, reasoning, learning, and decision-making

Change Management: Structured processes and communication strategies used to guide people, teams, and organizations through transitions caused by new technologies, workflows, or roles

Connectionist AI: An approach inspired by the brain that models computation with networks of simple units (neurons) which learn patterns from data through weighted connections

Data Bias: Systematic errors or skew in datasets that produce unfair or inaccurate model outcomes by underrepresenting or misrepresenting certain groups or scenarios

Data Pipeline: The sequence of processes and tools that collect, clean, transform, and deliver data from sources to systems or models for analysis and training

Data Quality: The degree to which data is fit for use, commonly measured by dimensions such as completeness, accuracy, relevance, timeliness, and consistency

Histogram: A chart that groups continuous numerical data into equal-width bins and displays the frequency of observations within each bin to show distribution

Fine-tuning: The process of adapting a pre-trained model to a specific task or domain by continuing training on task-relevant data

Explainability: The property of an AI model or system that provides understandable reasons or justifications for its outputs to improve transparency, trust, and accountability

First-party Data: Data collected directly by an organization from its own customers, products, or operations, typically considered the most reliable source for personalization and modeling

Generative AI: A class of AI models that create new content—like text, images, or audio—by learning the underlying patterns of training data and producing novel outputs often via natural language prompts

GPU (Graphics Processing Unit): Specialized hardware optimized for parallel numerical computation that significantly accelerates the training and inference of many AI and deep learning models

Human-in-the-loop: A system design in which humans are integrated into the AI decision process to review, correct, or approve outputs, especially where confidence is low or stakes are high

Machine Learning: A set of techniques that use statistical algorithms to enable computers to learn patterns from data and improve performance on tasks over time without being explicitly programmed for each task

MLOps: Practices, tools, and processes that operationalize machine learning model development, deployment, monitoring, retraining, and lifecycle management to ensure production reliability

Model Accuracy: A performance metric that measures the proportion of correct predictions or classifications made by a model relative to a labeled ground truth

Model Drift: The decline in a model's performance over time caused by changes in the underlying data distribution or relationships the model was trained on

Model: A computational representation learned from data that maps inputs to outputs and is used to make predictions, classifications, or generate content

Proof of Concept (POC): A small-scale implementation of an idea or solution built to demonstrate feasibility, validate assumptions, and gather evidence of potential value before a full-scale investment

Responsible AI: An approach to designing, developing, and deploying AI systems that prioritizes ethics, legality, fairness, transparency, security, and social impact throughout the lifecycle

Scalability: The ability of an AI solution, its infrastructure, and organizational processes to handle increased load, users, or data volume while maintaining performance and reliability

Symbolic AI: An early AI approach that represents knowledge explicitly with symbols and rules, enabling reasoning through logical manipulation rather than pattern learning

Third-party Data: Data acquired from external providers or partners that aggregate information across multiple sources to supplement an organization's first-party data

Upskilling: The process of training and educating employees to acquire new skills or improve existing ones—particularly AI literacy and tool proficiency—to maximize adoption and business value