SUSTAINABLEIT.ORG

AI General Glossary

IT Toolkit for Responsible & Sustainable AI: A Field Guide for Implementation at Scale

AI General Glossary

Terms in this glossary are relevant to various elements of IT Toolkit for Responsible and Sustainable AI by SustainableIT.org. Readers should refer back to this document for definitions as needed. The glossary is based on The Top 50+ AI/GenAI Terms Demystified, by Catherine Richards, Managing Partner, **Expera Consulting,** and adapted with permission.

A

AI (Artificial Intelligence): AI refers to smart technology that can analyze data, make decisions, and automate tasks.

AI, Agentic: A class of AI systems designed to exhibit agency—defined as the ability to make decisions, act independently, and learn over time through memory and feedback loops. These systems are modular and often include multiple specialized agents coordinated through an orchestration layer. This orchestrator integrates memory, decision-making, and interoperability across agents. Agentic AI is still emerging but already being tested in enterprise environments for tasks like onboarding, compliance audits, and cross-functional collaboration.

AI, Generative: GenAI is a specialized form of AI that goes beyond analysis. It actively creates new content, such as text, images, or even music, based on patterns it has learned. **The key difference between AI and GenAI** is that AI optimizes and automates tasks leveraging data patterns. Generative AI uniquely creates new content and media, customizable for personalization. While AI drives efficiencies, generative models power creativity.

AI Model: An AI model refers to a discrete AI system or unit of AI capability that has been trained on data to perform a specific task or set of tasks. Each model typically includes a trained algorithm, defined inputs and outputs, and operational parameters. In enterprise environments, the term is increasingly used as the central unit for AI governance, risk management, and lifecycle accountability, allowing organizations to track and manage AI capabilities as modular, auditable assets across their digital ecosystem.

AI Model Drift: Drift refers to a decline in an AI system's accuracy over time, caused by two main factors: data drift (when the input data changes) and concept drift (when the relationships the model learned no longer apply). It can lead to errors, operational risk, and compliance issues. Users help by flagging unusual results. Admins and technical teams address drift through monitoring, retraining, and governance.

AIOps: Short for Artificial Intelligence for IT Operations, AIOps is a discipline that applies machine learning, data analytics, and automation to improve and streamline IT operations. It enables real-time monitoring, analysis, and remediation of complex IT environments by ingesting large volumes of telemetry and event data from diverse sources.

AI Orchestration: The coordinated management of multiple AI agents and systems to perform complex, multi-step tasks. An orchestration layer often serves as the command center, handling user interaction, memory, task delegation, and agent-to-agent communication. Orchestration enables greater reliability, repeatability, and scalability in enterprise AI use cases.

AI Prompt: A prompt is a written instruction or question given to a GenAI system to generate a specific response.

AI Training Data: The dataset that is used to train an AI model. This data enables the model to recognize patterns and generate predictions. Biased or incomplete training data can lead to problematic outputs.

ALT Text (Alternative Text): A short, written description of an image that helps screen readers and search engines understand it.

Algorithm Bias: Systematic and repeatable errors in AI systems that lead to unfair, unethical, or discriminatory outcomes for certain user groups.

Answer Box: A direct answer shown at the top of search results or in an AI tool's response, pulled from a trusted source. It's also called a featured snippet.

Artificial General Intelligence (AGI): AGI, a term expressing concern from researchers, envisions highly autonomous AI systems surpassing humans in various economically valuable tasks, posing ethical considerations and implications. In essence, AGI = AI matching or topping human aptitude, and ASI = AI astronomically beyond all human cognition.

Artificial Superintelligence (ASI): A hypothetical futuristic form of AI theorized to greatly exceed human-level general intelligence.

Attribution AI: Applies machine learning to model touchpoint influence on conversions with enhanced accuracy over rules-based analytics.

B

Bias in AI: Bias in AI can arise from flawed data, model design, or societal inequities. Recognizing and mitigating it is essential to building ethical, inclusive AI solutions.

C

Chatbots: AI-powered systems that interact conversationally with users, enhancing engagement, qualifying leads, and collecting valuable analytics.

Conversational AI: The use of chatbots and virtual assistants for natural language processing, delivering personalized conversational experiences. **The key difference between chatbots and conversational** AI is that chatbots are one conversational AI use case in the overall field of conversational AI.

Context Window: The limit to how much information an AI model can remember in one prompt.

D

Data Mining: Applying AI and Machine Learning techniques to extract patterns and insights from datasets, informing decisions. **The key difference between data mining and machine learning** is that data mining applies algorithms to extract insights from data, while machine learning goes further by dynamically improving its analytic model over time through continuous learning from new information.

Deepfake: Synthetic media (often video) generated using AI to imitate real people, raising ethical concerns and misinformation risks.

Deterministic vs. Probabilistic AI: Deterministic AI gives the same answer every time. Probabilistic AI, like ChatGPT, may change the wording with each run.

E

Ethical AI: Ethical AI is the practice of designing and using AI systems in ways that are fair, transparent, accountable, and aligned with human values. It means considering who is impacted, how decisions are made, and whether outcomes can be understood and trusted. Delivering ethical AI requires human judgment, clear intent, and active responsibility, not just technical fixes.

Explainable AI (XAI): Making AI models' reasoning understandable to humans, fostering trust and oversight. The goal of XAI is to create transparent assistants rather than black-box oracles.

Experiential AI: Using AI in interactive technologies like AR and VR to evoke engagement in immersive digital experiences.

F

Fidelity: How precisely AI models capture the nuances of human behavior and judgment. In AI tools, higher fidelity indicates greater precision in emulating the nuanced richness of real language and expression.

Fine-Tuning: Customizing a pretrained model on new data to improve performance on specific tasks or match brand tone.

Folksonomy: User-created tags that describe content in their own words, outside of a formal system. Example: One teammate tags a blog post "customer win," another tags it "case study." AI tools can learn from these patterns to improve tagging, search, and recommendations, especially when building smarter content systems. (See also: Taxonomy, Ontology). **Taxonomy vs. Ontology vs. Folksonomy:** Taxonomy is the official filing system AI uses to know where things go. Ontology is the map of how ideas connect. AI uses it to understand relationships. Folksonomy is the messy list of tags people actually use AI learns from it to speak human.

G

Generative Adversarial Networks (GANs): A class of machine learning frameworks where two neural networks contest with each other to produce more accurate outputs.

Generative AI (GenAI): See AI, Generative.

H

Hallucination: Hallucination refers to AI systems generating fabricated predictions or content that seem believable but do not accurately reflect reality. This can lead to misguided strategic decisions or an inaccurate understanding of reality. To spot and deal with hallucinations, use diverse, unbiased data to train AI tools, always double-check data sources, and verify insights regularly to make sure decisions are based on accurate information.

Hyper-Personalization: Using AI for highly customized, individualized messaging and experiences based on stakeholder attributes and behaviors.

I

Image Generation: Creating images from text prompts using tools like DALL·E or Midjourney.

Inclusive Design: Mitigating issues of bias in AI systems by intentionally involving diverse voices in the development process and representing a wide range of human conditions across training data.

Intelligent Automation: The fusion of AI and automation to improve workflows and processes, often reducing manual labor and increasing efficiency.

Inference: Inference means an AI can gain insights on new data it hasn't seen before. It recognizes similar patterns from the examples it trained on earlier. The more training examples, the smarter AI gets at making accurate predictions and recommendations.

J

JSON (JavaScript Object Notation): A simple format used to organize and share data between systems. It looks like structured text and is often used behind the scenes in web development, APIs, and AI tools.

Jailbreak (AI context): When someone tries to trick an AI into breaking its own rules or filters, usually to get it to say something inappropriate, unsafe, or restricted.

K

Knowledge Base: Knowledge bases collect information into shared repositories and are found in GenAI tools like Jasper. AI uses a knowledge base to optimize. Generative AI uses them to create new things by expanding on what's already known.

L

Large Language Model (LLM): AI models trained on vast amounts of text data to understand and generate human-like language. Examples include OpenAI's GPT, Anthropic's Claude, and Google's Gemini.

> **Specialized LLMs:** Large language models fine-tuned on business-specific or domain-specific data to serve expert roles inside agentic ecosystems. These LLMs power intelligent agents that understand the nuances of enterprise tasks such as finance, legal, or compliance.

M

MCP (Model Context Protocol): A proposed interoperability protocol that allows agents to communicate with models and external systems in a consistent, structured way. MCP is foundational to enabling modular, agentic ecosystems where specialized agents can collaborate dynamically across tools and platforms.

Machine Learning (ML): A subset of AI where algorithms learn from data and improve performance over time without explicit programming.

Multi-Modal AI: AI models that can process and generate content across multiple types of data inputs—like text, image, audio, and video. These models enable richer and more dynamic content creation and user interactions across various platforms.

N

Natural Language Generation (NLG): Integrated into automation, NLG modes transform analytics data into written insights for diverse audiences.

Natural Language Processing (NLP): A field of AI focused on the interaction between computers and human language, enabling machines to understand and respond to text or speech inputs.

> **The key differences among LLMs, ML, and NLG** lie in their core functions. Large language models (LLMs) focus on generating original text content from scratch, machine learning (ML) extracts patterns from data to optimize decisions, and natural language generation (NLG) tailors analytics data into written narrations, providing distinct capabilities in the realm of content creation and decision optimization.

Neural Networks: Computing systems inspired by the human brain's network of neurons, used in machine learning to recognize patterns and solve complex problems.

O

Ontology: A structured system that connects related terms and ideas so AI tools understand how content fits together. Choosing one clear term—like "customer story" instead of switching between "case study," "testimonial," or "success profile"—helps AI tag and process information more accurately. (See also: Structured Content, Taxonomy.)

P

Predictive Analytics: Using data and AI algorithms to forecast future outcomes and trends, optimizing resource allocation.

Prompt Engineering: Crafting specific instructions for AI systems to generate desired content and optimize customization.

Q

Quantum Machine Learning (QML): A theoretical area with no current practical advantage over existing AI, requiring a pragmatic focus on proven analytics techniques.

R

RAG (Retrieval-Augmented Generation): A method that improves how AI answers by letting it look things up in real time. Instead of relying only on what it was trained on, the model pulls in current information.

Reinforcement Learning: An area of machine learning where agents learn to make decisions by performing actions and receiving feedback from the environment.

Responsible AI: Responsible AI refers to the broader governance and operational discipline that ensures AI systems are developed, deployed, and maintained in ways that align with legal, ethical, and organizational standards. It includes oversight practices, risk management, compliance, and lifecycle accountability across systems and teams. **The difference between Responsible AI and Ethical AI** and is that Ethical AI focuses on doing what is right, and Responsible AI ensures it is done right.

S

Search (in the AI era): AI-powered tools like Google SGE, Bing Copilot, and ChatGPT now give people direct answers instead of a list of links.

Schema: Code embedded in systems such as a website that labels information, such as “article,” “FAQ,” “event,” or “product.” Schema helps search engines and AI tools recognize the relevance of content. (See also: Structured Content.)

Semantic Search: Search that understands intent and meaning, not just keywords.

Sentiment Analysis: Using AI to determine subjective opinions in textual data, monitoring and analyzing feedback.

Small Language Model (SLMs): AI small Language Models are designed to understand and generate human language, similar to large language models (LLMs), but with fewer parameters, reduced computational requirements, and a smaller memory footprint. In enterprise contexts, SLMs are increasingly adopted as more sustainable, manageable, and cost-effective alternatives to LMs and large-scale generative AI.

Structured Content: Content that is organized clearly for both humans and machines. This includes things like headings, bullet points, summaries, and question-and-answer formats. (See also: Schema.)

Synthetic Data: Artificially generated data that mimics real-world datasets. It’s used to train AI models while reducing privacy risks and data sourcing constraints.

Sustainable AI: Sustainable AI refers to the design, development, deployment, and lifecycle management of AI systems in a manner that minimizes environmental impact, promotes social responsibility, and adheres to principles of ethical governance. For enterprise leaders, Sustainable AI serves as a framework for aligning AI innovation with enterprise sustainability goals, including carbon reduction, resource efficiency, equity, and transparency.

> The difference between Sustainable AI and Responsible AI is that Sustainable AI focuses on aligning AI with environmental and societal sustainability goals. Responsible AI provides the governance structure to make that possible. In practice, Responsible AI efforts often extend to and enfold the goals of Sustainable AI by embedding accountability, oversight, and compliance processes that ensure sustainability commitments are met across the AI lifecycle.

Sustainable AI focuses on *the impact on the planet and society.*
Example: Reducing energy use, promoting fairness, and aligning with sustainability goals.

Ethical AI focuses on *doing the right thing for people.*
Example: Making sure AI is fair, understandable, and aligned with human values.

Responsible AI focuses on making sure the right processes are in place.
Example: Setting up rules, checks, and teams to ensure AI is built and used safely and legally.

T

Taxonomy: A structured way to categorize content using consistent labels, like topics, formats, or audiences. A good taxonomy helps organize a content library, allowing AI tools and search engines to find and reuse assets more easily. (See also: Ontology, Structured Content.)

Text-to-Image Generation: Converting written descriptions into visually compelling images using AI.

Token: A chunk of text (like a word or part of a word) that large language models use to process and generate language. It affects how much content a model can handle in a single interaction.

U

Use Case: Use cases detail specific, real-world applications of GenAI capabilities addressing enterprise challenges or opportunities. Use cases can be generally categorized into prediction, language, and vision. Prediction is about forecasting future outcomes. Language relates to generating or comprehending text. Vision involves analyzing visual content.

V

Vision Recognition: Object detection, image categorization, and facial analysis for enhanced digital experiences and analytics.

W

Workflow Automation: Using AI to handle repetitive tasks like posting to social, sending emails, or moving content through approvals.

X

XML: A file format that helps search engines understand a website. XML is the sitemap that tells AI what pages exist. In contrast, schema is the label. It tells AI what each page is about. Structured content is how the label is written. It makes the message clear to both people and machines.

Y

YAML: A plain text format some AI tools use to organize instructions or templates. It looks like a simple checklist, like “Use a friendly tone” or “Keep responses under 150 words.” No coding required.

Z

Zero-click Search: When users get the answer they need directly in a search result or AI-generated response, without clicking a link.

Zero-shot Learning: When AI can perform a task it wasn’t specifically trained on by using general knowledge.

About SustainableIT.org

Vision

Advancing global sustainability through technology leadership.

Mission

Our mission is to unite the world's largest community of technology and sustainability leaders to define sustainability transformation programs, author best practices and frameworks, set standards and certifications for governance, provide education and training, and raise awareness for IT-centric ESG programs that make their organizations and the world sustainable for generations to come.

Mandates

Best Practices, Research and Standards

Identify sustainable digital transformation programs by industry. Research and define best practices, frameworks, and standards for all three pillars of sustainability (environmental, societal, and governance) for IT departments and organizations.

Global Awareness and Recognition

Promote sustainable digital transformation programs and advances in sustainability. Raise awareness through local, regional, and global awards, as well as through social media, publications, and public relations.

Community, Education and Training

Build local and regional communities for technology leaders to advance sustainability. Develop education and training programs for IT leadership and professionals for all three pillars of sustainability.

Transparency and Accountability

Set standards for metrics and reporting to enable transparency and accountability. Create certification programs for individuals and organizations with rights to use our sustainability emblem.

SUSTAINABLEIT .ORG

About SustainableIT.org

SustainableIT.org is a Delaware 501(c)(6) nonprofit, non-stock legal entity led by technology executives who will advance global sustainability through technology leadership. Our mission is to define sustainable transformation programs, author best practices and frameworks, set standards and certifications, provide education and training, and raise awareness for environmental and societal programs that make our organizations and the world sustainable for generations to come.