

Glossary

Note: this glossary has been compiled with the assistance of GPT4. It is by no means complete and other terms are used throughout the book.

Algorithm: A set of rules or instructions given to an AI, or a machine that it uses to solve problems or make decisions.

Artificial Intelligence (AI): The simulation of human intelligence processes by machines, especially computer systems. The processes include learning, reasoning, problem-solving, perception, and language understanding.

AutoML (Automated Machine Learning): The process of automating the end-to-end process of applying machine learning to real-world problems. AutoML makes it feasible for non-experts to effectively use machine learning models and techniques.

Autonomous Agents: In artificial intelligence, autonomous agents are systems capable of independent decision-making and action-execution based on the environment they are in, without continuous human guidance. For libraries, this could manifest in several ways, such as virtual assistants that help patrons navigate the digital library resources, autonomous robots that assist with book sorting or shelving, or intelligent systems managing queries and administrative tasks.

BERT (Bidirectional Encoder Representations from Transformers): A Transformer-based machine learning technique for natural language processing tasks. Developed by Google, BERT is designed to understand the context of a word in search queries, thereby helping data retrieval systems grasp the intent behind a user's query.

Bias in AI: Prejudice or unintended favoritism in AI decision-making. This bias occurs due to the data used to train the AI and can lead to unfair outcomes.

Big Data: Extremely large data sets that may be analyzed computationally to reveal patterns, trends, and associations, especially relating to human behavior and interactions.

Chatbot, aka Conversational Agent: A computer program that simulates and processes human conversation (either written or spoken), allowing humans to interact with digital devices as if they were communicating with a real person.

Collections-as-data: The implications of seeing GLAM collections as data.

Data Lake: A storage repository that holds a vast amount of raw data in its native format until needed. Libraries dealing with large sets of unstructured data can use data lakes for storage, which can then be processed by AI models for various tasks.

Data Mining: The process of discovering patterns and knowledge from large amounts of data. The data sources can include databases, data warehouses, the internet, and other information repositories.

Deepfakes: Use of generative AI techniques to create hard to detect false images or other content.

Deep Learning: A subset of machine learning where neural networks with many layers (hence “deep”) learn from large amounts of data. Deep learning methods are highly effective in discerning patterns and representations of data.

Digital Assistant (or Virtual Assistant): An application programme that understands natural language voice commands and completes tasks for the user.

Facial Recognition: A category of biometric software that maps an individual’s facial features mathematically and stores the data as a faceprint. It is used to identify individuals in photos, video, or real-time.

Federated Learning: A machine learning approach where the model is trained across multiple decentralised devices (or servers) holding local data samples and without exchanging them. This approach is useful in privacy-preserving and efficient decentralized systems.

Fine-Tuning: In the context of AI, fine-tuning refers to the process of making minor adjustments to a pre-trained model so that it can perform well on a new task. This process is essential when applying large language models to specific tasks or industries, including library science.

General AI: Also known as Strong AI. AI that works across multiple contexts, as opposed to narrow/weak AI.

Generative AI: AI that produces new content be that text, images or other media.

GPT (Generative Pre-trained Transformer): An advanced type of Transformer AI model used for natural language understanding and generation. GPT, developed by OpenAI, has had several versions, with GPT-3 being one of the most known for its large scale and versatility in handling various language tasks.

Hallucination: False information presented by AI, as in generative AI outputs.

Human-centred AI: An approach to AI development and use focused on benefiting humans.

Human in the Loop: The design of AI systems that retain a central role for human decision making, particularly in training and testing supervised machine learning.

Information Retrieval (IR): The activity of obtaining information resources relevant to an information need from a collection of information resources. Searches can be based on metadata or on full-text indexing.

Internet of Things (IoT): The network of physical objects – “things” – that are embedded with sensors, software, and other technologies for the purpose of connecting and exchanging data with other devices and systems over the internet.

Knowledge Graph: A form of knowledge base that uses a graph-structured model or network for semantic queries. It enhances search capabilities by understanding the context and relationships between different pieces of information.

Latent Dirichlet Allocation (LDA): LDA is a specific type of topic modelling technique. It is an unsupervised machine learning method that helps systems understand and identify topic structures within a text body, based on a set of observations. In a library context, LDA can be instrumental in identifying patterns and topics across different documents, such as research papers, books, and articles, even when those topics are not explicitly stated. This process enhances metadata tagging, making the search and retrieval process more efficient and user-friendly.

Large Language Models (LLMs): These are types of artificial intelligence models trained on vast amounts of text data. They are designed to understand language and generate coherent, contextually relevant text based on the input they are given. LLMs can answer questions, summarize content, translate languages, and even create content.

Machine Learning (ML): A subset of AI that provides systems the ability to automatically learn and improve from experience without being explicitly programmed. It focuses on the development of computer programs that can access data and use it to learn for themselves.

Metadata: This term refers to data that provides information about other data. In the context of libraries, metadata is crucial for the organization, discovery, and management of information resources. It can include details about the content, format, author, and creation date among other attributes of resources.

MLOps (Machine Learning Operations): A set of practices that combines machine learning, DevOps, and data engineering, which aims to automate and enhance the end-to-end lifecycle of ML models.

Multimodal AI: Multimodal AI refers to systems that can engage and integrate multiple types of data input or output, such as text, audio, visual, and sensory data, to make decisions or provide responses. In the context of libraries, multimodal AI can revolutionize user interactions by enabling more dynamic engagement. For instance, a system that understands speech, text, images, and gestures could provide more accessible services catering to diverse user needs, including those with disabilities. It could allow patrons to search for books through voice commands, receive information through interactive displays, or even direct them using augmented reality through their smartphones, creating an inclusive and interactive environment.

Narrow AI: Also known as weak AI. AI that solves a specific problem, as opposed to general/strong AI which could operate across multiple contexts.

Natural Language Processing (NLP): A field of AI that focuses on the interaction between computers and humans through language. It allows machines to read, decipher, understand, and make sense of human languages in a valuable way.

Opacity: A lack of transparency and explainability in AI decision making and outputs.

Optical Character Recognition (OCR): The electronic or mechanical conversion of images of typed, handwritten, or printed text into machine-encoded text, whether from a scanned document, a photo of a document, or text superimposed on an image.

Predictive Analysis: The use of data, statistical algorithms, and machine learning techniques to identify the likelihood of future outcomes based on historical data.

Prompt Engineering: The crafting of requests in natural language to shape the outputs from generative AI in desired directions.

Recommender System: A subclass of information filtering systems that seek to predict the “rating” or “preference” a user would give to an item. In libraries, this could be used to suggest books, articles, and other resources.

Reinforcement Learning (RL): A type of machine learning where an agent learns to behave in an environment by performing actions and receiving rewards for them. It is about taking suitable action to maximize reward in a particular situation, used in various sectors, including information retrieval and organization systems.

Reinforcement Learning with Human Feedback (RLHF): This is an approach combining reinforcement learning with human feedback to train models more efficiently or towards objectives that are hard to define formally using rewards alone. It's particularly relevant in interactive systems where human preferences or quality assessments guide the AI.

Responsible AI: A label applied to AI development which is ethical and trustworthy.

Robotics Process Automation (RPA): Technology that allows anyone to configure computer software, or a “robot,” to emulate and integrate the actions of a human interacting within digital systems to execute a business process.

Semantic Analysis: The study of meaning in language, covering the understanding of meaning in text or speech. In AI, this refers to machines understanding the meaning and nuances of human language.

Sentiment Analysis: Also known as opinion mining, it involves the use of NLP, text analysis, and computational linguistics to identify and extract subjective information from resources.

Strong AI: See General AI/

Synthetic Data: Data created artificially rather than from real world events. Sometimes used as training data.

TF-IDF (Term Frequency-Inverse Document Frequency): A numerical statistic that reflects how important a word is to a document within a collection or corpus. It's often used in information retrieval and text mining to evaluate and rank documents in relation to user queries.

Topic Modelling: In the context of AI, topic modelling is a type of statistical model used to uncover abstract topics within a large volume of text. For librarians, this is particularly useful in categorising and clustering large collections of documents or books based on their themes, significantly improving information retrieval processes. It allows for more intuitive searching capabilities, where users can search by topic or theme, rather than specific or exact keywords.

Transformer Models: A type of neural network architecture used in natural language processing. Unlike previous sequential models, Transformer models process words in relation to all other words in a sentence, making them more effective for understanding context within language.

Trustworthy AI: The European Union's key term to describe responsible, ethical AI.

Turing Test: A test of a machine's ability to exhibit intelligent behavior equivalent to, or indistinguishable from, that of a human.

Uncanny Valley: The negative emotional response of unease or disgusts when AI appears too human.

Weak AI: See Narrow AI.

Zero-shot Learning: A machine learning approach where the model can correctly make predictions about data it has not been explicitly trained on. This approach allows for more flexible and expansive applications of AI, as it requires less data-specific training.