

AI GLOSSARY

A plain English guide to key AI (Artificial Intelligence) terms
for the Australian legal profession



THE LAW SOCIETY
OF NEW SOUTH WALES



LexisNexis®



About the AI Glossary

Generative AI is driving a wave of technological progress across many industries, and it has the potential to transform the practice of law.

The Law Society of NSW saw an opportunity to provide AI educational material specifically for lawyers and has partnered with legal technology experts LexisNexis® Pacific to provide this guide for the Australian legal profession.

This AI glossary has been designed to help legal professionals navigate common AI terms with confidence. It provides straightforward explanations of key concepts – including large language models, machine learning, and natural language processing – and relates those concepts to legal practice where possible.

This glossary is a living document. As the technological landscape develops, we will continue to refresh and update it. Please visit the Law Society of NSW website to download the latest version.

We hope that you find this a useful reference.



Law Society of NSW
Brett McGrath
2024 President



LexisNexis Pacific
Greg Dickason
Managing Director, Asia & Pacific



Lexis® AI Insider

Become a Lexis AI Insider and keep up to date with the latest developments in legal AI, access exclusive AI education and insights, and help shape the future of AI in the legal profession.

REGISTER TODAY→

Contents

Artificial intelligence (AI).....	4
Large language model (LLM).....	4
Agentic AI.....	4
AI safety.....	4
AI system.....	5
Algorithm.....	5
Alignment.....	5
Application programming interface (API).....	5
Automation bias.....	5
Benchmark.....	6
Bias.....	6
Chatbot.....	6
ChatGPT.....	6
Claude.....	6
Cloud computing.....	6
Cognitive bias.....	7
Compute.....	7
Computer vision.....	7
Context window.....	7
Copilot for Microsoft 365.....	7
Data.....	8
Data science.....	8
Explainable AI (XAI).....	8
Fine-tuning.....	8
Foundation model.....	8
Generative AI.....	8
GPT.....	9
Guardrails.....	9
Hallucination.....	9
Human-in-the-loop.....	9
Machine learning.....	9
Modality.....	10
Model.....	10
Multi-modal.....	10
Natural language processing (NLP).....	10
Pre-training.....	10
Prompt.....	10
Prompt engineering.....	11
Reinforcement learning.....	11
Responsible AI.....	11
Retrieval-augmented generation (RAG).....	11
Supervised learning.....	12
Token.....	12
Training data.....	12
Transformer model.....	12
Unsupervised learning.....	12
Use case.....	12

Key Terms

Artificial intelligence (AI)

Broadly speaking, AI is the programming of machines to behave in ways that mimic human intelligence and human capabilities, including reasoning, decision-making, interacting, and perceiving. There are different sub-disciplines and techniques within the broad field of AI, including natural language processing (NLP) and machine learning.

AI has many applications, or use cases, in law, science, industry, and our day-to-day lives.

Related terms: machine learning; natural language processing (NLP); use case

Large language model (LLM)

A large language model, or LLM, is an AI model that is trained on vast amounts of text data to produce fluent responses to text drafted by the user (prompts). Anthropic's Claude and OpenAI's GPT-4 are examples of LLMs.

Because of their ease with language, LLMs are particularly well suited to use in fields like law. However, generic LLMs are not specifically trained for legal use cases, or for use in other professional fields. AI systems that are intended for use in a legal context offer more assurance of accuracy and usefulness.

Related terms: ChatGPT; Claude; foundation model; natural language processing (NLP); use case; prompt

Glossary

Agentic AI

Agentic AI is a term for AI systems that can comprehend nuanced instructions, autonomously make decisions, plan, pursue goals and understand workflows – like a human employee.

Related terms: AI system

AI safety

AI models are becoming increasingly sophisticated and powerful, carrying the potential for both beneficial applications and misuse. AI safety is a field of study focused on preventing the use of AI systems to cause harm and ensuring that AI behaves in accordance with desired human goals and values.

Related terms: alignment; bias; responsible AI

AI system

An AI system is an AI model or models combined with other elements that make the model output useful, including for example the design of the user interface.

Related terms: model

Algorithm

In computer science, an algorithm is a set of instructions given to a computer in order to achieve a task.

Related terms: machine learning; model

Alignment

In the context of AI, alignment means ensuring that AI systems behave in accordance with desired human goals and values and their intended objectives. One way of increasing alignment involves humans giving feedback to an AI system during the training process so that it learns to perform in intended ways – similar to the way in which humans learn through trial and error.

Related terms: AI safety; responsible AI; reinforcement learning

Application programming interface (API)

An application programming interface (API) is a set of pre-defined rules and protocols that allows software applications to exchange data in a standardised way, making them interoperable. APIs allow for the creation of products and services using a mix of the most appropriate data and technologies, including AI models.

For example, the Google Maps API allows software developers to include maps in their own applications.

Automation bias

Automation bias is a type of cognitive bias that may arise when a human interacts with an automated system. It refers to the human tendency to accept automated output despite evidence that the output is incorrect, and to ignore contradictory information even if correct, or to fail to act if the automated system incorrectly misses an issue.

A common example of automation bias in the context of driving is following GPS directions even though the driver suspects that they are wrong. An example of automation bias in a legal context is an administrative decision-maker uncritically accepting automated output when making a decision.

Related term: cognitive bias

Benchmark

A 'benchmark' in the context of AI is a standard against which the performance of an AI model on a particular task is evaluated. Benchmarks act as a reference point to facilitate comparison between different models and systematic evaluation of progress over time.

Bias

Bias is a broad and complex topic in AI. It refers to unfair distortions in AI output caused by discriminatory or unrepresentative training data or by errors in the algorithm itself. Bias is a significant risk in AI systems because it can cause these distortions to be reproduced on a large scale.

Related term: automation bias; cognitive bias; responsible AI

Chatbot

A chatbot is an application that simulates back-and-forth, human-like conversation. Chatbots are a common application of large language models (LLMs) and natural language processing (NLP). They are frequently used in customer support or in e-commerce.

Related terms: large language model (LLM); natural language processing (NLP)

ChatGPT

ChatGPT is a chatbot released in November 2022 by OpenAI. It is powered by a large language model (LLM) using generative pre-trained transformer (GPT) architecture. Transformer architecture allows for human-like back-and-forth conversations.

Related terms: chatbot; GPT; large language model (LLM); transformer model

Claude

Claude is a chatbot developed by the company Anthropic with substantial investment from Amazon and first released in March 2023. Like ChatGPT, Claude is based on a large language model (LLM).

Related terms: chatbot; large language model (LLM)

Cloud computing

Cloud computing allows organisations to pay a third party for the computing infrastructure that they require, instead of maintaining their own infrastructure (such as servers, software and databases) on their premises. Cloud computing providers include Amazon Web Services (AWS), Microsoft Azure, and Google Cloud.

Related terms: compute

Cognitive bias

Human beings tend to think and process information in certain ways that are not logical and do not reflect reality. These tendencies are called cognitive biases. These biases are often unconscious. Human decision-making is commonly affected by cognitive biases; similarly, AI systems may be affected by human cognitive biases, as AI ultimately depends on data generated by humans.

Examples of cognitive biases that may be familiar in a legal context include confirmation bias (interpretation of information in a way that confirms pre-existing beliefs) and anchoring bias (overreliance on information received early in a process).

Related terms: automation bias; bias

Compute

Compute is an umbrella term for computing resources, including hardware, software and other infrastructure, used to perform tasks such as processing data and training machine learning models. Current generative AI models require a large amount of compute – and therefore energy and other resources – to train and run.

Related terms: cloud computing

Computer vision

Computer vision is a field of AI that involves using machines to obtain and process visual data, such as photos and videos. One application of computer vision is autonomous driving. In the legal context, an important use case for computer vision is extracting text from handwritten or scanned documents.

Related terms: use case

Context window

A context window is the maximum amount of text (in the form of 'tokens') that a large language model (LLM) can take into account in one back-and-forth conversation. If the amount of text exceeds the model's context window, the model will 'forget' older text in the conversation and give less contextually relevant responses.

Related terms: large language model (LLM); token

Copilot for Microsoft 365

Copilot for Microsoft 365 is an enterprise-grade AI assistant and productivity tool for use in the Microsoft 365 applications.

Data

Data is a collection of raw information or values. In computer science terms, data is stored in sequences of binary digits (zeros and ones) called bits. The computer translates these bits into more meaningful information that a human can understand and contextualise, like numbers and words.

Related terms: data science

Data science

Data science is an interdisciplinary field of study focused on using data to help organisations make data-driven decisions. Data scientists may use AI techniques to derive meaning from data.

Related terms: data

Explainable AI (XAI)

The inner workings of AI models are generally opaque. Explainable AI (XAI) is a field of study that seeks to enhance the explainability or interpretability of AI models and their output.

Related terms: AI safety; responsible AI

Fine-tuning

Fine-tuning involves adapting a pre-trained model to specific use cases, using a smaller data set than the pre-training phase.

Related terms: foundation model; pre-training; use case

Foundation model

Foundation models are AI models that have been pre-trained on a large dataset and can be used to perform a range of different tasks across a variety of industries and fields, as opposed to models that have been trained to perform a specific task. A large language model (LLM) is a type of foundation model.

Related terms: large language model (LLM); pre-training

Generative AI

Generative AI models are models that are capable of producing new media (including text, images, video, code, and audio) based on very large amounts of training data. A large language model (LLM) is an example of a generative AI model.

Related terms: large language model (LLM)

GPT

GPT is an acronym that stands for 'generative pre-trained transformer'. A GPT is a natural language processing (NLP) model that is pre-trained on a large amount of text data to produce language fluently. ChatGPT is built using a GPT model.

Related terms: ChatGPT; natural language processing (NLP); model; pre-training; transformer

Guardrails

Guardrails are a safeguarding mechanism for generative AI models. They consist of limits and rules placed on a model to improve safety, reduce the risk of misuse, and ensure that the model provides an output that is aligned with its intended purpose.

An example of a guardrail on an AI application is a rule preventing the model from giving answers about topics for which it was not designed, like specific financial or legal advice.

Related terms: AI safety; alignment; responsible AI

Hallucination

Hallucination in the context of generative AI means the generation of incorrect or inaccurate but plausible-sounding information. Hallucinations may be caused by limitations in the underlying training data or in the training process, among other factors.

Hallucinations risk introducing factual inaccuracies and misinformation into decision-making processes. It is important that generative AI systems have guardrails and mechanisms in place to prevent or mitigate hallucination, particularly those intended for use in professional or high-risk contexts.

Related terms: AI safety; guardrails; training data

Human-in-the-loop

'Human-in-the-loop', or HITL, refers to an approach to AI systems in which humans are closely involved in the design, testing, and ongoing maintenance of the system. This approach helps to ensure that AI systems benefit from human judgment and expertise, ideally leading to more accurate output and higher levels of trust.

Related terms: reinforcement learning; supervised learning

Machine learning

Machine learning is a sub-field of AI. It is a field of study and research in which algorithms are trained to extract knowledge and insights from data without being specifically programmed, as in traditional computer programming. Training data is key to machine learning.

Related terms: algorithm; supervised learning; training data; unsupervised learning; reinforcement learning

Modality

A modality is a type or mode of data. Modalities include audio, video, text and images.

Related term: multi-modal

Model

An AI model represents what has been learned from applying an algorithm to training data. After training, the model can be applied to new, previously unseen data to generate an output.

Related terms: algorithm; foundation model; machine learning; training data

Multi-modal

A multi-modal AI model can generate content across multiple modalities, including text, images, video and speech.

Related term: modality

Natural language processing (NLP)

Natural language processing (NLP) is a field of AI that relates to understanding and generating human language in written and spoken form. Natural language understanding (NLU) and natural language generation (NLG) are subsets of NLP.

Large language models (LLMs) are very large AI models that are trained to perform NLP tasks. Chatbots are an example of an NLP use case.

Related terms: chatbot; large language model (LLM); model; use case

Pre-training

Pre-training is the initial phase of training an AI model on a large set of training data before fine-tuning it to suit more specific tasks.

Related terms: fine-tuning; foundation model; training data

Prompt

A prompt is the written instruction or query given as input to a generative AI model to cause it to produce a response. A simple example of a prompt in a legal research and drafting context is “Draft a clause for a privacy policy that deals with overseas disclosure of personal information.”

Related terms: generative AI; prompt engineering

Prompt engineering

Prompt engineering involves selecting and optimising the written instructions given to a generative AI model – for example, a large language model (LLM) – with the aim of producing the desired result for a particular task. Prompt engineering best practices include giving clear and concise instructions, providing any relevant context, letting the model know what role or persona to assume, and providing examples.

An example of a good prompt in a legal context is, “I am acting for the defendant in a negligence matter in the District Court of NSW. What are the requirements in relation to the giving of particulars in pleadings? Set out each requirement individually, along with the relevant provision.”

If a prompt is particularly complicated, it is good practice to break it down into smaller discrete sub-prompts that accomplish each step of a complex task. This process is called ‘prompt chaining’.

For tasks involving complex reasoning, it may also be useful to prompt the model to reason step by step or to ‘think before answering’ (‘chain-of-thought’ prompting). Research indicates that this prompting methodology leads to improved performance in tasks involving logical reasoning. An example of chain-of-thought prompting is, “Reasoning step by step, explain how the High Court applied the principles governing joint criminal enterprise in *IL v R*.”

Related terms: generative AI; large language model (LLM); prompt

Reinforcement learning

Like supervised and unsupervised learning, reinforcement learning is a machine learning technique. A reinforcement learning algorithm attempts to achieve its objective by interacting with the environment and learning from feedback in a process of trial and error – similar to the way in which humans learn.

Related terms: agentic AI; machine learning; supervised learning, unsupervised learning

Responsible AI

Responsible AI is a set of practices and principles intended to ensure that AI is beneficial, safe and ethical.

Related terms: AI safety; alignment; bias

Retrieval-augmented generation (RAG)

Large language models (LLMs) are trained on large amounts of text data to produce fluent language. Without fact-checking mechanisms, LLMs lack the ability to judge the accuracy of the text produced. Retrieval-augmented generation (RAG) enhances the accuracy, relevancy and authoritativeness of LLM-generated text by ensuring that the model creates its responses on the basis of trusted sources or sources from a pre-determined database.

Related terms: hallucination; large language model (LLM)

Supervised learning

In supervised learning, machine learning models are trained using training data that has been annotated. These annotations or labels assist the model in inferring rules, patterns and insights in the data.

Related terms: machine learning; training data; unsupervised learning

Token

In the context of large language models (LLMs), a token is a word or part of a word. LLMs have limits on the number of tokens that they can process at one time (their 'context window').

Related terms: context window; large language model (LLM)

Training data

Training data is the data used to teach a machine-learning model to extract patterns and insights. It is useful to think of training data as a large number of examples that will help the model to make predictions about data that it has not yet seen.

Related terms: machine learning; model; training data

Transformer model

A transformer model is a type of AI model that is particularly good at understanding context and the real-world connection between sequential data like words. Transformer models are the foundation of large language models (LLMs).

Related terms: GPT; large language model (LLM)

Unsupervised learning

In unsupervised learning, machine learning models are trained using training data that humans have not annotated. The model must infer rules and insights from the data on its own. In unsupervised learning, there is a risk of perpetuating biases that are present in the underlying training data.

Related terms: machine learning; supervised learning; training data

Use case

A use case is the application of a technology or product to achieve a specific goal. An example of a generative AI use case in the legal context is document drafting.