AI Glossary

This glossary is designed to provide lawyers with a shared language around artificial intelligence. It highlights terms most relevant to lawyers and legal professionals working with AI-enabled tools. It is not intended to be a comprehensive guide to computer science or advanced AI research.

This glossary is for professional reference only. Definitions are written in plain, non-technical language. It is not a substitute for IT, security, or compliance guidance.

Term	Definition
Core Concepts	
Artificial Intelligence (AI)	Technology that enables computers to perform tasks that normally require human intelligence, such as recognizing patterns, reasoning, or generating text.
Deep Learning	A subset of ML that uses multi-layered neural networks to process complex data.
Machine Learning (ML)	A branch of AI where systems learn from data and improve over time without being explicitly programmed.
Natural Language Processing (NLP)	AI's ability to understand, interpret, and respond to human language.
Large Language Model (LLM)	An AI system trained on vast amounts of text to generate human-like responses.
Token	A unit of text (like a word or part of a word) that AI systems process to understand and generate language.
Agentic AI	AI systems that can take actions or make decisions on their own to complete a task, rather than just responding to instructions.
Generative AI	
Generative AI	AI that can create new content such as text, images, or code.
Prompt	The input or instruction given to an AI system.
Prompt Engineering	The practice of crafting prompts in a clear, specific way to improve AI outputs.
Hallucinations	When AI generates information that sounds plausible but is factually incorrect.
Training Data	The information used to "teach" an AI model how to respond.
Context Window	The amount of information an AI system can "remember" at once when generating a response.
Guardrails	Built-in rules or limits that keep AI systems from producing harmful, misleading, or inappropriate results.

Retrieval Augmented Generation (RAG)	A method where AI pulls in relevant information from trusted sources before generating an answer, improving accuracy and reliability.
Risk & Compliance	
Bias	When AI produces results that are systematically unfair or skewed.
Explainability	The ability to understand how and why an AI system reached its output.
Metadata	Hidden information about a file, such as author, edits, or timestamps.
Data Residency	Where data is physically stored and processed.
Confidentiality Risk	The danger of exposing sensitive or privileged data when using AI.