

SPEAK THE LANGUAGE

Glossary of AI Terms

LLM	Large Language Model — a system that predicts the next chunk of text, one piece at a time.
Token	A word-piece; the unit an LLM reads and writes in ($\approx \frac{3}{4}$ of a word).
Next-token prediction	The one core operation: given the text so far, guess what comes next.
Parameters / weights	The billions of internal “dials” tuned during training.
Training / pretraining	Learning patterns by predicting hidden next words across huge amounts of text.
Knowledge cutoff	The date training stopped — the model doesn’t know events after it.
Context window	The text the model can “see” at once — its short-term memory, measured in tokens.
Prompt	The instructions and context you give the model.
System prompt	Standing instructions set once — the model’s role and rules.
Temperature	A dial for randomness: low = safe/consistent, high = creative/surprising.
Sampling	Choosing the next token probabilistically — why answers vary.
Hallucination	When the model states something false with confidence.
Fine-tuning	Extra training on good-answer examples to shape behaviour.
RLHF	Reinforcement Learning from Human Feedback — using human ratings to make it helpful.
Alignment	Teaching a model to be helpful, honest, and safe.

SPEAK THE LANGUAGE

Glossary of AI Terms

Jailbreak	A prompt that talks a model past its own safety rules.
Few-shot	Giving a few examples in the prompt so it pattern-matches off them.
Chain-of-thought	Reasoning step by step instead of jumping to an answer.
Inference	Running the model to generate an answer (vs training it).
Transformer	The neural-network design most modern LLMs are built on.
GPT	"Generative Pre-trained Transformer" — a common family of LLMs.
Embedding	Text turned into numbers that capture meaning.
Tokenizer	The component that chops text into tokens.
Prompt engineering	The skill of asking clearly to get better answers.
Tool use	Letting a model call a calculator, web search, or code runner to cover its gaps.
Agent	An LLM given tools and a loop so it can act, not just talk (see Course: AI Agents Explained).
Guardrails	Limits placed around a model to keep its behaviour safe.
Multimodal	A model that also handles images, audio, or video — not just text.
Open vs closed model	Whether the model's weights are publicly available or kept private.
Confidently wrong	The core risk: fluent, certain-sounding text that happens to be false.