



AI Literacy: Understanding LLMs

This course is part of the **AI Literacy** series and is designed to help meet the **AI literacy requirements** set out in **Article 4** of the **EU AI Act**, which **entered into force** on 2 February 2025.

By building the skills, knowledge, and understanding to interact effectively and responsibly with AI systems, you will not only improve your day-to-day results but also strengthen your organisation's compliance posture. The courses focus on practical, plain-language learning that benefits any role, from operational staff to senior leadership, while ensuring you can oversee, apply, and guide AI use with competence, compliance, and confidence.



MODULE 1: CORE CONCEPTUAL MODEL

1.1 What a Large Large Language Model (LLM) Really Is

- Cut through the hype and get a crystal-clear picture of what drives tools like ChatGPT, why it's neither a "search engine" nor a "knowledge vault," and what that means for your everyday use.

1.2 How LLMs Are Trained

- Learn the three-stage journey from raw internet text to a model that can write, reason, and follow instructions, and why understanding this process is the first step to mastering it.

1.3 How AI "Represents" Language

- Discover the hidden language inside the machine: how it breaks your words into units, maps them into meaning-rich spaces, and uses these maps to understand and respond.

1.4 How the Machine Thinks on the Fly

- Peek into the real-time process of how an LLM builds its answers, one piece at a time, and how small changes in your wording can change the entire outcome.

1.5 The Built-In Limits You Must Work Around

- Understand the root causes of AI's most talked-about flaws, hallucinations, bias, and outdated facts, so you can spot them instantly and prevent them from derailing your work.

MODULE 2: ARCHITECTURAL ESSENTIALS

2.1 Demystifying the Transformer

- Understand the breakthrough architecture that powers every modern LLM, and why the "attention mechanism" is the secret to its leap in language understanding.





AI Literacy: Understanding LLMs

MODULE 2: ARCHITECTURAL ESSENTIALS

2.2 The Scale Equation

- See how parameters, computing power, and data interact to shape a model's capability, and what trade-offs engineers make when building the AI you use.

2.3 From GPT-3.5 to GPT-5

- Explore how each new generation of GPT has evolved in reasoning, context handling, and efficiency, and how to pick the right one for your needs.

2.4 Speed vs. Frontier

- Learn the crucial difference between models built for fast, affordable production use and those designed for cutting-edge research, and when to choose each.



MODULE 3: BEHAVIOR & CONSTRAINTS

3.1 Controlling Creativity and Consistency

- See how a single setting can make the same AI either a rigid rule-follower or a free-flowing idea generator, and how to choose the right balance for your task.

3.2 The Sweet Spot in Prompt Size

- Learn why giving too little or too much information can backfire, and how to work within the model's "memory span" for maximum precision.

3.3 What the Model Knows and What It Guesses

- Understand the limits of the AI's training and why it sometimes improvises, so you can spot when it's reasoning vs. recalling.

3.4 The Two Layers of Every Conversation

- Discover how behind-the-scenes system prompts and your own instructions interact to shape every answer you get.

MODULE 4: RISKS & FAILURE MODES

4.1 Why AI Makes Things Up

- Uncover the mechanics behind "hallucinations" and why even the most confident answer can be completely wrong.

4.2 How Bias Creeps In

- See how an AI's training can tilt its answers and how those hidden influences can impact decision-making in sensitive areas.





AI Literacy: Understanding LLMs

MODULE 4: RISKS & FAILURE MODES

4.3 The Security Traps Few Users Spot

- Learn about prompt injection and data leakage risks before they catch you off guard.

4.4 From Risk-Free Fantasy to Governance Reality

- Shift your mindset from eliminating AI risks to actively managing them, using both technical safeguards and human oversight.



MODULE 5: APPLICATION

5.1 Turning Mechanics into an Advantage

- See how understanding the model's inner workings translates into more powerful, precise prompts.

5.2 Few-Shot vs. Zero-Shot

- Learn when to rely on the AI's built-in knowledge and when to guide it with tailored examples.

5.3 Decomposition

- Break down complex requests into smaller steps to improve accuracy and control.

5.4 Self-Critique

- Use the AI as its own editor to catch errors and refine outputs before you act on them.

5.5 Retrieval-Augmented Generation (RAG)

- Ground your AI's answers in real, verifiable sources to cut hallucinations and boost credibility.

5.6 Timeless Prompting Skills

- Discover why these techniques will stay valuable no matter how advanced future models become.

MODULE 6: GPT-5 PROMPTING & REASONING UPDATES

6.1 Understanding “Agentic” Behaviour

6.1.1 What agentic means in GPT-5

6.1.2 Controlling autonomy with “eagerness” settings

6.1.3 Why autonomy levels matter for oversight

6.2 Tool Preambles for Transparency

6.2.1 What they are and why they matter

6.2.2 How preambles improve trust and auditability





AI Literacy: Understanding LLMs

6.2.3 Designing an effective preamble

6.3 Instruction Hygiene & Prompt Optimizer

6.3.1 Why clarity matters more in GPT-5

6.3.2 Common prompt contradictions to avoid

6.3.3 Using an optimizer to improve reliability

6.4 New Parameters: Verbosity vs. Reasoning Effort

6.4.1 How verbosity controls output length

6.4.2 Allocating reasoning effort for complex tasks

6.4.3 Matching settings to your use case



The AI Literacy Series

If you want to take your prompting skills further into compliance-ready, governance-aligned competence, continue with [AI Literacy: Ethical Prompt Fluency](#), which builds on and extends the skills you gain here.

For those who would like to deepen their technical and conceptual understanding of AI systems, [AI Literacy: Technical Foundations for Non-Tech Roles](#) offer valuable context that will enhance your work with Prompt Fluency Toolkit, but it is not a prerequisite.

