

Understanding AI, Tools and Prompt Engineering

A practical one-day introduction for analysts and business professionals

FA-0708 · 1 day

AI for Leaders & Business · Agentic & Generative AI

Introduction

Explore AI foundations, compare selected tools and practise prompt and context design for everyday analytical work. This one-day introductory workshop combines short explanations, guided demonstrations and a small prompt-refinement exercise.

Compare capabilities, data access, reliability, cost and human oversight rather than assuming that every tool or subscription offers the same features. Examples use approved sample data; no coding or production deployment is required.

Learning outcomes

- Distinguish rule-based systems, machine learning, generative models and agentic workflows.
- Explain LLM training, adaptation and inference at an introductory level, including uncertainty and hallucination.
- Compare conversational assistants, Power Platform AI capabilities, agentic tools and source-grounded notebook assistants.
- Evaluate a proposed use case against accuracy, permissions, privacy, integration and cost constraints.
- Build and evaluate a reusable prompt with relevant context, examples, output constraints and human verification.

Prerequisites

- Comfort with business analysis, reporting or everyday office software.

- Basic familiarity with cloud applications; no programming or machine-learning background required.
- Access to the approved demonstration tools where available; account, licence and administrator restrictions are checked before practical work.

Training outline

1. Module 1 – AI foundations and practical limits

- Rule-based systems versus supervised, unsupervised and reinforcement learning; neural networks and deep learning in context.
- LLMs: an accessible overview of model architecture, training, fine-tuning and inference; fine-tuning is discussed conceptually, not promised as a hands-on service.
- Generative, predictive, prescriptive and conversational applications; distinguish an assistant response from an agent performing a sequence of actions.
- Potential help with routine tasks and insight generation; examine bias, hallucination, missing context and over-reliance.

2. Module 2 – Agents, tools and MCP

- Observe, plan, act and evaluate: trace a bounded multi-step workflow and identify the human approval points.
- Tool orchestration and Model Context Protocol: connecting an AI application to permitted data and tools through an open protocol.
- Interoperability does not grant data access or guarantee trustworthy actions; review authentication, permissions, auditability and action limits.

3. Module 3 – Compare AI tool categories

- ChatGPT for drafting, summarisation, questions and selected data or code assistance; check outputs against supplied evidence.
- Connected apps and plugins can add tools and reusable workflows, subject to account and workspace permissions; distinguish these end-user capabilities from developer API tool calling.
- Microsoft Power Platform: distinguish AI Builder, Copilot capabilities in individual products and Copilot Studio rather than treating them as one universal AI hub.
- Agentic assistants, including Manus as a comparison example: evaluate task scope, reliability and supervision without assuming a particular internal multi-agent architecture.
- Source-grounded notebook assistants from Google, including Gemini Notebook: use supported documents and web sources, follow citations and compare generated summaries with the originals. Public YouTube sources require captions; arbitrary video ingestion is not assumed.

4. Module 4 – Guided tool exploration

- Use two representative tasks, such as a document brief and a small reporting question, to compare available tools.
- Observe Power Platform automation or application examples and discuss connectors, setup, governance and feature availability.
- Refine a prompt, inspect an answer and identify where the assistant needs better context or a reliable source.
- Review an agentic workflow demonstration and a source-grounded research brief; record unsupported statements, missed evidence and supervision needs.
- Broader product comparisons are demonstrations or discussion, not a promise of hands-on accounts for every platform.

5. Module 5 – Prompt and context engineering

- Define the task, audience, constraints and output format; compare zero-shot instructions with a small set of relevant examples.
- Use roles and clear structure where helpful, but do not assume role labels improve factual accuracy. Adapt prompts to the selected model and test alternatives.
- Request concise explanations and verifiable evidence rather than hidden chain-of-thought; reasoning models do not generally need step-by-step prompting.
- Permit uncertainty and ask for citations where sources are available, then check the citations and claims; these instructions do not eliminate hallucination or bias.
- Select task-relevant documents and context, separate reference content from instructions, and manage stale or changing context. Distinguish supplying context from embedding-based retrieval.
- Refine one analyst prompt against sample cases using agreed quality criteria and human review; retain a reusable template with its assumptions and limits.
- Discuss iterative coding-assistant trends as context, not a coding lab or an attributed guarantee; choose a bounded next experiment.