Faysal Aziz

Demystifying AI: From Basics to Agents

A one-day conceptual overview of models, tools and responsible agent design

FA-0601 · 1 day

Agentic & Generative AI · AI for Leaders & Business

Understand how modern AI models and tool-using agents fit together, then review a responsible use-case proposal.

Introduction

This one-day overview explains AI foundations, machine learning, deep learning, language models and tool-using agents at a conceptual level. It connects prompting and MCP integration to business use cases, limitations and governance questions.

The emphasis is on informed discussion and a small responsible-agent proposal, not training models or deploying a production agent in one day. Practical examples are illustrative and do not imply client relationships, verified adoption rates or autonomous reliability.

Learning outcomes

The course teaches participants to:

- Distinguish AI capability labels, model-learning approaches and agent behaviour.
- Explain training and evaluation concepts and the roles of neural networks and language models.
- Frame prompts with clear context, constraints and review criteria.
- Describe MCP host/client/server roles and the difference between tool access and application orchestration.
- Analyse an agent use case with permissions, human oversight and failure handling.
- Discuss bias, privacy, accountability and the limits of explainability.

Prerequisites

- Basic familiarity with AI/ML terminology (e.g., “model,” “training,” “dataset”)
- Comfort with high-level concepts in programming, statistics, or data processing
- Openness to discuss real-world examples and ethical dilemmas

Training outline

1. Module 1 — AI Foundations

- Reactive vs. Deliberative vs. Learning agents
- Narrow AI vs. Artificial General Intelligence (AGI)

Distinguish agentic behaviour from model capability and deployment marketing; AGI is a concept, not an assumed capability of a current product.

2. Module 2 — Machine Learning and Deep Learning

- Overview: supervised, unsupervised, reinforcement learning
- Deep neural networks, architectures (CNNs, RNNs), backpropagation
- Use of Deep Learning in agentic systems

Training, validation, evaluation and data limitations through high-level examples rather than a full implementation lab.

3. Module 3 — Language Models and Prompting

- What makes LLMs 'large': data, transformer architecture, token prediction

Context windows, instructions, examples and iterative review; distinguish plausible text generation from verified knowledge.

Reasoning and chain-of-thought prompting as model-dependent concepts, not universal requirements or evidence of correctness.

Structured task requests and source checking before tools or downstream actions.

4. Module 4 — MCP and Tool Integration

- MCP host, client and server roles; tools, resources and prompts as supported primitives.
- Protocol/version and capability compatibility; tool access remains subject to application permissions.

- MCP exchanges context and tool capabilities; the AI application manages its model use and orchestration.
- Data permissions, untrusted tool content and prompt-injection risks; protocol use does not itself establish safety.

5. Module 5 — Agentic Systems and Use Cases

- Systems that use models to plan/select actions and interact with tools within defined boundaries.
- Autonomy, planning, context and escalation; distinguish automation from reliable independent judgement.
- Illustrative customer-service, finance, supply-chain and cybersecurity support scenarios.
- Evaluate business value, uncertainty, unauthorised actions and accountability without unsupported adoption or incident statistics.

6. Module 6 — Ethics and Governance

- Key principles: fairness, transparency, privacy, accountability
- Machine ethics taxonomy: impact, implicit, explicit, full agents
- Explainable AI (XAI): transparency in black-box systems
- Governance issues in agentic AI: oversight, human-in-the-loop, audit tracing
- Broader risks: job displacement, misinformation, alignment with human values

Explainability can aid understanding but does not prove fairness, safety or accountability; use human ownership and evidence-based review.

7. Module 7 — Synthesis and Discussion

- How ML/DL, LLMs, MCP, and ethics combine in end-to-end agent systems
- Group activity: draft a responsible agent proposal (use case, prompt, tools, safe-guards)

Discuss deployment and monitoring questions, remaining evidence and how to revise the proposal as technology changes.

© 2026 Faysal Aziz · Kuala Lumpur