Faysal Aziz

AI Essentials for Leadership

A one-day introduction to concepts, governance and adoption

FA-0561 · 1 day

AI for Leaders & Business · Agentic & Generative AI

A non-technical leadership workshop on AI capabilities, tool evaluation, prompting, oversight and a pilot-to-scale adoption roadmap.

Introduction

This one-day course gives leaders a practical vocabulary for discussing AI and making initial adoption decisions. It explains machine learning, deep learning, language models and tool-using agents without requiring programming.

Participants consider capabilities and limitations, evaluation criteria, governance responsibilities and a high-level adoption roadmap. The course supports informed questions and planning; it does not implement production systems or guarantee business impact.

Learning outcomes

The course teaches participants to:

- Explain AI, machine learning, deep learning and language models in accessible terms.
- Recognise variability, hallucinations, bias and unexpected behaviour in modern AI systems.
- Compare AI tool capabilities using cost, reliability, data, security and integration criteria.
- Guide prompt use with clear instructions, examples and output evaluation.
- Explain agent boundaries, tool access and human oversight.
- Identify responsible-use and governance responsibilities.
- Outline a measured pilot, evaluation and adoption roadmap.

Prerequisites

- No technical or programming background required
- Basic comfort with strategic, high-level discussions
- Openness to conceptual thinking and risk-based tradeoffs
- A desire to lead responsibly into the AI era

Training outline

1. 1 — AI from a manager's perspective

AI as a concept: beyond headlines

- Traditional automation vs adaptive intelligence
- Task-specific AI and general-intelligence concepts: distinguish demonstrated capability from speculation.

- Key dimensions: perception, decision, interaction
- Common misconceptions (AI will "think like humans," full autonomy too fast)

2. 2 — How modern AI works

- Data, pattern learning, predictions and decisions: distinguish model output from human or system action.
- Training vs inference (analogy: "learning phase" vs "doing phase")
- Tradeoffs: accuracy vs explainability, speed vs thoroughness
- Risk areas: overfit, bias, unpredictability

3. 3 — Machine learning, deep learning and LLMs

Machine learning: learning from historical data

- Supervised (input → label), unsupervised, reinforcement as business metaphors

Deep learning: layered "neural" abstractions

- Why deep learning is widely used for images, speech and unstructured data; suitability depends on the task and evidence.

Large Language Models (LLMs): generative text systems

- Tokenization and context (without math)
- Pretraining + fine-tuning logic
- Retrieval augmentation (hybrid systems)

Emergence & non-determinism

- Unexpected behaviours or capabilities: evaluate rather than assume intelligence or intent.
- Output variability: generation settings and system context can affect repeated responses; not every AI system is inherently non-deterministic.
- Risks: hallucinations, sensitivity to phrasing, hidden failure modes

4. 4 — Capabilities, tools and evaluation

- Core capability buckets: text, translation, summarization, question answering, image/audio, decision support
- Architectural pattern: tools + models + orchestration
- Evaluation criteria from a business lens: cost, latency, reliability, explainability, vendor stability
- Integration challenges: data pipelines, compliance, security, scaling

5. 5 — Prompting and output review

- Why prompts matter: your textual instructions are the "interface" to the model

  Core strategies:

  - Instruction style (clear directive)
  - Example style (show "good answers")
  - Request relevant evidence, assumptions or a concise rationale; use model-appropriate prompts rather than treating step-by-step reasoning prompts as universally necessary.
  - Iterative refinement (test, adjust, compare)
- Manager's role: setting guardrails for prompt use, reviewing prompt templates, monitoring output quality
- Prompt risk awareness: sensitivity, prompt injection, ambiguous instructions

6. 6 — Tool-using agents in business

What is an "agent"?

- Goal-oriented, multi-step, using tools, with a feedback loop

- Potential uses: research assistance, workflow orchestration and decision support, with defined limits.

  Key design decisions managers should know:

  - Scope boundaries and safeguards
  - Visibility, audit logs, human oversight
  - Fail-safe mechanisms, intervenability
- Risks to watch: objective misalignment, "drift," tool misuse

7. 7 — Governance and responsible oversight

- Why governance matters: risk, accountability, legal obligations and public trust.
- Core principles: fairness, transparency, accountability, safety, privacy

  Organizational levers:

  - Governance bodies or committees
  - Risk review, auditing, red teaming
  - Documentation, logging, incident response
  - Model version control, change tracking, risk thresholds
- Applicable laws, contractual requirements and voluntary frameworks: distinguish obligations from guidance.
- Governance in practice: balancing innovation vs control

8. 8 — A measured adoption roadmap

Personal adoption (for managers)

- Experiment with small use cases, build intuitive grounding
- Stay current, connect with practitioners, ask the right questions
- Develop a “critical mindset” toward AI outputs

Organizational adoption stages

- Identify high-value, low-risk pilot use cases
- Define metrics (value, error rate, cost, time saved)
- Pilot, evaluate and iterate; scale only when evidence and operational readiness support it.
- Build infrastructure, data capability, integration, maintenance
- Build governance, compliance, stakeholder alignment
- Training, change management, culture shift

Common project pitfalls and illustrative lessons

- Underestimating integration, data quality, ongoing monitoring
- Overhyping results, underpreparing for errors
- Model drift, versioning, operational debt
- Resistance from staff, misalignment of expectations

© 2026 Faysal Aziz · Kuala Lumpur