Faysal Aziz

AI for Managers & Leaders

From demystifying AI to governing it responsibly across the enterprise - 2 days

FA-0067 · 2 days

Agentic & Generative AI · AI for Leaders & Business · Data & Analytics

---

Introduction

AI is no longer a specialist topic sitting inside data science teams. It now affects how managers make decisions, redesign workflows, evaluate risk, improve customer experience, and lead teams through change. For an MNC, the challenge is not only understanding what AI can do, but also knowing where it is reliable, where it fails, how modern tools fit together, and how to adopt them without creating governance, security, or compliance problems. The current market is also shifting from standalone chatbots toward agentic systems, enterprise connectors, and tool-using models, which makes management literacy in AI a practical leadership requirement rather than a technical luxury.

This two-day manager-focused course is designed to build conceptual clarity first, then move into practical judgment: how AI works, what LLMs and agents are, what prompt engineering is and is not, why tokenization matters, how governance should be structured, and how managers should evaluate tools, vendors, risks, and use cases. The instructor brings over 30 years of industry experience and will teach with real industry-demanded content rather than an academic treatment of the subject.

Learning outcomes

By the end of the course, participants will be able to:

- Explain AI in plain business language, including core concepts such as machine learning, deep learning, generative AI, LLMs, agents, and MCP.
- Describe at a high level how modern AI systems work, including training, inference, context windows, retrieval, tokenization, and tool use.
- Distinguish between predictive AI, generative AI, copilots, assistants, and autonomous or semi-autonomous agents.
- Write and evaluate better prompts for managerial work, and understand when prompt engineering should give way to workflow design, evals, or context engineering.

- Assess major AI risks, including prompt injection, insecure output handling, training data poisoning, model denial of service, supply-chain vulnerabilities, and excessive agency.
- Identify governance requirements for enterprise AI adoption, including data controls, access boundaries, human oversight, auditability, and model or tool selection principles.
- Compare current tool categories and judge where tools such as NotebookLM, Manus, OpenClaw, Claude Cowork, and enterprise agent platforms fit into business workflows.
- Build a practical AI adoption lens for an MNC: where to start, what to avoid, what to pilot, and what management controls to put in place first.

Prerequisites

- No coding background required
- Prior people-management or functional-management experience in an enterprise environment
- Familiarity with standard business processes such as reporting, planning, operations, customer service, finance, HR, procurement, or risk management
- Basic exposure to tools such as ChatGPT, Copilot, Gemini, Claude, or similar is helpful but not required
- Willingness to engage with policy, governance, and change-management issues alongside the technology itself

Training outline

1. Why AI matters now for managers in an MNC

1. The shift from experimentation to enterprise adoption
2. Why AI literacy is now a management capability
3. Competitive, operational, and workforce implications
4. The difference between hype, capability, and business value
5. What senior leaders, middle managers, and functional heads each need to understand
6. Hands-on exercise: Identifying 3 high-impact AI opportunities within participants' own business units

2. Demystifying AI

1. What AI is and is not

   A simple taxonomy of AI

   1. Rules-based systems
   2. Machine learning
   3. Deep learning

4. Generative AI
5. Multimodal AI
6. Agentic AI

Common terms managers hear

1. Model
2. Algorithm
3. Training data
4. Parameters
5. Inference
6. Embeddings
7. Context window
8. Retrieval
9. Fine-tuning
10. Guardrails

4. Separating science fiction narratives from real enterprise capability
5. What AI can do well
6. What AI still does poorly
7. Hands-on exercise: Translate AI jargon into plain business language for a leadership audience

3. How AI works at a conceptual level

1. Pattern learning and statistical prediction
2. Training versus inference
3. Supervised, unsupervised, reinforcement, and self-supervised approaches
4. Neural networks in plain language
5. Why scale changed the field
6. Foundation models and frontier models
7. The role of data, compute, and feedback loops
8. Why outputs feel intelligent even when the underlying process is probabilistic
9. The limits of explanation and interpretability for managers
10. Hands-on exercise: Map a simple business problem to how an AI system would process it conceptually

4. From machine learning to generative AI

1. Traditional analytics and machine learning versus generative systems
2. Predictive AI versus generative AI

3. Structured data versus unstructured data
4. Classification, forecasting, recommendation, generation, and reasoning
5. Business examples across functions
6. Where generative AI complements analytics rather than replacing it
7. Hands-on exercise: Classify real business use cases into predictive vs generative AI

5. Understanding LLMs

1. What a large language model is
2. Tokens as the unit of processing

   Tokenization and why it matters

   1. Cost
   2. Latency
   3. Context limits
   4. Language effects
   5. Prompt design implications
4. Context windows and memory constraints
5. Pretraining, instruction tuning, and post-training
6. Hallucination, uncertainty, and confidence problems
7. Why LLMs are strong at language but uneven at truth
8. The role of retrieval, grounding, and source constraints
9. Multimodal expansion beyond text
10. Hands-on exercise: Analyze outputs from an LLM and identify strengths, weaknesses, and risks

6. Reasoning, tools, and action

1. The difference between answering and acting
2. Tool use in modern AI systems
3. Function calling and external actions
4. Retrieval-augmented generation
5. Knowledge-grounded systems
6. Workflows versus agents
7. When to use a chatbot, copilot, workflow engine, or agent
8. Human-in-the-loop versus human-on-the-loop operating models
9. Hands-on exercise: Design a simple AI-assisted workflow for a recurring managerial task

7. Agents, MCP, and connected AI systems

1. What an AI agent is
2. The rise of enterprise agents

   Agent capabilities

   1. Planning
   2. Tool invocation
   3. Context management
   4. File handling
   5. Web interaction
   6. Workflow orchestration

   Model Context Protocol fundamentals

   1. What MCP is
   2. Why it matters
   3. Clients, servers, tools, resources, and prompts
   4. Enterprise use cases for MCP
   5. Governance implications of connected agents
5. Excessive agency and operational risk
6. Patterns for safe deployment of connected AI
7. Hands-on exercise: Identify where MCP-style integration could connect systems in participants' organizations

8. Prompt engineering for managers

1. What prompt engineering is
2. What prompt engineering is not
3. Prompting as managerial communication with systems

   Core prompt patterns

   1. Role
   2. Objective
   3. Context
   4. Constraints
   5. Output format
   6. Quality criteria

   Prompting for executive use cases

   1. Summarization

    2. Decision support
    3. Scenario analysis
    4. Drafting
    5. Stakeholder communication
    6. Meeting preparation

   Improving output quality

    1. Better instructions
    2. Better context
    3. Better source material
    4. Better decomposition
7. Prompt chaining
8. Structured outputs
9. Reducing ambiguity
10. Knowing when prompt engineering is the wrong solution
11. Context engineering as the next step beyond prompting
12. Hands-on exercise: Rewrite weak prompts into structured, high-quality prompts for business scenarios

9. Context engineering

1. What context engineering is and why it matters
2. The shift from prompt-centric to context-centric design

   Types of context

    1. Static context
    2. Dynamic context
    3. Retrieved context
    4. User-specific context
    5. System and tool context

   Context assembly strategies

    1. Retrieval pipelines
    2. Document selection and ranking
    3. Chunking and embeddings
    4. Context compression and summarization
5. Managing context windows effectively
6. Tradeoffs between context size, relevance, cost, and latency
7. Designing reliable AI systems through better context rather than longer prompts

8. Context leakage and security considerations
9. Enterprise patterns for context management
10. Hands-on exercise: Design a context pipeline for a real business use case

10. Tokenization, cost, and performance management

1. What tokens are
2. Why token counts vary by wording and language
3. The relationship between tokens, context, cost, and response quality
4. Prompt bloat and context waste
5. Compression strategies for enterprise use
6. Long-context tradeoffs
7. Cost-awareness for managers procuring or scaling AI solutions
8. Throughput, latency, and user experience implications
9. Why governance should include token and usage discipline
10. Hands-on exercise: Optimize a prompt or workflow to reduce token usage while maintaining quality

11. Understanding model and tool choices

1. Frontier models versus open models
2. Hosted AI versus local-first and self-hosted approaches
3. General-purpose models versus specialist tools

   Tradeoffs

   1. Capability
   2. Security
   3. Privacy
   4. Control
   5. Speed
   6. Cost
   7. Vendor lock-in
5. Selection criteria for an MNC
6. Build, buy, integrate, or partner decisions
7. Hands-on exercise: Evaluate and select an AI solution based on a defined business scenario

12. Current AI tool landscape managers should know

Research and synthesis tools

1. NotebookLM and source-grounded knowledge work
2. Deep research style workflows

Enterprise agent platforms

1. OpenAI Frontier and managed enterprise agents
2. Claude Cowork and computer-assisted knowledge work
3. Manus and autonomous task execution

Open and local-first agent tools

1. OpenClaw and local-first multi-channel agent models

Collaboration and productivity integration

1. Connectors, shared context, and workflow embedding

Tool selection by use case

1. Executive productivity
2. Team productivity
3. Research
4. Knowledge management
5. Operations
6. Customer-facing workflows

Evaluation criteria for trending tools

1. Security posture
2. Permissions model
3. Data handling
4. Reliability
5. Auditability
6. Ecosystem fit
7. Geographic and regulatory considerations

7. Hands-on exercise: Compare two AI tools and recommend one for enterprise adoption

13. Business use cases across the MNC

1. Personal productivity for managers
2. Team productivity and knowledge reuse
3. Decision support and scenario planning
4. Sales and account management
5. Marketing and communications
6. Customer service and support operations
7. HR and learning

8. Finance and reporting
9. Procurement and supply chain
10. Legal and compliance support
11. IT service management
12. Product, strategy, and innovation functions
13. Cross-border collaboration and multilingual operations
14. Hands-on exercise: Build a use case portfolio for a department or function

14. Limits, failure modes, and managerial judgment

1. Hallucinations and fabricated detail
2. Overconfidence in fluent outputs
3. Bias and representational harm
4. Context loss
5. Inconsistent reasoning
6. Tool misuse and brittle automation
7. False precision in analysis
8. Domain-specific reliability limits
9. Why verification remains a management responsibility
10. Escalation and review thresholds for high-impact outputs
11. Hands-on exercise: Review AI-generated output and decide approval, revision, or rejection

15. Governance and enterprise controls

1. What AI governance means in practice

   Policy foundations

   1. Acceptable use
   2. Data classification
   3. Access controls
   4. Human oversight
   5. Approval workflows
   6. Logging and audit

   Governance roles

   1. Executive sponsors
   2. Risk and compliance
   3. IT and security

    4. Legal
    5. Business unit leaders
    6. End users
4. Model governance versus workflow governance
5. Vendor governance
6. Third-party risk and supply chain considerations
7. Documentation and traceability
8. Evaluation frameworks and red-teaming
9. Change management and workforce readiness
10. Measuring value without losing control
11. Hands-on exercise: Draft a basic AI governance checklist for participants' organizations

16. AI vulnerabilities and security risks

OWASP LLM risk categories

1. Prompt injection
2. Insecure output handling
3. Training data poisoning
4. Model denial of service
5. Supply chain vulnerabilities
6. Sensitive information disclosure
7. Excessive agency

2. Why connected tools expand the attack surface
3. Identity, permissions, and access sprawl
4. Data leakage paths
5. Jailbreaks and policy evasion
6. Malicious documents and retrieval poisoning
7. Security implications of agentic action
8. Sandboxing, containment, and monitoring
9. Observability for agent systems
10. Human approval checkpoints for critical actions
11. Hands-on exercise: Identify vulnerabilities in a sample AI workflow and propose mitigations

17. Risk, legal, and compliance considerations for MNCs

1. Cross-border data issues

2. Confidentiality and intellectual property
3. Record retention and discovery implications
4. Regulated workflows and explainability expectations
5. Contractual and procurement implications
6. Sector-specific risk thinking
7. Regional policy and standards awareness
8. Balancing innovation with compliance obligations
9. Hands-on exercise: Assess a use case for legal and compliance risks

18. Responsible use and workforce implications

1. Human accountability in AI-supported decisions
2. Transparency to staff and stakeholders
3. Skill shifts for managers and teams
4. New expectations for review, supervision, and delegation
5. Avoiding over-automation
6. Trust, adoption, and cultural readiness
7. AI literacy as part of leadership development
8. Hands-on exercise: Develop a communication plan for AI adoption within a team

19. Implementation roadmap for managers

1. How to identify high-value, low-risk starting points
2. Prioritizing use cases
3. Selecting pilot groups
4. Defining success metrics
5. Establishing governance before scale
6. Building reusable patterns
7. Avoiding fragmented tool sprawl
8. Partnering with IT, security, legal, and business stakeholders
9. Moving from pilots to operating model
10. Sustaining capability through training and review
11. Hands-on exercise: Create a phased AI adoption roadmap for a business unit

20. Manager toolkit and decision framework

1. Questions managers should ask before approving AI use

2. Questions managers should ask vendors
3. Questions managers should ask their own teams
4. A simple decision tree for use case suitability
5. A risk-value matrix for AI adoption
6. Red flags that require escalation
7. Practical checklist for safe and useful deployment
8. Hands-on exercise: Apply the decision framework to a real or hypothetical AI initiative

21. Latest tools and techniques spotlight

1. Source-grounded notebook workflows
2. Connected AI via MCP
3. Computer-use and browser-use patterns
4. Autonomous task agents
5. Local-first AI assistants
6. Shared-context enterprise agent platforms
7. Context engineering
8. Evals and benchmarking for business workflows
9. AI observability and agent monitoring
10. Emerging management implications of agentic work
11. Hands-on exercise: Explore and demonstrate one modern AI tool relevant to participants' roles

© 2026 Faysal Aziz · Kuala Lumpur