

AI for Executive Leadership and Governance

Conceptual literacy, oversight and an initial governance plan

FA-0564 · 1 day

AI for Leaders & Business · Agentic & Generative AI

A one-day executive workshop on AI capabilities, business value, risks, accountability, governance tools and an initial oversight charter.

Introduction

This one-day programme gives executives and board members conceptual AI literacy and a structured way to ask about value, risk and responsibility. It covers model and agent capabilities, deployment trade-offs and governance without mathematical or programming detail.

Topics are surveyed through illustrative scenarios and a focused governance-review exercise. Participants draft an initial charter or action plan for further organisational review. The workshop does not design a complete assurance programme, prove regulatory compliance or confer legal advice.

Learning outcomes

The programme teaches participants to:

- Explain key AI concepts and the conditions behind successful or unreliable use.
- Identify variability, drift, bias, hallucination, adversarial and accountability risks.
- Compare business opportunities, costs, pilot evidence and operational readiness.
- Explain tool-using and multi-agent systems and their control/oversight needs.
- Outline governance layers, decision rights and cross-functional responsibilities.
- Use a sample impact assessment, risk register and review checklist.
- Ask informed board/executive oversight questions and identify legal or specialist-review needs.
- Draft initial governance priorities and an executive charter.

Prerequisites

- Basic comfort with digital/IT concepts (data, software, systems)
- No deep technical background in AI or machine learning is required
- Familiarity with corporate governance, risk, compliance, audit processes

Training outline

1. 1 – Leadership context and AI concepts

Business context

- The macro AI landscape: acceleration, compute, data, talent, scaling
- Compare automation, augmentation and bounded agentic systems; these are choices, not an inevitable maturity ladder.
- Strategic questions: evidence of value, disruption assumptions and platform dependencies.
- Illustrative financial, supply-chain, customer-support and fraud-detection scenarios: benefits, failure modes and evidence requirements.
- Common misconceptions and hype traps

Definitions and system types

- Definitions: AI, machine learning, deep learning, symbolic AI, statistical AI
- Task-specific versus general-intelligence concepts and the uncertainty in capability claims.
- Generative models, reinforcement learning, agentic systems
- Hybrid models and neuro-symbolic models
- Relationship to data, feature engineering, architecture

2. 2 – How AI works and where it can fail

Conceptual foundations

- Training vs inference phases
- Model architectures: neural networks, transformers, LLMs (at a high level)
- Key enablers: compute, data scale, architectures, embeddings
- Overfitting, underfitting, regularization, generalization
- Distribution shift, concept drift, domain shift

- Adversarial inputs, robustness, sensitivity
- Interpretability, explainability, and “black box” limits
- Chain of components: data pipeline, model, inference, feedback loops

Risk and uncertainty overview

- Output variability and unexpected behaviour; not all systems are non-deterministic.
- Distribution shift, drift, feedback loops and compounding errors.
- Adversarial inputs, data poisoning, evasion and robustness.
- Bias, fairness, safety and objectives that reward the wrong behaviour.
- Hallucinations, overconfidence, spurious correlations and interpretability limits.
- Accountability gaps, ethical/legal/reputational risk and out-of-distribution failures.
- Low-probability/high-impact scenarios: distinguish evidence from speculation and select proportional mitigations.

3.3 — Business value and agentic systems

Value, readiness and adoption

- Use-case identification: process automation, predictive insights, decision augmentation, generative augmentation
- Value vs cost vs risk tradeoff
- ROI, metrics, pilots and POCs
- Readiness criteria and risks when considering pilot-to-production expansion.
- Data strategy: data governance, quality, features, sources, instrumentation
- Model release cycles, monitoring, retraining
- Infrastructure, MLOps, model ops, platform thinking
- Human-in-the-loop design, guardrails, fallback mechanisms
- Change management, adoption risks, cultural resistance
- Illustrative case reviews: what evidence would support a claimed success or failure?

Delegation and control

- Agentic AI: goal-oriented software using tools and feedback, including multi-agent designs; autonomy is bounded by system configuration.
- Delegation, trust, control: when do you give autonomy?
- Hierarchies vs peer agents, coordination, emergent behaviors

- Safety envelopes, constraint programming, kill-switch / override paths
- Simulations, sandboxing, red teaming
- Agentic risk: unintended incentives, goal conflicts
- Monitoring, oversight, rollback, escalation paths

4.4 — Governance and organisational design

- Governance goals: safety, robustness, transparency, accountability, alignment
- Layered governance: policy, process, technical guardrails, oversight bodies
- Compare organisational, process and technical governance layers; use an illustrative model rather than a prescribed universal architecture.
- Connect applicable obligations, organisational policy, review processes, technical controls and evidence.
- Adaptive governance for generative AI

Organizational roles and structures:

- AI leadership responsibilities, whether assigned to an existing executive or a dedicated role.
- AI ethics board, review committee, audit panels
- AI safety / assurance functions
- Cross-functional liaison (legal, risk, compliance, data, product)
- Embedding governance in existing structure (risk, compliance, internal audit)
- Lifecycle oversight: design review, impact assessment, continuous monitoring, incident response
- Escalation and decision rights: when human override, when board involvement
- Auditability, logging, traceability, versioning

5.5 — Practical review tools and case exercise

- AI impact assessment (privacy, fairness, safety, security)
- Risk register templates specific to AI
- Governance scorecards, maturity models
- Monitoring dashboards, anomaly detection, drift alerts
- Understand selected evaluation approaches such as fairness reviews, robustness tests and backtesting; specialist execution is outside this introductory course.
- Red-teaming, adversarial testing, scenario stress tests
- Incident reporting, root cause analysis, post-mortem
- Version control, rollback, fallback mechanisms
- Map relevant internal policies, external obligations and voluntary standards; seek appropriate specialist interpretation.

- Workshop: review a fictitious AI system and complete a short impact/governance assessment.

6. 6 — Board oversight and external requirements

Oversight questions and reporting

- The role of boards, risk committees, audit committees in AI oversight
- Questions executives / board members should ask
- AI inventory (catalog of AI systems in organization)
- Request proportionate documentation and visibility from development or procurement teams.
- Identify liability, fiduciary and director-responsibility questions for jurisdiction-specific legal review.
- Define escalation, intervention and shutdown/override decision rights where appropriate.
- Training and literacy programs for non-technical directors
- Reporting structures, dashboards, incident alerts

Regulation, frameworks and change

- Compare selected jurisdictional approaches, including the EU AI Act, at an awareness level; verify rules for the relevant context.
- Distinguish legal conformity/assessment requirements from voluntary assurance and certification schemes.
- International governance initiatives and proposals: distinguish recommendations from binding obligations.
- Trustworthy AI principles: transparency, fairness, robustness, privacy, explainability
- Standards, industry initiatives and assurance schemes: assess relevance rather than assume endorsement.
- Cross-border data flows, jurisdictional risk
- Adaptive governance — the need for co-evolving oversight in generative AI
- Scenarios: what future regulation or standard shifts might require adaptation

7. 7 — Priorities and initial charter

- Gap analysis: where does your organization stand?
- Prioritization: high-impact systems to govern first
- Pilot governance interventions, feedback loops
- Building capabilities: hiring, training, alliances
- Change management: embedding governance culture
- Monitoring and continuous improvement

- Scaling governance as AI adoption grows
- Capstone: draft an initial executive governance charter or action plan for further review.