Faysal Aziz

AI Governance - in depth

Building Trust, Control, and Accountability in the Age of Intelligent Systems

FA-0070 · one-day

Agentic & Generative AI · AI for Leaders & Business

Introduction

Understanding How AI Works, Why It Fails, and How Organizations Govern It Responsibly

Artificial Intelligence is no longer a futuristic concept reserved for research laboratories or technology giants. It now influences hiring decisions, financial approvals, healthcare recommendations, cybersecurity operations, customer engagement, and government services. Organizations are rapidly deploying AI systems to automate decisions, increase efficiency, and unlock new forms of value. Yet the same systems that create opportunity can also introduce serious risks: bias, misinformation, privacy violations, intellectual property concerns, security vulnerabilities, regulatory exposure, and loss of human oversight.

To govern AI effectively, professionals must first understand what AI actually is, how machine learning models function, why generative AI behaves unpredictably, and where the technology's limitations originate. Without that foundational understanding, governance becomes a checklist exercise instead of a meaningful risk and accountability framework.

This intensive one-day course provides a practical and business-oriented introduction to AI governance by first demystifying AI itself. Participants will explore how modern AI systems are built, trained, and deployed, how generative AI models such as large language models operate, and why these systems can produce inaccurate or harmful outputs. The course then transitions into governance, covering risk management, ethics, compliance, organizational controls, security, policy development, and emerging global governance frameworks.

The course is delivered from a real-world industry perspective by an instructor with over 30 years of industry experience, incorporating practical governance approaches, operational lessons, and enterprise implementation strategies rather than purely academic theory.

Learning outcomes

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

- Understand the foundational concepts behind Artificial Intelligence, Machine Learning, and Generative AI
- Explain how modern AI systems learn from data and generate outputs
- Identify the key limitations and risks associated with AI systems
- Recognize common AI failure scenarios including hallucinations, bias, misinformation, and model drift
- Understand why AI governance is becoming a critical business and regulatory priority
- Differentiate between AI ethics, AI risk management, and AI governance
- Identify core components of an AI governance framework
- Understand organizational roles and responsibilities in governing AI systems
- Evaluate AI-related operational, legal, security, and compliance risks
- Understand the importance of transparency, explainability, accountability, and human oversight
- Explore emerging global AI governance frameworks, standards, and regulations
- Develop practical approaches for governing AI adoption within organizations
- Understand governance considerations for Generative AI and Agentic AI systems
- Establish foundational governance controls for responsible AI deployment

---

Prerequisites

- Basic understanding of digital technologies and business operations
- General familiarity with modern software or cloud-based systems
- No programming or data science background required
- No prior AI experience required

Training outline

1. Introduction to Artificial Intelligence

Understanding Artificial Intelligence

- Definition of Artificial Intelligence
- Evolution of AI technologies
- Narrow AI versus General AI

- Predictive AI versus Generative AI
- AI versus traditional software systems
- Common AI terminology and concepts
- Real-world applications of AI across industries

Understanding How AI Works

- Data as the foundation of AI
- Machine Learning fundamentals
- Training data and learning patterns
- Algorithms and statistical modeling
- Neural networks and deep learning concepts
- Supervised learning
- Unsupervised learning
- Reinforcement learning
- Model training and inference
- AI model lifecycle

Understanding Generative AI

- What is Generative AI
- Large Language Models (LLMs)
- Tokens, embeddings, and transformers
- How AI generates text, code, images, and content
- Prompt engineering fundamentals
- Context windows and memory limitations
- Fine-tuning and Retrieval-Augmented Generation (RAG)
- AI agents and autonomous AI systems
- Multi-agent AI environments

Understanding AI Limitations and Risks

- Why AI systems make mistakes
- Hallucinations and fabricated outputs
- Bias in AI systems
- Data quality problems
- Incomplete and unrepresentative datasets

- Explainability limitations
- Black-box decision making
- Model drift and performance degradation
- Overreliance on automation
- Lack of contextual understanding
- AI security vulnerabilities
- Adversarial attacks against AI systems
- Intellectual property and copyright concerns
- Privacy and data leakage risks
- Ethical dilemmas in AI-driven decisions

---

2. Introduction to AI Governance

What is AI Governance

- Definition and purpose of AI governance
- Governance versus compliance versus ethics
- Why AI governance matters
- Business drivers for AI governance
- Risk management and organizational accountability
- Trust and transparency in AI adoption
- Balancing innovation with control

The Business Impact of Poor AI Governance

- Operational risks
- Legal and regulatory exposure
- Reputational damage
- Financial consequences
- AI incidents and governance failures
- Shadow AI and uncontrolled AI adoption
- Risks associated with public AI tools
- Governance challenges in Generative AI deployments

Core Principles of Responsible AI

- Fairness
- Accountability
- Transparency
- Explainability
- Privacy
- Security
- Human oversight
- Reliability and robustness
- Safety and resilience
- Ethical AI principles

---

3. AI Governance Frameworks and Standards

Global AI Governance Landscape

- Emerging global AI regulations
- Regulatory trends across regions
- Industry-specific governance requirements
- Cross-border AI governance challenges

AI Governance Framework Components

- Governance structures and committees
- Policies and standards
- AI risk classification
- AI inventory and system registration
- Model documentation requirements
- AI impact assessments
- Monitoring and auditing mechanisms
- Escalation and incident management processes

Governance Standards and Best Practices

- NIST AI Risk Management Framework

- ISO/IEC AI governance standards
- Responsible AI frameworks
- Enterprise AI governance models
- Risk-based governance approaches
- Governance maturity models

---

4. AI Risk Management

Identifying AI Risks

- Strategic risks
- Operational risks
- Compliance risks
- Ethical risks
- Cybersecurity risks
- Third-party AI risks
- Vendor and supply chain risks

AI Risk Assessment

- Risk identification methodologies
- Risk scoring and prioritization
- High-risk AI use cases
- Impact and likelihood analysis
- Governance controls mapping

Managing AI Risks

- Human-in-the-loop controls
- Human-on-the-loop oversight
- Access control and authorization
- Model testing and validation
- Bias detection and mitigation
- AI monitoring and observability
- Logging and audit trails
- Incident response planning

---

5. Governance for Generative AI and Agentic AI

Governance Challenges in Generative AI

- Unpredictable outputs
- Hallucinations and misinformation
- Data leakage risks
- Intellectual property concerns
- Unsafe content generation
- Prompt injection attacks
- Sensitive data exposure

Governing Enterprise AI Usage

- Acceptable use policies
- Employee AI usage guidelines
- AI procurement governance
- AI vendor evaluation
- AI model approval processes
- AI governance operating models

Agentic AI Governance

- Autonomous decision-making risks
- AI agents and delegated authority
- Identity and access management for AI agents
- Monitoring autonomous AI systems
- Governance for multi-agent ecosystems
- Accountability in autonomous workflows

---

6. Organizational AI Governance

Building an AI Governance Program

- Governance operating models
- Defining organizational roles and responsibilities

- Executive sponsorship
- AI governance committees
- Legal, risk, compliance, and security collaboration
- Data governance integration
- AI governance workflows

AI Policies and Controls

- AI acceptable use policies
- AI development standards
- Data handling and privacy controls
- Model approval and deployment controls
- Monitoring and review processes
- Governance documentation practices

AI Governance Across the AI Lifecycle

- Governance during data collection
- Governance during model development
- Governance during deployment
- Governance during ongoing operations
- Continuous monitoring and auditing
- Model retirement and decommissioning

---

7. AI Governance in Practice

Enterprise AI Governance Scenarios

- Governance for HR systems
- Governance for financial AI systems
- Governance for healthcare AI systems
- Governance for customer-facing AI
- Governance for cybersecurity AI tools

Practical Governance Challenges

- Balancing innovation with compliance

- Governance in fast-moving AI environments
- Managing shadow AI
- Governing third-party AI providers
- Governance for open-source AI models
- Scaling AI governance across the enterprise

Future of AI Governance

- Emerging trends in AI governance
- Governance for autonomous AI ecosystems
- Continuous AI assurance
- AI observability and monitoring
- The evolving role of regulators
- Governance as a business enabler

---

8. Course Wrap-Up

Key Takeaways

- Understanding AI before governing AI
- Governance as a strategic capability
- The importance of accountability and oversight
- Practical steps toward responsible AI adoption

9. Final Discussion and Q&A

- Organizational readiness considerations
- Governance priorities for different industries
- Building a roadmap for AI governance maturity

Disclaimer

This course outline is intended solely as a general training and discussion framework. The sequence of topics, scope of coverage, and specific subject matter may be modified, expanded, condensed, substituted, or otherwise adjusted by the instructor at any time based on participant background, organizational requirements, emerging industry developments, regulatory changes, time considerations, or

instructional judgment. The trainer reserves the right to amend the course content and delivery approach without prior notice in order to ensure relevance, practicality, and alignment with current industry practices.

© 2026 Faysal Aziz · Kuala Lumpur