BOARD INFINITY

COURSE 3 OF 3 · MANAGING AI SYSTEMS: DEVELOPMENT, DEPLOYMENT, AND GOVERNANCE

AI Risk and Compliance: Audit and Governance Foundations

Full course syllabus - every module, video lesson, reading and assignment, from red teaming and the EU AI Act to explainability, guardrails and PII protection.

Board Infinity
Instructor · Offered on Coursera

4	34	16	2 weeks
Modules	Video lessons	Assignments	At 10 hours a week

INSIDE THE COURSE

01	Adversarial Robustness & Red Teaming AI Systems	4 hours to complete
02	Regulatory Compliance: EU AI Act, NIST RMF & Copyright	4 hours to complete
03	Explainability (XAI) & System Transparency	4 hours to complete
04	Guardrails, PII Protection & Shadow AI Mitigation	4 hours to complete

Enroll on Coursera

https://www.coursera.org/learn/ai-risk-and-compliance-audit-and-governance-foundations

01 ABOUT THE COURSE

Govern, secure and audit enterprise AI.

This course takes a practical, end-to-end approach to governing, securing and auditing AI systems. You start with adversarial threats - jailbreaks, prompt injection, data leakage and manipulation - and practise structured red teaming, then work through the EU AI Act, the NIST AI RMF and the copyright risks in training data and outputs.

Next come explainability with SHAP, LIME and attention maps, and safety engineering with guardrails, Presidio PII protection and Shadow AI controls - 34 videos, 12 readings and 16 assignments in all. Basic AI/ML and Python familiarity helps.

WHO IT'S FOR

AI governance officers
Building governance and audit programmes for AI.

Risk & compliance teams
Applying the EU AI Act and NIST AI RMF.

AI & ML engineers
Red teaming and hardening the systems they build.

Policy makers
Understanding how AI risks are tested and controlled.

WHAT YOU'LL LEARN

01 **Red team AI systems** Module 1
Run adversarial scans with Giskard to find and prioritize AI vulnerabilities

02 **Apply AI regulation** Module 2
Classify systems under the EU AI Act and apply the NIST AI RMF across the lifecycle

03 **Explain model decisions** Module 3
Generate SHAP/LIME explanations and audit-ready transparency documentation

04 **Implement safety controls** Module 4
Guardrails, PII scrubbing with Presidio, and controls that mitigate Shadow AI

02 SYLLABUS

Four modules, from attack to audit.

2 weeks to complete at 10 hours a week. Flexible schedule - learn at your own pace.

01 Adversarial Robustness & Red Teaming AI Systems

Threat models, prompt injection, jailbreaks and Giskard

4 hours
8 videos

BY THE END OF THIS MODULE YOU CAN

- ✓ Identify LLM threats: jailbreaks, prompt injection, leakage and manipulation
- ✓ Distinguish benign failures from genuinely adversarial behavior
- ✓ Stress-test LLM defenses across architectures
- ✓ Run automated red teaming with Giskard and prioritize remediation

CONTENTS 8 videos (65 min) · 3 readings (90 min) · 4 assignments (105 min) · 1 plugin

TOPICS AI Threat Models | Prompt Injection & Jailbreaks | Automated Red Teaming with Giskard

You'll leave with: a red-team report on an AI system - vulnerabilities found, ranked and documented.

02 Regulatory Compliance: EU AI Act, NIST RMF & Copyright

Risk tiers, NIST functions and legal exposure

4 hours
7 videos

BY THE END OF THIS MODULE YOU CAN

- ✓ Classify AI systems under the EU AI Act risk tiers
- ✓ Explain high-risk obligations, conformity assessments and post-market monitoring
- ✓ Map the AI lifecycle to NIST AI RMF: Govern, Map, Measure, Manage
- ✓ Assess training-data licensing, output ownership and enterprise liability

CONTENTS 7 videos (58 min) · 3 readings (90 min) · 4 assignments (105 min)

TOPICS EU AI Act Risk Tiers | NIST AI RMF | Copyright & Legal Exposure

You'll leave with: concise compliance documentation that maps a system to the EU AI Act and NIST AI RMF.

03 Explainability (XAI) & System Transparency

SHAP, LIME, attention maps and transparency reports

4 hours
10 videos

BY THE END OF THIS MODULE YOU CAN

- ✓ Explain why explainability is difficult for LLMs
- ✓ Compare SHAP, LIME and attention maps, and choose the right one
- ✓ Present faithful, human-readable explanations to executives and regulators
- ✓ Design transparency workflows and document system and decision flows

CONTENTS 10 videos (66 min) · 3 readings (90 min) · 4 assignments (105 min)

TOPICS Interpreting Black Box Models · XAI for Stakeholders · Transparency for Governance

You'll leave with: SHAP or LIME applied to a text classifier, and an audit-ready transparency report.

04 Guardrails, PII Protection & Shadow AI Mitigation

Guardrails AI, NeMo, Presidio and governance controls

4 hours
9 videos

BY THE END OF THIS MODULE YOU CAN

- ✓ Build and test rule-based and semantic guardrails with Guardrails AI and NeMo
- ✓ Detect, mask and anonymize PII with Presidio, including in RAG pipelines
- ✓ Recognize Shadow AI and apply monitoring, policy and enforcement
- ✓ Design organization-wide governance controls

CONTENTS 9 videos (60 min) · 3 readings (90 min) · 4 assignments (105 min)

TOPICS Guardrails for Safety · PII Scrubbing with Presidio · Shadow AI Controls

You'll leave with: a capstone lab - a chatbot red teamed, guarded and scrubbed, with full evidence.

03 TOOLS & SKILLS

The tools you audit with.

Giskard

Red teaming

Automated vulnerability scans for AI and LLM systems.

In this course: Module 1, and again in the Module 4 capstone lab.

SHAP & LIME

Explainability

Explaining model decisions, alongside attention visualization.

In this course: Module 3, applied to a text classifier.

Microsoft Presidio

PII protection

Detecting, masking and anonymizing sensitive data.

In this course: Module 4, including RAG pipeline integration.

Guardrails AI & NeMo Guardrails

Safety

Rule-based and semantic guardrails for policy enforcement.

In this course: Module 4, tested against attacks.

SKILLS YOU LEAVE WITH

Risk & security

- AI security
- Vulnerability assessments
- Risk management
- Risk mitigation

Regulation

- Regulatory compliance
- Compliance auditing
- Legal risk
- Regulatory requirements

Governance & safety

- Governance
- Safety assurance
- Personally identifiable information
- Compliance reporting

04 CERTIFICATE & CAREER

Finish with a shareable certificate.

Add it to LinkedIn
Add this credential to your LinkedIn profile, resume, or CV.

Share it anywhere
Share it on social media and in your performance review.

Earned, not attended
Awarded on completion of the assignments and labs in all four modules.

Course 3 of 3
Completes the Managing AI Systems: Development, Deployment, and Governance Specialization.

ROLES THAT USE THESE SKILLS

AI Governance Officer · AI Risk Analyst · Compliance Officer · AI Red Teamer · AI Auditor



Board Infinity

Instructor · Offered on Coursera

Intermediate level · Taught in English · Flexible schedule

Ready to make AI safe to ship?

Enrollment, lessons, grading and the certificate all happen on Coursera.

Enroll on Coursera

Course page: coursera.org/learn/ai-risk-and-compliance-audit-and-governance-foundations

Want a coach alongside the course?

Board Infinity 1:1 coaching is a separate, optional Board Infinity offering - it is not part of the Coursera course. Details: https://www.boardinfinity.com/career-coaching/

Board Infinity is an independent education provider. Coursera and the Coursera logo are trademarks of Coursera, Inc. This independent educational resource is not affiliated with, endorsed by, sponsored by, or officially associated with any company, organization, regulator or certification body unless explicitly stated. All trademarks are the property of their respective owners.