MONARCH MANAGEMENT INSTITUTE

Open Course Network · Micro-Course 15

COURSE PROSPECTUS

Manage AI Risk Before It Manages You

A two-week, applied micro-course on the NIST AI Risk Management Framework — a practical method for governing, mapping, measuring and managing the risks of an organization's AI, worked on one AI system you carry all the way through, and taught for an executive audience.

2 wks	12	70%
Two weekly modules	Sections	To pass · unlimited attempts

At a glance

Duration	2 weeks — two weekly modules of six sections
Study time	≈ 14 directed hours (5–8 per week)
Level	Introductory; no prior AI or technical knowledge required · no prerequisites
Assessment	2 auto-graded scenario assessments · 70% to pass · unlimited attempts
Award	Monarch-endorsed Certificate of Knowledge
Format	Fully online, self-paced

What you will gain

By the end of this course you will be able to apply the full Govern-Map-Measure-Manage cycle to prioritize and address the risks of an organization's AI use, producing a defended risk response.

And what it counts toward

You earn the Monarch-endorsed Certificate of Knowledge and 2 MMI-Credits. MMI-Credits carry toward MMI program progression (PD-001).

What this course is

Manage AI Risk Before It Manages You is a fully online, self-paced micro-course on managing the risk of artificial intelligence. Most organizations did not decide to adopt AI — it arrived, in a vendor feature, a supplier's platform, a team pasting work into a public chatbot — and by the time leadership asks who owns the risk, AI is already shaping decisions that affect customers, staff and the balance sheet. This course teaches the discipline that brings that risk under control, through the most widely used public framework for it: the NIST AI Risk Management Framework (AI RMF 1.0), which organizes the work into four functions — Govern, Map, Measure and Manage (NIST, 2023).

You choose one real AI use in the opening section — your own organization's, a well-documented public example, or the course's running case — and carry it through all twelve sections, finishing with a defended, prioritized risk response for it: a document you can take to a board or a regulator.

The four functions you will learn to apply

- **Govern** — build the culture, policies, accountability, risk tolerance and inventory that make AI risk a managed responsibility, and oversee the third-party supply chain
- **Map** — establish the context of a specific AI system, categorize it, and identify its risks, benefits and impacts, and who bears them
- **Measure** — assess the system's trustworthiness with appropriate methods, metrics and independent TEVV, and track how its risk changes over time
- **Manage** — prioritize the measured risks against tolerance, respond proportionately, handle third-party and incident risk, and monitor so a treated risk stays treated

Who this course is for

- Executives and managers who find AI already at work in their organization and must be able to say it is under control
- Risk, compliance, legal and audit professionals extending their remit to the distinctive risks of AI
- Product and operations leaders who procure or deploy AI and answer for its outcomes
- Reflective managers and BBA, MBA or DBA candidates seeking a rigorous, applied grip on AI governance — no prior AI or technical knowledge assumed

What you will accomplish

- Explain why AI risk is socio-technical and distinct from ordinary technology risk, weighing both likelihood and magnitude of harm
- Apply the Govern function — establish policies, named accountability, a risk tolerance, an AI inventory and third-party oversight
- Apply the Map function — establish a system's context, categorize it, and identify its risks, benefits and impacts by who is affected
- Apply the Measure function — choose methods and metrics, assess the seven trustworthiness characteristics through TEVV, and track risk over time
- Apply the Manage function — prioritize against tolerance, respond, resource the controls, manage suppliers, and plan incident response
- Produce and defend a prioritized risk response for a real AI use, and name honestly what remains unmeasured

Your fortnight, section by section

Module 1 — Govern and Map the Risk

Section	Title
1	Why AI Risk Is Different
2	A Culture of AI Risk Management (Govern)

Section	Title
3	Risk Tolerance, Inventory and Third Parties (Govern)
4	Establishing Context and Categorizing the System (Map)
5	Identifying Risks, Benefits and Impacts (Map)
6	The Govern-and-Map Baseline (synthesis)

Module 2 — Measure and Manage the Risk

Section	Title
1	Choosing Methods and Metrics (Measure)
2	Assessing Trustworthiness, and TEVV (Measure)
3	Tracking Risk Over Time (Measure)
4	Prioritizing and Responding to Risk (Manage)
5	Resources, Third Parties and Incident Response (Manage)
6	The Full Govern–Map–Measure–Manage Cycle (synthesis)

Each section follows the same four steps — Explain, Show, Do, Feedback — worked through a running case (Rowan Cole's work at Wexler Insurance Group) and then on the AI use you choose.

The running case

Alongside your own AI use you follow Wexler Insurance Group, a fictional mid-sized insurer that has let AI spread ahead of its controls. Two systems recur through every section: an automated claims-triage model, built partly on a third-party vendor platform, that decides which claims are fast-tracked and which go to a human; and a customer-service chatbot built on a third-party large language model. You follow Wexler's newly appointed head of AI risk, Rowan Cole, against the foil of a product director whose refrain is "the model works — ship it," from an ungoverned start to a defended, board-ready response. Real incidents — an AI hiring tool, a contested risk-scoring system, an airline held liable for its chatbot — appear only as sourced illustrations of the ideas.

Assessment and certificate

Each module concludes with one auto-graded scenario assessment: a fifteen-item scenario drawn from a larger bank, presenting unfamiliar organizations and AI uses so it tests judgment rather than recall, with two formative checkpoints inside every section along the way. A score of **70%** passes, with unlimited untimed attempts, and a full explanation is shown for every item. On passing both module assessments you earn the **Monarch-endorsed Certificate of Knowledge**.

Taught with its limits

The NIST framework is a powerful instrument, but a bounded one, and the course teaches its limits alongside the method — that it is a voluntary way of organizing judgment rather than a compliance checklist; that third-party risk is often hard to measure because providers disclose little; that fairness cannot be reduced to a single number, since incompatible definitions cannot all hold at once; and that a first risk map will miss risks and must be revisited, not filed. You leave able to use the framework well and to say plainly what it does not settle.

> Choose one AI use, and learn to manage its risk before it manages you.