AI FOR BUSINESS AND NON-TECHNICAL AUDIENCES

Responsible AI and AI Ethics

Level: Foundation • 2 days (expandable to 3) • Virtual, In-person

Overview

AI systems now help decide who gets a loan, whose resume gets read, and what customers are told, which means their mistakes are no longer just technical bugs: they are unfair decisions, privacy breaches, and public failures with names attached. The hard part of responsible AI is not agreeing that fairness and transparency matter; everyone agrees. It is knowing what those principles mean in a specific use case, spotting the risk before deployment, and having the practical habits to act on concerns rather than admire them.

This is a hands-on, foundation course. It assumes no technical background and deliberately goes deep on a small set of principles and the practice of applying them, rather than surveying every ethics framework ever published. The gradient starts with why responsible AI matters in concrete business terms, builds the core principles one at a time (fairness, transparency, accountability, privacy), then turns them into practice: screening use cases, working through hard calls, and building a culture where concerns get raised. Every module includes a hands-on lab and builds on the one before.

Who Should Attend

- Business professionals and managers whose teams use or deploy AI
- Product, data, and technology staff who want a practical ethics grounding
- HR, legal, risk, and communications professionals increasingly pulled into AI questions

  Leaders responsible for building the formal governance machinery should continue to *AI Governance for Leaders*.

Prerequisites

- None. No technical background is assumed
- Basic familiarity with AI tools, at the level of *Generative AI for Every Employee*, is helpful

What You Will Learn

- Explain why responsible AI is a business necessity, using real failure cases
- Identify where bias enters AI systems and describe practical mitigations in plain language
- Apply transparency and accountability principles: explanation, disclosure, and clear ownership
- Assess privacy and generative AI risks, including hallucination and intellectual property
- Apply an ethical decision framework to hard, ambiguous AI use cases
- Build team-level practices that make responsible use routine rather than aspirational

Course Outline

Day one: the principles, grounded in real cases

- Why Responsible AI Matters
  - Real failures: biased hiring tools, wrongful denials, and public AI incidents

  - The business case: trust, regulation, and the cost of getting it wrong
  - Lab: analyze a public AI failure case and identify what went wrong and when it was catchable
- Fairness and Bias
  - Where bias comes from: data, design, and deployment
  - What mitigation looks like in practice, in plain language
  - Lab: spot the bias risks in a set of realistic AI use cases and propose responses
- Transparency and Accountability
  - Explainability and disclosure: what people deserve to know about AI decisions
  - Accountability: a human owns every AI-assisted outcome
  - Lab: draft disclosure language and an accountability map for an AI-assisted process

Day two: from principles to practice

- Privacy and the Risks Specific to Generative AI
  - Personal data, confidentiality, and what AI tools do with what you give them
  - Hallucination, intellectual property, and synthetic content risks
  - Lab: risk-screen a proposed generative AI use case and recommend go, fix, or stop
- Principles into Decisions
  - A practical ethical decision framework for AI use cases
  - Handling the genuinely hard calls, where principles pull in different directions
  - Lab: work a difficult, ambiguous scenario through the framework in teams and defend the call
- Building a Responsible AI Culture
  - Everyday behaviors: verification, disclosure, and raising concerns safely
  - Where culture meets structure, with a pointer to *AI Governance for Leaders*
  - Lab: draft a team-level responsible AI checklist you can put into use immediately

Extended Version

The three-day version keeps the same gradient and adds depth and practice:

- The regulatory landscape in more depth, including the EU AI Act and emerging rules
- Extended case-study work across industries, including attendees' own sectors
- Deeper practice with the decision framework on harder, multi-stakeholder scenarios
- A capstone in which teams build a complete responsible AI playbook and present it for critique