AI Fundamentals

Concepts, Risks, and Guardrails

A Course Offered by The RAI Fellow

Instructor: Iman Saleh, PhD

Course Overview

Artificial Intelligence is rapidly reshaping standard technical workflows. This course demystifies AI for anyone without requiring deep coding or heavy statistical knowledge.

Participants will trace the evolution of AI from classical rules-based automation to modern deep learning, generative and agentic AI systems.

Beyond foundational concepts, the curriculum explores critical failure modes, including model hallucinations, privacy leaks, systemic bias, and operational safety risks. Participants will gain a practical framework for deploying engineering guardrails, evaluating vendor tools, and managing risk when integrating AI into technical projects. They will also see AI in action by walking through a demo using Agentic AI to build a website.

Target Audience: Technical and non-Technical audience

Prerequisites: None. No prior programming or data science background required.

Course Objectives

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

- **Understand AI Core Terminology:** Explain the key distinctions between Artificial Intelligence, Deep Learning, Generative AI, and Agentic AI and articulate how AI processing differs from conventional deterministic software.
- **Classify AI Model Types & Applications:** Identify different paradigms of AI (supervised vs. unsupervised learning, discriminative vs. generative) and map appropriate model types to real-world engineering problem statements.
- **Trace the Technological Evolution of AI:** Summarize the historical trajectory of AI from early recommender systems to modern foundation models and recognize current technological trends shaping engineering industries.

© 2026 The RAI Fellow

- **Evaluate AI Operational Risks:** Recognize core AI risks across engineering contexts, including hallucinations, data privacy leaks, IP infringement, algorithmic bias, and physical safety failure modes in automated systems.
- **Apply Practical Guardrails & Mitigation Strategies:** Implement safeguards, such as human-in-the-loop oversight, input/output validation rules, vendor evaluation frameworks, and data governance policies to safely integrate AI tools into engineering workflows.
- **See Agentic AI in Action:** Walk through a demo that uses Agentic AI to build a personal website, applying the best practices on the way.

Instructor

Dr. Saleh is a Computer Scientist, Privacy Fellow, and Responsible AI leader with over 24 years of experience spanning academia and tech giants like Google, Intel, and Airbnb. Holding a PhD in Software Engineering, she is certified as a CIPP/US, CIPT, an AI Governance Professional (AIGP) and a Fellow of Information Privacy (FIP) by the International Association of Privacy Professionals (IAPP). Her technical background includes over 30 publications on big data, service-oriented architectures, and privacy-preserving solutions, along with an authorial book on formal data specification.

Throughout her career, Dr. Saleh has spearheaded cutting-edge research teams focused on Artificial Intelligence, Privacy, and Fairness. In industry, she specializes in building scalable AI governance programs from the ground up, embedding regulatory and ethical guardrails directly into technical infrastructures, and leading complex risk assessments.

© 2026 The RAI Fellow