Course Title: Digital Ethics & Responsible AI Governance

- **Duration:** 5 Days (approximately 35-40 instructional hours)
- **Course Goal:** To provide business leaders, technology strategists, and governance professionals with the essential frameworks, ethical principles, and practical tools to develop and implement a robust governance program for Artificial Intelligence (AI) and digital technologies. This course will enable participants to proactively identify and mitigate ethical risks, build trustworthy AI systems, and align their organization's innovation with core human values.
- **Core Focus:** Mastering the principles and practices of building a formal governance structure for AI that addresses fairness, accountability, transparency, and societal impact.

Learning Objectives:

Upon successful completion of this workshop, participants will be able to:

- **Define** digital ethics and articulate the business case for a proactive, responsible AI strategy.
- **Understand** the unique and amplified ethical risks posed by AI and machine learning (e.g., bias, opacity, autonomous decision-making).
- **Apply** established ethical frameworks to complex technology dilemmas.
- **Develop** a comprehensive Responsible AI Governance Framework, including principles, policies, and an oversight structure.
- **Establish** roles and responsibilities for AI governance, including an AI Ethics Committee or Review Board.
- **Conduct** an AI Risk & Ethical Impact Assessment for new and existing AI systems.
- **Implement** strategies to detect and mitigate bias in datasets and algorithmic models.
- **Champion** the principles of Explainability (XAI), Interpretability, and Transparency in AI systems.
- **Integrate** "Ethics by Design" principles into the AI development lifecycle.
- **Develop** a communication strategy to build stakeholder trust in the organization's use of AI.
- **Create** a roadmap for implementing a responsible AI governance program.

Target Audience:

- Chief Data Officers, Chief Analytics Officers, Chief Technology Officers (CTOs), and CIOs.
- Directors and VPs of AI, Data Science, and Analytics.
- Ethics, Compliance, Risk Management, and Legal Professionals.
- Product Managers and Innovation Leaders developing AI-powered products.
- Board members and senior executives responsible for corporate strategy and governance.
- Policymakers and regulators working on technology issues.

Prerequisites:

- A solid understanding of business strategy and risk management.
- Familiarity with data and analytics concepts is essential.
- No deep coding or data science expertise is required, but an understanding of what AI/ML does is important.

Teaching Methodology:

This is a highly interactive, strategic, and "gray area" focused workshop. The methodology includes:

- **Interactive Lectures & Frameworks:** Introducing concepts from ethics, law, computer science, and governance (e.g., NIST AI RMF, OECD AI Principles).
- **Intensive Case Study Deconstruction:** Deeply analyzing real-world AI ethical failures and successes from various industries.
- **Ethical "War Gaming" & Simulations:** The core of the course. Participants will act as an AI Ethics Committee, deliberating on complex, high-stakes scenarios.
- **Hands-on Workshops:** Participants will work in teams to draft AI principles, conduct risk assessments, and design a governance framework.
- **Group Debates & Discussions:** Exploring the challenging trade-offs between innovation, efficiency, and ethical responsibility.
- **Action Planning:** Developing practical roadmaps for implementation.

*(For **Online Delivery**, this is highly effective using breakout rooms for committee simulations, anonymous polls and Q&A for sensitive ethical questions, virtual whiteboards for mapping risks, and collaborative documents for drafting policies.)*

Materials Provided:

ADEPT TRAINING CENTER

- A comprehensive digital **"Responsible AI Governance Playbook"** with frameworks, checklists, and best practices.
- Templates for an AI Ethics Charter, an AI Risk & Impact Assessment, and an Implementation Roadmap.
- A detailed case study of a fictional company deploying multiple AI systems.
- A curated set of seminal articles, white papers, and resources on AI ethics and governance.
- Certificate of Completion.

Detailed Day-by-Day Curriculum

Day 1: The "Why" – The Ethical Imperative of AI

- **Module 1: The New Frontier: Defining Digital & AI Ethics (Approx. 2 hours)**
  - Welcome, introductions, and establishing a space for open and critical dialogue.
  - Moving beyond "data privacy": The broader scope of digital ethics.
  - What makes AI ethics different and more urgent? (Scale, speed, autonomy, opacity).
  - The business case for responsible AI: Trust as a competitive advantage, risk mitigation, and talent attraction.
- **Module 2: The AI Risk Universe: What Can Go Wrong? (Approx. 3 hours)**
  - A structured overview of the key ethical risks of AI:
    - **Bias & Fairness:** How AI can perpetuate and amplify societal biases.
    - **Transparency & Explainability:** The "black box" problem.
    - **Accountability & Responsibility:** Who is to blame when an AI system fails?
    - **Privacy:** New frontiers of data collection and inference.
    - **Safety & Security:** Adversarial attacks and unintended consequences.
    - **Human Agency & Oversight:** The risk of over-reliance on autonomous systems.
  - **Case Study Clinic:** Analyzing several real-world AI failures and mapping them to these risk categories.
- **Module 3: Foundational Ethical Frameworks for Technology (Approx. 2.5 hours)**
  - Applying classic ethical theories to AI dilemmas:
    - **Consequentialism/Utilitarianism:** Optimizing for the best outcome.
    - **Deontology/Duty-Based Ethics:** Following universal rules.
    - **Virtue Ethics:** The character of the developer/organization.
  - Introduction to modern principles-based frameworks (e.g., OECD AI Principles).
  - **End of Day 1 Recap & Q&A.**

Day 2: Building the Governance Structure

- **Module 4: Designing a Responsible AI Governance Framework (Approx. 2.5 hours)**
  - The goal: To create a system of authority and control for AI development and deployment.
  - Key components: Principles, Policies, Processes, People, and Technology.
  - The role of the Board of Directors vs. Executive Management.
- **Module 5: The AI Ethics Committee / Review Board (Approx. 3 hours)**
  - The purpose and mandate of an AI Ethics Committee.
  - Who should be on it? (A multi-disciplinary team: legal, compliance, tech, business, ethics, HR).
  - Developing a charter and operating model for the committee.
  - Decision-making processes and escalation paths.
  - **Workshop:** Participants work in groups to draft a charter for an AI Ethics Committee for the case study company.
- **Module 6: Defining Your Organization's AI Principles (Approx. 2 hours)**
  - Why a set of clear, public principles is the foundation of trustworthy AI.
  - Analyzing examples from leading tech companies.
  - The process of developing and embedding principles that are authentic to your organization's values.
  - **Activity:** Brainstorming and drafting a set of core AI Principles.
  - **End of Day 2 Recap & Q&A.**

Day 3: The Practice – Risk Assessment & Mitigation

- **Module 7: The AI Risk & Ethical Impact Assessment (Approx. 3.5 hours)**
  - A step-by-step process for proactively assessing new AI projects.
  - Key areas to assess:
    - Data: Source, quality, representativeness, privacy.
    - Model: Fairness, explainability, robustness.
    - Application: Intended use, potential for misuse, impact on users and society.
  - **Intensive Workshop:** Participants use a provided template to conduct a high-level Ethical Impact Assessment on a new AI initiative from the case study.
- **Module 8: The Problem of Bias: Detection & Mitigation (Approx. 3.5 hours)**

ADEPT TRAINING CENTER

- A deep dive into different types of bias (e.g., historical, measurement, representation).
- Technical tools for bias detection in datasets and models (overview).
- Process-based strategies for bias mitigation:
  - Diverse and inclusive development teams.
  - Careful data sourcing and pre-processing.
  - Regular model auditing and testing.
- Defining "fairness" for a specific use case – a complex trade-off.
- **Case Study Deep Dive:** Analyzing the data and outcomes of a biased hiring or lending algorithm.
- **End of Day 3 Recap & Q&A.**

Day 4: Ensuring Transparency, Accountability & Control

- **Module 9: The Demand for Explainability (XAI) & Transparency (Approx. 3 hours)**
  - The difference between transparency (of process) and explainability (of a model's decision).
  - When is a "black box" acceptable and when is it not? (Risk-based approach).
  - Techniques for model interpretability and explanation (SHAP, LIME - conceptual overview).
  - The importance of clear documentation for models and data (Model Cards, Datasheets for Datasets).
- **Module 10: Human-in-the-Loop: Accountability & Oversight (Approx. 2.5 hours)**
  - Designing effective human oversight for autonomous systems.
  - Different levels of human involvement (human-in-the-loop, human-on-the-loop, human-in-command).
  - Establishing clear lines of accountability when an AI system causes harm.
  - The importance of contestability and redress mechanisms for affected individuals.
- **Module 11: AI Governance Simulation: The Ethics Committee Meeting (Approx. 2 hours)**
  - **Intensive Simulation:** Participants take on roles in the case study's AI Ethics Committee.
  - They are presented with a proposal for a new, ethically challenging AI project (e.g., employee monitoring, predictive policing partnership, a sensitive health algorithm).
  - They must use their frameworks to debate the risks and benefits and come to a "Go / No-Go / Go with conditions" decision.
  - **End of Day 4 Recap & Q&A.**

Day 5: Implementation, Culture & Your Governance Roadmap

- **Module 12: "Ethics by Design": Integrating Ethics into the AI Lifecycle (Approx. 2.5 hours)**
  - Moving from a "gate-keeping" committee to an enabling function.
  - Embedding ethical checkpoints and reviews throughout the AI development process.
  - The role of training and empowering data scientists and engineers to be "ethical first-responders."
- **Module 13: Communicating Trust: Transparency with Stakeholders (Approx. 2 hours)**
  - Developing a communication strategy for your Responsible AI program.
  - How to talk about your use of AI with customers, regulators, and the public.
  - Creating transparency reports.
- **Module 14: Your Responsible AI Governance Roadmap (Capstone) (Approx. 2.5 hours)**
  - **Capstone Workshop:** Participants synthesize their learnings to create a phased, 12-month implementation roadmap for a Responsible AI Governance program at the case study company.
  - The roadmap will include key priorities, milestones, and success metrics for building out the program.
  - Groups present their roadmaps for peer and instructor feedback.
- **Module 15: Course Review & Commitment to Responsible Innovation (Approx. 1 hour)**
  - Review of the key components of a robust Responsible AI Governance program.
  - Open forum for final questions and sharing key commitments.
  - Course evaluation and closing.

ADEPT TRAINING CENTER