## Course Title: Ethical Leadership in Technology & AI: Navigating the New Frontier of Responsibility

**Duration:** 5 Days (approximately 35-40 instructional hours)

**Course Goal:** To provide technology leaders with the critical frameworks, foresight, and ethical reasoning skills to lead their teams and organizations in the responsible design, development, and deployment of technology and AI systems.

**Core Focus:** This is a leadership and foresight workshop. It is not about *how* to code an ethical algorithm, but about *how to lead* an organization that builds ethical technology. It tackles the "gray area" dilemmas at the intersection of innovation, profit, and societal impact, with a strong emphasis on building proactive governance rather than reactive damage control.

**Learning Objectives:**

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

- Articulate the unique ethical challenges posed by modern technology, particularly AI and data-driven systems.
- Apply multiple ethical frameworks (e.g., Consequentialist, Deontological, Justice) to complex tech dilemmas.
- Establish a robust ethical risk assessment framework for new products and features.
- Lead discussions on and develop strategies to mitigate algorithmic bias and promote fairness.
- Understand and implement principles of transparency, explainability (XAI), and accountability in AI systems.
- Foster a culture of "responsible innovation" and psychological safety where engineers and product managers feel empowered to raise ethical concerns.
- Develop a governance structure for AI ethics, including the role of ethics committees and review boards.
- Navigate the trade-offs between user privacy, personalization, and business objectives.
- Lead their organization through an ethical crisis caused by a technology failure.

**Target Audience:**

- **Technology Executives** (CTO, CPO, VP of Engineering/Product).
- **Startup Founders and CEOs.**
- **Directors and Senior Managers** in product, engineering, and data science.
- **Investors and Board Members** of technology companies.
- **Policy and Trust & Safety professionals** within the tech industry.

**Prerequisites:**

- **Significant leadership experience within the technology sector is essential.**
- A strong understanding of the product development lifecycle.
- A deep interest in the societal impact of technology.

**Teaching Methodology:**

This is a Socratic, case-based, and future-focused seminar.

- **Deep Case Study Method:** The curriculum is driven by in-depth, candid analysis of real-world tech ethics controversies (e.g., Cambridge Analytica, biased AI recruiting tools, social media's impact on mental health).

ADEPT TRAINING CENTER

- **Ethical "Red Teaming" Workshops:** A core feature where teams will be given a new product idea and tasked with identifying all the potential ways it could be misused or cause unintended harm.
- **Socratic Dialogue:** The facilitator will guide intense, challenging discussions with a focus on probing assumptions and exploring second- and third-order consequences.
- **"Future-Back" Thinking:** Exercises will challenge participants to imagine future ethical dilemmas and build systems to address them today.
- **Peer Consulting:** A confidential environment for leaders to discuss their own specific ethical challenges and get feedback from a trusted group of peers.

**Materials Provided:**

- **The Responsible Tech Leader's Playbook:** A briefing book with:
  - Frameworks for ethical risk assessment and decision-making.
  - An AI Ethics Principles checklist.
  - Templates for conducting an ethical "pre-mortem."
- A full set of all case study materials, including articles, post-mortems, and white papers.
- A curated reading list of seminal works on technology ethics.
- Certificate of Completion.

---

**Detailed Day-by-Day Curriculum**

**Day 1: The New Responsibility of Tech Leadership**

- **Module 1: Beyond "Move Fast and Break Things" (Approx. 2 hours)**
  - Welcome, introductions, and establishing a "brave space" for candid discussion.
  - The end of tech neutrality: Why "we're just a platform" is no longer a viable stance.
  - Defining the scope: From dark patterns to algorithmic bias to societal-level harms.
- **Module 2: A Leader's Toolkit for Ethical Reasoning (Approx. 3 hours)**
  - A rapid tour of key ethical frameworks applied to tech:
    - **Consequentialism:** "What is the net impact of this product?"
    - **Deontology:** "Are there certain rules we should never break, regardless of the outcome?"
    - **Virtue Ethics:** "What would a leader of integrity do?"
  - **Case Study Workshop:** Applying these frameworks to a real-world tech decision.
- **Module 3: Identifying Ethical Risks Before You Ship (Approx. 2.5 hours)**
  - The concept of the "Ethical Pre-Mortem."
  - A structured framework for identifying potential harms across different stakeholder groups.
  - **Workshop:** Conduct a rapid ethical risk assessment on a common tech feature (e.g., a "like" button, a friend recommendation algorithm).

**Day 2: The Challenge of Algorithmic Bias and Fairness**

- **Module 4: The Anatomy of Algorithmic Bias (Approx. 2.5 hours)**
  - Where does bias come from? (Biased data, flawed model assumptions, human bias in labeling).
  - The different types of bias and their real-world consequences (e.g., in hiring, lending, criminal justice).
  - **Case Study Deep Dive:** Analyzing the failures of a well-known biased AI system.

ADEPT TRAINING CENTER

- **Module 5: Defining and Measuring "Fairness" (Approx. 2.5 hours)**
    - Why "fairness" is a complex, socio-technical problem with no easy answers.
    - An overview of different mathematical definitions of fairness and their inherent trade-offs.
    - **Debate Workshop:** "Should an algorithm be optimized for equal opportunity or equal outcomes?"
- **Module 6: Strategies for Mitigation (Approx. 2.5 hours)**
    - Practical steps for mitigating bias: Data diversity, algorithmic auditing, human-in-the-loop systems.
    - The importance of diverse teams in building fairer products.
    - **Workshop:** Brainstorming mitigation strategies for a biased algorithm scenario.

**Day 3: Transparency, Privacy, and Trust**

- **Module 7: The Black Box Problem: Explainability and Interpretability (Approx. 3 hours)**
    - What is Explainable AI (XAI)?
    - The ethical imperative and business case for being able to explain *why* an AI made a certain decision.
    - Techniques for creating more interpretable models.
- **Module 8: The Privacy Paradox (Approx. 4 hours)**
    - Navigating the tension between data-driven personalization and user privacy.
    - The principles of Privacy by Design.
    - Understanding the spirit, not just the letter, of regulations like GDPR.
    - **Ethical Red Teaming Workshop:** Teams are given a new data-intensive product idea and must identify all the potential privacy risks and misuse scenarios.

**Day 4: Building a Culture of Responsible Innovation**

- **Module 9: AI Ethics Governance in Practice (Approx. 3 hours)**
    - Moving from principles to process.
    - The role and structure of AI Ethics Boards, Review Committees, and responsible innovation teams.
    - Creating clear channels for engineers and product managers to raise ethical concerns without fear of reprisal.
- **Module 10: The Role of the Leader (Approx. 4 hours)**
    - How leaders signal what truly matters.
    - Integrating ethical considerations into your product development lifecycle (e.g., adding an ethics gate).
    - Incentivizing responsible innovation: Aligning performance reviews and promotions with ethical behavior.
    - **Leadership Clinic:** A peer-consulting session where participants can discuss their own real-world challenges in fostering an ethical culture.

**Day 5: The Leader in the Arena**

- **Module 11: Leading Through an Ethical Crisis (Approx. 2 hours)**
    - What to do when your technology causes public harm.
    - The principles of a credible response: Acknowledgment, accountability, and a commitment to action.
    - The difference between a legal response and an ethical one.
- **Module 12: Capstone Simulation: The Launch Decision (Approx. 5.5 hours)**
    - **Final Project:** A complex, multi-faceted "Go/No-Go" decision simulation.

ADEPT TRAINING CENTER

    - **Scenario:** "You are the leadership team for a new generative AI product. It's incredibly powerful but internal testing has shown it can be used to create highly convincing deepfakes and misinformation. Your biggest competitor is rumored to be launching a similar product next month. The board is expecting a launch."
    - Teams must work through the ethical frameworks, assess the risks, consider mitigation strategies, and prepare a final recommendation to the "CEO" with a clear justification. Should they launch, delay, or kill the project?
- **Debrief & Personal Commitment (Approx. 0.5 hours)**
    - Debrief of the capstone simulation.
    - Participants will commit to one concrete change they will implement in their team or organization's process to improve ethical oversight.
    - Final Q&A.

ADEPT TRAINING CENTER