

THE AI PRECIPICE

An Executive Guide to Balancing AI Innovation and Safety

Academic Resource Package

Content is Suitable for both,
Undergraduate and
Graduate Programs



ISBN: 9781041020196 (hardcover) | 9781041038184 (paperback / ebook)

© 2026 Cathy Cobey and Dr. Cindy Gordon. All rights reserved.

The AI Precipice: An Executive Guide to Balancing AI Innovation and Safety is published by Routledge (Taylor & Francis Group), 2026.

This academic guide may be freely reproduced for non-commercial book club and educational use. For all other permissions, contact the publisher at permissions@tandf.co.uk.

Table of Contents

Academic Resource Package.....	1
Content is Suitable for both, Undergraduate and Graduate Programs.....	1
Table of Contents.....	3
Where to Get this Book.....	4
About the Authors.....	4
About This Package.....	5
Part 1: Course Alignment Guide.....	6
Programme Fit Matrix.....	8
Part 2: Student Reading Guide.....	9
Part 3: Discussion Questions.....	13

Where to Get this Book

Routledge: routledge.com/The-AI-Precipice

Amazon: amazon.com/dp/1041038186

Chapters/Indigo (Canada): indigo.ca — search “The AI Precipice”

BULK ORDERS

For bulk orders, please reach out to the publisher.

About the Authors

Cathy Cobey is EY 's Global Responsible AI Leader for Assurance and a leading AI advisor, strategist, and writer. She advises boards and executive teams on responsible AI adoption, governance, and risk management.

Dr. Cindy Gordon is CEO of SalesChoice, a prominent AI leader, speaker, researcher, and author of 15 books, with academic affiliations including the University of Toronto and George Brown College.

SPEAKING ENQUIRIES

For speaking or keynote bookings, please reach out to the authors. Additional contact information is available on the book website.

Dr. Cindy Gordon - cindy@saleschoice.com

Cathy Cobey - cathycobey@gmail.com

About This Package

This Academic Resource Package supports instructors and students using *The AI Precipice: An Executive Guide to Balancing AI Innovation and Safety* in undergraduate courses. The package is organized into five sections:

- **Part 1 – Course Alignment Guide:** A suggested 13-week syllabus and a programme fit matrix showing how the book maps to different undergraduate programmes.
- **Part 2 – Student Reading Guide:** A brief chapter-by-chapter orientation to help students engage with each chapter independently before class.
- **Part 3 – Discussion Questions:** Three discussion questions per chapter, suitable for seminars and tutorials. These are designed to be shared with students.
- **Part 4 – Instructor's Manual (INSTRUCTOR ONLY):** Learning objectives and teaching notes for each chapter. Not for student distribution.
- **Part 5 – Test Bank (INSTRUCTOR ONLY):** Multiple choice questions (with answer keys), short answer questions, and essay prompts for each chapter. Not for student distribution.

DISTRIBUTION NOTE

Parts 1, 2, and 3 are suitable for public distribution and may be posted on course websites or shared with students directly. Parts 4 and 5 are for instructor use only and should be hosted on a password-protected instructor portal.

Part 1: Course Alignment Guide

Suggested 13-Week Undergraduate Syllabus

The following syllabus maps the book's 16 chapters to a standard 13-week semester. Week 7 is reserved for a mid-term review and applied case study. Instructors may adjust pacing based on course level and student background.

Week	Topic	Chapters	Learning Focus
1	The AI Precipice: Why This Moment Matters	Preface + Chapter 1	Setting the stage; redefining the human-AI relationship
2	Responsible AI by Design	Chapters 2-3	Embedding ethics from day one; building executive AI literacy
3	Building AI Right: Lifecycle and Ecosystem	Chapters 4-5	Development lifecycle; AI as a broader ecosystem
4	AI Security and Risk Fundamentals	Chapter 6	Adversarial attacks; proactive and reactive security measures
5	AI Innovation vs. Safety: The Regulatory Landscape	Chapter 7	Global regulatory approaches; balancing innovation and safety
6	Corporate Governance and AI Strategy	Chapter 8	Aligning AI with corporate purpose; embedding responsible AI in culture
7	Mid-term Review and Applied Case Study	Chapters 1-8 review	Case analysis applying Sections I and II frameworks
8	AI Risk Management in Practice	Chapter 9	Risk prioritisation frameworks; building a risk-aware culture
9	Controls, Monitoring, and Assurance	Chapters 10-11	Cost-efficient controls; building stakeholder trust
10	Data-driven AI Monitoring	Chapter 12	Performance and ethics metrics; continuous improvement
11	Future-proofing: Scenario Planning for AI	Chapter 13	Scenario planning methodology; strategic agility
12	Utopia, Dystopia, and the Stakes of AI Leadership	Chapters 14-15	Positive and negative AI futures; governance responsibilities
13	Leading with a Thriving Mindset: Capstone	Chapter 16 + Synthesis	Leadership maturity framework; capstone reflection

Note: For a 12-week semester, Chapters 14 and 15 may be combined into a single session alongside Chapter 16.

Programme Fit Matrix

The AI Precipice is well suited to undergraduate and graduate courses across a range of programmes. The following matrix identifies the most relevant courses and key chapters by programme type.

Programme Type	Relevant Courses	Key Chapters
BBA / Bachelor of Commerce	Introduction to AI in Business, Business Ethics, Corporate Strategy	1, 2, 3, 7, 8, 16
MBA / Executive MBA	AI Strategy, Leadership & Governance, Risk Management	2, 7, 8, 9, 11, 13, 16
Master's in Data Science / AI	AI Ethics, Responsible AI, AI Governance	2, 3, 4, 5, 6, 9, 12
Public Policy / Law	Technology Law, AI Regulation, Digital Governance	7, 8, 14, 15
Risk Management / Audit	Enterprise Risk, IT Audit, AI Controls	6, 9, 10, 11, 12
Executive Education	AI for Executives, Boards and AI, Leading Digital Transformation	1, 3, 7, 8, 13, 16

Part 2: Student Reading Guide

This reading guide gives you a brief orientation to each chapter before you read it. Use it to focus your reading and prepare for class discussions.

Section I: The Courage Needed for Balanced AI

Chapter 1: Redefining our relationship with machines

Explores the profound shift in how humans interact with AI technologies, arguing that understanding both capabilities and limitations is essential to fostering a relationship that maximises benefits while ensuring ethical use. Education and transparency are positioned as the foundations of responsible adoption.

Chapter 2: Establishing responsible AI from day one

Introduces the concept of Responsible AI by Design — embedding ethical considerations, risk identification, and stakeholder involvement from the very beginning of an AI project. Provides practical frameworks for ensuring that responsible practices are built in rather than bolted on.

Chapter 3: Demystifying the magic of AI

Makes complex AI concepts accessible to executives by breaking down key model types and learning methods. Advocates for transparency in communicating AI capabilities and limitations as the foundation for building organizational trust and enabling informed governance decisions.

Chapter 4: Iterating AI from concept to reality

Provides a detailed overview of the AI development lifecycle from problem identification through to deployment and ongoing monitoring. Emphasises the iterative nature of AI development and the importance of embedding Responsible AI principles at each stage rather than treating them as a final checklist.

Chapter 5: AI is more than just models

Challenges the model-centric view of AI by mapping the full ecosystem of components — data collection, preprocessing, hardware, software, and network connectivity — that underpin any AI system. Emphasises that a holistic approach to AI governance must address all of these elements, not just the model itself.

Chapter 6: Battling AI breaches and adversarial attacks

Examines the security risks unique to AI systems, including adversarial attacks, data poisoning, and model inversion. Provides both proactive and reactive security strategies, drawing on case studies of real breaches to highlight the governance and reputational stakes of AI security failures.

Section II: Roadmap to achieving responsible AI

Chapter 7: The tug-of-war between AI innovation and safety

Examines the global tension between fostering AI innovation and ensuring safety and ethical standards. Surveys diverse regulatory approaches — including the EU AI Act and North American frameworks — and argues for governance structures that promote both rather than treating them as mutually exclusive.

Chapter 8: Aligning corporate profit and purpose

Argues that responsible AI and corporate profitability are complementary rather than competing goals. Provides practical strategies for integrating AI governance into corporate structures, aligning AI initiatives with strategic objectives, and embedding responsible AI in organizational culture through leadership and stakeholder engagement.

Chapter 9: Targeting the risks that truly matter

Provides frameworks for identifying and prioritizing AI risks by their potential impact. Covers risk assessment methodologies tailored to AI-specific threats — including model drift, algorithmic bias, and opacity — and describes what it takes to build a genuinely risk-aware organizational culture.

Chapter 10: Cost-efficient controls with maximum impact

Explores strategies for developing robust AI governance controls that are effective without being prohibitively expensive. Distinguishes preventative from detective controls, and provides practical guidance for selecting, implementing, and reviewing controls in proportion to the risks they are designed to address.

Chapter 11: Maintaining stakeholder trust and confidence

Examines the role of AI assurance — both internal and external — in building and sustaining stakeholder trust. Explores how organizations can demonstrate that their AI systems are operating responsibly, and provides strategies for communicating assurance outcomes to diverse stakeholder groups including employees, customers, regulators, and investors.

Chapter 12: The art of data-driven AI monitoring

Establishes continuous monitoring as a non-negotiable component of responsible AI governance. Covers the development of measurement frameworks that include both technical performance metrics and ethical compliance indicators, and advocates for data-driven governance interventions grounded in monitoring evidence.

Chapter 13: Future-proofing AI with scenario planning

Introduces scenario planning as a strategic tool for navigating AI-driven uncertainty. Explains how developing multiple plausible futures — rather than a single forecast — enables organizations to identify unintended consequences, test governance frameworks against different conditions, and build the strategic agility needed to thrive whatever the future brings.

Section III: Utopia or dystopia? The future of AI

Chapter 14: Building utopia

Envisions a future in which AI enhances human potential, promotes societal well-being, and fosters ecological sustainability. Outlines the principles of a positive AI future — including economic empowerment, inclusive governance, and ecological stewardship — and argues for the central role of purpose-driven businesses in helping realize this vision.

Chapter 15: Sliding into dystopia

Examines the pathways through which AI development could produce deeply harmful societal outcomes — including the exacerbation of inequality, erosion of privacy, AI-enabled surveillance, deepfake disinformation, and the undermining of human agency. Argues that these outcomes are not inevitable but require robust governance, ethical design, and multi-stakeholder oversight to prevent.

Chapter 16: Leading with a thriving mindset

Provides a capstone roadmap for executives navigating AI integration, organized around three maturity stages: survival, comfort, and thriving. Argues that the highest form of AI leadership requires balancing short-term business objectives with long-term societal responsibilities, investing in workforce transformation, and actively contributing to a positive AI future.

Part 3: Discussion Questions

The following questions are designed to support seminar and tutorial discussions. Three questions are provided per chapter, ranging from comprehension-level to critical and applied thinking. These questions may be shared with students.

Chapter 1: Redefining our relationship with machines

1. The authors argue that understanding AI's limitations is as important as understanding its capabilities. Do you agree? Provide an example from a field you are familiar with.
2. What cultural or organizational barriers might prevent people from adopting a balanced relationship with AI? How can leadership address these barriers?
3. The chapter highlights education and transparency as foundations for responsible AI. What would a meaningful corporate AI literacy programme look like in practice?

Chapter 2: Establishing responsible AI from day one

1. What does 'Responsible AI by Design' mean, and why does it matter more at the beginning of a project than at the end?
2. What are the risks of treating ethical considerations as an afterthought in AI development? Can you think of a real-world example where this caused problems?
3. How should organizations balance the pressure to move quickly in AI development with the need to embed responsible practices from the start?

Chapter 3: Demystifying the magic of AI

1. Why is it important for non-technical executives to understand, at a high level, how AI systems work? What governance risks arise from a lack of AI literacy at the leadership level?
2. The chapter argues for transparency in communicating AI capabilities and limitations. Who are the key audiences for this transparency, and what should each be told?

3. Choose one AI technology (e.g., machine learning, large language models, computer vision). What is its most significant business application, and what is one important limitation that decision-makers must understand?

Chapter 4: Iterating AI from concept to reality

1. Why is AI development described as 'iterative' rather than linear? What are the implications of this for how organizations plan and resource AI projects?
2. At which stage of the AI development lifecycle do you think Responsible AI principles are most difficult to apply, and why?
3. What role should business leaders play in the AI development lifecycle, and at which stages should they be most actively involved?

Chapter 5: AI is more than just models

1. The chapter argues that AI is an ecosystem, not just a model. What are the implications of this view for how organizations budget for and manage AI initiatives?
2. Why is data quality described as one of the most important factors in AI success? What can organizations do to improve the quality of data used in their AI systems?
3. What risks arise when organizations focus only on AI models while neglecting the broader infrastructure — data pipelines, hardware, network security — that supports them?

Chapter 6: Battling AI breaches and adversarial attacks

1. What makes AI systems uniquely vulnerable to security attacks compared to traditional software? What are the most significant types of attacks discussed in the chapter?
2. The chapter distinguishes between proactive and reactive security measures. Which do you think is more important for AI systems, and why?
3. If a company experiences an adversarial attack on its AI system, what are the likely business consequences beyond the immediate technical damage?

Chapter 7: The tug-of-war between AI innovation and safety

1. The chapter describes a 'tug-of-war' between AI innovation and safety. Do you think this tension is inherent, or can it be resolved? What conditions would need to be in place?
2. Compare the approach to AI regulation taken by two different jurisdictions. What are the advantages and disadvantages of each?
3. What role should industry leaders — as distinct from governments — play in establishing AI safety standards? What are the risks of self-regulation?

Chapter 8: Aligning corporate profit and purpose

1. Is there an inherent conflict between corporate profit and responsible AI? Make the case for why responsible AI can be a competitive advantage rather than a constraint.
2. What does it mean to 'embed responsible AI in organizational culture'? How would you know if a company had genuinely done this versus simply paying lip service to it?
3. The chapter discusses the role of leadership in AI governance. What specific behaviours should senior executives model to foster a responsible AI culture?

Chapter 9: Targeting the risks that truly matter

1. Why is it important to prioritise AI risks rather than treating all risks as equally important? What framework does the chapter offer for making these prioritisation decisions?
2. How does AI risk management differ from traditional enterprise risk management? What new risks does AI introduce that conventional frameworks may not fully capture?
3. What does it mean to foster a 'risk-aware culture' in the context of AI? What practical steps can organizations take to build this?

Chapter 10: Cost-efficient controls with maximum impact

1. What is the difference between a preventative control and a detective control in AI governance? Provide one example of each in the context of an AI system you are familiar with.

2. The chapter argues that AI controls must be cost-efficient. How should organizations decide which controls to invest in most heavily? What factors should guide this prioritisation?
3. Is it possible to over-control an AI system? What are the risks of excessive governance controls on organizational agility and innovation?

Chapter 11: Maintaining stakeholder trust and confidence

1. What is the difference between internal and external AI assurance, and why are both important for maintaining stakeholder confidence?
2. The chapter argues that trust in AI systems must be earned, not assumed. What does this mean in practice? How should organizations demonstrate trustworthiness over time?
3. How should organizations communicate about their AI governance practices to different stakeholder groups? Are different messages appropriate for different audiences?

Chapter 12: The art of data-driven AI monitoring

1. Why is ongoing monitoring of AI systems necessary even after they have been successfully deployed? What can change over time that makes monitoring so essential?
2. What metrics should organizations use to monitor an AI system's performance and ethical compliance? How are these different from traditional software performance metrics?
3. What is the risk of relying solely on technical monitoring metrics while ignoring qualitative or human-centred indicators of AI system health?

Chapter 13: Future-proofing AI with scenario planning

1. What is scenario planning, and why is it particularly valuable for organizations navigating rapid AI-driven change?
2. How does scenario planning differ from forecasting? Why might exploring multiple scenarios be more useful than a single best estimate of the AI future?
3. Develop a brief scenario for your industry: what would it look like if AI adoption in your sector accelerated dramatically over the next five years? What would be the key risks and opportunities?

Chapter 14: Building utopia

1. The chapter describes an 'AI utopia.' What are the core characteristics of this vision, and how realistic is it given current trajectories of AI development?
2. What role do businesses play in building a positive AI future? What specific commitments would a genuinely 'purpose-driven enterprise' need to make?
3. The chapter emphasises inclusive governance and ecological stewardship as components of an AI utopia. How do these values interact with the commercial imperatives of private sector AI development?

Chapter 15: Sliding into dystopia

1. The chapter identifies several pathways to a dystopian AI future. Which do you consider most dangerous, and why?
2. How does unchecked AI development threaten privacy and human agency? What governance measures could most effectively prevent this?
3. The chapter argues that surveillance and deepfakes represent critical societal risks. What responsibilities do technology companies bear for addressing these risks, and how far does individual corporate responsibility extend?

Chapter 16: Leading with a thriving mindset

1. The chapter describes three stages of AI leadership maturity: survival, comfort, and thriving. Where would you place most organizations today? What would it take to move to the 'thriving' stage?
2. What is a 'thriving mindset' in the context of AI leadership? How does it differ from a survival-oriented approach, in terms of both mindset and concrete behaviours?
3. The chapter argues that workforce transformation is essential for thriving AI leadership. What responsibilities do organizations have to employees whose roles are significantly affected by AI adoption?