



1. General Information

Course Subject	IIMT
Course Number	3671
Course Title	Responsible AI and Ethics
Academic Years	2026-2027
Grading Method	Letter

2. Instructors

Professor Bishop, David Lorin
Office: Room 1313 /F K.K. Leung Building
Email: dbishop@hku.hk
Office: 39174215
Subclasses: 2A

3. Teaching Assistants

Mr Jon Andre Pedersen
Email: jonp@hku.hk

4. Course Description

Course Description	The rapid integration of Artificial Intelligence into global business has powerfully reinforced the importance of ethical behavior, especially for managers deploying these novel technologies. Over the past few years, we have seen the incredible potential of AI, alongside stories of algorithmic bias, data misuse, and corporate misfeasance. Now more than ever, ethical standards regarding technology are being fiercely debated and enforced. It is imperative for future leaders to understand clearly what responsible AI entails, and have the tools to make ethical decisions when deploying these systems. The purpose of this course is to help you understand the role of ethics in developing, managing, and utilizing AI in a global environment. We will work together to develop the tools necessary to face complex ethical situations, where action often must be taken quickly regarding rapidly changing technology. The study of AI ethics is a vital conversation about human interaction, societal impact, and accountability. We will explore areas such as law, intellectual property, data privacy, algorithmic finance, and how traditional ESG (Environmental, Social, and Governance) frameworks must adapt to the digital age. This conversation is necessary to shape and counterbalance the relentless drive for technical innovation in business.
--------------------	---

5. Course Objectives

1. Students are constantly inundated with moral and ethical standards regarding technology, where the push for innovation often outpaces ethical consideration. Therefore, this course is largely to help you unlearn what you have learned about technology being purely objective or neutral.

By the end of the course you will:

1. be able to recognize ethical issues embedded in AI models, AI applications, and data systems, critically examining your own standards regarding technology use;
2. have developed the ability to critically analyze complicated business scenarios involving AI and make difficult ethical decisions based on what is 'right' rather than purely driven by efficiency;
3. appreciate the environmental and societal impacts of digital infrastructure, applying ESG materiality principles to technology;
4. develop a management perspective that includes an ability to formulate, analyze, and defend technological deployments and regulatory compliance (RegTech) in ethical terms.

6. Faculty Learning Goals

Goal 1: Acquisition and internalization of knowledge of the programme discipline

Goal 2: Application and integration of knowledge

Goal 3: Inculcating professionalism

Goal 4: Developing global outlook

Goal 5: Mastering communication skills

Goal 6: Cultivating leadership

7. Course Learning Outcomes

Course Teaching and Learning Activities	Aligned Faculty Learning Goals					
	1	2	3	4	5	6
CLO1. Gain the ability to explain ethical theories, legal frameworks, and ESG implications related to AI, both within Hong Kong and globally.	✓	✓				
CLO2. Critically assess the main concepts of responsible AI within the context of real-life business scenarios and simulations, and create appropriate solutions to technological ethical problems.		✓	✓		✓	
CLO3. Demonstrate the inter-relationship of developers, users, regulators, and the registered company, analyzing their roles in the context of ethical AI leadership and risk management.		✓	✓	✓	✓	
CLO4. Demonstrate effective communication skills, both in written and oral formats, particularly when defending or critiquing AI implementations.					✓	✓

8. Course Teaching and Learning Activities

Course Teaching and Learning Activities #	Expected Study Hours	Study Load (% of study)
T&L1. Lectures, quizzes, and in-class discussions	36	26.5
T&L2. Immersive Simulations & Preparation	35	25.7
T&L3. Final Writing Assignment	30	22.1

8. Course Teaching and Learning Activities		
T&L4. Independent Study	35	25.7
	Total: 136	Total: 100

9. Assessment Methods			
Assessment Methods	Description	Weight %	Aligned Course Learning Outcomes
A1. Final Discussion Assignment	Students will prepare for and engage in a comprehensive final discussion with the instructors. This session will test their ability to orally defend their perspectives on key course concepts, AI governance, and ethical deployments.	20%	1,2,3,4
A2. Immersive Simulations	Students will participate in intensive, realtime simulations throughout the semester (e.g., managing an AI crisis, pitching to an ethics board). They will be assessed on their preparation, role-play execution, and post-simulation critical analysis.	40%	2,3,4
A3. Final Writing Assignment	Students will be required to engage with a specific AI ethics topic, regulatory challenge, or proposed ESG governance framework for AI.	40%	2,3,4

Assessment Rubrics	
A1. Final Discussion Assignment	Students will prepare for a dedicated discussion session with the course instructor or teaching assistant at the end of the semester. This assessment replaces traditional class participation and evaluates the student's ability to articulate, defend, and debate complex ethical decisions regarding AI. Students will be judged on their depth of knowledge, critical thinking, and oral communication skills under direct questioning.
A+,A,A-	Demonstrates exceptional, original critical thinking and a profound depth of knowledge when articulating and defending complex ethical decisions. Handles direct questioning fluidly with excellent oral communication, rhetorical skill, and situational adaptability.
B+,B,B-	Demonstrates strong analytical ability and a solid grasp of AI ethics. Capably defends positions under direct questioning and shows strong oral communication skills, though arguments may rely more on course materials than highly original thought.
C+,C,C-	Shows a reasonable grasp of the subject matter, but responses under questioning are largely derivative. Demonstrates little evidence of deep critical thinking or ability to pivot when challenged, paired with fair oral presentation skills.
D+,D	Provides the bare minimum of information during the discussion. Oral defense is poorly organized, and the student struggles significantly to answer direct questions or apply ethical concepts.
F	Fails to demonstrate basic knowledge or understanding of AI ethics. Answers during the discussion are largely irrelevant, and the student is unable to articulate or defend a coherent position.
A2. Immersive Simulations	Students will be required to actively participate in and analyze immersive simulations relating to various AI ethics subjects discussed in the course. The grading for these simulations will cover: ● Preparation & Strategy: The quality of research and planning prior to the simulation. ● Execution: Active engagement, ethical decision-making under pressure, and adherence to assigned stakeholder roles during the simulation. ● Post-Simulation Briefs: Short executive summaries (no more than two pages) analyzing the outcomes, what went wrong/right, and the structural lessons learned.

Assessment Rubrics	
A+,A,A-	Exhibits outstanding research and strategic planning prior to the simulation. Actively drives the simulation forward with highly effective, ethical decision-making under pressure while perfectly maintaining their stakeholder role. The post-simulation brief is exceptional—concise, deeply analytical, and draws profound structural lessons.
B+,B,B-	Shows strong preparation and active engagement. Makes sound, defensible decisions under pressure and adheres well to their assigned role. The post-simulation brief is highly competent and well-organized, demonstrating clear analysis of outcomes, though it may lack highly original insights.
C+,C,C-	Preparation is adequate but lacks depth. Participates in the simulation but struggles with strategic foresight or ethical decision-making when placed under pressure. The post-simulation brief is more descriptive of what happened rather than critically analyzing the structural lessons.
D+,D	Demonstrates minimal preparation and poor execution, often breaking character or failing to engage with the ethical dilemma. The post-simulation brief is poorly organized, poorly digested, and lacks any real reflection on what went wrong or right.
F	Shows no evidence of preparation or strategy. Fails to participate meaningfully in the simulation or handle the scenario. The brief is irrelevant, fails to address the prompt, or is missing entirely.
A3. Final Writing Assignment	Students will be asked to engage with a specific AI ethics topic or proposed ESG governance framework for AI. More details concerning this assignment will be provided during the course.
A+,A,A-	Demonstrates exceptional original thought and deep critical analysis of the chosen AI ethics topic or ESG framework. The paper is impeccably organized, highly persuasive, and shows a masterful synthesis of complex regulatory and ethical concepts.
B+,B,B-	Shows strong critical and analytical thinking regarding the regulatory challenge or ESG framework, though the arguments may not be entirely original. The writing is well-organized, rhetorically sound, and clearly demonstrates a thorough grasp of the background reading.
C+,C,C-	Demonstrates a reasonable grasp of the chosen AI topic, but the analysis is mostly derivative and lacks deep critical evaluation of the governance frameworks. Organization and writing are fair but may lack flow or persuasive strength.
D+,D	Assembles only the bare minimum of information regarding the AI ethics topic. The writing is poorly digested, heavily disorganized, and shows no real evidence of critical thinking or understanding of the regulatory challenges.
F	Demonstrates a severe lack of understanding of the chosen topic. The assignment lacks coherence, contains largely irrelevant information, and fails to reach degree-level standards for written work.

10. Course Grade Descriptors	
A+,A,A-	Demonstrate evidence of original thought, strong analytical and critical abilities as well as a thorough grasp of the topic from background reading and analysis; should demonstrate excellent organizational, rhetorical, and presentational skills.
B+,B,B-	Demonstrate evidence of critical and analytical thinking but not necessarily original in their thinking; show adequate grasp of the topic from background reading and analysis; should demonstrate strong organizational, rhetorical and presentational skills.
C+,C,C-	Demonstrate evidence of a reasonable grasp of their subject but most of their information is derivative, with rather little evidence of critical thinking; should demonstrate fair organizational, rhetorical and presentational skills.
D+,D	Demonstrate evidence of being able to assemble the bare minimum of information, poorly digested and not very well organized in presentation. There is no evidence of critical

10. Course Grade Descriptors		
		thinking.
	F	Demonstrate evidence of poor knowledge and understanding of the subject, a lack of coherence and organization, and answers are largely irrelevant. Work fails to reach degree level.

11. Course Content and Tentative Teaching Schedule		
Topic/ Session	Content	
1	The Ethics of Innovation and Introduction to Responsible AI An overview of the tension between rapid technological advancement and ethical responsibility, challenging the assumption that technology is purely neutral or objective. This module will include: Approaching AI Philosophically (philosophy, epistemology, ethics)	
2	Transparency & Integrity - Algorithmic Bias and Data Equity Examining how AI models inherit and amplify human biases, with a focus on real-world business scenarios where skewed data sets have led to corporate misfeasance and discrimination.	
3	Accountability The ethical challenges of programming moral decisions into machines; the principles that should guide such decisions; and the legal and ethical frameworks for accountability when those decisions lead to harm.	
4	Fairness The trade-offs between technological efficiency and fairness; the ethical obligations to prevent bias in technological deployments; and the broader societal impacts of relying on flawed systems.	
5	Control - Data Privacy, Consent, and Surveillance Capitalism Analyzing the ethical boundaries of data collection and privacy in the digital age, mapping out the responsibilities of companies that gather user data to fuel AI systems. The balance between using technology for collective governance and the potential risks to personal liberty; the impact of algorithmic bias on such systems; and the meaning of user-centric data agency.	
6	AI in Finance and Algorithmic Markets A deep dive into algorithmic finance, exploring the ethical implications of high-frequency trading, automated credit scoring, and systemic risk in global markets.	
7	Sustainability - AI for Social Impact and ESG Adaptation Exploring how traditional Environmental, Social, and Governance (ESG) frameworks must evolve to account for digital infrastructure, and how organizations can actively leverage AI to create a shared, positive societal impact. We will also consider the environmental costs and benefits of deploying AI systems; and the ethical implications in aligning them with sustainable development goals.	
8	Intellectual Property and the Generative AI Era Navigating the complex legal and ethical landscape of copyright, ownership, and fair use when AI models are trained on, and subsequently generate, proprietary content.	
9	RegTech and Global AI Compliance Understanding the fragmented global regulatory environment and how companies can utilize Regulatory Technology (RegTech) to maintain compliance and proactively manage ethical AI risks.	
10	AI Leadership and Corporate Accountability Focusing on the executive perspective, this topic covers how leadership teams and ethics boards can formulate, analyze, and defend responsible AI deployments to external stakeholders.	
11	The Future of AI Ethics and Final Defenses A concluding synthesis of the course concepts, looking at the trajectory of AI technologies and preparing students for their final oral discussion assessments on ethical deployment.	

12. Required/Recommended Readings & Online Materials

Reading	Students will be provided with a course packet by the instructor. Moreover, because the landscape of artificial intelligence, regulation, and technology ethics is evolving rapidly, we frequently incorporate current news articles, breaking case studies, and recent policy updates into our teaching. Additional required readings from contemporary news sources will be distributed via Moodle throughout the semester to ensure our discussions reflect the most up-to-date, real-world business scenarios. Sample readings (subject to adjustment): · Policy on the AI Exponential · The Adolescence of Technology · Pope Leo XIV On Safeguarding the Human Person In the Time of Artificial Intelligence · The Ethics of Artificial Intelligence · Weapons of Math Destruction · The Alignment Problem
---------	--

13. Means / Processes for Student feedback on Course

✓	Conducting mid-term survey in additional to SETL around the end of the semester
✓	Online response via Moodle site
	Others

14. Course Policy

An orderly learning environment is extremely important for this course. Disruptive behaviors are inconsiderate to other students as well as to the instructor and are absolutely unacceptable. Talking during lectures, arriving to class late, and any other disruptions of mobile devices are not allowed; students who are responsible for any of these actions will be subject to academic penalty and will be asked to leave the classroom.

Any dishonesty such as cheating, false representation, plagiarism, etc. that comes to my attention will result in an F in the course.

Academic dishonesty includes cheating, plagiarism, unauthorized collaboration, falsifying academic records, and any act designed to avoid participating honestly in the learning process. Scholastic dishonesty also includes, but is not limited to, providing false or misleading information to receive a postponement or an extension on an exam or other assignment. The responsibilities of both students and faculty with regard to scholastic dishonesty are described in detail in the Disciplinary Committee Regulations. By teaching this course, I have agreed to observe all of the faculty responsibilities described in that document. By enrolling in this class, you have agreed to observe all of the student responsibilities described in that document. If the application of that policy statement to this class and its assignments is unclear in any way, it is your responsibility to ask me for clarification.

The Use of Generative AI in IIMT3671

The University of Hong Kong recognizes AI literacy as an essential competency alongside oral, written, visual, and digital communication. Because this course focuses explicitly on the responsible deployment of AI, you are actively encouraged to use Generative AI tools (e.g., ChatGPT, Claude, OpenAI developer tools, or AI coding assistants like Cursor) to brainstorm, structure your thoughts, or assist with research for your Final Writing Assignment and simulation preparations.

However, using these tools comes with strict ethical responsibilities regarding accountability and transparency:

- You are ultimately accountable for all work you submit. AI tools can hallucinate, produce biased outputs, or provide legally inaccurate information. You must critically verify all AI-generated content. Relying solely on AI without human judgment will result in a failing grade for the assignment.
- If you use GenAI tools to generate content, assist in writing, or analyze data, you must properly cite the tool used. In addition to a standard citation, you must include a brief "AI Methodology Declaration" at the end of your assignment. This declaration must state: Which tool(s) you used. How the tool was used (e.g., ideation, drafting, data analysis). How you exercised human judgment and oversight to arrive at the final product.
- Do not input sensitive, confidential, or proprietary data into public AI models. Ensure that any data you use complies with the University's privacy regulations.
- Passing off raw, AI-generated material as your own original thought without proper acknowledgement is a form of plagiarism and will be penalized in accordance with University regulations.

15. Additional Course Information

This course will utilize the Moodle framework.