

Spring 2025: **The Ethics and Governance of Artificial Intelligence** (PPHA 38850)

Background and Goals

This course critically examines the emergence of ethical frameworks in AI governance, exploring the norms, values, and political strategies that shape the development and oversight of AI systems across public policy, corporate practices, and international regulation. Students will analyze regional AI strategies and trace the origins and real-world applications of foundational AI principles such as fairness, accountability, and transparency.

At the heart of the curriculum is an interactive debate forum, where students will engage in rigorous discussions on the responsible use of AI across diverse case studies—spanning questions of human rights, sustainable development, and geopolitics. This format requires students to articulate, defend, and critique competing perspectives while sharpening both analytical precision and rhetorical agility.

By navigating the regulatory complexities of emerging technologies, this course serves as a training ground for future policymakers, advocates, and industry leaders—equipping them with the critical thinking and persuasive skills necessary to shape tech policy for the public good.

Prerequisites

This course welcomes students from all academic backgrounds. Although there are no formal prerequisites, participants should be comfortable with active engagement in lively discussions and collaborative group projects. A willingness to articulate perspectives, engage in constructive debate, and navigate multi-stakeholder dialogue is essential for success in this interactive learning environment.

How This Course Will Work

The course offers a dynamic blend of critical analysis and interactive debate. Each session begins with a mini-lecture and guided discussion of assigned readings, setting the stage for immersive debates where students rigorously examine and challenge diverse perspectives on AI governance through case studies. Active participation is expected—students should complete readings prior to class, attend all sessions, and engage fully in discussions and group exercises. All course materials, announcements, and assignment submissions will be managed through Canvas.

Course topics and schedule

Readings will be posted as digital copies on Canvas. No book purchases are required for full participation.

Week 1: AI, Power & Public Interest

How does AI shape society, and what are the key governance challenges?

- ⇒ The political and social stakes of AI governance; governance structures shaping the impact of AI on society; and the role of public institutions, corporations, and civil society in ensuring ethical AI.

Week 2: Transparency, Explainability & the Politics of AI

How can we evaluate and interpret AI decision-making?

- ⇒ The challenges of explainability of AI systems; the limits of interpretability in algorithmic decision-making; the “black box” problem in AI models; and the role of human oversight in building public trust.

Week 3: Fairness, Bias & Structural Inequalities

How do AI systems shape social and economic disparities, and what interventions are needed?

- ⇒ The ethical and political dimensions of algorithmic bias; trade-offs between efficiency and equity in AI policy; the ways AI can reinforce or mitigate systemic disparities; and regulatory responses to biased systems.

Week 4: Accountability, Responsibility & Moral Agency

How can legal and ethical frameworks address the consequences of AI-driven decisions?

- ⇒ The challenges of assigning responsibility for harms caused by AI systems; corporate, governmental, and individual obligations in AI oversight; and the debate over AI as a moral agent.

Week 5: Autonomy, Control & the Impact of AI on Freedom

To what extent does AI enhance or undermine human independence in decision-making?

- ⇒ The influence of AI on human choices and behaviors; algorithmic manipulation in politics and consumer decision-making; the risks of surveillance capitalism; and the ethics of AI nudging and behavioral control.

Week 6: AI, Democracy & Public Participation

How can different stakeholders influence AI governance, and what barriers exist?

- ⇒ The role of civil society in AI policy; challenges in ensuring broad participation in AI governance; strategies for inclusive and participatory AI policymaking; and regulatory approaches to democratic oversight.

Week 7: AI, Sustainability & the Political Economy of AI

What are the environmental and economic trade-offs of AI development and deployment?

- ⇒ The environmental cost of AI development; the impact of AI on labor and economic structures; global disparities in AI access and benefits; and alternative governance models for long-term sustainability.

Week 8: The Future of AI & the Governance of Emerging Technologies

How should AI governance evolve to address future risks and opportunities?

- ⇒ The long-term risks of AI; ethical dilemmas in AGI and superintelligence; geopolitical AI governance challenges; and different models for shaping the future of AI governance.

Week 9: Policy Memo Workshop

How can policy writing effectively communicate AI governance solutions?

- ⇒ Best practices for writing evidence-based policy memos; structuring arguments for clarity and impact; aligning recommendations with governance constraints; and peer review sessions to refine final projects.

Student Assignments

Please consult our course page on Canvas for detailed instructions on all assignments.

1. Reading Memos / 30% of final grade

Submit weekly reading memos articulating key themes, critical questions, and constructive critiques of assigned readings. These memos will be used to structure in-class discussions and collaborative activities.

2. AI Policy Review (Presentation) / 30% of final grade

Work in teams to evaluate the AI strategy of a selected regulatory region, assessing its policy frameworks, governance mechanisms, and alignment with global trends. Teams will analyze official policies, legislation, industry reports, and expert commentaries, considering both theoretical and practical challenges. The assignment culminates in a team presentation and Q&A session, where students must present a well-structured, evidence-based analysis.

3. AI Policy Memo / 40% of final grade

Individually develop a policy memo proposing a concrete solution to a critical AI challenge. The memo should define the problem, assess existing policies, and provide a well-researched and practical policy recommendation grounded in relevant evidence (e.g., legal texts, case studies, empirical research). Students must also outline a plan for measuring effectiveness and present their findings in a persuasive, policy-oriented format.

Classroom Policies

Grading

This course adheres to the University of Chicago's [Common Grade Policy](#), which uses a 4-point scale for quality grades. Each assignment will receive a letter grade corresponding to the point values outlined in the policy, as follows: 4.0 = A, 3.7 = A-, 3.3 = B+, 3.0 = B, 2.7 = B-, 2.3 = C+, 2.0 = C, 1.7 = C-, 1.3 = D+, 1.0 = D, F = 0. All assignments are graded qualitatively, and final grades will be calculated based on the weighted distribution of assignments outlined in the 'assignments' section of this syllabus. For further details on grading policies, students are encouraged to consult the University's Registrar website.

Attendance and Submission Deadlines

This course follows a strict attendance policy. Each unexcused absence will result in a one-third letter grade reduction (e.g., from A to A-). Similarly, late submissions of assignments will incur a one-third letter grade reduction (e.g., from A to A-) for each day they are overdue. Please contact the teaching assistant in writing to request an excuse in advance. Acceptable excuses for absences and late submissions include health issues, family emergencies, short-term professional obligations, or other documented reasons.

Electronics and AI Tools in the Classroom

Electronics are allowed in class, but only for purposes that enhance the learning experience (e.g., accessing readings or participating in discussions). Multitasking is prohibited, and electronic devices must remain focused on class activities. Similarly, students are permitted to use AI tools for assignments; however, full disclosure of AI usage is required. Students must document how and why AI tools were used in their work to ensure transparency and critical attribution.

Instructor Office Hours

Andre Uhl, PhD

uhl@uchicago.edu

Mondays, 10—11AM (Zoom) – sign up on Canvas

Teaching Assistant

TBA

General Resources Available to Students

- [Harris Academic Support Programs and Handbook](#)
- [Student Wellness](#)
- [UChicagoGRAD](#)

Harris School and University of Chicago Policies

- [Harris School Policies](#)
- [University General Policies](#)
- [University Academic Policies](#)