Ethics in AI Course Project Guidelines

For First-Year Artificial Intelligence Students

Non-coding, theory-based semester project

Purpose of the project: Students will study one ethical issue related to AI, analyze it using course concepts, discuss a real-world example, and suggest responsible solutions. The project does not require coding, model building, or software development.

Project Submission: 2026 May 31

1. Project Objectives

Your project should answer the following questions:

- What is the selected AI-related ethical problem?
- Why is this problem important?
- Who is affected by this problem?
- What are the possible benefits and risks?
- Which ethical frameworks help explain the issue?
- What responsible solutions can reduce the problem?

2. Required Project Components

- Written report
- Presentation slides
- Oral presentation

3. Recommended Report Length and Section Distribution

Total recommended report length: 6 to 8 pages, excluding references. Use font size 11 or 12 with 1.15 or 1.5 line spacing. If your example analysis is detailed, 8 to 10 pages is acceptable.

Section	Recommended Length	What to Include
Title Page	1 page	Project title, course name, student name or group members, department, and date.
Introduction	1 page	Introduce the selected topic and explain why it is important in AI ethics.
Background of the Topic	1 page	Explain the AI technology, system, or situation related to the topic in simple language.
Main Ethical Problem	1 page	Clearly describe the key ethical issue, such as bias, privacy, misinformation, cheating, job loss, transparency, or overreliance.
Real-World Example or Case Study	1 page	Discuss one real or realistic example. Explain what happened, who was affected, and why it matters ethically.
Ethical Frameworks and Analysis	1.5 to 2 pages	Apply at least 2 ethical frameworks from the list provided in this guideline. Strong projects may apply 3 frameworks.
Impact on People and Society	1/2 to 1 page	Explain how the issue affects individuals, students, workers, patients, organizations, or society.
Recommendations or Solutions	1 page	Suggest practical, realistic ways to reduce the ethical problem.
Conclusion	1/2 page	Summarize your main findings and give your final opinion.

References	As needed	Use at least 3 to 5 credible references. Lecture slides may also be used.

4. Suggested Report Structure

1. Title Page
2. Introduction
3. Background of the Topic
4. Main Ethical Problem
5. Real-World Example or Case Study
6. Ethical Frameworks and Analysis
7. Impact on People and Society
8. Recommendations or Solutions
9. Conclusion
10. References

5. Ethical Frameworks to Use in the Project

Students must apply at least 2 ethical frameworks. Do not only define the framework. Apply it directly to your selected topic.

Ethical Framework	How to Apply It
Utilitarianism	Focuses on consequences. Ask whether the AI system creates more benefit or more harm for society.
Deontology	Focuses on rules, duties, rights, and responsibilities. Ask whether the use of AI follows ethical rules and respects people's rights.
Virtue Ethics	Focuses on character and moral values. Ask whether the AI use reflects honesty, fairness, responsibility, and respect.
Justice and Fairness	Focuses on equal treatment. Ask whether the AI system treats different people or groups fairly.
Rights-Based Ethics	Focuses on human rights such as privacy, dignity, freedom, equal opportunity, and consent.
Care Ethics	Focuses on empathy, human relationships, and protection of vulnerable people such as students, patients, children, or elderly people.
Accountability and Responsibility	Focuses on who should be responsible if AI causes harm or makes a wrong decision.
Transparency and Explainability	Focuses on whether people can understand when AI is used and why the AI made a decision.
Privacy and Consent	Focuses on data collection, permission, data protection, and possible misuse of personal information.
Human-in-the-Loop Ethics	Focuses on keeping human supervision in important AI-supported decisions, especially in high-stakes areas.

6. How to Write the Ethical Analysis Section

Use the following simple format:

- Framework 1: Name the framework and briefly explain how it relates to your topic.
- Application: Explain what the framework reveals about the ethical problem.
- Framework 2: Repeat the same process with another framework.
- Comparison: Briefly explain whether the frameworks lead to the same or different conclusions.

Example: If your topic is deepfakes and misinformation, you may use utilitarianism to discuss social harm, rights-based ethics to discuss privacy and dignity, and accountability to discuss who should be responsible for harmful deepfake content.

7. Presentation Guidelines

Recommended presentation length: 8 to 10 slides. Recommended presentation time: 5 to 7 minutes per group.

Slide	Content
Slide 1	Title and group members

Slide 2	Introduction
Slide 3	Background
Slide 4	Main ethical problem
Slide 5	Case study
Slide 6	Ethical framework 1
Slide 7	Ethical framework 2
Slide 8	Impact on people and society
Slide 9	Recommendations
Slide 10	Conclusion and references

8. Example Project Topics and How to Analyze Them

The following examples show how to move from a simple topic to a proper ethical analysis. Students may use a similar structure for their own selected topics.

AI-Assisted Cheating in Education

Project focus: Study how students use AI tools for essays, homework, quizzes, coding tasks, summaries, or exam preparation. Clearly separate acceptable support, such as brainstorming or language improvement, from dishonest use, such as submitting AI-generated work as original.

Ethical analysis: This topic can be analyzed through deontology because academic rules and honesty matter even if AI helps students get higher marks. Virtue ethics is also useful because the issue is not only about grades, but also about student character, responsibility, and learning habits. Fairness should also be discussed because students who honestly complete their work may be disadvantaged if others misuse AI.

Affected stakeholders and risks: Students, teachers, universities, and future employers are affected. Main risks include unfair grades, weak learning, loss of trust, and difficulty in judging real student ability.

Possible recommendations: Recommend clear AI-use policies, required disclosure of AI assistance, oral defense of submitted work, in-class tasks, project-based assessment, and teacher guidance on responsible AI use.

Bias in AI Hiring Systems

Project focus: Study how companies use AI to screen CVs, rank applicants, analyze video interviews, or predict employee performance. Explain that AI may appear neutral but can learn unfair patterns from historical hiring data.

Ethical analysis: Justice and fairness are central because the system may treat applicants differently based on gender, race, age, university background, disability, or language style. Transparency is also important because applicants should know why they were rejected. Accountability matters because companies should not blame the AI system when unfair decisions occur.

Affected stakeholders and risks: Job seekers, employers, HR teams, and society are affected. Main risks include discrimination, loss of opportunity, hidden bias, and reduced diversity in workplaces.

Possible recommendations: Recommend fairness audits, diverse training data, human review before final decisions, clear explanation to applicants, appeal procedures, and regular monitoring of the system.

Deepfakes and Misinformation

Project focus: Study how AI can create fake videos, images, voices, or screenshots that look real. Explain how this can be used for entertainment, but also for scams, bullying, political manipulation, or reputation damage.

Ethical analysis: Utilitarianism can be used to compare possible benefits, such as education and creative media, with social harms, such as confusion, fear, and loss of trust. Rights-based ethics is important because deepfakes can violate privacy, dignity, and consent. Accountability is needed because creators, sharers, platforms, and sometimes organizations may all contribute to the harm.

Affected stakeholders and risks: Victims, social media users, voters, celebrities, students, families, and institutions can be affected. Main risks include fake evidence, damaged reputation, financial fraud, public panic, and reduced trust in real media.

Possible recommendations: Recommend content labeling, media literacy education, deepfake detection tools, platform reporting systems, legal responsibility for harmful use, and careful verification before sharing suspicious content.

Privacy and Surveillance in Smart Cities

Project focus: Study how AI is used in CCTV systems, facial recognition, traffic monitoring, smart buildings, and public safety systems. Explain the difference between using data for safety and using data for excessive monitoring.

Ethical analysis: Privacy and consent are central because citizens may be recorded or analyzed without clear permission. Utilitarianism can help compare public safety benefits with privacy risks. Rights-based ethics is also important because people should have dignity, freedom, and protection from unnecessary surveillance.

Affected stakeholders and risks: Citizens, students, visitors, police, city officials, and private technology companies are affected. Main risks include misuse of personal data, continuous monitoring, data leaks, unequal targeting, and loss of public trust.

Possible recommendations: Recommend limited data collection, public notice, strong data security, independent oversight, strict access control, deletion of unnecessary data, and clear rules on when facial recognition may be used.

AI in Healthcare Decision-Making

Project focus: Study how AI supports diagnosis, medical image analysis, treatment suggestions, hospital triage, or patient monitoring. Explain that AI can support doctors but should not replace professional medical responsibility in high-risk decisions.

Ethical analysis: Care ethics is important because healthcare decisions affect vulnerable people who need empathy and protection. Human-in-the-loop ethics is necessary because doctors should review AI outputs before action. Accountability matters because responsibility must be clear if AI gives wrong advice. Transparency is also important because patients and doctors need to understand the reason for AI-supported recommendations.

Affected stakeholders and risks: Patients, doctors, hospitals, families, insurance companies, and AI developers are affected. Main risks include wrong diagnosis, privacy leakage, overreliance, lack of explanation, and unclear liability.

Possible recommendations: Recommend doctor supervision, patient consent, explainable AI, clinical testing, privacy protection, clear responsibility policies, and a requirement that final decisions remain with qualified medical professionals.

AI and Job Displacement

Project focus: Study how AI and automation may change work in customer service, translation, design, accounting, manufacturing, education, transportation, or office administration. Explain that AI may not only remove jobs but also change the skills required for future work.

Ethical analysis: Utilitarianism can be used to discuss productivity benefits and economic growth, but these benefits must be compared with harm to workers. Justice and fairness are important because the negative impact may be stronger for low-income workers or people with fewer opportunities for retraining. Care ethics helps students discuss social responsibility toward affected workers and families.

Affected stakeholders and risks: Workers, students, companies, governments, families, and future graduates are affected. Main risks include unemployment, income inequality, stress, skill gaps, and unfair distribution of AI benefits.

Possible recommendations: Recommend reskilling programs, responsible company transition plans, government support, career counseling, lifelong learning, and education systems that prepare students for AI-supported workplaces.

AI Detection Tools and False Accusations

Project focus: Study how schools and universities use AI detection tools to check whether assignments were written by AI. Explain that these tools may be helpful but are not always accurate.

Ethical analysis: Justice and fairness are important because a false accusation can seriously harm an innocent student. Deontology is relevant because institutions have a duty to follow fair procedures and not punish students without strong evidence. Transparency and accountability are also needed because students should know how decisions are made and how to appeal.

Affected stakeholders and risks: Students, teachers, academic departments, and institutions are affected. Main risks include unfair punishment, anxiety, damaged reputation, and overreliance on detection scores.

Possible recommendations: Recommend that AI detection results should not be the only evidence, require human review, allow student explanation or oral defense, provide clear AI-use rules, and create an appeal process.

AI Chatbots and Hallucinated Information

Project focus: Study how AI chatbots sometimes give incorrect information in a confident style. Discuss examples from education, health advice, law, research, travel, or daily decision-making.

Ethical analysis: Transparency is important because users should understand that chatbot answers can be wrong. Responsibility matters because both developers and users have duties: developers should warn users and improve reliability, while users should verify important information. Virtue ethics can also be applied by discussing honesty, carefulness, and responsible use of AI-generated answers.

Affected stakeholders and risks: Students, teachers, researchers, patients, consumers, and general users are affected. Main risks include wrong decisions, fake references, academic mistakes, health risks, and spreading misinformation.

Possible recommendations: Recommend source checking, use of reliable references, warning labels, fact-checking habits, teacher guidance, and avoiding chatbot-only decisions in high-stakes matters.

Automation Bias in Decision Support Systems

Project focus: Study how people may accept AI recommendations without enough critical thinking in healthcare, education, banking, navigation, policing, or business decisions.

Ethical analysis: Human-in-the-loop ethics is central because humans must remain active decision-makers instead of blindly following AI. Accountability is important because someone must be responsible if a wrong decision is made. Care ethics is useful in cases where vulnerable people may be harmed by overreliance on automated decisions.

Affected stakeholders and risks: Decision-makers, service users, patients, students, customers, and organizations are affected. Main risks include wrong decisions, reduced human judgment, overconfidence in AI, and unclear responsibility.

Possible recommendations: Recommend decision checklists, human review, training on AI limitations, clear responsibility rules, second opinions for high-risk decisions, and system designs that show uncertainty.

AI-Generated Research and Data Fabrication

Project focus: Study how AI tools can be misused to create fake references, fake data, fake summaries, or misleading research text. Explain why this is harmful for academic integrity and scientific knowledge.

Ethical analysis: Virtue ethics is highly relevant because honest research requires truthfulness, discipline, and responsibility. Deontology also applies because researchers have a duty to follow academic rules. Accountability is important because authors, supervisors, reviewers, and institutions share responsibility for checking research quality.

Affected stakeholders and risks: Students, researchers, journals, universities, policymakers, and the public are affected. Main risks include false knowledge, wasted resources, damaged academic trust, and incorrect decisions based on fabricated findings.

Possible recommendations: Recommend citation verification, data checking, supervisor review, clear AI disclosure, research ethics training, and penalties for intentional fabrication.

Facial Recognition and Discrimination

Project focus: Study how facial recognition is used in airports, schools, phones, public security, attendance systems, or policing. Explain that inaccurate recognition can create serious ethical problems.

Ethical analysis: Rights-based ethics is important because facial data is personal and connected to identity. Justice and fairness are also central because some groups may experience higher error rates. Accountability matters because people should have a way to challenge wrong identification.

Affected stakeholders and risks: Citizens, students, travelers, minority groups, security staff, and institutions are affected. Main risks include privacy loss, misidentification, discrimination, false suspicion, and constant monitoring.

Possible recommendations: Recommend accuracy testing across different groups, consent rules, human confirmation, limited use in sensitive areas, clear appeal processes, and strong protection of biometric data.

AI Recommendation Systems and Social Media Manipulation

Project focus: Study how platforms recommend videos, posts, products, advertisements, or news. Explain how recommendation systems can influence attention, opinions, emotions, and behavior.

Ethical analysis: Utilitarianism can compare benefits, such as useful personalized content, with harms, such as addiction or misinformation. Transparency is important because users often do not know why content is recommended. Care ethics can be used to discuss protection of young users and vulnerable groups.

Affected stakeholders and risks: Students, children, social media users, advertisers, content creators, and society are affected. Main risks include filter bubbles, misinformation, manipulation, addiction, reduced attention, and polarization.

Possible recommendations: Recommend content labels, user control over recommendations, time-use warnings, digital literacy, transparency reports, and stronger protection for young users.

AI in Banking and Credit Scoring

Project focus: Study how AI may approve or reject loan applications, credit cards, insurance, or financial services. Explain that financial AI decisions can strongly affect people's life opportunities.

Ethical analysis: Justice and fairness are central because biased financial systems can exclude people unfairly. Transparency is also important because applicants should know why they were rejected. Rights-based ethics applies because people should have equal opportunity and a fair chance to appeal.

Affected stakeholders and risks: Loan applicants, banks, families, small businesses, and regulators are affected. Main risks include unfair rejection, hidden discrimination, lack of explanation, and limited access to financial services.

Possible recommendations: Recommend explainable decisions, fairness audits, appeal systems, human review, bias monitoring, and clear communication with customers.

AI and Creative Work

Project focus: Study how AI tools generate images, music, writing, videos, logos, or designs. Explain the debate about creativity, originality, copyright, and the use of artists' work in training data.

Ethical analysis: Rights-based ethics is important because creators may have rights over their work. Virtue ethics can discuss honesty in giving credit and not pretending AI-generated work is fully human-made. Fairness is relevant because artists and writers may lose income or recognition if their work is copied or used without permission.

Affected stakeholders and risks: Artists, writers, students, designers, companies, teachers, and audiences are affected. Main risks include copyright disputes, unfair use of creative labor, loss of income, and unclear authorship.

Possible recommendations: Recommend labeling AI-generated content, respecting copyright, giving proper credit, using licensed data, and setting clear rules for academic and professional creative work.

AI in Policing and Public Safety

Project focus: Study how AI may be used for crime prediction, surveillance, risk assessment, or resource allocation in policing. Explain why decisions related to safety and law enforcement require strong ethical control.

Ethical analysis: Justice and fairness are central because biased data may lead to unfair targeting of certain communities. Rights-based ethics is important because people have rights to privacy, dignity, and fair treatment. Accountability is required because authorities must justify decisions that affect citizens' freedom or safety.

Affected stakeholders and risks: Citizens, police departments, local communities, courts, and governments are affected. Main risks include discrimination, over-policing, privacy loss, lack of transparency, and reduced public trust.

Possible recommendations: Recommend independent review, public accountability, strict limits on use, bias testing, transparency reports, human oversight, and clear procedures for complaints.

AI Companions and Emotional Dependence

Project focus: Study how AI chatbots can act as friends, tutors, romantic partners, or emotional support companions. Explain that such systems may be useful but can create emotional and privacy risks.

Ethical analysis: Care ethics is highly relevant because emotionally vulnerable users may need protection. Privacy and consent are important because users may share sensitive personal information. Responsibility matters because companies should design these systems safely and avoid exploiting users' emotions.

Affected stakeholders and risks: Children, teenagers, lonely people, students, parents, educators, and technology companies are affected. Main risks include emotional dependence, harmful advice, privacy leakage, social isolation, and manipulation.

Possible recommendations: Recommend age-appropriate design, safety warnings, privacy protection, limits on sensitive advice, human support options, and clear information that the system is not a real human.

9. Dos and Don'ts

Do	Do Not
Understand your topic clearly.	Do not copy directly from websites or AI tools.
Use course concepts in your analysis.	Do not submit only general information without ethical analysis.
Include at least one real-world example.	Do not focus only on technology and ignore ethics.
Discuss both benefits and risks.	Do not use fake references or fake case studies.
Give your own opinion and practical recommendations.	Do not fill slides with long paragraphs.
Use at least 3 to 5 credible references.	Do not depend on only one group member.

10. Final Advice

Keep your project focused. Do not try to cover everything about AI ethics. A strong project explains one problem clearly, connects it to real life, applies ethical frameworks, and suggests realistic solutions.

A good project is not the one with the most difficult topic. A good project is one that shows clear understanding, ethical thinking, and responsible recommendations.