

Risk is a design question, not a panic button.

*Name risks precisely enough to make proportionate
teaching decisions — not sweeping ones.*

Four categories, unequal relevance.

01

IN SCOPE

System limitations and costs

Hallucination, bias, toxic content, privacy, environmental cost.

02

IN SCOPE

Misuse

Deliberate harm: disinformation, aiding prohibited activity.

03

IN SCOPE

Society-wide disruption

Competency loss, access gaps, ownership, labor, concentrated power.

04

ACKNOWLEDGED

Existential risk

Catastrophic, contested, and outside this course's actionable center.

This course concentrates on the risks faculty can directly influence.

A capability becomes a risk in context.

A COURSE HEURISTIC FOR SIZING RISK

Capability

What the model can fluently produce.



Exposure

How directly it touches a real decision.



Consequence

What is at stake if it is wrong.

LOW RISK

An AI recommendation used as the object of student critique.

HIGH RISK

The same recommendation used to guide a real patient or client decision.

*The capability is the same. You design the **exposure, oversight, and verification.***

Limitations require verification, review, and data boundaries.

Hallucination

Plausible, fluent, and false.

Design response: require named-source verification.

Bias

Defaults skew to dominant assumptions.

Design response: disciplinary review of outputs.

Data exposure

Inputs may be retained, depending on tool and settings.

Design response: prohibit sensitive uploads.

WHERE THE COURSE ACTS

Verification habits in Modules 3–4; the Access and Privacy dimensions of the Module 5 audit.

Misuse and overreliance need different responses.

INTENTIONAL MISUSE

Knowing violation or deliberate harm.

- Fabricating disinformation
- Aiding prohibited or unsafe activity
- Passing off AI work as one's own

RESPONSE

Boundaries, documentation, proportionate consequences.

OVERRELIANCE

Good-faith. Automation bias.

- Accepting fluent output without checking
- “Human in the loop” who rubber-stamps
- Skill that quietly stops being practiced

RESPONSE

Verification, adversarial review, independent checks.

Society-wide disruption enters the classroom.

Competency loss

Skills your field certifies may erode.

Unequal access

Paid tiers and time advantage the already-advantaged.

Ownership

Training data, attribution, and authorship blur.

Labor & environment

Human labor and computing costs sit behind model use.

Concentrated power

A few firms govern the defaults you inherit.

Course design can **amplify** these risks or **buffer** them. The Module 5 audit helps you make the call.

A mitigation can create a new risk.

Effectiveness

Does it actually reduce the risk?

Feasibility

Can you sustain it in a real course?

Tradeoffs

What equity or learning cost does it add?

APPLY IT TO FOUR COMMON MOVES

Blanket ban



new risk: may drive use underground and create uneven enforcement.

AI detection



new risk: may disproportionately flag multilingual writers.

Required AI use



new risk: can create access barriers or pressure data disclosure.

“Human in the loop”



new risk: can become a rubber stamp that adds no real check.

YOUR NEXT MOVE

One risk. One response. One tradeoff.

1. For one student task, name the risk that matters most.
2. Choose the design response you will actually use.
3. Name what it might worsen, burden, or leave unresolved.

These three lines become fields on your Module 1 worksheet — and feed your Discipline Statement.



ELABORATION

A defensible response names both the risk it reduces and the tradeoff it creates.