
Artificial Intelligence in the Age of AI

ECE 57000, Fall 2026

David I. Inouye

| **Doing** has become dramatically cheaper

AI can now produce code, prose, analyses, and prototypes that once demanded hours of skilled implementation.

The opportunity is not merely to do the old work faster.

| New tools are first forced into the **old classroom**

New tool	What became cheap	Naive pedagogical response
Calculator	Hand calculation	Ban it and preserve hand arithmetic
Word processor	Revising and reproducing text	Use the computer as a faster typewriter
Internet	Retrieving information	Treat the web as a larger encyclopedia
Generative AI	Producing an artifact	Police AI output or assign one-click artifacts

| Pedagogy eventually shifts the valuable skill **upstream**

New tool	What becomes more valuable
Calculator	Setup, estimation, interpretation
Word processor	Revision and deciding what is worth saying
Internet	Synthesis, source judgment, inquiry
Generative AI	Framing, choosing, critiquing, validating

| AI weakens the old link between **producing** and **learning**

Before generative AI

- Deriving, coding, and writing created productive friction.
- A polished artifact was imperfect evidence of learning, but usually meaningful.

With generative AI

- Strong artifacts can appear without the same underlying mental model.
- The product alone no longer tells us what was learned.

| The redesign must answer **two different questions**

1. How do we **induce learning** when AI removes productive friction?
2. How do we **measure learning** when the finished product is no longer enough?

| The course mission is to teach **principles and discernment**

Teach the principles of AI - and the **discernment** to work effectively in the age of AI.

| **Content** and **context** reinforce each other

Content: Principles of AI

- Enduring ideas organize the course.
- Methods remain vehicles for technical depth.
- Technical intuition is the deliverable.

Context: *The Age of AI*

- Discernment becomes a central human contribution.
- Learning and doing require different evidence.
- Process and verification matter more than polish.

| Methods will serve the **principles**, not the reverse


- We will still study machine learning, deep learning, CNNs, transformers, generative models, reinforcement learning, and agents.
- A method earns time by the intuition it builds and the decisions it clarifies.
- Durable concepts matter more than a fixed list of methods.

| Technical intuition connects **math, assumptions, behavior, and evidence**

When you meet a new AI method, you should be able to ask:



- What problem has been formalized?
- What assumptions and constraints create its behavior?
- What evidence would distinguish success from a convincing failure?
- What changes when the data, objective, or environment changes?

| Has this training process converged?




Discuss with a partner: Has training converged? Why or why not?

| A linear scale can hide orders of magnitude of progress




The objective did not stop near 10^{-1} ; it continued toward 10^{-7} .

| Now the training objective has **actually** flattened



Discuss again: Will further training help, or only waste computation?

| Training convergence does not settle **generalization**



- A flat training objective does not imply that test performance has stopped changing.
- Epoch-wise double descent is a counterexample: in some regimes, test error improves again with more training.

Conceptual illustration after Nakkiran et al., “Deep Double Descent,” ICLR 2020.

| A new learning target requires **new evidence**

Principles, intuition, and discernment must become observable in more than one way.

- Oral or written quizzes make current understanding visible.
- Assignments create focused practice and evidence.
- Exams test cumulative individual reasoning.
- The AI Product tests whether judgment creates stakeholder value.

| The assessment portfolio separates learning from doing

Evidence	Weight	What it is mainly for
AI Product	40%	Creating and validating stakeholder value
Written or oral quizzes	10%	Making the student's mental model observable
Assignments	10%	Practicing and connecting forms of evidence
Exams	40%	Cumulative individual technical reasoning

| We will pilot **oral quizzes** at scale

- Explain reasoning aloud, not only submit a polished answer.
- A fresh question tests whether the mental model transfers.
- The syllabus walkthrough covers the authoritative policies and deadlines.

LIVE DEMONSTRATION

Open the current oral-quiz interface.

| The AI Product asks whether your work **matters to someone**

- Build an **individual** product whose primary use case materially depends on AI.
- Develop a clear stakeholder value claim and test it with evidence.
- Choose one of two **full-credit** paths:
 - an encouraged **real-stakeholder path**; or
 - a **virtual-stakeholder path** using three distinct AI personas and critical synthesis.
- Value, evidence, and judgment matter more than complexity or feature count.

| The AI Product turns **judgment into a process**

1. **Ideation and stakeholder framing:** identify who cares, why, the product idea, a tentative value claim, and the stakeholder path.
2. **Preliminary product:** build enough to learn, gather evidence, obtain structured feedback, and revise.
3. **Final product:** demonstrate value honestly, name limitations, and choose a defensible next step.

Within the AI Product grade: **Milestone 1 = 10%**, **Milestone 2 = 25%**, and the **final deliverable = 65%**.

| The two pilots reveal **different evidence**

Oral quizzes

- Is the student's mental model actually present?
- Can the student explain and reason without assistance?

AI Product

- Can the student identify and create stakeholder value?
- Can the student direct AI, test claims, and revise responsibly?

One makes **learning** visible; the other makes effective **doing** visible.

| Exams remain **cumulative individual reasoning**

- Midterm 1: **Wednesday, October 7, 2026** - 12%
- Midterm 2: **Wednesday, November 11, 2026** - 12%
- Final exam: **December 14-19, 2026**, time and location TBD - 16%
- Each exam permits one 8 1/2 x 11 inch page, both sides, of physically handwritten notes.

| We will now walk through the **syllabus and schedule**

- Open Brightspace for the required course syllabus
- Open the tentative course schedule

The syllabus is **required reading**. Quiz questions may assess syllabus content.

Evacuate to the **assembly area across North Street**


Once out of the building, occupants shall meet at the following
Evacuation/Emergency Assembly Area (EAA) locations:

1. Primary EAA Location (should be outside, in an area away from the building):

Gather on sidewalk away from the building and building entrances across the street by parking garage.

2. Secondary EAA Location (should be inside a nearby building in case of inclement weather):

Purdue Memorial Union or Grissom Hall




Shelter from severe weather in UC B014

SHELTER IN PLACE

During a severe weather incident (such as a tornado warning), occupants can take shelter in a safe location, such as an interior room with no windows, ideally in the lower level of the building.

Severe weather shelter-in-place options in this building include, but are not limited to:

Go to lowest level possible – study lounge in the basement UC B014.



| Complete these **first-week** actions

1. Read the syllabus and tentative schedule.
2. Confirm access to Piazza and Gradescope; check Piazza regularly.
3. Complete the **ungraded prerequisite quiz by Monday, August 31**.
4. Maintain Claude Pro or **ChatGPT Plus** access—the approximately \$20/month plans; wait to buy Colab Pro until it is needed.

| Our shared target is **learning that transfers**

Understand the principles. Make good judgments. Use AI effectively. Build something that creates real value.

The course will challenge all four - and make the reason for each challenge visible.