

AI Fluency Quizzes: Student Guide

ECE 4424 / CS 4824: Machine Learning: Learn by Building

Fall 2026

Instructor: Ming Jin

Virginia Tech

What this quiz format is

An AI Fluency quiz asks you to evaluate a short AI-generated answer about the technical material in the course. The answer will usually mix good observations with a claim that is incomplete, unsupported, or too confident. Your job is not to declare that “AI is bad.” Your job is to decide exactly what the evidence supports and what it does not.

The posted sample quiz is **not for credit**. It shows the structure we may use on future in-class quizzes.

How the two-pass quiz works

1. **Cold audit.** Read only the front. Mark one strong claim, one claim you would verify, and the most consequential leap. Record an initial disposition.
2. **Course instruction.** During the lecture, learn the relevant technical ideas. Keep your original marks visible.
3. **Evidence audit.** Turn the sheet over when instructed. Use the new evidence to audit one claim, complete a short technical check, and request a useful verification artifact.
4. **Revision.** Give a final disposition. You may change your mind, or keep the same verdict for a better reason. Either can demonstrate learning.

A useful answer structure

For the main audit, move through four questions:

Claim	What, exactly, is the AI asserting? Restate it in your own words.
Evidence	Which number, assumption, definition, or missing test bears on that claim?
Action	Does the evidence justify only a comparison, a threshold choice, a pilot, or a real deployment?
Next step	What concrete artifact would reduce uncertainty, and what result would make you stop?

What the quiz assesses

- **Foundational machine learning:** definitions, assumptions, confusion-matrix rates, threshold tradeoffs, calibration, conditional probability, validation, and data scope.
- **Technical judgment:** choosing the right denominator, separating a model score from a real-world guarantee, and recognizing when evidence is too narrow for the proposed action.
- **AI fluency:** tracing claims to evidence, noticing confident but unsupported transitions, asking for falsifiable checks, and retaining responsibility for the final decision.

Practical expectations

- Bring something to write with; a second ink color is helpful but optional. Turn the page only when instructed.
- AI tools may not be used during an in-class quiz unless the instructor explicitly authorizes them or an approved accommodation applies.
- Treat the scenario as a technical decision exercise. The mushroom example is fictional and must never be used for real identification.

AI Fluency Quizzes: Student Guide

ECE 4424 / CS 4824: Machine Learning: Learn by Building

Instructor: Ming Jin

Fall 2026

Virginia Tech

What strong work looks like

A strong response is specific and economical. It usually does four things:

- names a particular claim rather than criticizing the entire answer;
- cites packet evidence or completes the relevant calculation correctly;
- explains why that evidence changes the proposed action; and
- asks for a test or artifact with a clear stop condition.

You do **not** need to diagnose every sentence. Prioritize the flaw that could cause the most consequential decision error.

Mini-example

Suppose an AI writes: “The classifier is 95% accurate, so it is ready to deploy.” A useful audit would not merely say “unsupported.” It would explain that overall accuracy hides which class was harmed, request confusion counts at the proposed operating threshold, and state what high-cost error would block deployment.

That response connects a foundational concept to an action. It also makes the next AI interaction more productive because the follow-up request is concrete and testable.

How we grade

The main focus is your **thinking process**: whether you identify a consequential claim, connect it to evidence, explain what that evidence authorizes, and make a defensible decision. Open-ended judgment questions may have more than one strong answer. You do not need to match a preferred conclusion if your reasoning is technically sound and well supported.

When a question has a clear right or wrong answer—for example, a calculation, conditional probability, or definition—the answer and supporting work are graded for correctness.

The quiz follows the course’s short in-class 2 / 1 / 0 policy:

- **2 points**: clear, evidence-linked reasoning and correct work when a question has a single answer;
- **1 point**: relevant progress with a missing reasoning link or a technical error; and
- **0 points**: absent or materially blank.

Your initial cold judgment is not graded for being “right.” What matters is making your reasoning visible and improving it when the evidence warrants revision.

Why the quiz is designed this way

AI answers can be fluent even when the evidence-to-conclusion step is weak. The two-pass design captures your first reaction, then lets you use the lecture’s technical ideas to reassess it. This checks understanding more directly than recognizing an option in a multiple-choice list, and it treats revision as evidence of learning rather than as a mistake.

The paper format also matters: your original marks, later evidence, and final decision remain visible together. That creates a small record of how your reasoning changed.

Best preparation: understand the lecture’s core concepts well enough to ask, “What evidence would have to be true before I could rely on this conclusion?”