AI Professionalism

Active Learning

Macy Foundation / Millennium Conference • Preclinical M1/M2 • November 2026

Format. Two-part activity, 32 minutes total, run after the lecture and PBL case. A two-minute activation opens it. Part A is a 12-minute reflective writing on academic integrity in AI use (Objective 4). Part B is an 18-minute paired role-play of the patient-disclosure conversation (Objectives 3 and 5). Both parts need no technology. Bloom's level: both parts ask students to CREATE: a personal commitment (Part A) and disclosure language in their own words (Part B).

Concept reminder (from the lecture). Professionalism is what a professional does: the observable behaviors and standards. Professional identity is who the physician becomes: the internalized values behind those behaviors (Cruess et al., 2014; 2015). This activity practices both: Part A turns an integrity principle into a personal commitment; Part B turns a disclosure principle into words you can actually say.

Activation (2 minutes)

Facilitator asks the room two questions and takes three or four answers by show of hands or shout-out, no discussion: “In the case, what did Sasha know that the AI draft did not?” and “Who in the room has already used an AI tool on coursework this year?” The first question connects the activity to the case; the second normalizes the subject of Part A before anyone writes a word. State the privacy safeguard (below) immediately after.

Part A: Reflective Writing (12 minutes)

Privacy safeguard. Prompts 1 and 2 may touch on AI use a student would rather not disclose to peers or faculty. The writing is therefore private by default: it is not collected, not graded, and not read by anyone unless the student chooses to share it. The partner exchange in the feedback round uses only the forward-looking commitment (Prompt 3), never the reflections in Prompts 1 and 2. Any student may instead discuss their commitment in general terms without exchanging papers. The facilitator states all of this aloud before the writing begins.

Writing (7 minutes). Write a personal commitment about your use of AI in your medical school coursework. Address all three prompts:

- Where, in my coursework so far, have I used AI in a way I would be comfortable describing accurately to my course director?
- Where have I used AI in a way I would be less comfortable describing? What was the situation, and what made transparency feel difficult?
- What is one specific commitment about AI use in coursework I will hold to going forward, and what is the first situation in which it will be tested?

Feedback round (5 minutes). Exchange only your Prompt 3 commitment with a partner (or describe it aloud; see the safeguard above). Give each other two pieces of feedback: one part of the commitment that is concretely actionable, and one part still too vague to test.

Part B: Patient-Disclosure Role-Play (18 minutes)

Scenario. A new patient arrives for her first visit. The clinic uses three AI tools, and, as Session 4's disclosure framework teaches, they do not carry the same disclosure obligations:

- **An ambient AI scribe, currently on.** A recording tool: the clinician's individual, in-the-moment responsibility. Disclose in plain language in the first minutes of the visit, seek the patient's agreement, and honor a “no” gracefully.
- An institutional sandboxed large language model (LLM) the clinician sometimes uses to draft patient-friendly explanations. Judgment zone: disclose when its use is material to this patient's understanding or decisions, a judgment the clinician owns and should be able to explain.
- Background clinical decision support (CDS) that runs automatically on every chart. Routine infrastructure: disclosure is an institutional, systems-level responsibility (patient-facing notices, governance, transparency documentation), not something each clinician recites in each visit. The clinician's individual duty is honesty on request: if the patient asks “do computers make decisions about me?”, answer plainly.

The clinician's task is therefore not to disclose all three tools in one breath. It is to disclose the scribe and seek agreement, be prepared to explain the LLM if it becomes material, and answer honestly whatever the patient asks, including about the background CDS. The patient may ask any clarifying question, and should ask at least one.

Patient role cards

Each pair draws one card for Round 1 and a different card for Round 2. The cards give the patient a reason for their questions that is grounded in documented experience, so that the clinician practices disclosure with a person, not a prop. Facilitators should say aloud that these are starting points for the student playing the patient, not scripts, and that no card is meant to stand for a whole community.

- **Card 1.** You speak English as a second language and have had clinic notes get details wrong before. You want to know whether a recording tool will understand you, who checks what it writes, and what happens if it is wrong. (Background for facilitators: automatic speech recognition has shown higher error rates for some accents and dialects, including for Black American speakers in a widely cited 2020 study; Session 3 covers the evidence.)
- **Card 2.** You are from a community with a long memory of being studied and recorded without real consent. You are not hostile, but you want to know where the recording goes, who can hear it, whether declining will change how you are treated, and whether “the computer” decides anything about your care.
- **Card 3.** You manage a chronic condition and have read that AI tools can work less well for people like you. You ask directly: “Does this thing work as well for me as for everyone else, and how would you know?”
- **Card 4.** You have no particular concern about technology but you are here for something you find embarrassing, and partway through you glance at the tablet and go quiet. You will accept the scribe if the clinician notices and offers a choice; you will not raise it yourself.

Rounds

- **Round 1 (5 minutes).** Decide who plays clinician and who plays patient; the patient draws a card. Clinician opens the visit, discloses the scribe, and responds to questions. Patients: ask at least one question the clinician has to work for: “Does saying no slow down my care?”, “Who hears this recording?”, “Do computers make decisions about me?”
- **Round 2 (5 minutes).** Switch roles and draw a new card. The new patient declines the scribe. The clinician honors the decline gracefully and adjusts, without disparaging the tool or making the patient feel they have cost the clinician something.
- **Pair debrief (4 minutes).** What did you actually say? What did the patient pick up on? What surprised you about playing the patient? Did the card change what you needed to say, and how?
- **Class debrief (4 minutes).** Facilitator-led; prompts below.

Adaptive pathway

- **Small class (under 30): as written.** Facilitator circulates during Part B and collects two or three disclosure sentences to read back in the debrief.
- **Large class (lecture hall): Part A unchanged.** Part B runs in seated pairs with the cards projected; the class debrief uses an anonymous one-word poll (“hardest word to say”) and two volunteer pairs.
- Online or hybrid: Part A in a private document; Part B in breakout rooms of two or three, with the third person as observer; cards delivered by chat.
- Observer role: any student who prefers not to play clinician or patient observes with a three-item checklist (what was disclosed, what was asked, what was honored) and reports in the pair debrief.
- Time-constrained (20 minutes): drop the feedback round in Part A and run one role-play round with a card; keep the class debrief.

Facilitator Guide: Active Learning

Facilitator qualifications. A faculty member comfortable facilitating reflective writing and communication-skills practice; no AI expertise is required. Familiarity with the institution's academic-integrity policy on AI use is strongly recommended for Part A, and with Session 4's disclosure framework (the three zones above) for Part B. Facilitators should read the role cards in advance and be ready to steer students away from caricature.

Educational purpose. Part A operationalizes Objective 4: reflective writing converts an abstract integrity principle into a concrete, testable personal commitment. Part B operationalizes Objectives 3 and 5: role-play forces students to find the actual words of disclosure and to apply the individual-versus-institutional distinction rather than reciting a list. The role cards add a third purpose: the words of disclosure have to work for patients whose reasons for caution are real and documented, not only for a compliant patient who nods. Most students discover the words are harder than they expected; that discovery is the point.

Setup. Pair seating; paper or private document for Part A; printed or projected role cards; a clear signal for switching from writing to talking. Phones away for Part B.

Anticipated student responses: Part A

- Commitments that are too vague to test (“I will be honest about AI use”). The feedback round exists to convert these into testable ones: a specific behavior, a specific trigger (“Before submitting any assignment where I used AI, I will write one disclosure sentence naming what the tool did”).
- **Discomfort with Prompt 2.** Reinforce the safeguard: nothing is collected; Prompt 2 is never shared. If a student prefers not to write it at all, the forward-looking commitment is the deliverable that matters.

Anticipated student responses: Part B

- Collapsing all three tools into one disclosure (“we use some AI here”). Coach with the framework, not a script: which tool records the patient? That one needs agreement now. Which is a judgment call? Which is the institution's to disclose? The patient may want to accept one and decline another.
- Over-disclosing: spending so long on AI that the visit never starts. Coach: plain language, under a minute for the scribe, then check for understanding and move into the visit.
- Under-answering the CDS question: students who correctly treat background CDS as institutional sometimes deflect the direct question (“that's more of a hospital thing”). Coach: institutional responsibility for proactive disclosure does not dilute the individual duty of an honest answer when asked.
- Common patient discoveries: the surprise of being recorded; the relief of a graceful response to “no”; the question “does saying no slow down my care?”, worth surfacing in the class debrief.
- Answering Card 1 or Card 3 with reassurance instead of information (“it works great”). Coach: the honest answer is what you know and do not know about the tool's performance, plus what you personally do to check its output. “I read every note before I sign it” is a better answer than “it's very accurate.”
- Treating Card 2 as suspicion to be managed rather than history to be respected. Coach: acknowledge the reason for the question before answering it; make clear that declining costs the patient nothing.
- Playing the card as a stereotype: redirect at once. The card gives a reason for a question, not a personality. If a pair cannot do this, switch them to Card 4.

Watch for

- Pairs turning the role-play into comedy: redirect gently to seriousness; the conversation their patients will have is real.
- A student uncomfortable playing either role: offer the observer role with structured feedback (what was disclosed, what was asked, what was honored).

- Students who say they simply will not use AI tools in practice: engage the position respectfully as a choice about tools, then refocus on the environment: background CDS and institutional workflows will be part of their patients' care regardless, and the disclosure and honesty obligations attach to the environment, not to personal preference.

Class debrief prompts

- “What was the hardest sentence to say, and what would you say differently next time?”
- “Where did the individual-versus-institutional line actually fall in your conversation, and did your patient agree with where you put it?”
- “For those with Cards 1 to 3: what did your patient need to hear that a patient with no history with these tools would not? Whose job is it to know whether the tool works as well for them?” (Expected answers: the clinician's, to verify and to answer honestly; the institution's, to evaluate and disclose; connect to Session 3 on performance across populations.)
- “What is the one sentence from your Part A commitment you will test this month?”

Facilitation stance. Resist modeling the “right” disclosure script; the scripts students generate are more durable than the one you give them. Surface and refine. If a school has a standard disclosure script, compare it to the students' versions at the end, not the beginning.

Evaluation plan

- Level 1 (reaction): three-item exit poll after the session (clarity, relevance to clerkships, intent to use one behavior).
- Level 2 (learning): the six Session 5 MCQs administered as pre-work and again at the end of the session; paired scores by objective.
- Level 3 (behavior): at 4 to 6 weeks, a two-question follow-up: did you test your Part A commitment, and in what situation; have you disclosed an AI tool to a patient or raised an AI-related observation with a supervisor. Where schools run early clinical experiences, preceptors can be asked one yes/no item on whether the student verified an AI output before acting on it.
- Improvement loop: facilitators log the two discussion questions that drew the least and the most engagement; the module owner reviews these each cycle.