

THE AI SAFETY PAUSE™

Faculty Implementation Toolkit

Integrating AI Competency Into Nursing Education

Three Sessions. No New Course Required.

Developed by

Susan Deane, EdD, MSN, CNE

Co-founder, Nurses for AI™

Based on AI Competency Domains for Nurses:

A Framework for Education, Practice, and Leadership

© 2026 Nurses for AI™. For non-commercial educational use.

PURPOSE OF THIS TOOLKIT

This toolkit is designed for nursing faculty who want to integrate AI competency into their existing courses without adding a new course to an already-packed curriculum.

It provides everything you need to deliver three structured sessions that introduce nursing students to the AI Safety Pause™ — a clinical reasoning tool for evaluating AI-generated recommendations at the point of care.

The sessions are designed to thread into courses you already teach: clinical reasoning, patient safety, evidence-based practice, or professional nursing. No AI expertise is required to facilitate them.

What This Toolkit Includes

- Three ready-to-teach session plans with learning objectives, facilitation guides, and timing
- Five clinical scenarios with AI Safety Pause application guides
- Student handout: The AI Safety Pause quick reference card
- Discussion prompts and reflection questions for each session
- Assessment rubric for evaluating student AI reasoning
- Faculty preparation guide (what you need to know before Session 1)
- Data and evidence brief for department chairs and curriculum committees

Design Principles

- **Integration, not addition:** These sessions fit inside existing courses, not alongside them
- **Clinical reasoning first, technology second:** The AI Safety Pause extends clinical judgment — it doesn't replace it with technical knowledge
- **Faculty-friendly:** You do not need to be an AI expert. If you can teach clinical reasoning, you can teach this
- **Scaffolded:** Sessions build from awareness (Session 1) to application (Session 2) to integration (Session 3)

FACULTY PREPARATION GUIDE

Before facilitating Session 1, spend 30–45 minutes reviewing this section. It covers what you need to know about AI in clinical nursing workflows — not to become an expert, but to feel confident facilitating student discussion.

What Faculty Need to Know

AI is already in clinical workflows. Students will encounter AI-powered tools in their first clinical placements: sepsis prediction alerts, clinical deterioration early warning scores, fall risk assessments, pressure injury prediction models, and ambient documentation tools. These are not future technologies — they are current practice.

The problem is not the technology. The problem is that no one is teaching nurses how to evaluate AI output critically. Most nurses treat AI-generated recommendations the way they treat physician orders: follow it, document it, move on. The AI Safety Pause provides a structured alternative.

You do not need technical AI knowledge. The AI Safety Pause is a clinical reasoning framework, not a technology framework. If you can teach students to question a lab value that doesn't match the clinical picture, you can teach them to question an AI alert that doesn't match the patient in front of them. The cognitive process is identical.

Key Data Points for Context

- Only 41% of nurses use AI for work, compared to 57% of doctors (Elsevier, 2026)
- Only 42% of nurses trust AI tools today, yet 61% believe AI will improve care in the next decade (Elsevier, 2026)
- 68% of clinicians report inadequate AI training (Elsevier, 2026)
- 71% of nurses want more involvement in AI tool design and implementation (McKinsey, 2026)
- Trust is the #1 barrier to AI adoption among nurses (McKinsey, 2026)
- The ANA's first AI in Nursing report (2026) calls for nurses to participate in AI governance, evaluation, and implementation

SESSION 1: THE AI SAFETY PAUSE — INTRODUCTION

Recommended Course: Foundations of Nursing, Clinical Reasoning, or Professional Nursing

Time Required: 50–60 minutes

Format: Lecture/discussion with student reflection

Learning Objectives

- Describe at least two ways AI-generated recommendations currently appear in clinical nursing workflows
- Explain why AI output requires the same critical evaluation as any other clinical data source
- Apply the three AI Safety Pause questions to a clinical scenario

Session Outline

Opening (10 min): Ask students: “If a machine told you your patient was going into sepsis, what would you do?” Allow open discussion. Most students will say they would initiate the sepsis bundle. Some will say they would assess the patient first. Use this to introduce the tension: AI recommendations are not orders, but most clinicians treat them that way.

Content (20 min): Introduce the three AI Safety Pause questions with brief explanation of each:

- Question 1: Does this make clinical sense? — Trust your training. If something feels off, it probably is.
- Question 2: What might the AI be missing? — Context, history, the patient in front of you.
- Question 3: Who is accountable for this decision? — The algorithm doesn’t have a license. You do.

Share 2–3 data points from the Faculty Preparation Guide to establish why this matters now. Emphasize that AI tools are already in clinical settings students will enter.

Application (15 min): Present Scenario 1 (Sepsis Alert). Walk through the AI Safety Pause as a class. Model the reasoning process aloud.

Reflection (10 min): Students write a brief reflection (1 paragraph): “Describe a situation in clinical practice where you might need to question what a technology tells you. What would you do?”

Facilitation Tips

- Normalize uncertainty. Tell students: “You are not expected to know how AI works. You are expected to know how to think critically about what it tells you.”
- Avoid framing AI as good or bad. Frame it as a clinical data source that requires evaluation, like any other.

- If students ask technical questions you can't answer, say: "That's a great question. The point of the AI Safety Pause is that you don't need to know how the algorithm works to evaluate whether its recommendation makes clinical sense."

SESSION 2: APPLYING THE AI SAFETY PAUSE IN CLINICAL SCENARIOS

Recommended Course: Medical-Surgical Nursing, Critical Care, Pharmacology, or any clinical course

Time Required: 50–60 minutes

Format: Small group case-based learning

Learning Objectives

- Apply the AI Safety Pause to evaluate AI-generated recommendations across multiple clinical contexts
- Identify what contextual information AI models typically lack in clinical decision-making
- Articulate the chain of accountability when AI is part of a clinical decision pathway

Session Outline

Review (5 min): Brief recap of the three AI Safety Pause questions. Ask 2–3 students to paraphrase each question in their own words.

Small Group Work (30 min): Divide students into groups of 4–5. Each group receives a different scenario from the Scenario Bank (Scenarios 2–5). Groups work through the AI Safety Pause for their scenario and prepare a brief report-out:

- What did the AI recommend?
- What did you decide, and why?
- What was the AI missing?
- Who is accountable for this decision?

Report-Out and Discussion (15 min): Each group presents their scenario and reasoning. Faculty facilitates cross-group discussion: What patterns did you notice across scenarios? Where did the AI get it right? Where did it fall short? What role did clinical context play?

Closing (5 min): Introduce the concept of “distributed governance” — the idea that every nurse who interacts with an AI tool is part of the governance loop, not just the people on the AI committee. The AI Safety Pause is how governance reaches the bedside.

Facilitation Tips

- Resist the urge to tell groups the “correct answer.” The goal is clinical reasoning, not a single right response. Multiple reasonable decisions are possible.
- If a group decides to follow the AI recommendation, that’s fine — as long as they can articulate why it made clinical sense, what it might be missing, and who is accountable.

- Encourage students to disagree with each other. The discussion is the learning.

SESSION 3: FROM CLASSROOM TO BEDSIDE — INTEGRATION AND ADVOCACY

Recommended Course: Leadership, Professional Issues, Community Health, or Capstone

Time Required: 50–60 minutes

Format: Discussion, reflection, and action planning

Learning Objectives

- Evaluate how organizational AI governance structures include or exclude nursing perspectives
- Develop a personal plan for applying the AI Safety Pause in clinical practice
- Articulate why nursing's participation in AI governance is a patient safety issue, not just a professional development issue

Session Outline

Opening Discussion (15 min): Present the Elsevier data: 41% of nurses say their views are rarely or never part of organizational AI decisions, compared to 19% of doctors. Ask: Why does this gap exist? What are the consequences? What would need to change?

The Six Competency Domains (15 min): Briefly introduce the six AI Competency Domains for Nurses as the broader framework the AI Safety Pause is built on:

- **AI Literacy** — knowing what a sepsis prediction model can and cannot do
- **Data Stewardship** — recognizing when incomplete documentation affects AI accuracy
- **Clinical Decision Support** — evaluating whether to act on, question, or override an AI recommendation
- **Ethics and Legal** — understanding informed consent when AI contributes to care decisions
- **Quality and Safety** — reporting patterns of false-positive alerts as a systems issue
- **Leadership and Advocacy** — requesting a seat on the organization's AI governance committee

Personal Action Plan (15 min): Students write a brief plan (half page) answering three questions:

- How will I use the AI Safety Pause in my first clinical position?
- What is one question I will ask about AI tools during a job interview?
- What would I do if I disagreed with an AI recommendation and my supervisor told me to follow it?

Closing (10 min): Facilitated discussion on the third question — this is where clinical judgment, professional advocacy, and organizational culture intersect. End with: “The algorithm doesn’t have a license. You do. That’s not just a line — it’s your professional responsibility.”

CLINICAL SCENARIO BANK

Each scenario is designed for use with the AI Safety Pause. Scenarios are mapped to common clinical courses and can be adapted to your program's specific clinical contexts.

SCENARIO: The 3 AM Sepsis Alert (Med-Surg / Critical Care)

A 68-year-old patient is post-operative day 2 following a hip replacement. At 3 AM, the AI sepsis prediction model fires a high-risk alert. The patient's temperature is 38.1°C, heart rate is 94, and WBC is slightly elevated. The patient is resting comfortably and had an uneventful post-op course.

Q1 — Does this make clinical sense? Post-surgical inflammatory markers are expected. The vital signs are borderline but not alarming. The clinical picture does not yet match sepsis.

Q2 — What might the AI be missing? The AI does not know the patient just had surgery, that these labs are trending from a higher baseline, or that the patient appears clinically well.

Q3 — Who is accountable? If the nurse initiates a sepsis bundle based solely on the alert, who ordered it? If the patient suffers harm from unnecessary antibiotics, who bears responsibility?

Teaching Point: The AI Safety Pause does not tell the nurse to ignore the alert. It creates space for the nurse to assess, contextualize, and decide whether to act, monitor, or escalate — rather than reflexively initiating a protocol.

SCENARIO: The Fall Risk Override (Med-Surg / Geriatrics)

An 82-year-old patient is flagged as low fall risk by the AI-powered risk assessment tool. However, the nurse notices the patient is unsteady when transferring to the chair, has a new medication (sedative) added to their MAR this morning, and seems confused about where they are.

Q1 — Does this make clinical sense? The AI says low risk, but the nurse's assessment says otherwise. Clinical observation trumps algorithmic scoring.

Q2 — What might the AI be missing? The AI risk model may not have updated for the new sedative. It cannot observe gait instability or assess orientation in real time.

Q3 — Who is accountable? If the patient falls and the nurse relied solely on the AI assessment, the nurse's professional judgment is what will be evaluated — not the algorithm's.

Teaching Point: This scenario illustrates that AI tools can be wrong in clinically significant ways. The nurse's bedside assessment is the safety net the system depends on.

SCENARIO: The Deterioration Score That Doesn't Match (Critical Care / ICU)

A patient in the ICU has a clinical deterioration score that reads “stable.” But the nurse has noticed subtle changes over the last two hours: the patient is slightly more anxious, their skin is clammy, and they’re asking for water more frequently. Vitals are technically within normal limits.

Q1 — Does this make clinical sense? The numbers say stable. The nurse’s pattern recognition says something is changing. Early deterioration often presents with subtle signs before vital signs change.

Q2 — What might the AI be missing? The AI model relies on documented vital signs and lab values. It cannot sense anxiety, clamminess, or behavioral changes that experienced nurses recognize as early warning signs.

Q3 — Who is accountable? If the nurse dismisses their own clinical instinct because the algorithm says “stable,” and the patient deteriorates, the failure is not technical — it is a failure to preserve clinical judgment.

Teaching Point: This is the Benner novice-to-expert concept in action. Expert nurses recognize patterns before data confirms them. The AI Safety Pause protects that expertise from being overridden by algorithmic confidence.

SCENARIO: The Medication Recommendation (Pharmacology)

An AI clinical decision support tool recommends a dosing adjustment for a patient’s anticoagulant based on recent lab values. The recommended dose is within standard parameters. However, the nurse knows the patient has been non-compliant with dietary restrictions (eating large amounts of vitamin K-rich foods) and that the lab was drawn shortly after a missed dose.

Q1 — Does this make clinical sense? The dose adjustment is technically correct based on the lab value. But the lab value itself may not reflect the patient’s true therapeutic status.

Q2 — What might the AI be missing? The AI does not know about dietary non-compliance or the missed dose. It is calculating based on a data point that the nurse knows is incomplete.

Q3 — Who is accountable? If the dose is adjusted and the patient experiences a bleeding or clotting event, the nurse who administered it will need to account for their clinical reasoning.

Teaching Point: AI models optimize based on available data. When the nurse knows the data is compromised, the AI Safety Pause creates the space to intervene before the recommendation becomes an action.

SCENARIO: The Discharge Readiness Score (Community Health / Transitions of Care)

An AI tool calculates a patient’s discharge readiness score as “high” based on clinical criteria: vital signs stable, labs normalized, ambulating independently. The nurse, however, knows the patient lives alone, has no reliable transportation to follow-up appointments, expressed

confusion about their new medications, and has limited health literacy.

Q1 — Does this make clinical sense? Clinically ready to discharge is not the same as safe to discharge. The AI measures physiological readiness but not psychosocial readiness.

Q2 — What might the AI be missing? Social determinants of health, health literacy, medication comprehension, and support systems at home — none of which are captured in the AI's scoring model.

Q3 — Who is accountable? If the patient is readmitted within 48 hours because they could not manage their care at home, the discharge decision will be scrutinized — not the algorithm that said they were ready.

Teaching Point: This scenario extends the AI Safety Pause beyond clinical parameters into the holistic assessment that defines nursing practice. It demonstrates why nursing judgment encompasses dimensions AI models do not measure.

STUDENT HANDOUT

Print and distribute to students. May also be posted in simulation labs and clinical sites.

THE AI SAFETY PAUSE™

Before you act on AI output, ask:

1 Does this make clinical sense?

Trust your training. If something feels off, it probably is.

2 What might the AI be missing?

Context, history, the patient in front of you.

3 Who is accountable for this decision?

The algorithm doesn't have a license. You do.

Clinical judgment was never meant to be outsourced.

ASSESSMENT RUBRIC: AI CLINICAL REASONING

Use this rubric to evaluate student performance on AI Safety Pause application in case studies, simulation debriefs, or written reflections.

Criteria 1: Clinical Sense-Making (Question 1)

Exceeds Expectations: Student identifies specific clinical indicators that support or contradict the AI recommendation and articulates a reasoned position.

Meets Expectations: Student recognizes whether the AI recommendation aligns with the clinical picture and provides basic rationale.

Developing: Student accepts or rejects the AI recommendation without clinical reasoning.

Criteria 2: Contextual Awareness (Question 2)

Exceeds Expectations: Student identifies multiple contextual factors the AI cannot access (social, behavioral, historical, environmental) and explains how they affect the recommendation's validity.

Meets Expectations: Student identifies at least one factor the AI might be missing and explains its relevance.

Developing: Student does not consider what information the AI lacks.

Criteria 3: Accountability Reasoning (Question 3)

Exceeds Expectations: Student articulates the professional, legal, and ethical dimensions of accountability when AI contributes to clinical decisions, including documentation and escalation pathways.

Meets Expectations: Student identifies who is accountable and recognizes that the nurse retains professional responsibility regardless of AI input.

Developing: Student attributes accountability to the AI system or does not address accountability.

Criteria 4: Decision Quality

Exceeds Expectations: Student arrives at a well-reasoned clinical decision that integrates AI output with clinical assessment, contextual factors, and professional judgment.

Meets Expectations: Student makes a reasonable clinical decision and can articulate the reasoning behind it.

Developing: Student follows or rejects the AI recommendation without demonstrating clinical reasoning.

DISCUSSION PROMPTS AND REFLECTION QUESTIONS

Session 1 Discussion Prompts

- Have you ever seen a technology give a recommendation that didn't match what you observed clinically? What did you do?
- What's the difference between following a physician's order and following an AI recommendation? Should there be a difference?
- If an AI tool tells you a patient is low risk but your gut says something is wrong, what do you do? Why?

Session 2 Discussion Prompts

- In your scenario, what would have happened if the nurse had simply followed the AI recommendation without pausing?
- What kinds of patient information does AI typically have access to? What kinds does it not?
- If you override an AI recommendation and the patient has a bad outcome, how do you defend that decision? What if you follow the AI recommendation and the patient has a bad outcome?

Session 3 Discussion Prompts

- Why do you think nurses are less likely than physicians to be included in technology decisions at their organizations?
- What would you want to know about an AI tool before using it on your patients?
- If your first employer has an AI tool you've never been trained on, what steps would you take before using it?

Written Reflection Prompts

- Describe a scenario where the AI Safety Pause would change the outcome for a patient. Walk through all three questions.
- In 250 words, explain why "the algorithm doesn't have a license" matters for nursing practice.
- Compare how you would evaluate a lab value that doesn't match the clinical picture with how you would evaluate an AI recommendation that doesn't match. What's similar? What's different?

EVIDENCE BRIEF FOR DEPARTMENT CHAIRS AND CURRICULUM COMMITTEES

Use this one-page summary to make the case for integrating AI competency into your nursing program.

The Problem

AI-powered clinical decision support tools are now embedded in nursing workflows across health systems nationwide. Nursing students entering practice will encounter sepsis prediction algorithms, deterioration scoring models, fall risk assessments, and ambient documentation tools from their first clinical rotation. Yet nursing curricula have not kept pace.

The Evidence

- Only 41% of nurses use AI for work, compared to 57% of doctors (Elsevier, 2026)
- Only 42% of nurses trust AI tools today (Elsevier, 2026)
- 68% of clinicians report inadequate AI training (Elsevier, 2026)
- 71% of nurses want more involvement in AI tool design and implementation (McKinsey, 2026)
- The ANA's 2026 AI in Nursing report calls for nurses to participate actively in AI evaluation, governance, and implementation
- A 2026 systematic review recommends integrating AI competencies throughout nursing curricula rather than standalone AI courses (SAGE, 2026)
- AACN has called for curriculum redesign to include AI competency development (2026)

The Solution

The AI Safety Pause™ Faculty Implementation Toolkit provides a structured, evidence-informed approach to integrating AI competency into existing nursing courses. It requires three sessions, no new course development, no faculty AI expertise, and no additional technology. It builds on clinical reasoning skills students are already developing and extends them to AI-generated recommendations.

The Framework

The toolkit is grounded in the AI Competency Domains for Nurses (Deane & Koyfman, 2026), the first structured competency framework designed specifically for nursing practice in AI-enabled environments. The book has held the #1 New Release position in Nursing Trends on Amazon since its April 2026 launch.

The Ask

Integrate the three AI Safety Pause sessions into one existing clinical or professional nursing course as a pilot. Evaluate student outcomes using the included assessment rubric. Share findings with the curriculum committee to inform broader integration across the program.