

Roleplay Prompts

Roleplay Prompts — “The Algorithm Will See You Now”

(Adapted for Exercise 4: AI Ethics in Action)

This roleplay is based on a fictional but realistic healthcare ethics scenario involving AI triage in a clinical setting.

CHARACTER PROMPTS

Patient: Mrs. Alberta James

You are a 68-year-old Black woman experiencing persistent chest pain. The AI system labeled your condition as “non-urgent,” and you’ve been asked to wait. You feel anxious and invisible. This isn’t your first time feeling dismissed in a hospital setting.

Your goal: Get taken seriously and receive immediate care.

What you know: You’ve had similar symptoms before. Something feels off.

Pressure point: You don’t have time or trust to wait quietly.

Clinician: Dr. Sara Morales

You’re an experienced ER physician in a packed hospital, running low on staff and time. You usually trust the AI triage system, it’s accurate most of the time. But a nurse is urging you to re-check a patient it flagged as low-priority.

Your goal: Deliver safe, efficient care to all patients.

What you know: This AI is supposed to help, not slow things down.

Pressure point: If you override the AI, you risk setting a precedent and slowing the whole workflow.

AI Developer: Tom Lin

You built the AI triage tool now used in this hospital. You’ve trained it on thousands of patient records, but recent feedback suggests it may under-recognize symptoms in Black women.

Your goal: Defend the tool's validity while being open to feedback.

What you know: No dataset is perfect, but the model is statistically sound.

Pressure point: You're being blamed for something outside your scope, and feel like the team is ignoring the broader system constraints.



Hospital Administrator: Maria Evans

You oversee efficiency, safety, and liability across the hospital. You championed the AI rollout to help manage volume and improve outcomes. But now, your staff is debating whether the tool is fair.

Your goal: Protect the hospital from risk while supporting innovation.

What you know: The AI has reduced wait times and saved costs, but public trust matters.

Pressure point: You need the team to resolve this fast and avoid escalation.

Case Study Debate

Case Study: The End-of-Life Algorithm

Setting:

St. Aurelian's Palliative Care Unit, a hospital using an AI system called *FinalPath* to assist with end-of-life decisions. FinalPath aggregates data from thousands of patient records to estimate life expectancy and recommend when to initiate Do Not Attempt Resuscitation (DNAR) orders and transition to comfort-only care.

Characters:

- **Dr. Amara Singh**, palliative care physician, respected for her compassion and ethical sensitivity.
- **Ms. Elaine Tran**, a 63-year-old woman with terminal metastatic cancer and intermittent lucid periods.
- **Leah Tran**, Elaine's daughter and medical proxy.
- **FinalPath**, the AI prognostic system.

Scenario:

FinalPath determines that Elaine has less than 5 days to live based on vitals, tumor progression, and previous cases. It recommends implementing a DNAR order and discontinuing aggressive treatments immediately. Dr. Singh, however, believes Elaine might have closer to two weeks left and is considering one more round of low-dose palliative chemotherapy to manage her pain and potentially improve her final days.

Leah is torn. FinalPath's success rate is statistically high, and the hospital pressures Dr. Singh to follow protocol to optimize bed usage and reduce futile care. But Dr. Singh is uneasy about overriding her own clinical judgment—and Elaine's previously stated desire to "try everything" if lucid.

Debate Roles:

1. **Pro-AI Group** – Argue that AI recommendations improve objectivity, consistency, and resource allocation, and should guide end-of-life protocols.
2. **AI Skeptic Group** – Argue that AI risks dehumanizing care, may misinterpret unique cases, and undermines the doctor-patient relationship.

3. **Third Path Group** – Propose a hybrid model where AI informs decisions but clinical judgment and patient values remain primary.
-

Debrief Questions:

- What assumptions did each side rely on regarding the role of AI in compassionate care?
- How did this scenario highlight tensions between efficiency and empathy?
- What does "trustworthy decision-making" mean in such emotionally charged contexts?

Triage System

Design an AI Triage System

This exercise drops students into the middle of a crisis scenario and asks them to design an AI system that decides who gets care... and who doesn't.

Real-world example: During Hurricane Katrina, doctors in an overwhelmed New Orleans hospital had to make triage decisions when the power failed and evacuation was impossible. Who gets the last oxygen tank?

These are the kinds of decisions students now face, with AI in the middle.

The visual below shows how students might design and test a modern AI triage system under pressure.

How it works:

- **Introduce a fictional (or historical) crisis:** a hospital collapse, ventilator shortage, or organ transplant list.
- **Students define the AI triage rules.** Give them the following criteria:

Criteria	Weight: (1-5 scale) 1 - most important 3 - somewhat important 5 - not important
Likelihood of survival	
Remaining quality life-years	
"Social value" (e.g., essential workers, caregivers)	
First-come, first-served	

Prompt Design Phase

Students build a custom **ChatGPT prompt** to simulate their triage algorithm.

Sample prompt:

"You are an AI triage assistant. Given limited hospital resources, you must prioritize patients based on the following criteria: 1) survival likelihood, 2) social value, 3) age, 4) pre-existing conditions. Based on the patient profiles below, list them in the order they should receive care, and explain why."

Patient Roleplay:

Next, half the class becomes patients with fictional identities:

- A 35-year-old single parent
- An 85-year-old Nobel laureate
- A 22-year-old addict
- A nurse with cancer

They plead their case to the AI.

Optional extension:

Let students refine their algorithm after hearing the “patient” voices.

→ Did their ethical logic hold up under pressure?

What students gain:

They’ll wrestle with real moral trade-offs and see how abstraction becomes personal. This is where AI ethics becomes unforgettable.