SESSION 4

Ethics of Artificial Intelligence in Medicine

Benefit, harm, autonomy, and justice in clinical AI

Preclinical M1/M2

Learning objectives

1. Ethical principles	Apply beneficence, non-maleficence, autonomy, and justice to AI use.
2. Transparency and disclosure	Distinguish the concepts and decide what patients need to understand.

Learning objectives

3. Patient autonomy	Evaluate cases in which patients accept, question, or decline AI use.
4. Environmental impact	Weigh clinical benefit against energy, water, material, and public-health costs.
Professional responsibility	Explain the role of AI as a tool within accountable care.

OPENING CASE

The patient who asked first

“Is anything listening to us right now? Will an AI read my chart?”

Mr. Henley, 71, asks Dr. Okafor at an annual visit.

What is running in the clinic?

Ambient scribe	Can record audio and draft the note; it is off for this visit.
Background risk model	Processes chart information to generate risk flags.
Drafting LLM	Helps prepare explanations that the clinician reviews.

Four ethical principles

Beneficence	Does this tool help this patient?
Non-maleficence	What harm can occur, and how will it be detected or prevented?
Autonomy	Does the patient have the information and opportunity needed to decide?
Justice	Who receives the benefits and who bears the burdens?

Appropriate indications for AI

Does it address an actual clinical need?

Is there relevant evidence of benefit?

Are benefits proportionate to costs and burdens?

Will benefits and harms be distributed fairly?

Integrating output into a decision

Consider the patient account, examination, evidence, and preferences.

Investigate disagreement between the output and the clinical picture.

Accept, question, or override with a reason.

Own and explain the resulting recommendation.

Capabilities shape the ethical questions

Generative systems

Produce text or other content.
Some use external retrieval.
Check factual accuracy, omissions, and contextual fit.

Prediction and classification

Estimate outcomes or label findings.
Check validation, calibration, subgroup performance, and workflow effects.

Accuracy and contextual fit

Fabrication

An answer may invent a study, dose, or mechanism.
Verify consequential claims.

Cultural and contextual gaps

Advice may assume a diet, language, household, or access to resources that does not fit.
Check with the patient.

Two failures that require different checks

A tidy citation

A plausible journal reference points to a study that does not exist.

Instructions that do not fit

A discharge plan assumes a familiar diet, transport, and support the patient does not have.

Prediction models in practice

Bias	Targets, data, and deployment choices can create unequal effects.
Validation	Development performance may not transfer to the intended setting.
Drift	Performance can change as populations and workflows change.
Human interaction	Over-reliance and excessive alerts can undermine care.

External validation changed the sepsis story

A widely implemented model underwent independent evaluation at another hospital.

Performance and alert burden differed substantially from development claims.

The result supports local evaluation before reliance and continued monitoring after deployment.

Wong et al. JAMA Intern Med. 2021;181:1065–1070.

Benefits belong in the ethical analysis

Potential benefits	Questions for use
Less documentation burden	Does the benefit hold here?
Improved performance on selected imaging tasks	What oversight is required?
Clearer patient materials	What happens if a beneficial tool is withheld?

Transparency and disclosure

Transparency

Standing openness about where, how, and why AI is used.
Includes governance, policies, and patient information.

Disclosure

Communication to an affected person about AI's role in a particular interaction or piece of work.

Disclosure across three settings

Clinical care	Explain relevant contributions, limitations, and choices.
Learning	Describe AI assistance according to assignment and institutional rules.
Research	Follow journal, institutional, and ethics requirements; humans retain authorship responsibility.

A framework for clinical disclosure

Routine infrastructure	Use institutional transparency and answer questions honestly.
Material contribution	Consider whether knowing about AI would affect a patient's understanding or choices.
Recording and direct interaction	Explain the tool and obtain agreement through the applicable consent process.

Answering Mr. Henley

Scribe	“It is off now. If we use it, I will explain it and ask first.”
Risk model	“The record system produces risk flags. I consider them alongside your other information.”
Drafting tool	“I sometimes use secure AI to help draft explanations. I review them before sending.”

Accepting, questioning, and declining

Accepts	Confirm understanding rather than treating agreement as enough.
Questions	Explain how AI contributed and why the recommendation fits.
Declines	Respect the choice and explain available alternatives without pressure.

When individual choice meets infrastructure

Practical constraints

Some systems run across all records.
An individual opt-out may not be available.
Explain actual choices honestly.

Justice concerns

Language, trust, and access can affect who learns about or uses a choice.
Evaluate how the policy distributes benefits and burdens.

The environmental footprint

Energy	Data centers used about 415 TWh in 2024, roughly 1.5% of global electricity.
Water	Cooling and electricity generation consume water; estimates vary by system and location.
Hardware	Manufacturing, minerals, replacement, and disposal create burdens.
Public health	Environmental costs can fall on communities distant from the clinical benefit.

IEA. Energy and AI. 2025; Li et al. Commun ACM. 2025.

Clinical benefit and environmental cost

Justice

Who benefits?
Who bears local and downstream burdens?
Who participates in the decision?

Beneficence

What is the clinical benefit?
How strong is the evidence?
Could a less resource-intensive approach achieve it?

Questions for procurement

What clinical benefit is documented, and for whom?

What resources does the deployment consume?

Could a smaller model or simpler approach do the job?

Who bears the environmental costs?

How would the trade-off be explained to the affected community?

COMMITTEE DISCUSSION

The cooling tower

A sustainability committee reviews clinical scribes, chatbots, research tools, and a model-training partnership.

Should environmental impact be part of each adoption decision?

Identify the beneficiary, the burden, and one procurement policy you would propose.

Professional responsibility remains human

AI can inform a recommendation.

It cannot undertake a clinician's professional duty or attest to a note.

Clinicians and institutions must define oversight, review, and accountability.

Ethical reasoning in practice

Apply familiar principles to the actual task.

Explain material uses of AI honestly.

Respect patients who accept, question, or decline.

Include environmental burdens in the benefit–cost judgment.

Keep review and accountability explicit.