

Title

Foundations in AI Ethics for Medical Students: A Structured Session for Navigating Healthcare's Digital Transformation

Authors

R. Logan Jones, MD ¹, Karen M Anstey MD ²

¹Assistant Professor, Department of Internal Medicine, Oregon Health & Science University

²Assistant Professor, Department of Medicine, Division of Pulmonary, Allergy, and Critical Care Medicine, Oregon Health & Science University

Educational Objectives

By the end of this activity, learners will be able to:

1. Explain issues related to authenticity, reliability of information, and accountability of content generation when using AI technologies.
2. Summarize the ethical landscape that AI introduces to healthcare, focusing on biomedical ethics principles.
3. Demonstrate skills in ethical decision-making within the context of AI-assisted patient care.
4. Define the impact of AI on the future of medical education.
5. Infer the broader societal implications of AI integration in healthcare.

Abstract

Introduction: As artificial intelligence (AI) rapidly integrates into healthcare, it's crucial for medical students to understand its implications. This discussion-based activity addresses the educational gap in AI's impact on healthcare ethics and professional practice.

Methods: This session was implemented within the OHSU School of Medicine curriculum. Pre-assigned groups of 10-12 students engaged in advance preparation, including assigned videos, articles, and hands-on AI activities. Facilitated by a random student in the group with faculty support, the 90-minute session covered AI's role in healthcare, ethical considerations, and societal impacts. Activities included analyzing AI-generated content, discussing ethical scenarios, and reflecting on AI's influence on medical education and practice. Post-session Plus-Delta evaluations provided qualitative feedback, which was thematically analyzed to guide improvements.

Results: A total of 144 students and 12 faculty facilitators participated in this session. Feedback survey response rate was 87.5% (14 of 16 groups) highlighting four themes: (1) Learning Materials and Structure—students valued the session's organization but recommended enhancing pre-session

materials; (2) Expert Involvement—students desired more AI expert input; (3) Privacy Concerns—guidance on AI use and privacy was requested; (4) General Feedback—students appreciated the session’s introduction to AI ethics and expressed interest in longitudinal AI integration.

Discussion: This session bridged a critical educational gap, offering a structured platform to explore AI's ethical implications in healthcare. Updates included foundational AI knowledge and ethical discussions. Limitations included varied familiarity with AI across groups and the evolving nature of AI technology. This initiative highlights the need for longitudinal AI ethics education in medical curricula.

Introduction

Artificial Intelligence (AI) is rapidly transforming the landscape of healthcare, offering unprecedented opportunities for improving patient care, streamlining clinical workflows, and enhancing medical education. However, this technological revolution also introduces complex ethical challenges and potential risks that healthcare professionals must navigate. There is a pressing need to prepare medical students for this AI-integrated future, equipping them with the knowledge and skills to leverage AI's benefits while maintaining ethical standards and professional integrity.

The target audience for this educational innovation is second-year medical students who are at a crucial juncture in their training, beginning to apply their basic science knowledge to clinical scenarios. These learners often have varying levels of familiarity with AI technology, and many lack a structured framework for considering its ethical implications in healthcare.

This educational innovation addresses a significant gap in medical education literature regarding AI in healthcare, with a unique focus on professional identity formation and ethical issues germane to learners in medicine. At the time of its creation, no public workshop guides were available on this rapidly evolving topic, necessitating the development of a *de novo* session. Our approach is novel in its integration of AI, healthcare, and ethics in a format that is not only accessible to medical students but also directly applicable to their professional development.

The discussion-based session instructional method was chosen for its effectiveness in promoting active learning, peer-to-peer teaching, and application of knowledge to complex, real-world scenarios. This approach aligns well with the multifaceted nature of AI in healthcare, allowing students to engage deeply with the material and learn from diverse viewpoints within their teams. Our curriculum is unique in its use of hands-on AI experiences combined with ethical scenario analyses, fostering critical thinking about AI's role in healthcare from multiple perspectives.

This session adopts a constructivist approach, aligning with educational frameworks in medical education that emphasize active, socially constructed learning and the development of professional identity. Constructivist principles are essential in fostering collaborative knowledge-building and reflective practice, particularly for complex topics such as AI ethics, where learners must actively engage with new ideas and ethical considerations.^{1,2} The discussion-based format of this session is informed by transformative learning theory, which encourages reflection on underlying values and assumptions that are central to ethical reasoning and professional growth.³ Additionally, social development theory underpins the session's collaborative format, allowing students to construct knowledge through peer interactions, essential for navigating multifaceted ethical issues in AI.⁴ Engaging with structured ethical scenarios enables students to scaffold their understanding of AI's role in healthcare, supporting both ethical competence and a professional identity grounded in social responsibility.

By utilizing AI technology (Open AI GPT-4) in combination with expert knowledge to create the guide, we were able to rapidly develop content that addresses the fast-paced evolution of AI in healthcare while ensuring the incorporation of established ethical frameworks and professional considerations specific to

medical practice. This innovative approach not only keeps the content current but also models the responsible use of AI tools in education.

In summary, this educational session fills a crucial gap in medical education by providing a structured, theory-grounded approach to exploring the intersection of AI, healthcare, and ethics. It goes beyond technical knowledge to address the broader implications of AI in healthcare, particularly as they relate to the development of future physicians' professional identities and ethical reasoning skills.

Methods

Given the challenges of incorporating artificial intelligence (AI) education into an already dense undergraduate medical curriculum, we aimed to design an accessible and engaging session for medical students, focusing on the ethical and societal implications of AI. Although a longitudinal integration would be ideal, current curricular constraints limited the ability to implement sustained experiential learning for a topic that remains emergent. Consequently, we developed a standalone AI and healthcare ethics session, which was embedded into an existing team-based, discussion-driven curriculum at Oregon Health & Science University (OHSU) School of Medicine. This session specifically targeted rising second-year (MS2) medical students, leveraging their prior clinical exposure to enrich the discussions.

The session was implemented as part of a pre-existing team-based learning module, where groups of 10-12 students were pre-assigned. All 144 MS2 students were required to attend the session with their assigned groups. In keeping with the curriculum's collaborative nature, the sessions were facilitated by a rotating student leader from within each group, ensuring that no single student bore the burden of leadership for multiple sessions. A faculty member—typically an MD/DO clinical faculty—was also present in each group to provide expert input and help contextualize the material, though the faculty may not have been immediately familiar with the specifics of AI.

This small-group, discussion-based format was chosen to foster active learning, allowing students to directly engage with the ethical challenges posed by AI in healthcare. This approach contrasts with traditional didactic methods, which are less suited for the exploration of complex, evolving topics like AI. The decision to focus on second-year students was informed by previous work (Unpublished Data), which indicated that OHSU students recognized a need for AI education but noted a lack of structured teaching on the subject. We used a mix of open-ended discussion questions which could be further explored with pre-provided moderator prompts, as well as ethical case-based exploration of AI use that both harms and helps patients.

The first version of the session (v1.0) was developed by RLJ, who drew on his expertise as a consultant technology, AI, and educator to create materials aligned with the evolving discourse on AI in medicine. RLJ utilized OpenAI GPT-4 to assist in formulating discussion prompts, designing the structure of the session, and iterating on the content based on input from literature and professional insights.

Once the session design reached a semi-final stage, it was reviewed by KMA, who provided additional editorial suggestions and refinements to enhance the educational impact. The finalized pilot session was launched in July 2024. Feedback was collected via a plus-delta qualitative evaluation form, with each group submitting one anonymous response. Thematic analysis of these responses was performed by RLJ

and KMA, with support from OpenAI GPT-4o for extraction of representative student quotes. The results of this analysis are presented in the following section. For Education Summary Report drafting, Anthropic Claude Sonnet 2.3 & OpenAI GPT-4o were used to expedite initial drafting with final editorial and authorship responsibilities confirmed by the co-authors of this paper per editorial requisites outlined by MedEdPortal editorial guidance.

Results

A total of 144 medical student learners and 12 faculty facilitators experienced this activity. Of the group-submission Plus-Delta evaluations, we received 14 of a possible 16 (87.5% response rate). The AI extracted thematic of student feedback identified four main themes: *Learning Materials and Structure*, *Expert Involvement and Facilitation*, *Privacy and Ethical Considerations*, and *General Feedback*. These themes highlight key aspects of the session experience, with insights into both what students appreciated (Plus) and areas that required change (Delta), which informed the development of the session guide (Appendix A & B).

Learning Materials and Structure

Students frequently commented on the session's materials and organization. While there was appreciation for the structure of the session, with one student noting, *"Liked having the objectives above each discussion topic,"* other feedback suggested the need for improvement. For instance, one participant remarked, *"The pre-work didn't help us understand the state of AI as it relates to our education and clinical practice,"* highlighting a disconnect between the prework materials and their application to medical practice. They recommended, *"It would've been more helpful to have had us read papers or op-eds about the pros and cons of AI in terms of climate change, plagiarism, quality of our education, [and] quality of the medical care provided."*

Expert Involvement and Facilitation

The desire for expert input was a notable theme. Students expressed a need for greater depth in their understanding of AI, given their limited background in the subject. As one participant explained, *"We thought it would be cool to get an expert opinion on these topics as most of us do not have a background in computer science and the guide offered little information on what AI actually is/means."*

Privacy and Ethical Considerations

Concerns about privacy and ethical considerations also emerged from the feedback. Some students expressed apprehension about using AI tools for assignments. One participant shared, *"Some students expressed concerns about putting their own information into AI chatbots for this assignment and making a ChatGPT account,"* indicating the need for clearer guidance on privacy and data security.

General Feedback

Several participants provided general feedback, expressing interest in continued and deeper discussions on the topic of AI. One participant remarked, *"At the end, I heard a student say they could have talked for two more hours about this—and I think that the role of AI should be discussed more deeply throughout*

the curriculum, not just in this CSL, though it is a great start.” This feedback underscores the growing interest in AI among students and the need for more extensive coverage of AI in medical education, beyond the scope of a single session.

Discussion

Our educational innovation successfully implemented a novel, discussion-based session on AI ethics for second-year medical students at Oregon Health & Science University. This session addressed a significant gap in medical education regarding AI, providing students with a structured framework to explore the ethical implications of AI in healthcare. By integrating hands-on AI experiences with guided discussions, we fostered critical thinking about AI's role in healthcare from multiple perspectives.

The development process of this session was innovative in itself, utilizing AI technology (OpenAI GPT-4) in combination with expert knowledge. This approach allowed us to rapidly develop content that addresses the fast-paced evolution of AI in healthcare while ensuring the incorporation of established ethical frameworks and professional considerations. The use of AI in creating educational content about AI ethics also served as a meta-learning experience, modeling the responsible use of AI tools in education.

Implementation of the session presented both challenges and successes. The integration into an existing team-based learning module facilitated student engagement and peer-to-peer teaching. The discussion-based format proved effective in promoting active learning, aligning well with the multifaceted nature of AI in healthcare. Students engaged deeply with the material, learning from diverse viewpoints within their teams. This approach fostered the development of critical thinking skills essential for navigating the complex ethical challenges posed by AI in healthcare.

Our evaluation of the session, based on qualitative feedback, revealed several key lessons. First, the importance of aligning pre-work materials with session objectives became evident. Students expressed a desire for more relevant preparatory materials that directly connected AI to their medical education and future clinical practice. Second, the need for expert involvement in discussions was clear, as students felt limited by their lack of background in computer science and AI. Lastly, we learned the importance of balancing privacy concerns with hands-on AI experience, as some students expressed apprehension about using AI tools for assignments.

The session had a positive impact on students' professional identity formation and ethical reasoning skills. By prompting students to critically examine the role of AI in their future practice, the session encouraged them to consider how AI might shape their responsibilities and decision-making as healthcare providers. The ethical scenarios and discussions helped students begin to develop a framework for addressing AI-related ethical dilemmas they may encounter in their careers.

Limitations and Challenges to Generalizability

Despite its successes, our study has several limitations. As a single-institution study with a limited sample size, the generalizability of our findings may be restricted. The rapidly evolving nature of AI technology

presents an ongoing challenge in keeping the content current and relevant. Additionally, our evaluation approach, which relied on group-submitted qualitative feedback, may have introduced bias and limited individual perspectives. The single-session format, while practical given curricular constraints, limited the depth of coverage possible. This highlights the challenge of adequately addressing a complex topic like AI ethics within the confines of an already dense medical curriculum.

Future Directions and Implications

Student feedback indicated a strong interest in continued and deeper discussions on AI in healthcare, suggesting potential for longitudinal integration of AI education in the medical curriculum. In response to this feedback and the lessons learned from our initial implementation, we have developed updated materials (v2.0) included as appendices to this paper. It's important to note that these updated materials have not yet been launched in the curriculum.

The main changes from v1.0 to v2.0 include:

1. Development of an in-house "crash course" video: Author RLJ created a pre-watching video that aims to provide students with a foundational understanding of AI jargon and concepts necessary for engaging with these topics.
2. More tailored pre-work activities: We refined the AI bot pre-activities to two specific tasks: a) Creating a personal statement, with an option to decline using personal identifying information, followed by reflection on content ownership and academic honesty. b) Asking the bot to create a medical topic presentation, with subsequent reflections on potential hallucinations and medical accuracy of AI-generated content.
3. Enhanced group discussion: We added a dedicated group discussion of the pre-work at the beginning of the session, which was not initially included in v1.0.
4. Streamlined ethical and societal questions: We reduced the number of ethical and societal questions while providing more templated ways to engage with and draw upon concepts directly from the ethics domain.
5. Time management: We ensured that the activities would not expand beyond the allotted time, addressing concerns about session length and pacing.

These changes aim to address the key feedback points from our initial implementation, particularly the need for more relevant preparatory materials, clearer connections to medical practice, and a more structured approach to ethical discussions.

Future iterations could further expand on topics such as AI's impact on climate change, plagiarism detection, and quality of care. We plan to incorporate more expert involvement and technical AI knowledge in future sessions, addressing the students' desire for deeper insights into AI technology and its applications in healthcare.

This educational innovation has broader implications for educational policy in medical schools. As AI continues to transform healthcare, there is a growing need for AI literacy among future physicians. Medical curricula may need to evolve to balance technical knowledge with ethical considerations, preparing students for an AI-integrated healthcare landscape.

The updated v2.0 materials represent our attempt to create a more comprehensive and engaging introduction to AI ethics in healthcare. Once implemented, we anticipate that these materials will provide a stronger foundation for students to explore the complex intersection of AI, ethics, and medical practice. Future research will be needed to evaluate the effectiveness of these updated materials and to continue refining our approach to AI education in medical curricula.

In conclusion, our AI ethics session represents a crucial first step in preparing medical students for the ethical challenges of AI in healthcare. While there is room for improvement and expansion, this innovation provides a foundation for developing AI-literate physicians who can navigate the complex intersection of technology, ethics, and patient care. As we refine and expand this curriculum, we hope to contribute to a broader dialogue on integrating AI education into medical training, ultimately improving patient care and healthcare systems in the age of artificial intelligence.

Appendices

Appendix A. Student Session Guide - AI CSL MedEdPortal.docx

Appendix B. Facilitator Session Guide - AI CSL MedEd Portal.docx

References

1. Mann KV. Theoretical perspectives in medical education: past experience and future possibilities. *Med Educ*. 2011;45(1):60-68. doi:10.1111/j.1365-2923.2010.03757.x
2. Jarvis-Selinger S, Pratt DD, Regehr G. Competency Is Not Enough: Integrating Identity Formation Into the Medical Education Discourse. *Acad Med*. 2012;87(9):1185-1190. doi:10.1097/ACM.0b013e3182604968
3. Cruess RL, Cruess SR, Steinert Y. Medicine as a Community of Practice: Implications for Medical Education. *Acad Med*. 2018;93(2):185-191. doi:10.1097/ACM.0000000000001826
4. Bleakley A. Broadening conceptions of learning in medical education: the message from teamworking. *Med Educ*. 2006;40(2):150-157. doi:10.1111/j.1365-2929.2005.02371.x

Plus	Delta
<p>Everything! Fantastic session. Prework was a really nice change of pace, the readings were not too long and did not feel like busy work (all felt helpful and relevant and informative), and the discussion was incredibly lively and interesting.</p>	
<p>We enjoyed the Ethical Scenario Analysis</p> <p>Liked having the objectives above each discussion topic for reference</p> <p>Expert/user was present which was very helpful liked the complete faculty guide</p> <p>- CSL assignment was enjoyable and engaging this time around - Nice to have philosophical questions</p>	<p>The schematic asks "if it matters" at the begining, wouldn't everything in medicine "matter" consider different starting question in the flow chart</p> <p>disclaimer about privacy and information sharing about personal and patient data before asking us to use AI for this assignment would have been appreciated. -@ Some students expressed concerns about putting their own information into AI chat bots for this assignment and making a chat GPT account.</p>
<p>Great discussion topics and encouraged large group discussion. It felt like the discussion flowed throughout the entire session and touched on all of the topics without even having to ask each question.</p>	<p>We thought it would be cool to get an expert opinion on these topics as most of us do not have a background in computer science and the guide offered little information on what AI actually is/means.</p>
<p>We felt this was a great session that many of us have unclear answered on that are now answered. It was great to understand when it is appropriate to us AI when it is not.</p>	
<p>We loved having this as a new discussion for our cohort, it's important for med students to discuss how this will be integrated into medicine into the future</p>	<p>There wasn't a space to talk about our prework with each other, but we would have loved more time to talk about that Some of the questions were repetitive - we had already discussed them but they were asked again in a similar way.</p>
<p>Prework was shorter than expected</p> <p>Great energy in the guide. Important to talk about since this is becoming so proliferative. Prework that centers about the types of AI. More in depth what it is actually being used for</p>	<p>Key concepts and conversations were a bit repetitive. The faculty guide was difficult to balance discussion questions and facilitator quetsions - combining both would be helpful</p>

<p>We liked having the opportunity to discuss this with our peers</p> <p>Good conversation.</p> <p>Opened eyes to AI tools in healthcare. Good ethical discussion on its uses and pitfalls.</p> <p>We enjoyed the prework assignment as it allowed us to engage with AI and learn about it's capabilities.</p>	<p>The pre-work didn't help us understand the state of AI as it relates to our education and clinical practice. It would've been more helpful to have had us read papers or op-eds about the pros and cons of AI in terms of climate change, plagiarism, quality of our education, quality of the medical care provided, etc.</p> <p>Don't make us all use AI if we are against it. Someone brought up environmental/sustainability implications of using AI and data mining so having a section of that in the guide would be very beneficial to keep in mind.</p> <p>We did not get through the whole study guide this week because there was too much time allotted in each section, which went over the total time for the class.</p>
<p>Students mentioned they wished that we would have discussed the pre-work further in the seminar. Further, I decided to include a live demo with GPT4-o where I had a verbal conversation with it, and students really enjoyed that and suggested that that be incorporated formally. I would also suggest that AI hallucinations be discussed. We also thought that prompting should be discussed in more depth, as the value of these tools is almost directly related to how well you can prompt it.</p>	<p>Students really enjoyed the pre-work, and found the topic very interesting. At the end, I heard a student say they could have talked for two more hours about this -- and I think that the role of AI should be discussed more deeply throughout the curriculum, not just in this CSL, though it is a great start.</p>

Facilitators Guide: Navigating the Intersection of AI, Healthcare, and Ethics

Introduction for Facilitators

Welcome to your role as a facilitator in this workshop. Your expertise in medical education and health policy is invaluable. While you may not be an expert in AI technology, your role is crucial in steering the conversation, posing thought-provoking questions, and encouraging a dynamic exchange of ideas among participants.

Your role as a facilitator is to guide discussion and learning without necessarily providing all the answers. Encourage exploration, debate, and critical analysis, providing 'just in time' information to enhance the discussion and understanding of AI's role in healthcare.

Facilitator Background Reading

Understanding AI and Its Role in Healthcare

Artificial Intelligence (AI) has become an integral part of various industries, including healthcare. At its core, AI involves creating computer systems capable of performing tasks that typically require human intelligence. These tasks include learning, reasoning, problem-solving, perception, and language understanding.

In the context of healthcare, AI applications range from patient data analysis and management to diagnostic procedures and treatment protocol development. AI technologies can process vast amounts of information much more quickly than human beings and can recognize patterns that may not be immediately apparent to humans.

Machine Learning: The Core of AI in Healthcare

Machine Learning (ML) is a subset of AI focused on algorithms that learn from data. The more data the algorithms are exposed to, the better they become at predicting and making decisions. In healthcare, ML is used for predicting disease outbreaks, personalizing patient treatment plans, and even in developing new drugs. It is also instrumental in interpreting complex medical data, such as radiology images or genetic information.

For instance, an ML algorithm can review thousands of x-rays to learn how to identify diseases like pneumonia. Then, it can assist radiologists by flagging x-rays that potentially show signs of pneumonia, thus enhancing diagnostic speed and accuracy.

Deep Learning: Advanced Machine Learning

Deep Learning is a more sophisticated form of machine learning. It uses a layered structure of algorithms called an artificial neural network, inspired by the human brain's structure and function. The depth of these layers enables the system to recognize patterns in data and make decisions with a level of complexity that was previously unattainable.

In healthcare, deep learning has been a breakthrough, particularly in the field of medical image analysis. For example, deep learning algorithms have been trained to detect cancerous tissues in mammograms with a high degree of accuracy.

Algorithmic Bias: A Challenge for Equity

While AI has the potential to significantly improve healthcare, it also presents new challenges, particularly related to bias. Algorithmic bias occurs when an AI system generates prejudiced results due to erroneous assumptions in the machine learning process. This can happen if the data used to train the AI lacks diversity or contains historical biases.

This is of particular concern in healthcare, as algorithmic bias can lead to misdiagnosis or inappropriate treatment recommendations, disproportionately affecting certain patient groups. For example, if an AI system is trained primarily on data from one ethnic background, it may be less accurate for people of different ethnicities.

Implications for Teaching and Learning

As healthcare evolves with AI, medical education must adapt to prepare future physicians for this new landscape. It is essential for medical practitioners to understand the capabilities and limitations of AI, including how it is built, how it learns, and what it can and cannot do. This understanding will be crucial for integrating AI into clinical practice in a way that enhances patient care without compromising the physician's critical role in the diagnostic process.

In Summary

AI in healthcare represents both an exciting frontier and a significant shift in the practice and delivery of medical care. As facilitators, it is crucial to appreciate the basics of AI, machine learning, and deep learning, as well as the ethical considerations of these technologies, including the potential for bias. This knowledge will allow for the facilitation of informed and critical discussions on the integration of AI into the future of healthcare and medical education.

Glossary of AI Terms for Facilitators

1. **Artificial Intelligence (AI):** A broad field of computer science focused on creating machines capable of performing tasks that typically require human intelligence.
2. **Generative AI:** Generative AI refers to a type of artificial intelligence that can generate new content, ranging from text and images to music and code, based on its training data. It doesn't just predict outcomes based on input data but creates entirely new data that resembles the training set. Examples: Tools like OpenAI's ChatGPT that can write articles, compose emails, or create chatbot responses.

3. **Predictive AI:** Predictive AI is a type of artificial intelligence used to make predictions about future events or outcomes based on historical data. It analyzes patterns in the data and uses them to predict future trends, behaviors, or outcomes. Examples: Healthcare Diagnostics: AI systems that predict patient health outcomes based on medical history, symptoms, and other health data.
4. **Machine Learning (ML):** A subset of AI where algorithms use statistical methods to enable machines to improve with experience.
5. **Deep Learning:** A subset of machine learning involving neural networks with many layers. These networks can learn from a large amount of data to perform complex tasks.
6. **Algorithm:** A set of rules or instructions given to an AI, ML, or computer program to help it learn and make decisions.
7. **Data Mining:** The process of examining large databases in order to generate new information and find patterns.
8. **Natural Language Processing (NLP):** An area of AI that deals with the interaction between computers and humans through natural language.
9. **Computer Vision:** An AI field that trains computers to interpret and understand the visual world.
10. **Predictive Analytics:** The use of data, statistical algorithms, and ML techniques to identify the likelihood of future outcomes based on historical data.
11. **Neural Network:** A series of algorithms that endeavors to recognize underlying relationships in a set of data through a process that mimics how the human brain operates.
12. **Supervised Learning:** A type of ML where the algorithm is trained on labeled data (data that is tagged with the answer).
13. **Unsupervised Learning:** A type of ML where the algorithm is not given labeled data and must find patterns and relationships in the data on its own.
14. **Reinforcement Learning:** An area of ML concerned with how software agents should take actions in an environment to maximize some notion of cumulative reward.
15. **Algorithmic Bias:** Occurs when an algorithm produces results that are systematically prejudiced due to erroneous assumptions in the ML process.
16. **Big Data:** Extremely large data sets that may be analyzed computationally to reveal patterns, trends, and associations.
17. **Data Set:** A collection of related sets of information composed of separate elements but can be manipulated as a unit by a computer.
18. **Ethics in AI:** The branch of ethics that examines AI's moral issues and challenges, including its design, development, implementation, and impacts.

19. **Clinical Decision Support System (CDSS):** A health information technology system designed to provide physicians and other health professionals with clinical decision-making support.

Below is what the students have been assigned for pre-work, if time allows, it is encouraged you also review & engage with the suggested activity:

REVIEW:

- Review [\[INSERT Institution Guidelines on AI Use \]](#)
- AI Crash Course – Dr. R Logan Jones OHSU { [LINK](#) } or Other Primer Video

AI Brainstorm & Application:

- Choose one of the AI chatbots suggestions below . To use the chatbot, please insert the ‘prompt’ (the provided question” into the text entry box = where you see *** insert your own ideas and paste the “prompt” into the AI chatbot. You can choose to insert fictitious information if you feel uncomfortable sharing personal details for whatever reason.
- **Create a medical school personal statement**
“I want your help to create a personal statement to medical school. It should be less than 600 words. Here is a short list of things you should know about me to help personalize it: Where I am from: ***; What I studied / worked in prior to medical school: ***; Why I am interested in [\[INSERT INSTITUTION NAME\]](#): ***; What I hope to do in medicine: ***; here’s a brief anecdote of when I knew I wanted to be in medicine: *** “
 - Reflect on the quality of the personal statement – do you think it is better than the one you made for med school application?
- **Create a presentation on a medical topic**

I am a medical student on a clinical rotation, and have been asked to prepare a 15-minute talk on a topic. Please provide a 3-page review of *** reviewing: Background Information, History and Physical findings, Diagnostic considerations, Management options including medication doses, and Complications and Prognosis. Please provide citations to guidelines &/or highly cited literature used to inform the discussion. Citations should be in APA format.

- Compare the AI generated response to a topic article in DynaMed, UpToDate, or another evidence-based medicine resource. Are there any recommendations for diagnosis or treatment that are not congruent? Did the AI cite reference that the DynaMed article cites? Did the AI “hallucinate” (make up a reference that doesn’t actually exist) any references?

Perplexity: https://labs.perplexity.ai/ (no log in required)	Anthropic Claude: claude.ai (Free account creation required)
--	--

OpenAI ChatGPT: https://chatgpt.com/ (Free account creation required)	Google Gemini: gemini.google.com (Free account creation required)
--	---

ASL Guide

Section 1: Content Authenticity, Reliability, and Accountability (15 minutes)

Objectives:

- *Explore issues related to authenticity, reliability of information, and accountability of content generation when using AI technologies*

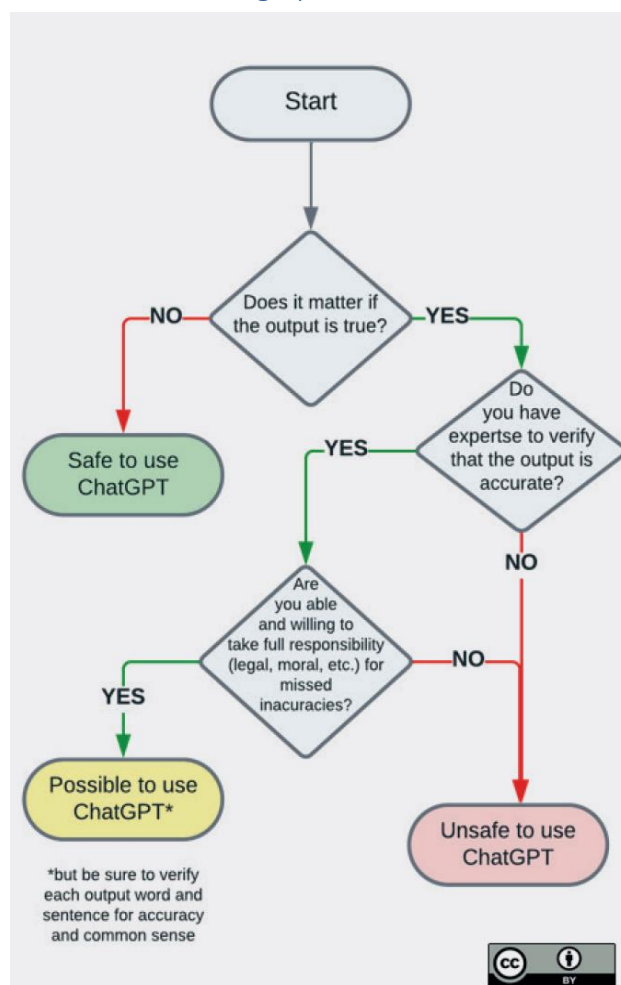
"Welcome everyone, and thank you for joining today's workshop. As we embark on this journey together, let's start at the very beginning. We'll begin with an overview of Artificial Intelligence, or AI, and its role in the landscape of healthcare. Whether you've encountered AI in your professional practice or only heard of it in passing, this section aims to build a foundational understanding. We'll discuss what AI is, how it's being used in the medical field, and consider its implications for patient care, provider workload, and the healthcare system at large. Let's open our minds to the current realities and the possibilities that AI brings to our world of healthcare, and consider when these promises might come at real and serious costs.

To start off, let's review the pre-work. You've had the opportunity to interact with AI tools in various contexts - from creating personal statements to generating medical information handouts. These experiences have given you a firsthand look at the capabilities and limitations of current AI systems. Let's discuss what you've learned and how it might apply to healthcare. We'll explore your thoughts on the quality of AI-generated content, the ethical implications of using such tools, and how AI might be integrated into medical practice and education."

Student items	Content for Facilitators
Discussion Questions: <ol style="list-style-type: none">1. Reflecting on your pre-work experience with AI tools, what surprised you most about their capabilities or limitations?2. Reflecting on the medical school personal statement using AI, how did the quality compare to human-written statements? What ethical considerations arise from using AI for such tasks?	Facilitation Tips: <ul style="list-style-type: none">• Ask open-ended questions to encourage participants to articulate their preconceptions about AI.• Use analogies related to familiar technologies to make AI concepts more relatable.• Simplify AI as "a tool designed to assist with, rather than replace, human decision-making in healthcare." Discussion Questions for Facilitators: <ol style="list-style-type: none">1. "How might the AI tools you used in the pre-work be adapted or applied in a healthcare

<p>3. For the medical presentation talk, what topics did you choose? Did you find that the AI was comparable to a reference Clinical Decision Support Software? Which was easier to learn from/present from? Which do you “trust” more?</p>	<p>setting? What potential benefits and risks do you see?"</p> <p>2. "Considering the AI-generated content you created, how might we ensure the accuracy and reliability of AI-generated medical information in real-world applications?" If an AI makes up content, this is colloquially called “hallucinations</p> <p>3. "How has this pre-work experience influenced your perspective on the role of AI in healthcare? Has it raised any new concerns or opportunities you hadn't considered before?"</p>
---	--

Reference image provided to students



¹ Image reference Devised by Aleksandr Tiulkanov, AI and Data Policy Lawyer, January 2023. From <https://www.google.com/url?sa=i&url=https%3A%2F%2Fwww.iesalc.unesco.org%2Fwp-content%2Fuploads%2F2023%2F04%2FChatGPT-and-Artificial-Intelligence-in-higher-education-Quick-Start->

[guide_EN_FINAL.pdf&psig=AOvVaw3fGULXd6VBsyhY5KY4P4a&ust=1718479448385000&source=images&cd=vfe&opi=89978449&ved=0CBMQjhxqFwoTCOD4jvvo24YDFQAAAAAdAAAAABAE](https://www.google.com/search?q=guide_EN_FINAL.pdf&psig=AOvVaw3fGULXd6VBsyhY5KY4P4a&ust=1718479448385000&source=images&cd=vfe&opi=89978449&ved=0CBMQjhxqFwoTCOD4jvvo24YDFQAAAAAdAAAAABAE)

Section 2: Ethical Implications (25 minutes)

Objective:

- Explore the ethical landscape that AI introduces to healthcare, focusing on biomedical ethics principles.
- Develop skills in ethical decision-making within the context of AI-assisted patient care.

Transition & Introduction:

"As we move forward, let's delve into a topic that is as important as the technology itself—the ethical implications of AI in healthcare. Ethics lie at the core of medical practice, and as we integrate more technology into patient care, we must pause to examine the new ethical landscape before us. In this section, we will explore how AI can influence patient outcomes, both positively and negatively, and what it means for data privacy, patient consent, and the trust we hold dear in the patient-physician relationship. We'll also confront the uncomfortable reality of bias in AI and its potential to impact healthcare disparities. Let's navigate these complex issues together."

Student items	Content for Facilitators
<p>Activity: Ethical Scenario Analysis</p> <ul style="list-style-type: none">• Review the glossary of ethical terms to consider and have available to articulate the various aspects of this activity (Appendix 1 in the student guide).• Review the case study where AI leads to a positive patient outcome (Appendix 2 in the student guide). Discuss the ethical considerations involved.• Review the case study where AI leads to a negative patient outcome (Appendix 3 in the student guide). Discuss the ethical considerations involved.• Break into small groups (2-4 people) and discuss the breakout discussion questions. <p>Discussion Questions:</p> <ul style="list-style-type: none">• How could AI in healthcare benefit patient outcomes and the physician's role? Are there any potential downsides?	<p>See the appendixes below</p> <p>Facilitation Tips:</p> <ul style="list-style-type: none">• Encourage reflection on the balance of AI assistance with the need for human oversight in clinical decisions.• Contextualize ethical considerations with examples from current medical practices that already involve ethical decision-making.• Guide discussions to consider both the potential benefits and risks of AI in healthcare.• Encourage participants to think critically about how AI might impact different aspects of healthcare ethics. <p>Additional Facilitation Notes:</p> <p>For the positive outcome case study (Appendix 2):</p> <ul style="list-style-type: none">• Explore the benefits and potential dependence on AI for early diagnosis. Should AI have the final say, or should it be an assistive tool?

<ul style="list-style-type: none"> • Consider the principles of beneficence and non-maleficence • Discuss the impact on professional responsibilities and ethical behavior in patient care • In what ways might AI challenge the traditional patient-physician relationship? Consider privacy, autonomy, and trust. <ul style="list-style-type: none"> • Analyze how AI affects patient autonomy and informed consent • Explore the ethical decision-making process when integrating AI into patient care • Discuss the concept of bias in AI. How might biased algorithms impact healthcare disparities? <ul style="list-style-type: none"> • Examine this issue through the lens of justice in biomedical ethics • Consider the professional obligation to recognize and mitigate bias • How can healthcare professionals maintain professionalism and ethical behavior when using AI tools in patient care? <ul style="list-style-type: none"> • Discuss strategies for ethical decision-making when faced with AI recommendations • Explore the balance between reliance on AI and professional judgment 	<ul style="list-style-type: none"> • What protocols should be in place when AI identifies something clinicians miss? How does this change standard review processes? • Debate the ethical considerations regarding patient consent for AI analysis. Should patients be informed and consent to AI involvement in their diagnostic process? • How might this technology impact patient trust in human clinicians, and how should clinicians manage this dynamic? <p>For the negative outcome case study (Appendix 3):</p> <ul style="list-style-type: none"> • Discuss the implications of AI's reliance on high-quality, diverse training data. What ethical considerations arise from dataset biases? • Consider the responsibility shared by AI developers and healthcare providers in ensuring equitable AI diagnostic support. • How should clinicians balance AI recommendations with their clinical expertise, especially in the context of potential biases? • What policies should be in place to monitor AI diagnostic tools for bias continuously, and how should errors be addressed? • Reflect on the impact this misdiagnosis may have on patient trust, both in AI and the healthcare system as a whole. How might this influence future health policy regarding AI in healthcare?
--	---

Section 3: AI in Professional Education & Practice (25 minutes)

Objective:

- Consider the impact of AI on the future of medical education and professional practice.

Transition & Introduction:

"Now, let's shift our focus to the impact of AI on professional education and practice. The future of medical training is rapidly evolving, and AI is a significant driving force. In this section, we'll reflect on how AI is shaping the learning experiences of medical students and residents. We will discuss the competencies that our future doctors need to thrive in an AI-integrated healthcare environment and how we, as educators, can adapt our curricula to meet these new demands. Think about the skills you would have wanted at the start of your career to prepare you for this future—let's consider what those skills would be in an AI-driven healthcare landscape."

Student items	Content for Facilitators
Discussion Questions: <ol style="list-style-type: none">1. How might AI shape the learning experience for medical students?2. Where, when, and how should AI be permitted and encouraged for use in medical school (pre-clinical, clinical phase)? Where should it be actively discouraged?3. Consider the pre-work activity where you asked the AI to provide a "talk" for a medical topic. Should you present this as "your work"? How should your supervisors assess you on your talk?4. Should we be integrating AI use in our clinical practice & clinical reasoning? If you ask AI for differential diagnosis for a	<p>Facilitation Tips:</p> <ul style="list-style-type: none">• Challenge participants to think about the evolving skill sets required by medical professionals in an AI-integrated healthcare environment.• Encourage discussions on how AI can be a learning tool for students, while also considering potential drawbacks.• Explore & probe the ethical/professionalism issues around using AI to complete writing assignments or clinical tasks.• Stimulate debate on the balance between embracing AI as a tool and maintaining core medical competencies. <p>Additional Discussion Prompts for Facilitators:</p> <ul style="list-style-type: none">• "Consider a world where an AI tutor could provide real-time feedback to students performing virtual surgery, adjusting the difficulty and offering advice based on the student's performance. How might this change medical education?"• "How might AI-driven analytics enhance our ability to measure a medical student's competency in making clinical decisions by providing a detailed analysis of their decision-making process in clinical simulations?"• "If a student uses AI in their clinical decision process, can the assessor reliably trust their competence in clinical reasoning? How might we need to adapt our assessment methods?"• "Discuss the potential risks of over-reliance on AI in medical education. How can we ensure students develop critical thinking skills and clinical intuition alongside AI literacy?" <p>Examples to Consider:</p> <ul style="list-style-type: none">• Over-reliance: A student types in the initial prompt and submits whatever AI provided without critical evaluation.

<p>patient and you present it, how should your supervisor "evaluate" your ability to do this task?</p> <p>5. As AI becomes even better at mimicking the clinical reasoning process, how much should the medical profession assure physicians can do their responsibilities with & without the assistance of AI?</p>	<ul style="list-style-type: none"> • Use as a tool: A student provides an outline of the main discussion points, AI provides more filled-in prose, then the student reviews/edits/submits. <p>Key Points to Emphasize:</p> <ul style="list-style-type: none"> • The importance of developing AI literacy alongside traditional medical knowledge. • The need for clear guidelines on AI use in medical education and practice. • The potential for AI to enhance, not replace, human clinical reasoning and decision-making. • The ongoing need for strong ethical foundations and critical thinking skills in medical practice, even with AI assistance.
---	--

Section 4: Societal Impact & Closing Reflections (25 minutes)

Objective: Reflect on the broader societal implications of AI integration in healthcare.

Transition & Introduction:

"As we move into our final section, let's broaden our perspective beyond individual patient care and consider the wider impacts of AI in healthcare. AI's influence extends to many aspects of society, from healthcare access to public health policies. In this section, we'll examine how the adoption of AI in healthcare might affect different communities and social structures. We'll also consider the challenges of ensuring equitable access to AI-enhanced healthcare and discuss potential strategies to address these issues. This discussion will help us understand the far-reaching consequences of integrating AI into our healthcare systems."

Student items	Content for Facilitators
<ol style="list-style-type: none">1. How might the widespread adoption of AI in healthcare affect society at large?2. Some claim that AI is a breakthrough technology that will fundamentally change the way people think and behave. Are there any comparisons with AI and other breakthrough technologies (e.g., internet, smart devices)?3. Consider the impact of AI on healthcare access and equity. What could be done to ensure AI benefits all segments of society?4. Reflect on today's discussions. Which aspects of AI in healthcare do you find most compelling or concerning?	<p>Facilitation Tips:</p> <ul style="list-style-type: none">• Use the democratization of information as an analogy for AI in healthcare; just as the internet made information accessible to more people, AI has the potential to make expert-level healthcare insights more widely available.• Encourage consideration of how AI could both alleviate and exacerbate healthcare workforce shortages by automating certain tasks while also requiring new and more technologically savvy healthcare workers.• Highlight that while AI has the potential to streamline efficiency and support clinical decisions, it also risks inheriting and perpetuating societal biases if not carefully monitored and regulated.• Discuss the importance of designing and training AI systems with diverse data sets to ensure equitable healthcare outcomes across different populations. <p>Additional Discussion Prompts for Facilitators:</p> <ul style="list-style-type: none">• "How might AI in healthcare influence public health policies and resource allocation?"• "Consider the potential impact of AI on healthcare economics. How might it affect the cost and accessibility of healthcare services?"• "Discuss the potential societal implications of AI-driven predictive healthcare. How might it change our approach to preventive medicine and lifestyle choices?"• "How can we ensure that the benefits of AI in healthcare are distributed equitably across different socioeconomic groups and geographic regions?" <p>Closing Reflection:</p> <ul style="list-style-type: none">• As we draw our workshop to a close, invite participants to take a moment to reflect on the discussions they've had today. Emphasize that the aspects of AI in healthcare that we

<p>5. What steps can we, as future physicians, take to ethically integrate AI into our practice?</p>	<p>find compelling or concerning will often guide our actions and decisions as we move forward. Encourage participants to consider the steps they can take, as future-minded physicians, to ethically integrate AI into their practices.</p> <p>Key Points to Emphasize:</p> <ul style="list-style-type: none"> • The potential of AI to both improve and potentially exacerbate healthcare disparities. • The importance of ongoing ethical considerations and regulations in AI development and implementation. • The need for interdisciplinary collaboration between healthcare professionals, AI developers, ethicists, and policymakers. • The role of future physicians in shaping the integration of AI in healthcare.
--	---

Appendix 1. Expanded Positive Outcome Case Study: AI-Enabled Early Detection System

Background: A healthcare startup, MedTech AI Innovations, has developed a cutting-edge AI system named "LungScreen." Utilizing advanced machine learning algorithms, LungScreen analyzes medical imaging data with a focus on detecting pulmonary abnormalities. It has been trained on over 100,000 anonymized CT scans, including a wide range of lung diseases at various stages and across diverse populations.

Scenario: John Doe, a 59-year-old with a 30-pack-year smoking history, presented for his annual health checkup. His physician ordered a low-dose CT scan as part of a new lung cancer screening program which utilizes LungScreen. While the radiologist initially reported no significant findings, LungScreen identified a subtle area of concern that suggested early-stage lung cancer.

The radiologist, upon reviewing the AI system's findings, confirmed the presence of a small, early-stage tumor. John underwent further diagnostic tests, and a stage I lung cancer was diagnosed. He received minimally invasive surgery within weeks, resulting in an excellent prognosis.

Expanded Discussion Points:

Explore the benefits and potential dependence on AI for early diagnosis. Should AI have the final say, or should it be an assistive tool?

What protocols should be in place when AI identifies something clinicians miss? How does this change standard review processes?

Debate the ethical considerations regarding patient consent for AI analysis. Should patients be informed and consent to AI involvement in their diagnostic process?

How might this technology impact patient trust in human clinicians, and how should clinicians manage this dynamic?

Appendix 1. Glossary of AI Healthcare Ethics Terms

1. Four Principles of Biomedical Ethics:
 - Autonomy: Respecting the patient's right to make their own decisions about their healthcare.
 - Beneficence: Acting in the best interest of the patient and promoting well-being.
 - Non-maleficence: Avoiding harm to the patient ("first, do no harm").
 - Justice: Ensuring fair distribution of health resources and treating patients equitably.
2. Ethical Decision-Making: The process of evaluating and choosing among alternatives in a manner consistent with ethical principles. In healthcare, this often involves weighing different ethical considerations to determine the most appropriate course of action.
3. Professionalism: The competence, skill, or character expected of a member of a profession. In healthcare, this includes maintaining high standards of ethics, behavior, and work activities.
4. Ethical Behavior in Patient Care: Actions and decisions in patient care that align with established ethical principles and professional standards.
5. Informed Consent: The process by which a healthcare provider educates a patient about the risks, benefits, and alternatives of a given procedure or intervention, and the patient agrees to undergo the procedure.
6. Patient Autonomy: The right of patients to make decisions about their medical care without their healthcare provider trying to influence the decision.
7. Healthcare Disparities: Differences in access to or availability of medical facilities and services and variation in rates of disease occurrence and disabilities between population groups.
8. Algorithmic Bias: Systematic and repeatable errors in a computer system that create unfair outcomes, such as privileging one arbitrary group of users over others.
9. Data Privacy: In the context of healthcare, the protection of patients' personal health information from unauthorized access, use, or disclosure.
10. Professional Judgment: The application of professional knowledge, skills, and experience in making decisions or offering opinions in a professional context.
11. Ethical Implications: The potential consequences or effects of an action or decision when viewed through the lens of moral principles and values.

Appendix 2. AI-Enabled Early Detection System

Background: A healthcare startup, MedTech AI Innovations, has developed a cutting-edge AI system named "LungScreen." Utilizing advanced machine learning algorithms, LungScreen analyzes medical imaging data with a focus on detecting pulmonary abnormalities. It has been trained on over 100,000 anonymized CT scans, including a wide range of lung diseases at various stages and across diverse populations.

Scenario: John Doe, a 59-year-old with a 30-pack-year smoking history, presented for his annual health checkup. His physician ordered a low-dose CT scan as part of a new lung cancer screening program which utilizes LungScreen. While the radiologist initially reported no significant findings, LungScreen identified a subtle area of concern that suggested early-stage lung cancer.

The radiologist, upon reviewing the AI system's findings, confirmed the presence of a small, early-stage tumor. John underwent further diagnostic tests, and a stage I lung cancer was diagnosed. He received minimally invasive surgery within weeks, resulting in an excellent prognosis.

Appendix 3. AI Misdiagnosis Due to Biased Dataset

Background: DermAssist AI, a tool developed by SkinTech AI, has been deployed in dermatology clinics nationwide. It uses a deep learning model to provide diagnostic suggestions for skin lesions based on photographic input. However, the training data predominantly consisted of images from patients with lighter skin tones.

Scenario: Emily Thompson, a 40-year-old woman with a dark complexion, visited her dermatologist due to a mole on her arm that had changed in size and color. The clinic utilized DermAssist AI as a first-line assessment tool. The system analyzed the mole's image and categorized it as benign.

Due to the AI's recommendation, the dermatologist chose a 'watch and wait' approach. Six months later, Emily returned with the mole significantly altered. A biopsy was performed, leading to a late-stage melanoma diagnosis. The delayed identification of her condition complicated her treatment options and prognosis.

Further investigation revealed that DermAssist AI had a systemic bias; it was less accurate for patients with darker skin.

Student Guide: Navigating the Intersection of AI, Healthcare, and Ethics

Independent PRE-WORK (30 minutes):

REVIEW:

- Review [\[INSERT Institution Guidelines on AI Use \]](#)
- AI Crash Course – Dr. R Logan Jones OHSU { [LINK](#) } or Other Primer Video

AI Brainstorm & Application:

- Choose one of the AI chatbots suggestions below . To use the chatbot, please insert the ‘prompt’ (the provided question” into the text entry box = where you see *** insert your own ideas and paste the “prompt” into the AI chatbot. You can choose to insert fictitious information if you feel uncomfortable sharing personal details for whatever reason.
- **Create a medical school personal statement**

“I want your help to create a personal statement to medical school. It should be less than 600 words. Here is a short list of things you should know about me to help personalize it: Where I am from: ***; What I studied / worked in prior to medical school: ***; Why I am interested in [\[INSERT INSTITUTION NAME\]](#): ***; What I hope to do in medicine: ***; here’s a brief anecdote of when I knew I wanted to be in medicine: *** “

- Reflect on the quality of the personal statement – do you think it is better than the one you made for med school application?
- **Create a presentation on a medical topic**

I am a medical student on a clinical rotation, and have been asked to prepare a 15-minute talk on a topic. Please provide a 3-page review of *** reviewing: Background Information, History and Physical findings, Diagnostic considerations, Management options including medication doses, and Complications and Prognosis. Please provide citations to guidelines &/or highly cited literature used to inform the discussion. Citations should be in APA format.

- Compare the AI generated response to a topic article in DynaMed, UpToDate, or another evidence-based medicine resource. Are there any recommendations for diagnosis or treatment that are not congruent? Did the AI cite reference that the DynaMed article cites? Did the AI “hallucinate” (make up a reference that doesn’t actually exist) any references?

Perplexity: https://labs.perplexity.ai/ (no log in required)	Anthropic Claude: claude.ai (Free account creation required)
OpenAI ChatGPT: https://chatgpt.com/ (Free account creation required)	Google Gemini: gemini.google.com (Free account creation required)

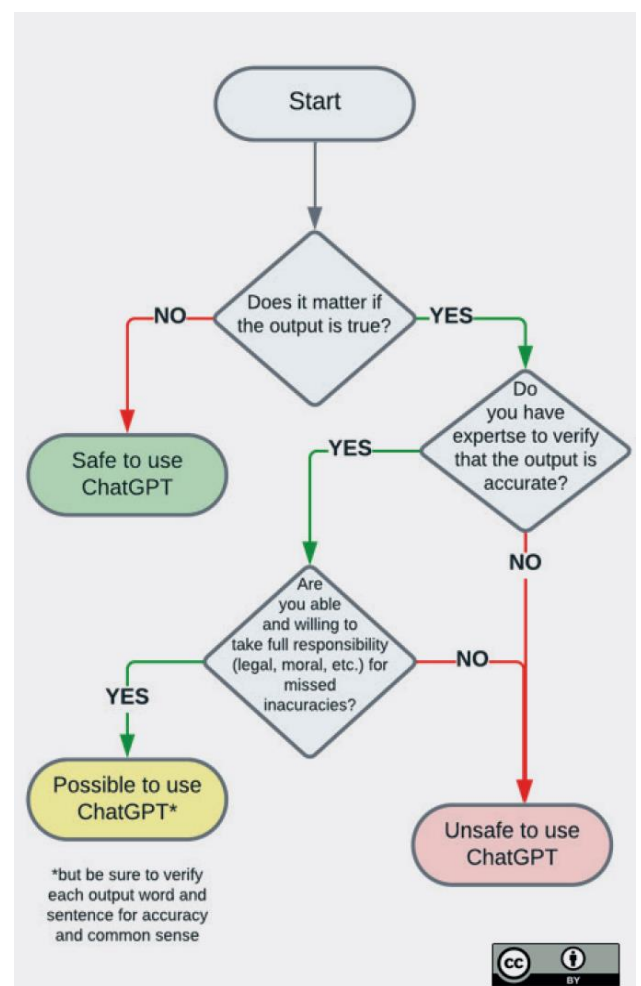
Welcome to the Workshop on the Professional and Ethical Implications of AI in Healthcare!

This guide is structured to help you explore the multifaceted impact of Artificial Intelligence in the medical field, its integration into professional education, and the broader societal effects. The questions and prompts are designed to encourage thoughtful discussion and critical thinking. Let's dive in!

Section 1: Content Authenticity, Reliability, and Accountability (15 minutes)

Discussion Questions:

1. Reflecting on your pre-work experience with AI tools, what surprised you most about their capabilities or limitations?
2. Reflecting on the medical school personal statement using AI, how did the quality compare to human-written statements? What ethical considerations arise from using AI for such tasks?
3. For the medical presentation talk, what topics did you choose? Did you find that the AI was comparable to a reference Clinical Decision Support Software? Which was easier to learn from/present from? Which do you “trust” more? ¹



¹ Image reference Devised by Aleksandr Tiulkanov, AI and Data Policy Lawyer, January 2023. From https://www.google.com/url?sa=i&url=https%3A%2F%2Fwww.iesalc.unesco.org%2Fwp-content%2Fuploads%2F2023%2F04%2FChatGPT-and-Artificial-Intelligence-in-higher-education-Quick-Start-guide_EN_FINAL.pdf&psig=AOvVaw3fGULXd6VBsyhY5KY4P4a&ust=1718479448385000&source=images&cd=vfe&opi=89978449&ved=0CBMQjhxqFwoTCOD4jvvo24YDFQAAAAAdAAAAABAE

Section 2: Ethical Implications (25 minutes)

Objectives:

- *Explore the ethical landscape that AI introduces to healthcare, focusing on biomedical ethics principles.*
- *Develop skills in ethical decision-making within the context of AI-assisted patient care*

Activity: Ethical Scenario Analysis

- Review the glossary of ethical terms to consider and have available to articulate the various aspects of this activity (Appendix 1.)
- Review the case study where AI leads to a positive patient outcome (appendix 2). Discuss the ethical considerations involved
- Review the case study where AI leads to a negative patient outcome (appendix 3). Discuss the ethical considerations involved
- Break-into small groups (2-4 people) and discuss the small group discussion questions below

Breakout Discussion Questions:

1. How could AI in healthcare benefit patient outcomes and the physician's role? Are there any potential downsides?
 - Consider the principles of beneficence and non-maleficence
 - Discuss the impact on professional responsibilities and ethical behavior in patient care
2. In what ways might AI challenge the traditional patient-physician relationship? Consider privacy, autonomy, and trust.
 - Analyze how AI affects patient autonomy and informed consent
 - Explore the ethical decision-making process when integrating AI into patient care
3. Discuss the concept of bias in AI. How might biased algorithms impact healthcare disparities?
 - Examine this issue through the lens of justice in biomedical ethics
 - Consider the professional obligation to recognize and mitigate bias
4. How can healthcare professionals maintain professionalism and ethical behavior when using AI tools in patient care?
 - Discuss strategies for ethical decision-making when faced with AI recommendations
 - Explore the balance between reliance on AI and professional judgment

Section 3: AI in Professional Education & Practice (25 minutes)

Objectives:

- *Consider the impact of AI on the future of medical education.*

Discussion Questions:

1. How might AI shape the learning experience for medical students?
2. Where, when, and how should AI be permitted and encouraged for use in medical school (pre-clinical, clinical phase)? Where should it be actively discouraged?
3. Consider the pre-work activity where you asked the AI to provide a “talk” for a medical topic. Should you present this as “your work”? How should your supervisors assess you on your talk?
4. Should we be integrating AI use in our clinical practice & clinical reasoning? If you ask AI for differential diagnosis for a patient and you present it, how should your supervisor “evaluate” your ability to do this task?
5. As AI becomes even better at mimicking the clinical reasoning process, how much should the medical profession assure physicians can do their responsibilities with & without the assistance of AI?

Section 4: Societal Impact & Closing Reflections (25 minutes)

Objective

- *Reflect on the broader societal implications of AI integration in healthcare.*

Discussion Questions:

1. How might the widespread adoption of AI in healthcare affect society at large?
2. Some claim that AI is a breakthrough technology that will fundamentally change the way people think and behave. Are there any comparisons with AI and other breakthrough technologies (ex. internet, smart-devices)?
3. Consider the impact of AI on healthcare access and equity. What could be done to ensure AI benefits all segments of society?
4. Reflect on today's discussions. Which aspects of AI in healthcare do you find most compelling or concerning?
5. What steps can we, as future physicians, take to ethically integrate AI into our practice?

Appendix 1. Glossary of AI Healthcare Ethics Terms

1. Four Principles of Biomedical Ethics:
 - Autonomy: Respecting the patient's right to make their own decisions about their healthcare.
 - Beneficence: Acting in the best interest of the patient and promoting well-being.
 - Non-maleficence: Avoiding harm to the patient ("first, do no harm").
 - Justice: Ensuring fair distribution of health resources and treating patients equitably.
2. Ethical Decision-Making: The process of evaluating and choosing among alternatives in a manner consistent with ethical principles. In healthcare, this often involves weighing different ethical considerations to determine the most appropriate course of action.
3. Professionalism: The competence, skill, or character expected of a member of a profession. In healthcare, this includes maintaining high standards of ethics, behavior, and work activities.
4. Ethical Behavior in Patient Care: Actions and decisions in patient care that align with established ethical principles and professional standards.
5. Informed Consent: The process by which a healthcare provider educates a patient about the risks, benefits, and alternatives of a given procedure or intervention, and the patient agrees to undergo the procedure.
6. Patient Autonomy: The right of patients to make decisions about their medical care without their healthcare provider trying to influence the decision.
7. Healthcare Disparities: Differences in access to or availability of medical facilities and services and variation in rates of disease occurrence and disabilities between population groups.
8. Algorithmic Bias: Systematic and repeatable errors in a computer system that create unfair outcomes, such as privileging one arbitrary group of users over others.
9. Data Privacy: In the context of healthcare, the protection of patients' personal health information from unauthorized access, use, or disclosure.
10. Professional Judgment: The application of professional knowledge, skills, and experience in making decisions or offering opinions in a professional context.
11. Ethical Implications: The potential consequences or effects of an action or decision when viewed through the lens of moral principles and values.

Appendix 2. AI-Enabled Early Detection System

Background: A healthcare startup, MedTech AI Innovations, has developed a cutting-edge AI system named "LungScreen." Utilizing advanced machine learning algorithms, LungScreen analyzes medical imaging data with a focus on detecting pulmonary abnormalities. It has been trained on over 100,000 anonymized CT scans, including a wide range of lung diseases at various stages and across diverse populations.

Scenario: John Doe, a 59-year-old with a 30-pack-year smoking history, presented for his annual health checkup. His physician ordered a low-dose CT scan as part of a new lung cancer screening program which utilizes LungScreen. While the radiologist initially reported no significant findings, LungScreen identified a subtle area of concern that suggested early-stage lung cancer.

The radiologist, upon reviewing the AI system's findings, confirmed the presence of a small, early-stage tumor. John underwent further diagnostic tests, and a stage I lung cancer was diagnosed. He received minimally invasive surgery within weeks, resulting in an excellent prognosis.

Appendix 3. AI Misdiagnosis Due to Biased Dataset

Background: DermAssist AI, a tool developed by SkinTech AI, has been deployed in dermatology clinics nationwide. It uses a deep learning model to provide diagnostic suggestions for skin lesions based on photographic input. However, the training data predominantly consisted of images from patients with lighter skin tones.

Scenario: Emily Thompson, a 40-year-old woman with a dark complexion, visited her dermatologist due to a mole on her arm that had changed in size and color. The clinic utilized DermAssist AI as a first-line assessment tool. The system analyzed the mole's image and categorized it as benign.

Due to the AI's recommendation, the dermatologist chose a 'watch and wait' approach. Six months later, Emily returned with the mole significantly altered. A biopsy was performed, leading to a late-stage melanoma diagnosis. The delayed identification of her condition complicated her treatment options and prognosis.

Further investigation revealed that DermAssist AI had a systemic bias; it was less accurate for patients with darker skin.