# The Physicians' Charter for Responsible AI

A Practical Guide to Developing, Testing, and Using AI Tools in Clinical Practice

## Executive Summary

physicianscharter.ai

organized by MD+CALC


# 01. About the Charter

**The origin of our Physicians' Charter stems from a growing concern among the physician creators of MDCalc about the rapid pace of AI and how it will be implemented in healthcare. In the absence of an existing resource that provided practical, clear guidance using real-world clinical scenarios and authored by frontline physicians, we assembled a diverse group of experts to create one. This document is the collective effort of practicing physicians across numerous medical specialties. We all share enthusiasm for AI's potential in medicine, and are steadfast in our commitment to its ethical, fair, and patient-focused implementation.**

**Delivering care to patients every day provides us with a unique understanding of the intricacies of healthcare, a perspective we consider essential in guiding AI's integration into our field. There is a true urgency from physicians to set safe boundaries and demand high expectations from AI in the clinical environment.**

**We hope this charter offers a practical, understandable, and accessible framework to guide all stakeholders. As physician leaders in this era of AI evolution, we must always prioritize the values and the welfare of our patients above all. This charter is our pledge to ensure AI in medicine is effective, ethical, and fundamentally patient-centric.**

Since the first physician saw the first patient, doctors have always searched for, developed, and adopted new tools to improve patient care. From the advent of the stethoscope in the early 19th century to the development of advanced medical imaging like the ultrasound during the 20th, physicians have integrated technology to diagnose diseases more accurately and improve patient outcomes. Similarly, clinical scores and algorithms have enabled physicians to make better, more informed decisions through research and evidence-based care. Physicians have endeavored to use these medical advancements with a focus centered on patient welfare and trust.

Today, we stand at the threshold of a new era. Artificial intelligence, which we will define as "software or systems that are capable of performing tasks by learning from and analyzing language and medical data" — is increasingly present in society and could even become an essential component of medicine.

As authors of this document, we wish to emphasize our enthusiasm and excitement for the promises AI brings. We are not creating this charter out of fear. Instead, our motivation stems from our passion and belief in AI's transformative potential. If developed and deployed responsibly, we believe AI can help to re-center the practice of medicine around the patient — not the computer screen — bringing physicians back to the bedside in a way reminiscent of how medicine was practiced centuries ago.

The promise of AI is significant: improved disease diagnosis, personalized treatment plans, and overall enhancements to patient care. However, with these opportunities come numerous challenges and ethical considerations. How do we ensure that AI systems offer transparency and respect patient privacy? How can we reduce algorithmic bias to prevent exacerbating health disparities? How do we maintain the irreplaceable humanity in medicine while leveraging the benefits of AI?

As we grapple with these and other numerous ethical questions, our charter seeks to provide a collective response. Our position is unequivocal: **patients must continue to be the central focus of medical care, and the sanctity of the patient-doctor relationship must be upheld**. AI is a tool designed to support and augment the capabilities of the healthcare professional - not replace them. AI tools can amplify our skills, acting as a co-pilot, enabling us to work more efficiently and smartly and ultimately allowing us to focus more on what truly matters: our patients.

In this document, we discuss our mission statement and vision for how AI tools — from generative chatbot models to machine learning and neural networks — should be used meaningfully, ethically, and to the benefit of the patients we serve. We base these recommendations on the ethical principles physicians have relied upon for centuries. We have also given great thought as to how these new technologies may require us to focus on several new principles: transparency, accountability, equity, and most importantly, human-centered care.

We then present our 10 Rules of the Road for integrating AI into medical practice. These practical guidelines are designed to be easily understandable by those who work in healthcare and include examples from clinical practice, drawing parallels with our past as we navigate the future.

As you explore this outline, we hope it catalyzes thoughtful discussions and informed actions around integrating AI into clinical practice. We believe in a future where AI, like the stethoscope or ultrasound, becomes an indispensable, supportive tool that enhances our ability to provide the best care for our patients, with the doctor-patient relationship always at its core.

# 02. Our Core Values

**Our core values build on the four pillars of medical ethics:** autonomy, beneficence, non-maleficence, and justice.

**Human-Centered Care:** The priority is always to serve the patient's needs and preferences, focused on their values. This includes respecting patient autonomy and involving patients in decision-making; it also includes focusing on outcomes, diagnoses, and treatments that are relevant and impactful to patients and their lives. AI should be used to enhance — not replace — the patient-doctor relationship and help physicians provide more personalized, effective, and efficient care.

**Transparency:** Clear and open communication about how AI tools function, how they were developed, and what data they use. This includes sharing how these tools impact clinical decisions and patient care.

**Privacy and Security:** Safeguarding sensitive patient information by employing robust data protection measures and ensuring adherence to relevant laws and regulations.

**Equity:** Ensuring AI tools do not exacerbate health disparities but instead work to promote equitable care and outcomes. It means recognizing and acknowledging bias that exists today, prioritizing bias mitigation, demanding diverse data representation, and deploying AI equitably.

**Collaboration:** Engaging a variety of stakeholders in AI development and use, such as physicians, other healthcare professionals, data scientists, ethicists, and patients themselves. This collaboration facilitates a multidisciplinary approach and a more holistic view of patient care.

**Accountability:** Responsibility for AI's implementation and the outcomes it generates. This includes the need for regulatory oversight, accountability of systems not individuals, malpractice clarification and reform, adherence to privacy laws, and safety and error management. Additionally, we must consider the downstream effects of AI implementation: Are resources allocated differently? What are the unintended consequences and impact on patient care?

**Continuous Learning and Improvement:** Embracing a growth mindset, in which there is continuous monitoring of AI tools, ongoing validation, and an open environment for learning and adaptation. This includes both the AI systems and the healthcare professionals using them.

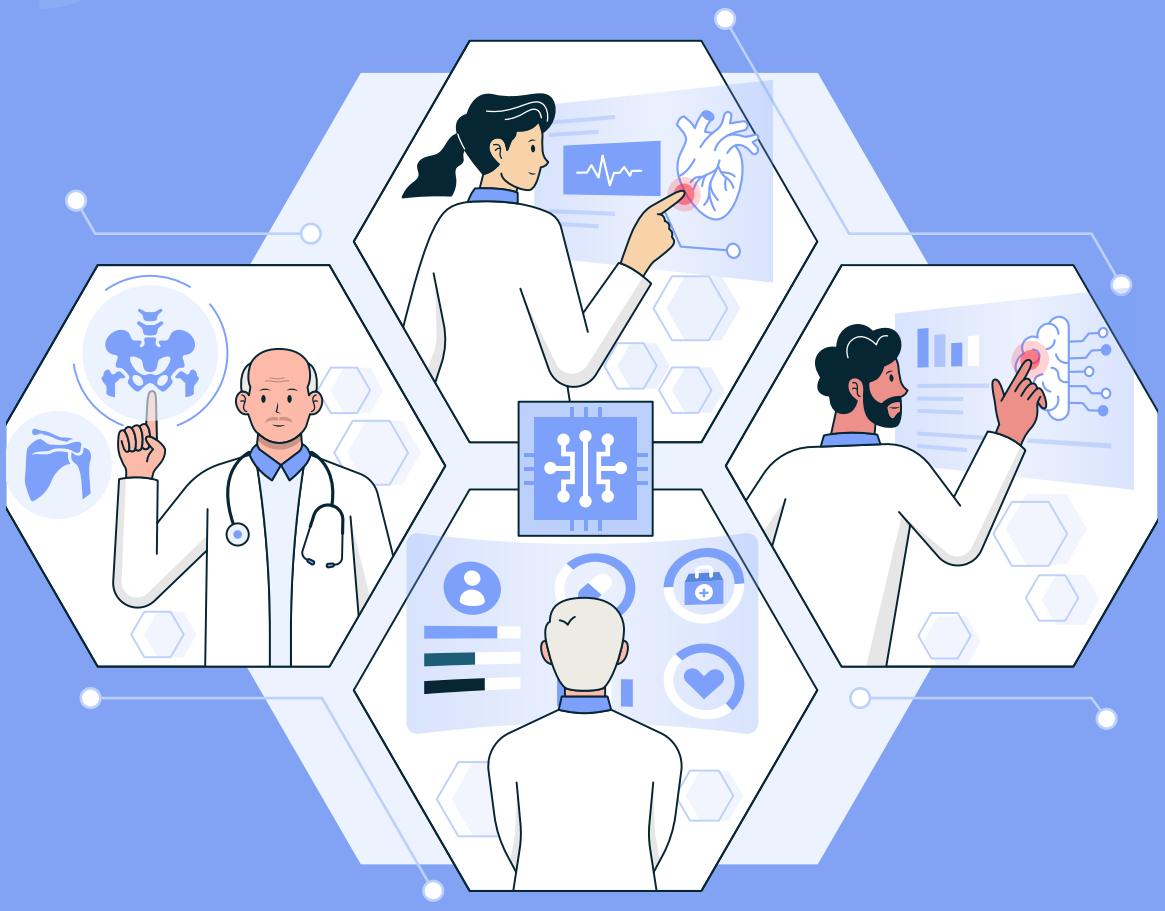

# 03. The 10 Rules of the Road for AI Implementation

| Topic and Description | Examples in Clinical Practice |
| --- | --- |
| 01<br>  <br>**Human-Centered Design and Engagement**<br>Keep the patient-doctor relationship central, involve patients and doctors early in the development process and inform them about how AI is utilized in their care. | • In developing an AI tool for assessing depression, feedback from patients and psychiatrists is included in development from the outset to ensure the tool is both clinically useful and user-friendly.<br>• During an AI-guided surgery, the surgeon discloses and explains to the patient the role of AI in assisting but not performing the procedure.<br>• An AI diagnostic tool is used by a physician during the patient visit, so that the doctor can explain and discuss the tool with their medical expertise. |
| 02<br>  <br>**Data Quality and Privacy**<br>Prioritize high-quality, diverse, and geographically relevant data for training AI models. Respect patient privacy and foster responsible data interpretation. | • An AI tool for predicting disease progression should use diverse datasets representing different geographic and demographic cohorts, ensuring its efficacy across a broad patient population.<br>• EHR data used to train AI models for outcome prediction should be de-identified and encrypted end-to-end to maintain patient privacy. |

| Topic and Description | Examples in Clinical Practice |
| --- | --- |
| 03<br>  <br>**Ethics, Bias Mitigation, and Their Implications**<br>Expect, monitor and mitigate biases in AI algorithms and consider potential ethical implications in AI deployment. | • In developing an AI tool for skin cancer detection, the model is trained on a diverse dataset representing various skin types to minimize bias.<br>• When deploying an AI tool for prioritizing patient referrals, consider its impact on access to care to ensure it does not inadvertently favor or disadvantage certain patient groups. |
| 04<br>  <br>**Trust: Transparency, Explainability, and Accountability**<br>Encourage a “glass box” approach to AI, provide clear information about its workings, and establish a robust framework for trust and accountability. | • When using an AI model for predictive analytics, both patients and physicians are provided with clear, understandable explanations of how the model works, what data it uses in its analysis, and how it makes predictions (some AI models make this not possible).<br>• An accountability framework is implemented so that in case of misdiagnosis by an AI tool, there are mechanisms for addressing the error and preventing recurrence through rapid feedback from clinician to AI developer. |

| Topic and Description | Examples in Clinical Practice |
| --- | --- |
| 05<br>  <br>**Continuous Validation, Feedback, and Improvement**<br>Ensure that AI tools are evaluated on and iterate from formal, objective evaluations of their utility as well as in everyday use; models require continuous review. Encourage feedback and provide clear paths for users to share their experiences and insights, keeping the tools effective, safe, and up-to-date. | • An AI tool for diagnosing diabetic retinopathy is repeatedly and regularly validated on diverse, independent datasets. Its performance is closely monitored over time using standardized benchmarks. Any erroneous suggestions from this AI tool are directly reported by physicians via a dedicated feedback system, leading to the tool's refinement and improvement.<br>• An AI system built for heart disease diagnosis is not just validated once, but regularly checked against standard performance measures and monitored for drift in accuracy. Doctors using the system can report any inaccuracies they find, helping make the system better over time. |
| 06<br>  <br>**Collaborative Approach and Workflow Integration**<br>Promote a collaborative, compensated, multidisciplinary approach to AI development, focusing on AI tools that integrate seamlessly into healthcare workflows. | • In developing an AI tool for radiology, radiologists, data scientists, ethicists, and patients are all involved, with compensation structures in place for the time required to review these models.<br>• An AI tool for analyzing CT scans is designed to integrate directly into a hospital's existing imaging and EHR systems, providing insights within the existing workflow. |

| Topic and Description | Examples in Clinical Practice |
| --- | --- |
| 07<br>  <br>**Regulatory Compliance and Safety**<br>Adhere to regulatory guidelines for AI development and implementation. Implement robust safety measures to protect patient safety. | • AI-based diagnostic tools are developed in accordance with guidelines and regulations, ensuring their safety and efficacy.<br>• In AI-assisted surgery, backup safety measures are considered and ready for use to prevent potential harm from AI-induced errors. |
| 08<br>  <br>**Education and Support**<br>Provide comprehensive education and training to healthcare providers about AI, and support them in their roles as primary interpreters of AI outputs. | • In a hospital deploying an AI tool for radiology interpretation, a comprehensive training program is provided, offering radiologists extensive knowledge about the tool, its use cases, and how to interpret and verify its outputs.<br>• A healthcare organization provides an ongoing support program for clinicians, providing regular updates, resources, and direct communication lines with the AI development team for queries and feedback. |

| Topic and Description | Examples in Clinical Practice |
| --- | --- |
| 09<br>  <br>**Patient-Centered Outcomes and Value in Healthcare**<br>Develop clinically meaningful AI tools that enhance healthcare value, reduce overdiagnosis and overtreatment, and provide better outcomes at the same or lower cost. | • An AI tool for lung cancer screening is trained to accurately differentiate between benign and malignant nodules, thus reducing unnecessary invasive procedures and patient anxiety.<br>• An AI system for detecting pulmonary emboli takes into consideration that some emboli may be clinically insignificant (or even false positives). It includes the risk of anticoagulation and PE treatment into its model, recognizing that patient-important outcomes are the overall goal, not just detection of clot. |
| 10<br>  <br>**Understand the Limits of AI**<br>Recognize that while AI can augment and improve healthcare delivery, it is not a panacea and cannot solve every problem in our complex, fragmented healthcare system. We must understand its limitations, understand when human intervention is needed, and find a balance between technological assistance and human action to provide optimal care to patients. | • AI models can assist in predicting disease progression based on extensive data sets, but these predictions are purely statistical and do not account for individual patient responses and differences. Clinicians must interpret these predictions while considering their personal understanding of the patient's condition and unique circumstances.<br>• An AI algorithm may be capable of sorting and prioritizing patient referrals based on their medical data, but it cannot wholly substitute the human touch in empathizing with patient fears and anxieties. Clinicians are needed to communicate comfort and provide the caring human interaction that patients often require. |

# 05. The Physicians' Charter Conclusion

As we conclude our guide, a new era dawns — one brimming with potential. Just as antibiotics and X-rays have improved the practice of medicine, so too will artificial intelligence. AI's potential to redefine the practice of medicine for patients, physicians, healthcare workers, policymakers, and ethicists is enormous. There are countless ways that it could help improve diagnostics, deliver better education, offer new treatment recommendations, and even allow physicians more time with their patients. Yet it brings with it a critical obligation to ensure ethical, safe, and respectful deployment — and physicians are essential experts who can help navigate this space.

And as we delve deeper into this brave new world, it has become clearer that introducing AI into healthcare is not a mere plug and play scenario. Rather, it requires continuous, meticulous effort in calibrating, validating, and updating the AI systems, which in turn requires time and resources for ensuring the reliability, safety, accuracy, and trustworthiness of these tools.

Throughout this document, we've discussed our 10 Rules of the Road for AI Implementation, each one serving as a beacon to guide us through the complexities of incorporating AI into healthcare. But as we write these rules in the summer of 2023, we acknowledge the need for adaptability. Rather than rigid edicts, they are dynamic guideposts, capable of evolving in step with the expanding and transformative world of AI. Our framework, therefore, must be as fluid and adaptable as the technology it seeks to shepherd. And when in doubt, our values must serve as our compass, steering our decisions and molding the development and implementation of AI.

The trust given to us by patients is the sacred bedrock upon which all healthcare is built. It is equally susceptible to erosion from indifference or malpractice as it is from unethical or misguided AI deployments. Missteps — no matter the intent — have the potential to echo far beyond individual patients, beyond the confines of one doctor or one institution, and beyond the boundaries of medical specialties. These reverberations could indeed stall or even halt further advancement of AI in healthcare.

AI offers immense promise, from improved patient outcomes to enhanced provider experiences. Let us embrace AI as a valuable ally and co-pilot — supporting the patient-physician relationship — and enabling us to help, counsel, and guide our patients through their lives. Our journey towards responsible AI usage is a path to better healthcare for all. We invite you to join us on this journey — guided by firm principles yet adaptable in our approach, and always anchored by our dedication to our patients.

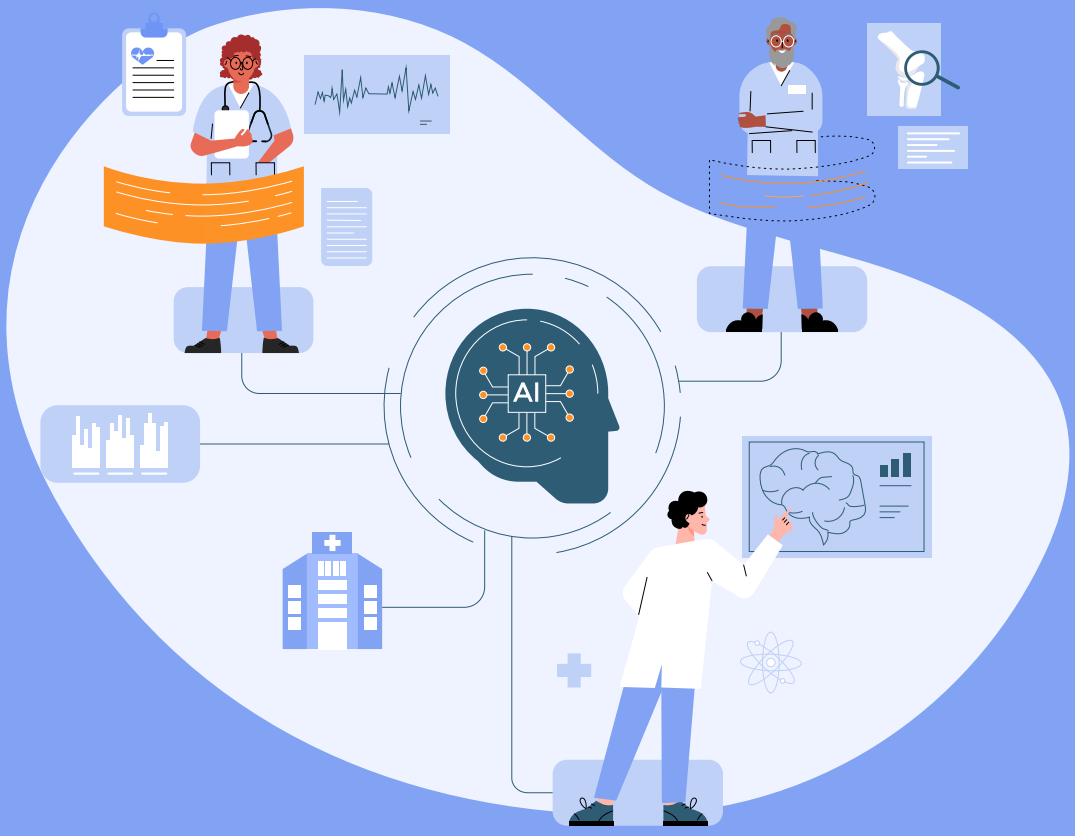


# 06. About the Authors

**Anthony Cardillo, MD** is a pathologist and Clinical Informatics fellow at NYU Langone. His primary interests are in medical cybersecurity and the digital transition of pathology. He presently serves on two national committees involving artificial intelligence and ethics in the College of American Pathologists and the American Medical Informatics Association. More recently, Dr. Cardillo was recognized in The Pathologist's Power List in 2021 and 2022, and in 2023 placed in the US and UK Summit for Democracy competition to develop secure AI models.

**William "Will" Collins, MD** is a hospital medicine physician and Clinical Assistant Professor at the Stanford School of Medicine. He is also the current president of the Society of Hospital Medicine San Francisco Bay Area Chapter. He has been captivated by both the potential and the risk of AI applications in medicine. From his experience in clinical research, he is interested in designing rigorous trials to assess AI interventions to show meaningful outcomes for patients and medical providers.

**Dustin Cotliar, MD MPH**, brings significant care delivery expertise that comes from over eight years of clinical practice and studying healthcare policy and management at Columbia University. He has served as a clinical consultant with the Kaiser Family Foundation where he published health system research that has been cited in articles by the NY Times, VOX, politico, and others. A recent first-place winner at MIT's Hacking Medicine, one of the largest clinical hackathons in the country, Dr. Cotliar is passionate about building innovative clinical products, especially those rooted in artificial intelligence and machine learning.

**Carly Eckert, MD**, MPH is a physician technologist located in Chapel Hill, NC. She is double-boarded in preventive medicine and clinical informatics. Carly has led clinical and data teams within healthcare startups for nearly a decade. Her areas of focus include AI governance, ethics, and bias. She also enjoys teaching physicians and other healthcare providers on the topics of practical and applied AI solutions and how to communicate with technical teams.

**Sarah Gebauer, MD** is an experienced hospital leader and healthcare technology consultant with a background in clinical informatics. She's passionate about physician engagement with artificial intelligence and founded Machine Learning for MDs, a free online community providing education, training, and networking for physicians in the AI space.

**Raouf Hajji, MD**, PhD is an Assistant Professor of Internal Medicine, Medicine Faculty of Sousse, Tunisia. With his expertise in clinical practice, biomedical research, and academia, he is the author, reviewer, and editor of many peer-reviewed medical journals and book chapters. He is Co-founder and Medical Lead of International Medical Community (IMC), an international initiative working as an Innovation Health technologies Hub with the main scope of advancing international cooperation and creating a link to cutting-edge technologies for the healthcare sector worldwide. You can join him on LinkedIn where he publishes weekly medical newsletter: Healthcare Present & Future with Updates on Biomedical Research, Academia, Clinical Practice and Emerging Technologies in Healthcare.

**S. Morgan Jeffries**, MD is a neurohospitalist and physician informatician at Geisinger, where his work focuses on quality measures, workflow improvements, and AI strategy. He's also an assistant professor at the Geisinger Commonwealth School of Medicine and a member of Epic's Adult Neurology Specialty Steering board. He's interested in the similarities and differences between human and AI minds, AI safety and alignment, and AI evaluation. He occasionally writes on LinkedIn and less frequently on X (née Twitter).

**Matt Sakumoto, MD** is a virtualist primary care physician in San Francisco, and Adjunct Clinical Professor at UCSF focusing on virtual care and clinician efficiency tools for the EHR. With prior industry experience at multiple telehealth startups and as a clinician-advisor to many early-stage companies, he is passionate about exploring and expanding the digital health landscape.

**William Small, MD**, MBA is a hospital medicine physician and clinical informatics fellow at NYU Langone Health who is focused on the impact of communication technologies on the clinician experience with the EHR and patient outcomes. He is dedicated to understanding how best to evaluate outputs of generative AI and is a key member of the team evaluating the effects of generative AI chatbot integration into patient-provider Inbasket communications on provider efficiency and satisfaction.

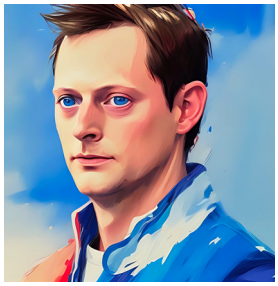

**Graham Walker, MD** is an emergency physician and clinical informaticist in San Francisco, California with The Permanente Medical Group (TPMG) and enjoys working at the intersection of technology and medicine and created and organized the Physicians' Charter for Responsible AI. He also built MDCalc and theNNT, two free online resources that have allowed millions of clinicians from around the world to incorporate evidence-based decision-making into their medical practice. You can find him at LinkedIn writing about medicine and technology and referring to himself in the third person.

## Disclaimer

The views and opinions expressed by the authors do not necessarily reflect the views of our employers or other organizations with which we're affiliated. While every effort has been made to ensure the accuracy and currency of the information, the rapidly evolving nature of AI in healthcare means that some details may change over time; we hope to review this document annually. This document was created by individual physicians in our free time. © August 17, 2023.

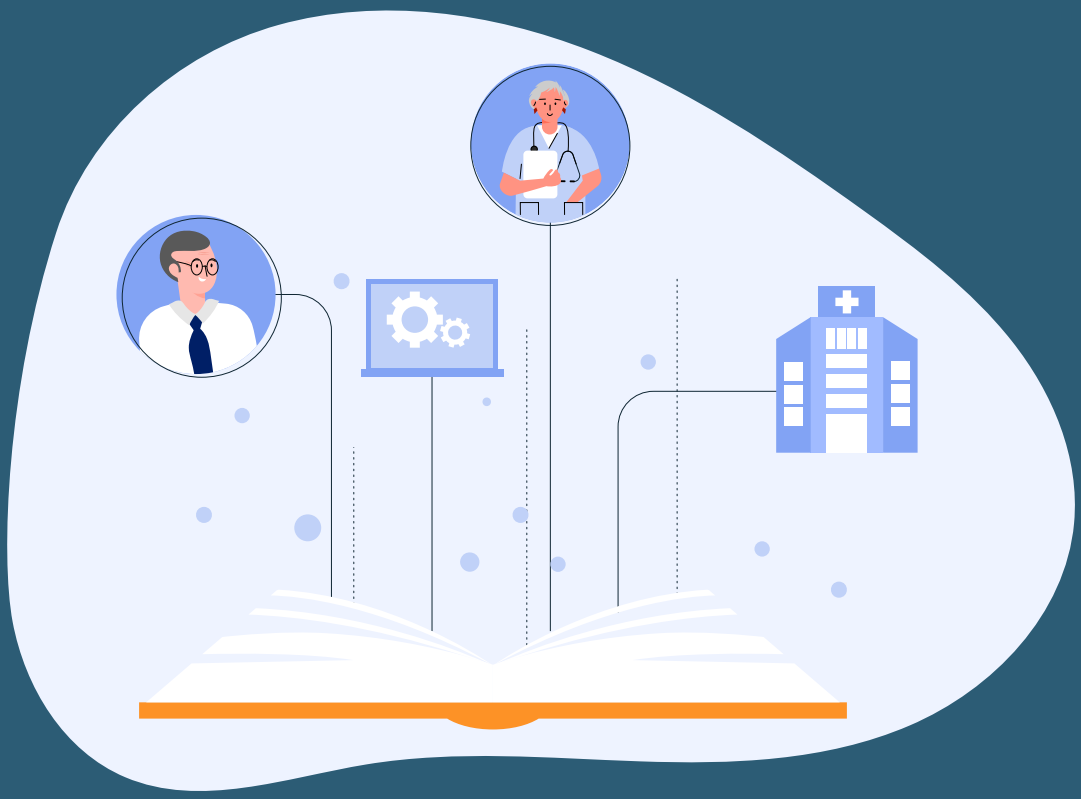

# 07. Acknowledgements

- Many thanks first to Dr. Scott Campbell, an emergency physician, founding member of the American Board of AI in Medicine, and incredibly gifted AI expert who helped develop the concept for this charter in the first place.

- Thank you to Dr. Sarah Gebauer for creating the Machine Learning for MDs online community, without which this paper would have struggled to find diverse authorship, as well as Dr. Shoreh Irani for her Physician-Led AI in Medicine Facebook group.

- Appreciation to Drs. Basil Kahwash, Jane Wang, Oluseyi Fayanju for their edits and suggestions during the writing process, along with Dr. Sarah Gebauer and Dr. Raouf Hajji for their review and editing as well. Graham Walker's parents (Susan and Don Walker) also provided editing advice, as they have been doing ever since Graham was old enough to write.

- Thank you to Drs. Edward Yap and Charulata Ramaprasad for providing their opinions as the project was just starting out.

- Gratitude to Dr. Justin Norden for his leadership in the intersection of AI and medicine, and for providing feedback about the charter as it developed.

- We look forward to other organizations and partnerships building on and collaborating with this charter — and recommend that we continue to work with expediency in this area, as the world of AI is moving extremely fast, and at a much faster pace than medicine would normally progress. We were inspired by the Coalition for Health AI (CHAI) at https://www.coalitionforhealthai.org/ and the Health AI Partnership at https://healthaipartnership.org/. We also look forward to the work of the National Academy of Medicine's Health Care Artificial Intelligence Code of Conduct, and hope it will start to provide expedited recommendations as well.


*Finally, this document was born out of the discussions the physician creators of MDCalc, both hopeful-yet-concerned after seeing the power of generative AI in late 2022. MDCalc supported and this project's work and allowed it to go from vague concept to illustrated, edited, completed execution.*

# 08. Further Reading and Supporting Research

## Introduction, Values, Mission, and Vision

a. This is a wonderful piece on Dr. Rene Theophile Hyacinthe Laënnec, a French physician and the inventor of the stethoscope in the early 1800s.

b. In a similar vein, we have supporting histories of the PET scan and medical ultrasound.

c. Stanford's Human-Centered Artificial Intelligence has a great brief overview of different AI terms and definitions.

### 1. Human-Centered Design and Engagement

a. The New England Journal of Medicine provides a thorough summary of Artificial Intelligence and Machine Learning in Clinical Medicine, including history, ML, and chatbots, with proposals for research standards as well.

b. Stanford's HAI (Institute for Human-Centered Artificial Intelligence) is an outstanding resource for all-things AI, and its Values section aligns well with this Charter's vision, and its Humanity section aligns perfectly with our first chapter.

c. This paper from San Diego comparing ChatGPT responses to Reddit physician responses made headlines and sparked controversy when it suggested that OpenAI's tool provided more empathetic responses than the physician users on Reddit.

d. Want to learn more about Human-Computer Interaction? Ben Shneiderman's book entitled Human-Centered AI is a great place to get started.

e. And for more HCI information, look to HCI International's book series and conference.

### 2. Data Quality and Privacy

a. This 2023 article from Computers in Biology and Medicine delves into the barriers of AI adoption in healthcare, focusing on various privacy and data concerns, and presents an overview of advanced privacy-preserving techniques like Federated Learning and Hybrid Techniques.

b. A team in China provides an excellent review on federated learning and privacy-preserving algorithms as solutions to data fragmentation and privacy challenges in healthcare AI.

c. This is a summary of a roundtable discussion by the US Department of Health and Human Services (HHS) on the opportunities, challenges, and strategies for using data to train AI models in healthcare, offering recommendations for HHS and stakeholders to further AI advancements.

3. **Ethics and Bias Mitigation**
    a. This NEJM paper discusses the use of race in predictive algorithms, the problems that arise when using race, and highlights the importance of knowing what goes into algorithms.
4. **Trust: Transparency, Explainability, and Accountability**
    a. Carnegie Mellon's Violet Turri has an outstanding piece on "What is Explainable AI?"
    b. This famous paper from Microsoft on explainable models revealed an issue with a neural network that was predicting that patients with asthma had a lower likelihood of mortality from pneumonia (when in actuality they have a higher mortality); the model was making technically accurate conclusions, but made these conclusions because asthmatic patients were more often managed in the ICU, lowering their mortality due to more aggressive, intensive care.
    c. This paper from PLOS is an outstanding review of the ethical, theoretical, and practical concerns around AI models and tools — specifically focusing on how emergency dispatch operators did not adopt a tool that predicted which emergency calls were for a cardiac arrest case because they did not trust or understand it.
    d. Epic's Sepsis model is discussed in this paper and is unfortunately a good example of a model failing "in the wild."
5. **Continuous Validation, Monitoring, and Improvement**
    a. This Lancet paper suggests concerns around generalizability of models in healthcare and explains the reasons that models may not be as generalizable as we would like to think.
    b. This NEJM correspondence (in particular, its Table 1) provides an overview of approaches to recognizing and addressing dataset shift.
6. **Collaborative Approach and Workflow Integration**
    a. Authors from Ohio State provide a roadmap for the integration of AI into Radiology workflows specifically, from the Journal of Medical Imaging.
    b. European Radiology reviews the challenges — and offers solutions to them — in this piece, again focusing around AI in radiology.
    c. This article from the UK discusses advances in Human-Computer Interaction, breaking the paper up into 6 categories: Interfaces, Visualization, Electronic Health Records, Devices, Usability, and Clinical Decision Support Systems.

7. **Regulatory Compliance and Safety**
    a. The FDA provides guidance for AI and ML in Software as a Medical Device applications.
    b. The FDA also has a helpful navigator to help developers determine if their software is a medical device.

8. **Education and Support**
    a. This paper from Health Education UK argues that the healthcare workforce in the UK will need education and training — and the creation of an educational framework — to use AI successfully.
    b. This article interviews 45 physician champions and discusses what they felt were critical to the adoption of a new EHR and what challenges they faced.
    c. Health Affairs discusses 7 lessons from EHR implementation — including hands-on training.
    d. Here are 10 more lessons learned from an academic medical center that adopted an EHR for its 6 hospitals, 2 campuses, and 46 outpatient sites.

9. **Patient-Important Outcomes and Value in Healthcare**
    a. This paper reviews the very concept of patient-important outcomes, and acknowledges that medicine doesn't often ask patients what's important to them as an outcome.
    b. Even in research today, we don't focus nearly enough on patient-important outcomes — in diabetes and critical care as just two examples.

10. **Understanding the Limits of AI**
    a. Thinking, Fast and Slow is a book by psychologist Daniel Kahneman, who describes two systems that humans use when thinking; a fast, instinctive system, and a slow, deliberative thought process.
    b. The complementarity-driven deferral to clinicians (CoDoC) system proposes a model that could even help clinicians decide when to rely on AI tools and when to defer to clinician judgment.
    c. This paper demonstrates how AI can be helpful to humans — by re-ordering CT scan reading queues — without replacing physician interpretation.
    d. The tragic crash of flight AF447 is an example of the devastating consequences of automation bias.
    e. Automation bias is hard to overcome, even when humans are educated and warned that it exists.