



Navigating the Responsible and Ethical Incorporation of Artificial Intelligence into Clinical Practice

Adopted by FSMB House of Delegates, April 2024

EXECUTIVE SUMMARY

Artificial Intelligence (AI) holds tremendous potential to aid healthcare providers in diagnosis, treatment selection, clinical documentation, and other tasks to improve quality, access, and efficiency. However, these technologies introduce risks if deployed without proper “guardrails” and understanding which may impact considerations in clinical practice as well as regulatory processes of state medical boards. By taking a proactive and standardized governance approach anchored in ethical principles, state medical boards can promote safe and effective integration of AI, in its various forms, while prioritizing patient wellbeing.

This report summarizes expert opinion and proceedings to develop guidance from the FSMB Ethics and Professionalism Committee to aid physicians and state medical boards in navigating the responsible and ethical incorporation of AI centered on (1) education, (2) emphasizing human accountability, (3) ensuring informed consent and data privacy, (4) proactively addressing responsibility and liability concerns, (5) collaborating with experts, and (6) anchoring AI governance in ethical principles.

Clinical systems and processes making use of AI must be continually monitored and refined. This should not occur in a vacuum but should be the focus of collaborative efforts among physicians, health systems, data scientists, and regulatory agencies, *including state medical boards*. By thoughtfully addressing the opportunities and challenges posed by AI in healthcare, state medical boards can promote the safe, effective, and ethical use of AI as a tool to enhance, but generally not replace, human judgment and accountability in medical practice. In fulfilling their missions to ensure that patients benefit from and are not harmed by applications of AI in their care, it is essential that state medical boards avoid over-regulation and regulatory overreach by attempting to regulate that which is not in their purview. With focused efforts on the current and future state of the use of AI by licensees, state medical boards may sustain regulatory efficiency, achieve consistency across jurisdictions in the regulation of AI in clinical practice, help secure the benefits of AI, and proactively safeguard patients while upholding professional standards.

Section I. Background

The rapid development of Artificial Intelligence (AI) technologies, a subset of which is sometimes referred to as “Augmented Intelligence,” presents new opportunities to improve healthcare quality, access, and efficiency. However, it also poses ethical challenges regarding accountability, transparency, equity, and patient safety. State medical boards have an important role in systems of governance and oversight to ensure the safe, effective, and ethical use of AI in clinical practice within the scope of their regulatory duties and oversight of licensees. Although medical boards are limited from directly regulating AI as a healthcare device or tool, they are the only regulatory agency with the explicit authority to regulate physicians who use AI to provide care.

Recognizing the important role played by state medical boards to uphold ethical and professional standards, as well as the centrality of medical professionalism to the responsible application of AI in patient care, FSMB Chair, Jeffrey D. Carter, MD, tasked the FSMB’s Ethics and Professionalism Committee, chaired by Mark B. Woodland, MS, MD, with analyzing AI from a medical regulatory perspective. The Committee was asked to identify ethical principles that will guide the FSMB’s approach to developing an understanding of AI and help inform medical regulatory considerations for state medical boards as they encounter the application of AI in the clinical practice of licensees.

This report summarizes key discussion findings and provides guidance to state medical boards for oversight of utilization of AI in medical practice to promote patient safety, quality care, equity, and accountability. Recommendations center on education for clinicians, emphasizing human accountability, ensuring informed consent and data privacy, proactively addressing liability concerns, collaborating with experts, and anchoring AI governance in ethical principles.

The statements and recommendations offered in this report apply existing and well-established regulatory considerations to new technologies and tools in the provision of health care. Professional responsibilities and expectations of medical licensees remain the same; how they are fulfilled may differ based on the AI application utilized in delivering care.

Defining AI

To fully grasp the impact of AI in healthcare, it is essential to understand its foundational components, such as algorithms, data analytics, and machine learning. AI systems in healthcare leverage complex algorithms and advanced data analytics to make predictions or decisions. AI’s capability to automate routine tasks, provide diagnostic support, and enhance physician cognitive functions presents a significant shift in healthcare practices.

The training of AI is conducted through a process called ‘machine learning’, which involves feeding large amounts of data into a computer system, allowing it to learn patterns, make predictions, or create decisions based on that data. As the system processes more data, its ability to make predictions or decisions improves. This learning can be either ‘supervised’ or ‘unsupervised’ depending on the objective. Specifically, ‘supervised learning’ is the creation of a prediction from labeled training data that is weighted for a specific purpose. This could be something as simple as a spam email filter (i.e., classification) or predicting future housing prices (i.e., regression). In ‘unsupervised learning’ the system creates predictions by analyzing data sets that are not labeled or weighted to achieve a specific outcome. For example, unsupervised algorithms designed to recognize and decode human speech are fed thousands of hours of spoken language to foster the ability to learn, understand, and interpret different words and phrases. In essence, the key difference is that supervised learning works with labeled data and

aims to predict an output, while unsupervised learning works with unlabeled data and aims to understand the underlying structure of the data.

Multiple AI applications are currently in use in healthcare. Some common examples of such applications include:

- Computer vision systems to analyze medical images,
- Natural language processing to review clinical notes,
- Predictive algorithms and advanced data analytics to forecast clinical trends,
- Voice recognition to support clinical documentation,
- **“Chatbots”** to provide patient education and triage.

These tools analyze large datasets to identify patterns, classify information, and make predictions to support clinical decision-making.

Recent developments in unsupervised machine learning have led to more sophisticated content generation, commonly referred to as **“generative AI”** (GAI). GAI can generate new content or data that is similar to human-generated content. Unlike traditional AI systems that are designed to analyze data and make predictions or decisions, GAI focuses on the creation of new, original outputs. Key characteristics of GAI include the ability to utilize advanced data analytics to learn patterns, styles, or structures from large datasets to produce novel creations that don’t simply replicate the input data but exhibit some form of creativity or innovation. GAI is used in a wide range of applications, including creating artwork, composing music, and generating realistic human voices. In the healthcare setting, GAI is being used in drug discovery, personalized medicine, medical imaging, and predictive analytics for epidemic outbreaks.

The most recent breakthroughs in AI have been in the field of GAI, specifically with what are known as **“Large Language Models”** (LLMs). LLMs are created using artificial neural networks – algorithms inspired by the structure and function of the human brain – which are fed vast amounts of large datasets consisting of a wide variety of text sources such as books, articles, websites, and more, allowing the LLM to learn language patterns, grammar, and context. The 'large' in LLMs refers to the size of the neural networks used, specifically regarding the number of parameters they contain. If you were to think of parameters in the context of medicine, they would be analogous to human neurons in the brain. For example, the human brain comprises approximately 86 billion neurons,¹ while the largest LLM comprises approximately 175 billion parameters.² These parameters can be adjusted as the LLM ingests more information to improve accuracy of predictions and generation of coherent and contextually relevant text. Once trained, LLMs can perform a variety of language-related tasks such as translating languages, answering questions, summarizing texts, and even creating content like stories or articles.

AI in Healthcare

In healthcare, AI is already being applied to all aspects of a physician’s workflow, from purely administrative tasks such as patient scheduling and clinical documentation, all the way to clinical decision support. Some experts forecast that artificial intelligence will replace as much as 80% of

¹ Northwestern Medicine 2019, <https://www.nm.org/healthbeat/healthy-tips/11-fun-facts-about-your-brain#:~:text=Research%20suggests%20the%20human%20brain,can%20combine%2C%20increasing%20storage%20capacity>

² Amazon Web Services, <https://aws.amazon.com/what-is/large-language-model/>

what doctors currently do,^{3,4} thereby allowing physicians to realign their focus on patient care rather than administrative tasks. AI applied to the front-end and ongoing documentation of a patient encounter can accomplish much of the clerical and administrative work and refocus a physician's attention towards delivery of care. This presents enormous potential for reducing physician burnout by eliminating redundancies and systemic waste. However, it is important to recognize the need for verification of AI-generated clinical information for accuracy. As such, it is important to state that these tools have generally not been designed to, nor are they yet capable of, replacing a physician's professional judgment, ethical responsibilities, or accountability to state medical boards.

A key professional responsibility in medical practice has always been the assurance that diagnoses, clinical decisions, and recommendations are not biased. As with any other tool or differential used to diagnose or treat a condition, medical professionals are responsible for ensuring accuracy and veracity of evidence-based conclusions. AI systems encumbered by false or inaccurate information may carry a bias that can be detrimental to providers and harmful to patients. Physicians should therefore make reasonable efforts to identify and address such biases before using AI systems in patient care.

Regulatory Landscape

The rapid development, deployment, and wide-spread utilization of AI has left regulators across the globe struggling to identify ways to regulate AI through existing structures. This has resulted in a regulatory framework for AI, including but not limited to government statutes, industry guidance, and professional opinion, that is complex and yet still largely underdeveloped, both in the United States and globally. As a result, physicians and state medical boards should anticipate significant, continual change as key players such as legislatures, government agencies such as the Food & Drug Administration (FDA), the Federal Trade Commission (FTC), and industry groups explore the extent of their authority over AI systems, and the legal and regulatory considerations that arise in response. The reaction of medical professional associations such as the American Medical Association, American Osteopathic Association, and specialty societies, as well as those responsible for medical education and its accreditation, will further influence the regulatory work of state medical boards.

The result is a confusing landscape of AI tools and resources, and regulatory patchworks promising to transform healthcare in meaningful ways while providing little insight or guidance into exactly how this outcome will be achieved. Therefore, it is critical that clinicians become educated about what AI is, how it can be used, what are its limitations, and what are the clinician's role and responsibilities in its use.

FSMB Experience in AI

Recognizing the trajectory of AI and its impact on medical systems and clinical care, the FSMB co-hosted a symposium with leading health law firm McDermott Will and Emery in November 2019. This symposium served as an introduction between regulators and industry leaders and a collaborative effort to explore themes and concepts emerging from the initial uses of AI in a

³ Longoni C and Morewedge CK, AI Can Outperform Doctors. So Why Don't Patients Trust It? *Harvard Business Review*, October 30, 2019, <https://hbr.org/2019/10/ai-can-outperform-doctors-so-why-dont-patients-trust-it>

⁴ Curry R, The A.I. revolution in health care is coming. CNBC July 12, 2023, <https://www.cnbc.com/2023/07/12/the-ai-revolution-in-health-care-is-coming.html>

healthcare setting. Resulting from this symposium, the FSMB convened a taskforce that monitored developments and advised the Board of Directors on how best the FSMB could play a role in shaping the future AI ecosystem. This taskforce concluded that AI should be thought of as a tool for medical practice and recommended that initial policy guidance address the ethical and professional responsibilities of physicians choosing to employ it in the delivery of care. The content of this current report is consistent with the taskforce's recommendation.

Section II. Education

A physician has the duty to maintain the requisite skill and knowledge to provide safe and effective health care. As AI is continually utilized and integrated into existing healthcare infrastructures, it is imperative that physicians remain attuned to developments in AI and strive to understand the benefits and risks it poses. Underappreciation of the ability of AI to improve healthcare delivery may restrict a physician from practicing to the top of their license and may result in a physician not taking full advantage of the tools that can improve patient outcomes. At the same time, over-reliance on AI can lead to real harms in independent clinical thinking and critical decision making such as misdiagnosis, medical errors, dependence, and skill degradation. This risk of harm increases in situations of algorithmic bias or where misinformation is present.

Accordingly, medical education, at all levels, should include an emphasis on advanced data analytics and use of AI in a clinical setting. Consistent with their duties under the principles of justice, beneficence and non-maleficence, physicians should regularly engage in accredited continuing medical education programs designed to improve competence in understanding the application, benefits, and risks of AI and its implications on patient care.

Section III. Accountability

State medical boards do not regulate tools or technologies, only the licensed physicians that use those tools. Consistent with the prevailing standards for any tool used in the delivery of healthcare, the physician is ultimately responsible for the use of AI and should be held accountable for any harms that occur. The extent to which a physician will be held accountable by the state medical board will depend on the relationship between the AI being used and risk that the tool may either create patient harm or otherwise impact the professional obligations of the physician. As **Figure 1** illustrates, as AI tools perform functions that more closely model the practice of medicine, the risk to patients of their application generally increases. The appropriate level of regulatory scrutiny and accountability to the regulator by the licensee using the tool should increase commensurately.




Fig. 1 - Modeling Risk v. Function

Decision Support

Physicians may consider AI as a decision-support tool that assists, but does not replace, clinical reasoning and discretion. Physicians should understand the AI tools they are using by being knowledgeable about their design, training data used in its development, and the outputs of the tool in order to assess reliability and identify and mitigate bias. Once a physician chooses to use AI, they accept responsibility for responding appropriately to the AI's recommendations. For example, if a physician chooses to follow the course of treatment provided by an AI-generated response, then they should be prepared to provide a rationale for why they made that decision. Simply implementing the recommendations of the AI without a corresponding rationale, no matter how positive the outcome may be, may not be within the standard of care. Alternatively, if the physician uses AI and then suggests a course of treatment that deviates from one delineated by AI, they should document the rationale behind the deviation and be prepared to defend the course of action should it lead to a less than optimal or harmful outcome for the patient. Generally, the reason a physician provides for disagreeing with an AI's recommendation should be because following that recommendation would not uphold the standard of care. As with any tool, once it produces a result, the outcomes cannot be ignored; there must be documentation reflecting how it was or will be utilized by the physician in the care provided. While the expanded use of AI may benefit a physician, failure to apply human judgement to any output of AI is a violation of a physician's professional duties.

The use of AI in medical practice may present challenges to the provision of a rationale for following or ignoring an AI tool's advice in situations where algorithms informing the tool's recommendation are too complex for humans to understand. Such "black box" algorithms may still hold tremendous benefit for patients and should not be disallowed outright. While a licensee may not be able to explain in step-by-step fashion precisely how an AI tool arrived at a clinical recommendation, they should still be expected to offer a reasonable interpretation of how the AI arrived at a particular output (i.e., recommendation) and why following or ignoring that output meets the standard of care.

Medical Records

Disciplinary data from state medical boards indicate that failure to maintain adequate medical records served as the basis for 6% of all disciplinary actions from 2015-2019. AI is growing in use to serve as a medical scribe and interact with electronic medical records to automate this component of medical practice. This application of AI holds great promise in reducing a recognized cause of burnout. However, physicians should be aware that the use of AI to generate medical records, without proper oversight, may lead to inaccurate documentation and subsequent patient harm for which the physician will likely be accountable.

Part of the use of AI in documenting medical care requires these tools to access and review personal health information (PHI). Physicians should be aware of what security measures are in place to ensure the PHI provided to AI systems remains secure and in compliance with existing state and federal laws, as well as the patient's preferences. Physicians retain their duty to review records created with AI to ensure that the data captured is accurate and properly managed.

Section IV. Informed Consent and Data Privacy

One of the primary goals of the informed consent process is to ensure patient autonomy in clinical decision making. This is accomplished both by informing patients about diagnosis and treatment planning *and* safeguarding patient privacy.

For informed consent to be valid a patient must be adequately informed about their diagnosis and treatment options, the risks and benefits involved, and reasonable alternatives. These duties under the principle of autonomy apply in all clinical encounters, including those that use AI to inform diagnosis and treatment plans. A physician must be able to independently explain components of diagnosis and treatment options in order to fulfill their professional responsibilities relating to the informed consent process. Informed Consent is not a list of AI-generated risks and benefits, but instead a meaningful dialogue and shared decision-making between the physician and patient. AI may be used to assist in this process but the ultimate responsibility rests with the physician.

Because data received during a patient encounter may be input into AI tools, physicians should receive a patient's consent prior to application of a tool to a patient's care. Physicians should disclose to patients when and how AI is used in their care and clearly communicate about the capabilities and limitations of their tools to the patient, including how they use and share any patient data obtained during a patient encounter. Physicians should also be prepared to disclose how they used AI in their diagnosis and treatment planning, discuss the continued role and responsibilities of the physician, and describe any safeguards that have been put in place to ensure reliability of the AI's output. A lack of transparency regarding the role that AI has played in the delivery of care and the inability of the physician to communicate with the patient can

undermine trust and may serve to highlight the physician's lack of understanding of how the AI tool works.

Section V. Equity and Bias

As noted earlier in this document, AI systems encumbered by false or inaccurate information may carry a bias that can be detrimental to providers and harmful to patients. The principle of justice dictates that physicians have a professional responsibility to identify and eliminate biases in their provision of patient care, including those that may arise through biased AI algorithms.

AI also poses an opportunity to expand access to care for populations historically marginalized and otherwise disadvantaged. Efforts must be made to ensure that all patients have equitable access to the benefits of AI and that existing disparities are not further exacerbated.

FSMB recognizes that it and its member medical boards have an interest in assisting other regulatory agencies and responsible developers of AI to promote systemic standards that require disclosure of information about the training data set, such as race/ethnicity breakdown, and further information about potential biases and risks related to the use of the tool. Because biased training data incorporated into AI tools may ultimately impact patient care and because of the potential that generative AI could perpetuate, rather than eliminate, bias in healthcare, the FSMB should join with other interested parties to understand and resolve the issue of algorithmic bias.

Section VI. AI Governance Through Ethical Principles

Because of the rapidly evolving nature of AI, attempting to regulate its specific applications in healthcare will prove ineffective as the regulatory process will not be able to keep pace with AI's technological advancement. As such, medical boards should instead focus on governing the use of AI through established ethical principles, including respect for patient autonomy, non-maleficence, beneficence, and justice, that have served as the foundation of professional expectations and demonstrated applicability in a variety of situations, regardless of treatment modalities or technology involved.

The following principles and accompanying recommendations are offered by the FSMB to state medical boards and other relevant parties to support the responsible and ethical regulation of clinical care that incorporates AI:

1. *Transparency and Disclosure:*

- Licensees should be required to maintain transparency about the use of AI in healthcare.
- State medical boards should develop clear guidelines for licensees about the disclosure of AI usage to patients that contribute to patient and physician understanding but do not create unnecessary administrative burden.
- FSMB should develop documentation detailing the capabilities and limitations of the most commonly used AI tools to assist medical boards in their role as regulators.
- FSMB should develop a frequently asked questions and best practices document to serve as a resource for medical boards and licensees regarding transparency and use of AI in the provision of care.

2. *Education and Understanding:*

- FSMB and its partners in the medical education community should identify structured educational resources for physicians, medical boards, and patients about AI in healthcare. Such programs should include resources to help understand how AI works, its benefits, potential risks, and implications for patient care.
- FSMB should collect resources, recommendations, guidelines and commentary regarding responsibility and accountability with AI and the medical regulatory process.

3. *Responsible Use and Accountability:*

- Developers should provide agency to physicians assist in their ability to know when and how to use the AI tool in patient care.
- Hospital systems, insurers, or others who select AI tools to support clinical decision making should provide physicians with education about AI tools, access to performance reports of the individual tools, and should design a process for regular review of the efficacy of the tools.
- AI tools should be designed in a manner which would provide state medical boards the ability to audit and understand them, in order to appropriately assess whether a physician who relied upon a tool's output has deviated from standard of care.
- FSMB should support state medical boards in interpretation of responsible and accountable use of AI by clinicians.

4. *Equity and Access:*

- Efforts should be made to ensure equitable access to the benefits of AI for all patients.
- FSMB and state medical boards are committed to the principle that care provided by licensed physicians, physician assistants and other health care professionals is equitable and not influenced by bias based on race, ethnicity or other forms of discrimination.
- FSMB should join with other interested parties to understand and resolve the issue of algorithmic bias.

5. *Privacy and Data Security:*

- Developers of AI tools must implement rigorous safeguards to protect patient data used in the development and evaluation of AI.
- Licensees should generally be informed about how patient data will be used and be prepared to convey this to patients.
- FSMB, along with industry stakeholders, should create policies for the use and dissemination of patient data by AI systems, including minimum data protection measures for patient data used in AI development or evaluation. Where possible both state and federal regulators should coordinate to ensure any policies are not duplicative.
- FSMB should support state medical boards in developing clear patient information materials about patient rights with respect to acceptable use of their data and the role of regulators in this space, both at the state and federal levels.

6. *Oversight and Regulation:*

- State medical boards must retain the authority to discipline physicians for the inappropriate application of AI tools in the delivery of care. This includes considering issues of accountability, particularly as AI systems become more autonomous.
- State medical boards should examine how the “practice of medicine” is legally defined in their jurisdiction for purposes of ensuring continued regulatory oversight of those who provide healthcare, human or otherwise.
- FSMB should explore and pilot ways in which AI can aid medical boards in decision-making, with the potential to shift from a reactive to a proactive system.
- FSMB should work with state medical boards to help develop policies that address the use of AI systems by licensees, particularly as AI systems become more autonomous.

7. *Continual Review and Adaptation of Law and Regulations:*

- State medical boards, with the support of the FSMB, should continually review and update guidelines and regulations related to AI as it continues to evolve.
- Policy makers should consider the impact of AI on fundamental legal principles such as the definition of practice of medicine and the impact of AI on the corporate practice of medicine.
- FSMB should establish a dedicated team for the ongoing review and adaptation of AI guidelines and regulations.

Section VII. Conclusion

The incorporation of AI in medical practice presents tremendous benefits to patients and physicians alike. It also presents significant risk of harm to patients and physicians if it is developed and used irresponsibly. A sensible approach to the regulation of AI by state medical boards and its incorporation into practice by licensees holds greater promise of realizing AI’s benefits while minimizing potential harms. Adherence to traditional professional expectations for the provision of medical care will help achieve the patient safety goals of physicians and state medical boards.

FSMB Ethics and Professionalism Committee

2023-2024 Members

Mark B. Woodland, MS, MD (Chair), Pennsylvania State Board of Medicine

Eric Eskioğlu, MD, MBA, The Health Management Academy (Subject Matter Expert)

Sara Gerke, Dipl-Jur Univ, MA, Penn State Dickinson Law (Subject Matter Expert)

Leonard Green, MPS, Rhode Island Board of Medical Licensure and Discipline

Jade James-Halbert, MD, Missouri Board of Registration for the Healing Arts

W. Reeves Johnson, Jr., MD, Tennessee Board of Medical Examiners

Christine Khandelwal, DO, North Carolina Medical Board

Heidi Koenig, MD, Kentucky Board of Medical Licensure

Hemesh Patel, DO, MS, Osteopathic Medical Board of California

Sarvam TerKonda, MD, Florida Board of Medicine

Christy Valentine Theard, MD, MBA, Louisiana State Board of Medical Examiners

Ex Officio Members

Jeffrey Carter, MD, FSMB Chair

Katie Templeton, JD, FSMB Chair-elect

Humayun “Hank” Chaudhry, DO, MACP, FSMB President and CEO

FSMB Staff

Eric Fish, JD, Chief Legal Officer

Frank Meyers, JD, Deputy Legal Counsel

Mark Staz, MA, Chief Learning Officer

State Medical Board or organizational affiliations are presented for purposes of identification and do not imply endorsement of any draft or final version of this report.