VIEWPOINT

AI AND CLINICAL CARE

Software as a Medical Practitioner—Is It Time to License Artificial Intelligence?

Eric Bressman, MD, MSHP; Carmel Shachar, JD, MPH; Ariel D. Stern, PhD; Ateev Mehrotra, MD, MPH

Editorial

The Healthy Technology Act of 2025, a bill being considered by Congress, would permit artificial intelligence (AI) systems to prescribe medications without human sign-off. While allowing AI to practice independently may still be a stretch, large language models are already being used to support diagnosis and treatment.[1] These applications challenge the current regulatory framework for clinical AI, which has focused on software as a medical device (SaMD). The US Food and Drug Administration (FDA) has cleared more than 1000 AI tools using this framework, mostly for narrow, well-defined tasks. However, generative models can be applied across domains, straining the SaMD framework. In this Viewpoint, we explore the application of a licensure paradigm to clinical AI systems, arguing that this offers a more reliable regulatory framework.

Limitations of SaMD

Within the FDA's SaMD framework, an algorithm that meets the statutory definition of a medical device may be evaluated through premarket approval, de novo, or 510(k) pathways, depending on risk and existence of a predicate (a previously FDA-cleared comparator). The process is modeled on how the FDA assesses hardware devices like pacemakers and powered wheelchairs. More than 95% of all AI authorizations to date have followed 510(k), by demonstrating "substantial equivalence" to one or more predicates. Sponsors can now preauthorize limited SaMD revisions through predetermined change-control plans that specify allowable model updates. This framework works for algorithms with narrow use cases (eg, an automated image classifier that labels lung lesions), for which metrics are well defined and a predicate likely exists.

Several assumptions of this framework are challenged by the emergence of generative AI models. First, current guidance assumes a static or predictably updated algorithm that is amenable to 1-time review, followed by well-defined change control plans. However, generative models are frequently updated across training and deployment cycles and their behavior can change via fine-tuning, retrieval, and policy updates; this results in a mismatch for 1-time review and predetermined change plans. Second, SaMD presumes a narrow indication label (eg, "classify high-risk lung lesions"). A general-purpose model resists such labeling because of its broad potential scope. A single model might interpret computed tomography scans, dose vancomycin, and counsel on fertility, making it impossible to prespecify every use case. Third, SaMD assumes single manufacturer control, whereas modern models may be built on open-weight foundations, fine-tuned locally, and extended via third-party plug-ins, leaving ownership (and therefore regulatory accountability) uncertain.

Licensure Framework

These shortcomings highlight the need for a different form of oversight that is fit for today's emergent AI tools.[2] We believe that the concept of licensure, the method used for regulating clinicians, can be adapted for AI. Medical licensure in the US emerged in the late nineteenth century amid public concern over quack remedies and variable training. The concerns now attached to generative AI, such as hallucinations, performance drift, and data poisoning, may be unique, but they mirror long-standing worries about clinician competence and misconduct. Licensure is designed to handle such concerns by coupling entry to practice standards with ongoing surveillance and education, thus sustaining public trust.

Regulation of AI tools used by clinicians can be built on the precedent of licensure models for supervised and collaborative practice, such as those for physician assistants and nurse practitioners.[3] Generative AI models could be licensed within a restricted practice framework, requiring physician oversight for each class of function (eg, note writing, medication prescribing), or a reduced practice framework, in which certain lower-risk functions can be handled autonomously.

The **Table** outlines an approach to applying the components of licensure to generative AI. The developer would seek a core license for the base model, and the licensing authority would define a range of competencies relevant to the model's intended scope of practice. Similar to the human practitioner, the model would need to meet minimum performance thresholds in each of these domains in the preclinical setting and then undergo a period of supervised training and practice in an accredited clinical environment.[4] Even after licensure has been granted, there would be multiple mechanisms of supervision in place. This would include regular retesting and a designated discipline board for processing complaints swiftly. Institutional credentialing, like for human practitioners, should play an essential role in local implementation and oversight. By anchoring oversight in these familiar levers (competency testing, maintenance of certification, and disciplinary transparency), licensure can provide the continuous governance that static device regulation was never designed to deliver.

Challenges and Limitations

Governance Architecture

Who would oversee this? A single, novel entity, such as a federal digital licensing board, would be ideal, but less realistic. A pragmatic so-

Table. Components of Licensure

Licensure concept	Human clinician	Generative AI
Prelicensure requirements	Coursework part of accredited degree Passing examinations (eg, USMLE) Supervised period of clinical training	Technical validation for predefined competencies (AI board examinations) Supervised pilot in accredited "implementation centers" (AI "residency")
Scope of practice	Delineation of approved medical services, in which populations, with degree of autonomy Collaboration or supervision agreements for PAs/NPs	Delineation of approved functions (eg, image interpretation), in which populations, with degree of autonomy Guidance on supervising clinician oversight for each function
Institutional credentialing	Hospitals: Credential to perform specific procedures Review outcomes and can suspend privileges for safety concerns	Local health systems' AI governance committees: Vet site-specific implementation Determine local privileges within the licensed AI's scope of practice Monitor local quality metrics, can revoke a privilege or deactivate a model if thresholds not met
Continuing oversight	Continuing education requirements Periodic knowledge assessments (for maintenance of board certification)	"Digital CME/MOC": annual rerun of updated benchmarks for each competency Reporting of clinical performance measures for review by board
Discipline	State medical boards investigate complaints Can fine, suspend, or revoke license, or mandate retraining Actions reported to the National Practitioner Data Bank	Digital boards receive and process complaints Can place AI system on probation; require model patches or additional guardrails; suspend or revoke license Maintain a public database of disciplined models and corrective action plans

Abbreviations: AI, artificial intelligence; CME, continuing medical education; MOC, maintenance of certification; NP, nurse practitioner; PA, physician assistant; USMLE, US Medical Licensing Examination.

lution could entail contributions from existing federal and state bodies, provided each invests in dedicated AI expertise. The FDA could retain its role in premarket assessments, with the broadened approach outlined previously, to prevent developers from needing to submit to 50 state licensing authorities. Designated health systems with recognized AI expertise could function as implementation centers; others have proposed roles for independent nonprofit organizations.[5] Continuing oversight and discipline could be an adjunctive function of state medical boards (eg, a digital division), ideally with a federal coordinating body to harmonize standards.

Liability

Licensure only works if accountability is clear when harm occurs. We envision liability scaling with autonomy: within a restricted license, or for higher-risk functions, liability rests with the supervising clinician and institution; when models are granted a degree of autonomy for lower-risk functions, developers must assume responsibility, including in instances of malpractice. This raises challenges, but paves the way for functioning AI insurance markets.[6]

Conclusions

Generative AI systems are evolving rapidly, and existing regulatory frameworks are no longer sufficient to provide oversight for their broad clinical skill set. General-purpose large language models have been left unregulated even though they are being used in practice and likely meet the definition of medical devices outlined in the 21st Century Cures Act.[7] Amid uncertainty around the AI regulatory environment (as the current administration has signaled a lighter touch), there may be an opportunity to craft a more agile, forward-thinking approach for clinical AI. A licensure framework may help ensure that innovation scales with accountability and not ahead of it.

ARTICLE INFORMATION

Author Affiliations: Department of Medicine, University of Pennsylvania, Philadelphia (Bressman); Leonard Davis Institute of Health Economics, University of Pennsylvania, Philadelphia (Bressman); Health Law and Policy Clinic at Harvard Law School, Cambridge, Massachusetts (Shachar); Hasso Plattner Institute, University of Potsdam, Potsdam, Germany (Stern); School of Public Health, Brown University, Providence, Rhode Island (Mehrotra).

Corresponding Author: Eric Bressman, MD, MSHP, Department of Medicine, University of Pennsylvania, 423 Guardian Dr, 12th Floor, Philadelphia, PA 19104 (eric.bressman@pennmedicine.upenn.edu).

Published Online: November 17, 2025. doi:10.1001/jamainternmed.2025.6132

Conflict of Interest Disclosures: Dr Shachar reported grants from Viiv, the BMSF Foundation, Gilead, AbbVie, Commonwealth Fund, Massachusetts Department of Public Health, and Kaiser Permanente outside the submitted work. Dr Stern reported personal fees from Peterson Health Technology Institute and Roche Diagnostics; nonfinancial support from the German Society for Digital Medicine; and a role as a research fellow of the International Collaborative Bioscience Innovation & Law Programme, which is supported by Novo Nordisk Foundation grant NNF23SA0087056, outside the submitted work. Dr Mehrotra reported personal fees from Black Opal Ventures outside the submitted work. No other disclosures were reported.

REFERENCES

1. Blease CR, Locher C, Gaab J, Hägglund M, Mandl KD. Generative artificial intelligence in primary care: an online survey of UK general practitioners. *BMJ Health Care Inform*. 2024;31(1):e101102. doi:10.1136/bmjhci-2024-101102

2. Babic B, Glenn Cohen I, Stern AD, Li Y, Ouellet M. A general framework for governing marketed AI/ML medical devices. *NPJ Digit Med*. 2025;8(1):328. doi:10.1038/s41746-025-01717-9

3. Morrell W, Shachar C, Weiss AP. The oversight of autonomous artificial intelligence: lessons from nurse practitioners as physician extenders. *J Law Biosci*. 2022;9(2):lsac021. doi:10.1093/jlb/lsac021

4. Longhurst CA, Singh K, Chopra A, Atreja A, Brownstein JS. A call for artificial intelligence implementation science centers to evaluate clinical effectiveness. *NEJM AI*. 2024;1(8):AIp2400223. doi:10.1056/AIp2400223

5. Shah NH, Halamka JD, Saria S, et al. A nationwide network of health AI assurance laboratories. *JAMA*. 2024;331(3):245-249. doi:10.1001/jama.2023.26930

6. Stern AD, Goldfarb A, Minssen T, Price II WN. AI insurance: how liability insurance can drive the responsible adoption of artificial intelligence in health care. *NEJM Catalyst Innovations in Care Delivery*. 2022;3(4):CAT.21.0242. doi:10.1056/CAT.21.0242

7. Weissman GE, Mankowitz T, Kanter GP. Unregulated large language models produce medical device-like output. *NPJ Digit Med*. 2025;8(1):148. doi:10.1038/s41746-025-01544-y