


Health AI Principles for State Policymakers

Position Paper



About Connected Health Initiative

The Connected Health Initiative (CHI) is a multistakeholder coalition dedicated to advancing the responsible development, deployment, and use of innovative technologies and digital solutions in healthcare. As a collective voice for stakeholders across the healthcare ecosystem, including technology companies, healthcare providers, payers, patient advocacy groups, and research institutions, CHI works to promote policies that enable seamless connectivity, improved care delivery, and better health outcomes. With a focus on leveraging technologies like artificial intelligence (AI), telehealth, and remote patient monitoring, CHI seeks to foster collaboration, innovation, and the responsible use of technology in healthcare.

The United States faces projected workforce shortages of approximately 141,000 physicians, 355,000 nurses, and 423,000 behavioral health professionals by 2038. Clinicians, patients, consumers, and caregivers increasingly rely on digital health tools, including those with AI components such as clinical decision support, ambient documentation, predictive analytics, and generative AI assistants, to deliver higher-quality, more accessible, and more cost-effective care. As adoption accelerates, clinicians, health systems, and patients benefit from clearer information about how to evaluate these tools for safety, efficacy, real-world performance, and impact on care.

The role of state governments in facilitating responsible health AI is meaningful but should be carefully calibrated; a growing patchwork of state laws risks creating compliance burdens that fall disproportionately on small and mid-sized developers, delaying patient access to beneficial tools without measurably improving safety. These principles give state policymakers a framework to evaluate interventions related to health AI across key areas including quality assurance and oversight, thoughtful design, access and affordability, transparency, ethics, and workforce development.

The Quadruple Aim, which serves as the guiding star for these principles, captures four interrelated objectives: improving the patient experience of care, augmenting population health management, reducing costs, and supporting the healthcare workforce. CHI believes that responsibly designed and deployed health AI, and digital health tools more broadly, can and must contribute to achieving these aims. State policy should be evaluated against whether it accelerates or impedes that contribution.

We urge state legislatures, executive agencies, and other policymakers to refer to these principles as a guide as they consider the role of state law and policy in the use of AI in the delivery and management of healthcare. CHI is committed to being a resource for state legislatures to achieve our shared goals for the responsible development and deployment of healthcare tools with AI components and we look forward to engaging in working groups, task forces, and roundtable discussions.

Ten Principles for Evaluating State Health AI Policy

- 1 A Coordinated Approach to AI in Healthcare
- 2 Quality Assurance and Oversight
- 3 Thoughtful Design and Oversight
- 4 Access and Affordability
- 5 Transparency and Research
- 6 Modernized Privacy and Security Frameworks
- 7 Bias Mitigation
- 8 Ethics
- 9 Workforce and Education
- 10 Intellectual Property



1 A Coordinated Approach to AI in Healthcare

A wide range of federal, local, state, and international laws prohibit harmful conduct in health-related contexts regardless of whether the use of AI is involved. Generally, policymakers should leverage existing authorities first and adopt AI-specific requirements to address identified gaps. For example, all state consumer protection statutes prohibit a wide range of unfair or deceptive acts or practices in some form or another. A growing number also have more specific requirements with respect to privacy and data security. The use of AI does not shield companies from these prohibitions; some state privacy laws even require algorithm users to honor consumer requests to exclude personal information about them in algorithm training. However, state agencies must approach the applicability of these laws in AI contexts thoughtfully and with great sensitivity to the novel or evolving risks AI systems present. State legislators and other policymakers must first understand how existing laws apply to activities involving AI, including in healthcare contexts, to avoid creating sweeping new authorities or agencies that awkwardly or inconsistently overlap with current policy frameworks.

Relevant federal anchors include U.S. Food and Drug Administration (FDA) oversight of AI-enabled software as a medical device, Office of the National Coordinator for Health's (ONC) certification and information-blocking rules, the Federal Trade Commission (FTC) Act, and the National Institute for Standards and Technology (NIST) AI Risk Management Framework. Where states do act, they should look to harmonize with these federal and consensus frameworks to reduce state-to-state fragmentation that disproportionately burdens small developers and slows patient access.

States should also recognize their distinct leverage as purchasers, through Medicaid, state employee health plans, public hospitals and universities, and corrections health to advance responsible practices through procurement standards rather than broad regulation. State grant programs offer a similar opportunity to effect responsible health AI policy through funding choices and program design rather than prescriptive mandates. As a prime example, the Rural Health Transformation Program (RHTP), through which the Centers for Medicare & Medicaid Services (CMS) is distributing \$50 billion to states between FY2026 and FY2030, gives states significant discretion in their approved plans to direct funds toward technology-driven solutions, including AI-enabled tools that expand access, support rural workforces, and modernize care delivery. State plans across other grant programs (workforce development, behavioral health, maternal health, and broadband-enabled care) can and should advance the responsible deployment of AI in healthcare through state funding conditions.



2 Quality Assurance and Oversight

Policy frameworks should utilize risk-based approaches to ensure that the use of health AI aligns with any relevant recognized standards of safety and efficacy. Risk-based oversight should be calibrated to the AI tool's intended use and clinical context, the clinical significance of the decision being supported, the consequence of error, the level of autonomy of the tool, and the degree of human oversight in the workflow. Oversight should additionally distinguish among FDA-regulated AI-enabled medical devices, non-device clinical decision support, administrative and operational AI, and patient-facing tools, avoiding duplicative obligations on tools already subject to federal premarket review or post-market surveillance requirements.



All stakeholders will benefit from understanding the distribution of risk and liability in building, testing, and using AI tools. Policy frameworks addressing liability should ensure the appropriate distribution and mitigation of risk and liability. Specifically, those in the value chain with the ability to minimize risks based on their knowledge and ability to mitigate should bear the burden of doing so. Some recommended areas of focus include:

- Ensuring AI is safe and efficacious and providing transparent, auditable data demonstrating this safety and efficacy
- Encouraging AI developers to consistently utilize rigorous procedures and enabling them to document their methods and results.
- Encouraging those developing, offering, or testing AI systems intended for use in the healthcare value chain to provide truthful, meaningful and easy-to-understand representations regarding intended use, risks, and limitations that would be reasonably understood by those intended and expected to use the AI solution. Such representations should also address the responsible design choices that shape the system's behavior and its performance across relevant dimensions, which may include fairness, explainability, reliability and robustness, safety, controllability and human oversight, privacy and security, transparency, and governance across the AI supply chain.
- Updating/patching solutions to address known, demonstrated vulnerabilities and harms.

3 Thoughtful Design and Oversight

Policy frameworks should encourage design of health AI systems that are informed by real-world workflows, human-centered design and usability principles, and end-user needs. Tools should integrate with existing health technology infrastructure and interoperability standards so that AI enhances care. AI systems should facilitate a transition to changes in the delivery of goods and services that benefit consumers and businesses. The design, development, and deployment of AI should leverage collaboration and dialogue among users, AI technology developers, and other stakeholders to have all perspectives reflected in AI solutions.

AI tools should be enabled to support the judgment of licensed clinicians and other accountable decision-makers. For higher-risk uses, such as diagnoses, treatment recommendations, and coverage or utilization management decisions that restrict or deny access to care, meaningful human review is and will remain vital while other lower risk uses that can affirm or expedite access may not require the same case-by-case human review. Oversight requirements should reflect that lower-risk administrative or documentation tools should not be subject to the same level of review as tools that drive clinical or coverage decisions.



4 Access and Affordability

Policy frameworks should enable health-related products and services that involve AI systems to be accessible and affordable. Even where digital health tools and AI-enabled services demonstrably improve outcomes and reduce total cost of care, legacy coverage and payment frameworks often fail to recognize them, leaving patients, clinicians, and health systems without a sustainable path to use them at scale. States should modernize Medicaid coverage policies, state employee health plan benefit design, and state-level payment models to reimburse digital health and AI-enabled services and support outcomes-based and value-based arrangements that share the savings these tools can generate.

5 Transparency and Research

Policy frameworks should support and facilitate research and development of healthcare AI by prioritizing and providing sufficient funding while also maximizing innovators' and researchers' ability to collect and process data from a wide range of sources. State disclosure regimes should align transparency expectations with established consensus frameworks, including the FDA's labeling and transparency expectations for AI-enabled devices and the NIST AI Risk Management Framework. Transparency obligations should be calibrated to risk and to the appropriate audience.

Research on the costs and benefits of transparency in AI should also be a priority and involve collaboration among all affected stakeholders to develop a better understanding of how and under which circumstances transparency mandates would help address risks arising from the use of AI systems.



6

Modernized Privacy and Security Frameworks

The many new AI-driven uses for data, including sensitive personal health information and data used to train, fine-tune, or prompt AI systems, raise privacy questions but also offer the potential for more powerful and granular privacy controls for consumers. Policy frameworks must be scalable and assure that an individual's data is properly protected, both respecting an individual's right to privacy and maintaining robust cybersecurity protections, while also allowing the flow of information and responsible evolution of AI. A balanced privacy and security framework should ideally rest on a comprehensive federal floor that applies consistently across state lines, which would lessen entity compliance burden, avoid undue barriers to data processing and collection while advancing reasonable data minimization, data use transparency, consent, and consumer/patient rights frameworks.

7

Bias Mitigation

Biases in datasets are, and will remain, a pressing issue with AI systems, particularly those that utilize machine learning techniques. Data provenance and bias issues present in the development and uses of AI solutions should be mitigated to avoid harm to users or consumers of products or services involving AI, including through unlawful discrimination. Beyond examination at the development stage, developers and deployers should continue to evaluate model performance across relevant subpopulations, monitor for performance drift in real-world use, and document known limitations so that clinicians and patients can rely on AI tools appropriately. State oversight should focus on real-world testing and on whether AI tools are used consistently with their validated intended use. Bias in datasets that results in harm to users or consumers should remain actionable under existing civil rights and consumer protection law.

8 Ethics

The success of AI's use in healthcare depends on an ethical approach. A policy framework must promote many of the existing and emerging ethical norms for broader adherence by AI technologists, innovators, computer scientists, and those who use such systems. Where appropriate, state authorities may need to set and enforce standards (e.g., professional conduct and the standard of care for licensed clinicians) about AI's use in care. Relevant ethical considerations include:

- Applying ethics to each phase of an AI system's life, from design to development to use.
- Ensuring that AI solutions benefit consumers across socioeconomic, age, gender, geographic origin, and other groupings.
- Appreciating that AI tools may reveal sensitive and private information about a user and ensure that laws require the protection of such information.

9 Workforce and Education

Policy frameworks should support education for the advancement of AI, promote examples that demonstrate the success of AI, and encourage stakeholder engagements to keep frameworks responsive to emerging opportunities and challenges. States should support the current healthcare workforce in using AI tools well, recognizing that the successful integration of AI into care delivery depends on clinicians and administrative staff who have the time, training, and institutional support to use AI tools effectively and safely.

States, licensing boards, and continuing-education accreditors should play a constructive role in setting expectations for AI literacy across professions.

- Academic education should include curricula that will advance the understanding of and ability to use AI solutions.
- The healthcare workforce should have access to in-service training, continuing education, and protected time to develop competency with AI tools as they are introduced into clinical and administrative workflows.
- Workforce development grants, apprenticeships, and rural and safety net workforce investments should be designed to build AI fluency where it is most needed and to reduce administrative burden on clinicians and care teams.

10 Intellectual Property

The protection of IP rights is critical to the evolution of health AI. In developing approaches and frameworks for health AI governance, state policymakers should be mindful of how current legal protections apply in circumstances involving AI and ensure that compliance measures and requirements do not undercut IP or trade secrets, including training data.



ConnectedHealthInitiative

