

PROPOSED POLICY POSITIONS: COMPREHENSIVE APPLICATION OF HUMAN OVERSIGHT FOR AI IN CLINICAL DECISION SUPPORT TOOLS

DEFINITIONS

- *Human oversight* is the process whereby a qualified individual or defined group of individuals oversees AI systems.¹
 - Key to effective and ethical AI adoption in healthcare are humans with the appropriate knowledge and skillset making final clinical choices, using AI as a support tool, to ensure patient safety.^{2,3}
- *Human-in-the-loop* is a concept that is widely endorsed but poorly defined.⁴ As such, AMIA suggests considering the following specific loops.⁵
 - *The design loop* establishes a continuous human-in-the-loop process by engaging clinicians and patients in human-centered design to ensure AI-enabled CDS is co-designed, implemented, and refined through ongoing human feedback, clinical expertise, and person-centered evaluation.
 - *The clinical loop* is where the AI-enabled CDS meets point-of-care decision-making: who initiates the prompt, who receives and interprets the output, how the output is presented, and whether clinicians accept, modify, or override it in practice.⁶
 - *The governance loop* governs approval and deployment: what AI-enabled CDS gets deployed where, who approves it, what evidence is required, and what trade-offs are acceptable.⁷
 - *The learning loop* oversees monitoring and model updates after deployment: who monitors performance, detects drift, audits accuracy and equity, manages updates, and bears responsibility when the model changes.⁸
- *Clinical decision support (CDS)* is a computerized tool that provides timely and person-specific information, intelligently filtered or presented at appropriate times, to enhance patient outcomes and quality of care. It can be AI-enabled.
 - CDS is used to drive care quality, safety, efficiency, and effectiveness.
- Anyone may use *patient-centered CDS*, which help to boost engagement by involving consumers, patients, or caregivers with their clinicians and as active participants in the decision-making process.⁹ This may include wearable devices, chatbots, care coordination

¹ <https://artificialintelligenceact.eu/article/14/>

² van Voorst R. Challenges and Limitations of Human Oversight in Ethical Artificial Intelligence Implementation in Health Care: Balancing Digital Literacy and Professional Strain. *Mayo Clin Proc Digit Health*. 2024 Sep 7;2(4):559-563. doi: 10.1016/j.mcpdig.2024.08.004. PMID: 40206522; PMCID: PMC11976012.

³ <https://www.lexisnexis.com/blogs/en-ca/b/legal-ai/posts/ethical-consideration-ai-adoption-human-oversight>

⁴ *NEJM AI* 2026;3(3) DOI: 10.1056/Ale2600084.

⁵ See Figure 1.

⁶ *Ibid.*

⁷ *Ibid* 4.

⁸ *Ibid* 4.

⁹ Clinical Decision Support. Office of the National Coordinator of Health Information Technology. April 1, 2026. <https://healthit.gov/clinical-quality-and-safety/clinical-decision-support/>.

tools, preliminary explanations of results, symptom checkers, triage apps, and gamified digital therapeutics.

- CDS are used at any point-of-care or point-of-need to drive care quality,¹⁰ safety, efficiency, and effectiveness.

AMIA POLICY POSITION: AMIA supports these definitions for human oversight, human-in-the-loop (design loop, clinical loop, governance loop, and learning loop), AI-enabled clinical decision support.

ISSUE AREAS

I. MEANINGFUL HUMAN OVERSIGHT IS INTEGRAL TO AI HEALTHCARE POLICY

1. “Meaningful” oversight requires more than just a human in the loop at the point of decision-making.

Keeping humans in the loop is not a sufficient safeguard unless the appropriate humans with the relevant expertise to the decision at hand are empowered at the right points in the system.¹¹ Oversight in the form of “a clinician reviews the output,” only works if clinicians can understand and appropriately use the system. Clinicians undergo years of medical training, not computational training. They may have varying understanding of how AI-enabled CDS algorithms work, yet they are expected to understand algorithmic decision-making and override recommendations at the point-of-care.

“[T]he problem [is] that the increase of algorithmic technology is accompanied by an unrealistic expectation for professional caretakers to fully understand it, so they can serve as effective human overseers [sic] in human-nonhuman (algorithmic) decision making. The oversighting role of humans is necessary condition for labeling AI as ‘ethical,’ where the idea is that the algorithm provides a calculation or advice to the human physician and it is then up to that human physician to either follow or deviate from that advice.”¹²

Proper oversight requires at a minimum:

- Transparency concerning the source, scope, and characteristics of training data and model design
- Clear communication of intended use, limitations, and risks
- Agreed coding standards like drug labeling with indications and contraindications
- Clinical validation, post-deployment monitoring, and updating the tool with learnings from the monitoring

¹⁰ Quality is STEEP (Safe, Timely, Effective, Efficient, Equitable, Patient-centered). Crossing the Quality Chasm: A New Health System for the 21st Century (2001).

¹¹ NEJM AI 2026;3(3); DOI:10.1056/Ale2600084

¹² van Voorst R. Challenges and Limitations of Human Oversight in Ethical Artificial Intelligence Implementation in Health Care: Balancing Digital Literacy and Professional Strain. Mayo Clin Proc Digit Health. 2024 Sep 7;2(4):559-563. doi: 10.1016/j.mcpdig.2024.08.004. PMID: 40206522; PMCID: PMC11976012.

- A system that provides explanation and documentation that is reviewed by humans, not exclusively a secondary AI.

2. Safe and effective use of AI-enabled CDS tools requires that the higher the potential risk of harm, the greater the need of human oversight.

The degree of human supervision determines how risky an AI-enabled CDS tool is.

- Systems that are fully autonomous or act directly based exclusively on patient data should be subject to the highest level of regulatory scrutiny.
- By contrast, CDS mediated by clinicians (human-in-the-loop) is generally seen as lower risk.

AMIA POLICY POSITION: AI-based decision support requires ongoing human engagement with clinicians and patients at any point-of-care, and with the appropriate person with relevant subject matter expertise empowered to impact the design, clinical, governance, and learning loops.

II. FEDERAL OR STATE POLICY ON HEALTHCARE AI MUST INCLUDE PROPER OVERSIGHT ON SEVERAL LEVELS, IN ADDITION TO BASELINE HUMAN OVERSIGHT

1. Relying solely on industry or clinician oversight is inadequate.

- a. Gaps in oversight can expose patients to bias and safety risks.
- b. Systems with limited opportunity for human intervention should receive stronger regulatory control.
- c. Apps that are intended or marketed to diagnose, suggest treatment, analyze health data, and are patient-facing, should clearly inform users if they have been submitted to the FDA for review and if they are: Class I (exempt), Class II (moderate risk) and cleared, Class III (high risk) and approved.
- d. Apps developed through the CMS Health Tech Ecosystem and accepted into the Medicare App Library offer a promising model of meaningful human oversight. They undergo independent review through either the DiMe Seal or the CARIN Alliance pathway, which includes privacy, security, clinical evidence, usability, equity, and technical interoperability criteria.¹²

2. A patchwork of State policymaking is likely to prove confusing, leading to inconsistencies in decision making in comparable situations in different states.

- a. Compliance may also be compromised, as efforts to satisfy one set of State rules may conflict with similar efforts in another State.

3. Human oversight must be embedded in a broader governance framework.

Oversight is not merely the responsibility of one or a collection of individuals. Rather, effective oversight requires:

- A clear statement of the risks in the use of AI and/or CDS with a view to mitigation,

- An oversight plan must be meaningful, informed, and supported by transparency,
- Oversight within a broader governance and regulatory framework, especially for adaptive AI,
- Institutional governance bodies exercising authority over AI use,
- Ongoing monitoring, evaluation, and maintenance of models,
- Lifecycle management from design → deployment → post-market performance.

4. Oversight is especially critical for adaptive (learning) systems

Adaptive CDS continues to learn from the cases it encounters in actual practice, so its underlying algorithms change after deployment. This creates new risks:

- Performance can drift or change after deployment,
- Biases may emerge or worsen over time, and
- Inconsistencies may arise in decisions at different times or in different settings.

AI in clinical decision support should augment—not replace—clinicians, with accountable humans in the loop, but only within a system of transparency, monitoring, and institutional oversight. “Clinician checks output” is not enough given the limited information a clinician has in an interaction—oversight includes organizational accountability and continuous supervision.¹³

AMIA POLICY POSITION: Policymakers should require that oversight means: not relying solely on industry or clinical judgement, consistency and enforceability in federal law rather than a patchwork of state laws, and application of a broad governance framework with a special emphasis on adaptive learning systems.

III. DEVELOP A PUBLIC-PRIVATE PARTNERSHIP TO STUDY, TEST AI IN HEALTHCARE

1. Federal policymakers should create a public-private partnership of healthcare stakeholders to thoroughly vet and recommend AI health care policies.

The partnership must include key HHS government agencies (ONC, CMS, AHRQ, HRSA), industry, professional societies, health systems, academia, and other patient and provider stakeholders.

2. AI tools intended to target patient engagement should be produced and/or evaluated in partnership with the patient population on which they are used and developed through public-private partnerships.

- Patient-facing technologies have greater reach and impact when developed in collaboration with people who are living with or at risk for the conditions the technical tool is targeting.

¹³ *Journal of the American Medical Informatics Association*, Volume 28, Issue 4, April 2021, Pages 677–684, <https://doi.org/10.1093/jamia/ocaa319>

- Human-centered design is an evidence-based approach that is focused on building solutions that people find both usable and useful^{14,15}. In the development of AI tools, human-centered design can mitigate bias by validating the assumptions a developer makes about the values and needs of people who are engaging with an AI tool for health.¹⁶
- Model testing is critically important. Policymakers should create a sophisticated CMMI model to study the relationship between humans and AI in the provision of patient care. Model testing, at a minimum, should include careful review of:
 - Patient data privacy,
 - Documentation and cognitive burden,
 - Standardized measurement approaches for data centers, transparency requirements, regulatory frameworks to incentivize sustainable AI deployment, and
 - Fact-checking AI hallucinations.

AMIA POLICY POSITION: Policymakers must clarify accountability, liability, and financial models by better understanding the role of humans and the role of AI in the provision of patient care through testing and analysis that deliberately includes broad stakeholder involvement.

Figure 1: Human-in-the-Loop

¹⁴ Yen, P. Y., & Bakken, S. (2012). Review of health information technology usability study methodologies. *Journal of the American Medical Informatics Association*, 19(3), 413-422.

¹⁵ Papoutsis, C., Wherton, J., Shaw, S., Morrison, C., & Greenhalgh, T. (2021). Putting the social back into sociotechnical: case studies of co-design in digital health. *Journal of the American Medical Informatics Association*, 28(2), 284-293.

¹⁶ Amershi, S., Weld, D., Vorvoreanu, M., Fourney, A., Nushi, B., Collisson, P., ... & Horvitz, E. (2019, May). Guidelines for human-AI interaction. In *Proceedings of the 2019 chi conference on human factors in computing systems* (pp. 1-13).

Human in the Loop Clinical Decision Support

AI Supports Better Decisions. **Humans remain responsible.**

