WHITE PAPER

Why FDA Clearance Does Not Clear Providers from "High-Risk" AI Liability

Dangerous Misunderstanding of the Post-Deployment Monitoring Gap, Emerging State Legislation, and the Imperative for Continuous AI Governance in Healthcare

A White Paper from Asher Informatics Leadership
February 2026

Authors:
Charlotte Kalafut, CEO & Co-Founder
John F. Kalafut, PhD, CTO & Co-Founder

Executive Summary

Healthcare systems across the United States are rapidly adopting AI-enabled medical devices and software, with the FDA now listing over 1,250 authorized AI/ML-enabled products. A pervasive and dangerous assumption has taken root among healthcare delivery organizations: that FDA clearance provides comprehensive risk coverage for these AI tools throughout their operational life. **It does not.**

FDA 510(k) clearance, the pathway through which over 96% of AI medical devices reach market, establishes that a product is substantially equivalent to a predicate device for the purpose of *marketing permission*. It does not validate long-term clinical efficacy, guarantee freedom from algorithmic bias, or ensure sustained performance in real-world clinical environments. A 2025 JAMA study of 691 FDA-cleared AI/ML devices found that 46.7% of decision summaries failed to report study designs, 95.5% omitted demographic information, and fewer than 2% cited randomized clinical trials.[1] The FDA's updated Clinical Decision Support (CDS) Guidance (January 6, 2026) reinforces that classification boundaries do not eliminate deployer responsibility.

Meanwhile, landmark research published in *Scientific Reports* (Nature Publishing Group) demonstrated that 91% of machine learning models experience temporal quality degradation after deployment, a finding with profound implications for clinical AI systems making diagnostic and treatment decisions.[2] When these two realities intersect with emerging legislation like the Colorado Artificial Intelligence Act (SB 24-205) for example, which holds *deployers* legally liable for the AI systems they operate, healthcare organizations face a compounding risk that many have not yet recognized.

The critical gap: FDA clearance addresses how a product can be marketed. It does not address how that product performs after it enters your clinical environment. The deploying organization bears that responsibility, and emerging state laws are codifying it - regardless of how broader federal AI policy continues to evolve.

1. What FDA Clearance Actually Means, and What It Does Not

1.1 The 510(k) Pathway: Marketing Permission, Not Clinical Validation

The FDA's 510(k) pre-market notification pathway is the dominant route to market for AI/ML-enabled medical devices. Under this pathway, a manufacturer demonstrates that its product is "substantially equivalent" to a legally marketed predicate device in terms of intended use and technological characteristics. If the FDA concurs, the device receives *clearance* to be marketed, a designation that is explicitly **not the same as approval.**[3]

The Bipartisan Policy Center's 2025 analysis of FDA oversight for health AI tools affirms this distinction clearly: 510(k) clearance does not imply endorsement or validation of clinical benefit; it only means that the device meets the statutory standard of safety and effectiveness relative to its predicate.[4] This is a critical nuance that many healthcare organizations overlook when making procurement and deployment decisions.

1.2 The Evidence Gap in Cleared AI Products

The pre-market evidence submitted for most AI/ML devices is far thinner than healthcare leaders might expect. A cross-sectional study published in JAMA Health Forum in 2025 analyzed all 691 AI/ML devices cleared through July 2023 and found alarming gaps:[1]

- **46.7%** of FDA decision summaries did not describe the study design
- **53.3%** omitted training sample size
- **95.5%** did not report demographic information
- **Only 6 devices (1.6%)** cited a randomized clinical trial
- **Fewer than 1%** reported actual patient health outcomes
- **Only 28.2%** of device summaries mentioned any safety assessment

A 2025 systematic review published in JAMA Network Open examining FDA-authorized radiology AI devices confirmed these patterns: of 717 radiology devices with submission documentation, only 5% underwent prospective testing, 8% included a human-in-the-loop evaluation, and 29% incorporated any form of clinical testing.[5] Just15 devices across the entire dataset employed both prospective and clinical testing.

1.3 Post-Market Surveillance: The Manufacturer's Responsibility That Falls Short

While the FDA expects manufacturers to monitor real-world performance and report adverse events, the system relies heavily on voluntary compliance. A PMC-published review titled "The Illusion of Safety" found significant inconsistency in post-market surveillance guidelines, leading to industry-wide variability that makes it difficult to compare safety and performance across different AI/ML-enabled medical devices.6 The review recommended that post-deployment evaluation should prioritize patient outcomes, health system performance, and workforce impact, metrics that are not currently standardized.

Key Insight: FDA clearance tells you a product met a pre-market standard for marketing authorization. It tells you nothing about how that product will perform in your specific patient population, with your infrastructure, and over time as clinical data distributions evolve.

1.4 FDA's CDS Guidance Clarifies Classification, Not Operational Accountability

The January 6, 2026, update to the FDA's CDS Software Guidance clarifies when certain CDS functions are excluded from the definition of a medical device under the 21st Century Cures Act.

This clarification is often treated as a product classification determination. In practice, it establishes a boundary condition.

To qualify as Non-Device CDS, software must meet four statutory criteria. The fourth requires that a healthcare professional be able to independently review the basis for the recommendation and not rely primarily on the software to make a clinical decision. That requirement is frequently interpreted as a design feature. It is more accurately understood as an operational condition.

Whether independent review is meaningfully preserved depends on how the system is deployed - within time-constrained environments, alongside other AI systems, and within real-world reliance patterns. The Guidance does not assess those environmental factors. It defines classification.

Accordingly, FDA status does not eliminate deployer accountability. Where AI systems are embedded across shared workflows, preserving independent clinical judgment becomes an enterprise governance responsibility, not a product-level determination.

2. The 91% Problem: AI Model Degradation Is Not a Theory

2.1 The Landmark Study

A seminal study published in *Scientific Reports* (Nature Publishing Group) by researchers from MIT, Harvard, the University of Monterrey, and Cambridge provides the most comprehensive analysis to date of AI temporal quality degradation.[2] The researchers developed a testing framework and applied it to 32 datasets across four industries (healthcare, transportation, finance, and weather) using four standard ML model types (Linear Regression, Random Forest, XGBoost, and Multilayer Perceptron Neural Network).

The finding: across all 128 model-dataset pairs tested, temporal model degradation was observed in 91% of cases.

The study identified several critical degradation patterns that are directly relevant to clinical AI deployments:

- **Gradual degradation:** Model accuracy erodes steadily over time, even when underlying data distributions remain relatively stable.
- **Abrupt breakage points:** Models perform acceptably for extended periods and then experience sudden, dramatic performance collapse, often without warning.
- **Increasing error variability:** Even when median model error remains acceptable, the gap between best-case and worst-case predictions widens significantly, creating unreliable outputs that mask themselves behind average performance metrics.

2.2 Why This Matters for Clinical AI

In a clinical context, these degradation patterns carry direct patient safety implications. Consider an AI triage system that prioritizes radiology worklists based on suspected pathology. If that system's sensitivity gradually declines by even a few percentage points over 12 months (a pattern well within the study's findings), critical findings may be delayed without any single dramatic failure event to trigger investigation.

Supplementary research in *Nature Communications* (2024) reinforced these concerns with experiments on real-world medical imaging data, concluding that monitoring model performance alone is not a reliable proxy for detecting data drift, and that institutional guidance on appropriate steps after drift detection remains lacking.[7]

A 2025 JAMA Network Open prognostic study of 143,049 patients across seven hospitals found that significant data shifts caused by changes in demographics, hospital types, admission sources, and laboratory assays substantially degraded clinical AI performance, particularly during the COVID-19 pandemic.[8] The study demonstrated that proactive monitoring and continual learning strategies were essential to maintaining safe and equitable AI deployment.

The clear implication: an AI model that was clinically validated at the time of FDA clearance is statistically likely to degrade in your environment after deployment. Without continuous monitoring, you will not know when or how severely this degradation is affecting patient care.

3. The Colorado AI Act: A Case Study in Deployer Liability

The FDA's Clinical Decision Support (CDS) Guidance defines a classification boundary. That boundary depends on meaningful independent clinical review in practice. Where independence erodes, exposure does not disappear - it shifts to the deploying organization. Emerging state-level frameworks reflect this broader shift from product regulation to deployment liability.

3.1 Why Colorado Matters for Every Healthcare Organization

The Colorado Artificial Intelligence Act (SB 24-205), enacted May 17, 2024, with enforcement beginning June 30, 2026, is the most comprehensive state-level AI regulation in the United States.[9] It is not merely one state's policy experiment. It is the template that other states are watching, modeling, and in some cases directly copying. For healthcare organizations, Colorado's Act demands particular attention for three reasons.

First, the Act explicitly designates healthcare services as a domain where AI systems make "consequential decisions", meaning decisions that have a material legal or similarly significant effect on consumers.[10] This means that **any AI system involved in clinical triage, diagnostic prioritization, treatment recommendations, or even administrative decisions that affect access to care is automatically classified as high-risk** under the Act. There is **no opt-out**, **no size threshold** for this classification, and **no exemption** based on FDA clearance status.

Second, the Act applies to any entity **"doing business in Colorado"**, a phrase that Colorado interprets broadly to include organizations that solicit business from or serve Colorado residents, even without a physical presence in the state. Multi-state health systems, telehealth providers, and organizations with Colorado-based patients should assume the Act applies to them.

Third, violations are treated as deceptive trade practices under the Colorado Consumer Protection Act, carrying penalties of **up to $20,000 per violation**, with exclusive enforcement authority held by the Colorado Attorney General (AG).[9]

3.2 The Developer/Deployer Distinction: Where the Risk Falls

The Act draws a critical distinction between two roles that most healthcare organizations will occupy simultaneously:

Role	Definition	Primary Obligations
Developer	An entity that builds or substantially modifies an AI system	Duty of care; provide model documentation, dataset cards, performance evaluations to deployers; disclose risks to AG within 90 days
Deployer	An entity that uses a purchased or internally developed high-risk AI system	Risk management program; impact assessments (initial + annual); consumer notification; correction/appeal rights; AG disclosure; public transparency statement

Most healthcare delivery organizations are deployers. However, a hospital that builds its own clinical decision support tool, triage bot, or predictive algorithm becomes a **developer** under the Act, triggering additional reporting requirements to the Attorney General and intensive dataset documentation obligations. Many organizations will find themselves occupying both roles simultaneously.

3.3 Specific Deployer Obligations Under the Colorado AI Act

The Act imposes a structured set of obligations on deployers of high-risk AI systems. For healthcare organizations, these obligations represent a significant operational lift that cannot be satisfied by relying on vendor assertions or FDA clearance alone:

- **Risk Management Policy and Program:** Deployers must implement a formal, documented program to govern the deployment of every high-risk AI system in their organization. This is not a one-time exercise; it must be an ongoing, operationalized program.
- **Impact Assessments:** Deployers must complete an initial impact assessment before or upon deployment, conduct annual reviews thereafter, and update the assessment within 90 days of any substantial modification to the AI system. Assessments must evaluate risks of algorithmic discrimination, describe the data processed, document the AI system's outputs and limitations, and identify transparency measures. These assessments must be retained for three years.
- **Consumer Notification:** Before a high-risk AI system makes, or serves as a substantial factor in making, a consequential decision concerning a patient, the deployer must notify that patient. In a healthcare context, this means informing patients when AI is materially involved in their triage prioritization, diagnostic workup, or treatment pathway.
- **Correction and Appeal Rights:** Patients must be given the opportunity to correct any incorrect personal data that the AI system processed, and to appeal adverse consequential decisions through human review (where technically feasible).
- **Attorney General Disclosure:** If the deployer discovers that a high-risk AI system has caused algorithmic discrimination, it must notify the Colorado Attorney General within 90 days of discovery.
- **Public Transparency Statement:** Deployers must publish and maintain a statement on their website describing the types of high-risk AI systems deployed, how they manage known or foreseeable discrimination risks, and the nature, source, and extent of information collected and used.

3.4 Purchased AI vs. Custom-Built AI Under Colorado

The distinction between purchased (off-the-shelf) AI and custom-built (bespoke) AI creates different compliance pathways, but neither path eliminates the deployer's obligations:

Purchased (Off-the-Shelf) AI

Even when a healthcare organization purchases an AI tool from a third-party vendor, the deploying organization remains legally responsible for conducting its own impact assessment. The deployer must verify that the vendor (developer) has provided the required documentation: a general statement of intended use, performance metrics, dataset descriptions, known limitations, and information necessary to complete the deployer's impact assessment. Standard vendor contracts often lack these governance clauses. Healthcare organizations should proactively negotiate these requirements into procurement agreements.

Critical point for purchased FDA-cleared AI: While FDA clearance provides a clinical safety baseline and may satisfy some documentation requirements, it does *not* satisfy

Colorado's legal standard for "algorithmic discrimination" assessment or "consumer notification" regarding high-risk decisions such as triage prioritization or diagnostic bias. FDA-cleared AI still requires a separate Colorado AI Act compliance program.

Custom-Built (In-House) AI

When a hospital or health system builds its own AI tool (a triage bot, a sepsis prediction model, a scheduling optimizer), it becomes a **developer** under the Act. This triggers the full developer obligation set: providing documentation, conducting bias evaluations, and reporting algorithmic discrimination to the Attorney General if discovered. Combined with the deployer obligations that still apply, custom-built AI represents the highest compliance burden under the Act.

3.5 A Healthcare Scenario: What Colorado Compliance Looks Like in Practice

Consider a Colorado hospital system that has deployed three AI tools: (1) an FDA-cleared radiology triage algorithm from Vendor A that prioritizes stroke-suspect imaging studies; (2) a commercially licensed sepsis early warning system from Vendor B; and (3) an internally developed readmission risk model used by care coordinators. Under the Colorado AI Act, this organization must:

- Classify all three systems as high-risk, since each influences consequential healthcare decisions.
- Implement a risk management program covering all three systems, with documented policies, roles, responsibilities, and escalation procedures.
- Conduct separate impact assessments for each system, evaluating algorithmic discrimination risk specific to their patient population, and repeat these assessments annually.
- Verify that Vendors A and B have provided the required developer documentation (model summaries, dataset descriptions, known limitations, performance evaluations). If the vendors have not provided this documentation, the deployer is still on the hook.
- Notify patients when these AI systems are a substantial factor in decisions affecting their care, including when the triage algorithm reprioritizes their imaging study or when the sepsis system triggers a clinical alert.
- For the internally developed readmission model, fulfill both developer and deployer obligations, including dataset documentation and AG notification requirements.
- Publish a public transparency statement describing all three deployed systems and the organization's risk management approach.
- Monitor all three systems for algorithmic discrimination on an ongoing basis and report any discovered discrimination to the Attorney General within 90 days.

None of these obligations are satisfied by FDA clearance status. The FDA clearance tells you the product can be marketed. Colorado tells you that you, as the deployer, must govern it.

3.6 The Affirmative Defense: Your Best Protection

The Act provides a meaningful incentive for proactive governance. A deployer has a **rebuttable presumption of reasonable care** if it can demonstrate compliance with a recognized AI risk management framework, most notably the NIST AI Risk Management Framework (AI RMF 1.0) or a substantially similar framework.[11] Additionally, the Act provides an **affirmative defense** against Attorney General enforcement if the deployer discovers and cures violations through its own governance processes.

This means the difference between a defensible governance posture and a vulnerable one may come down to whether your organization has: (a) formally adopted and operationalized NIST AI RMF; (b) mapped your AI governance policies to the framework's controls; and (c) documented your ongoing compliance activities. The affirmative defense effectively rewards organizations that invest in governance infrastructure and exposes those that don't.

The NIST AI RMF alignment requirement creates a direct, actionable pathway: organizations that can demonstrate systematic framework compliance gain legal protection. Organizations that cannot are exposed to up to $20,000 per violation with no defense.

3.7 Key Dates and Enforcement

Date	Milestone
May 17, 2024	Colorado AI Act (SB 24-205) signed into law by Governor Polis
February 1, 2026	Original effective date for developer and deployer obligations; public transparency statements and impact assessments required
June 30, 2026	Enforcement effective date (delayed from February 1 by SB 25B-004). Full compliance mandatory; AG enforcement authority begins
Ongoing (Annual)	Annual review of each deployed high-risk AI system for algorithmic discrimination
Within 90 Days	Impact assessment update required after any substantial AI system modification
Within 90 Days	AG notification required upon discovery of algorithmic discrimination

3.8 The ONC/HTI-1 Exemption: A Narrow Carve-Out

The Act does include an exemption for deployers and developers of high-risk AI systems that are in compliance with standards established by the Office of the National Coordinator for Health Information Technology (ONC), specifically the HTI-1 Final Rule.[12] This exemption is designed to avoid regulatory duplication for certified health IT developers and providers already complying with ONC's information disclosure and risk management requirements.

However, this exemption is narrower than many organizations assume. Not all healthcare AI systems are certified health IT, and not all AI deployed in clinical settings falls under ONC's regulatory scope. Standalone AI diagnostic tools, third-party triage algorithms, and internally developed predictive models typically do *not* qualify for this exemption. Organizations should not assume the ONC carve-out protects them without a careful, system-by-system analysis.

3.9 The Small Business Exemption: Limited Relief

Deployers with fewer than 50 employees are exempt from several requirements, including the public transparency statement, impact assessments, and risk management policy, provided they do not train AI models using their own data.[9] However, larger health systems, hospital networks, and regional medical centers receive no such exemption. These organizations must manage the full compliance burden across potentially dozens of AI systems operating at varying risk levels.

3.10 FDA Clearance Does Not Satisfy Colorado's Legal Standard

This is the central message of this white paper, and it bears repeating with specificity: FDA clearance and the Colorado AI Act address fundamentally different risk domains that do not overlap.

Dimension	FDA 510(k) Clearance	Colorado AI Act
Purpose	Marketing authorization	Consumer protection from algorithmic discrimination
Focus	Substantial equivalence to predicate device	Ongoing equity in consequential decisions
Who Bears Liability	Manufacturer (developer)	Deployer (healthcare organization)
Bias Assessment	Not required (95.5% omit demographic data)	Required: impact assessment for discrimination risk
Ongoing Monitoring	Manufacturer's post-market responsibility	Deployer must annually review each AI system

Patient Notification	Not required by clearance	Required before AI-influenced consequential decisions
Performance Drift	Not addressed at clearance	Deployer responsible for detecting and remedying
Enforcement	FDA warning letters and recalls	AG enforcement: up to $20,000 per violation

An AI diagnostic tool may carry FDA 510(k) clearance while simultaneously violating the Colorado AI Act if the deploying organization has not conducted impact assessments, implemented risk management programs, or monitored for algorithmic bias in its patient population. A vendor's assertion that their product is "FDA cleared" provides zero protection to a healthcare organization under the Colorado AI Act. The two frameworks are complementary, not substitutes.

The bottom line for Colorado healthcare organizations: On June 30, 2026, your FDA-cleared AI tools will still need a separate, documented governance program. If you have not built one, you are exposed to enforcement actions that FDA clearance cannot shield you from.

3.11 Beyond Colorado: A National Trend

Colorado's legislation is not an isolated development. The EU AI Act (enacted March 2024) establishes comparable high-risk AI requirements across the European market. Multiple U.S. states are pursuing similar legislation, with Colorado's framework serving as a reference model. New York City's Local Law 144 already regulates automated employment decisions, and California's CPPA draft regulations address AI profiling. Healthcare organizations operating across state lines should anticipate a growing patchwork of deployer obligations. The organizations that invest in Colorado compliance infrastructure now will have a transferable foundation for multi-state readiness.

Recent federal executive direction has signaled interest in reducing fragmented state-by-state AI regulation in favor of a more unified national posture. However, executive action does not displace existing state statutes or near-term deployer obligations. In a period of regulatory variability, governance maturity becomes the stabilizing control layer: organizations must be able to demonstrate disciplined oversight regardless of how the federal-state balance evolves.

Separately, regardless of regulation there is still risk and liability and the best way to mitigate that risk is through comprehensive AI governance done, locally and independently.

4. The Compounding Risk: Where These Gaps Converge

When we examine the three realities outlined above, the limitations of FDA clearance, the near certainty of model degradation, and the legal obligations now being placed on deployers, a compounding risk picture emerges that demands immediate attention.

Risk Domain	Common Assumption	The Reality
Clinical Safety	FDA clearance ensures our AI is safe	96%+ of AI devices use 510(k); fewer than 2% cite Randomized Controlled Trials (RCT)s
Performance Stability	The AI works the same way over time	91% of Machine Learning (ML) models degrade after deployment
Bias & Equity	The vendor tested for bias	95.5% of cleared devices omit demographic data
Legal Liability	The vendor bears the risk	Colorado AI Act holds deployers liable for ongoing governance
Regulatory Compliance	FDA clearance covers our obligations	State laws impose separate, additive requirements on deployers

Healthcare organizations that rely solely on their vendors' FDA clearance status as evidence of AI governance are exposed across every dimension in this table. The gap between assumption and reality is where patient safety incidents, regulatory penalties (up to $20,000 per violation under the Colorado AI Act), and reputational damage emerge.

5. What Healthcare Organizations Must Do Now

Addressing the convergence of these risks requires healthcare organizations to build institutional capabilities for clinical AI governance that extend well beyond procurement and initial validation. The following priorities should guide immediate action:

5.1 Establish a Formal AI Risk Management Program

Align your governance framework with NIST AI RMF 1.0 and prepare for updates. This is not only best practice; it is the foundation for the affirmative defense provided by the Colorado AI Act and similar legislation. Your risk management program should define roles, responsibilities, escalation procedures, and decision-making authority for all deployed AI systems.

5.2 Conduct Impact Assessments for All Deployed AI

Inventory all AI systems currently in clinical or operational use. For each system, document its intended use, the consequential decisions it influences, the data it processes, and its known limitations. Conduct formal impact assessments that evaluate algorithmic discrimination risk and identify monitoring gaps.

5.3 Implement Continuous Performance Monitoring

Healthcare organizations need the ability to track model performance metrics, detect distributional drift in input data, and identify emerging equity concerns, on an ongoing basis, not just at deployment. Given the 91% degradation rate documented in peer-reviewed research, post-deployment monitoring is not optional. Monitoring should also include **reliance signals** (e.g., override/acceptance patterns and workflow time-to-decision) to validate that meaningful independent clinical review remains operationally realistic where AI materially influences decisions.

5.4 Negotiate Vendor Transparency Requirements

The Colorado AI Act requires developers to provide deployers with documentation including model summaries, dataset descriptions, known limitations, and performance evaluations - including information sufficient to support meaningful independent clinical review where applicable. Healthcare organizations should proactively build these requirements into vendor contracts and procurement processes, before legislation mandates it.

5.5 Prepare for Multi-State Compliance

Organizations operating in multiple states should anticipate a growing body of AI governance legislation and build governance infrastructure that is adaptable to varying requirements. A framework-aligned approach (e.g., NIST AI RMF, ISO 42001) provides the most durable and scalable foundation for multi-jurisdictional compliance.

6. How Asher Informatics Can Help

Asher Informatics PBC was founded specifically to address the governance and oversight gap that this white paper describes. As a Public Benefit Corporation, our mission is to democratize access to clinical AI governance tools, particularly for the healthcare organizations that need them most: rural hospitals, community health centers, critical access facilities, and resource-constrained health systems that lack the specialized expertise to build governance infrastructure from scratch.

The AshMatics AI Governance Studio

Our platform, the AshMatics AI Governance Studio, provides an integrated solution for the full clinical AI governance lifecycle:

- **Policy & Process Studio:** Agentic configuration engine that generates organization-specific governance policies mapped to NIST AI RMF, ISO 42001, and emerging state legislation including the Colorado AI Act. Move from framework to operational policy in days, not months.
- **Value & Use Case Studio:** Structured evaluation tools for assessing AI solutions against your organization's clinical needs, patient population characteristics, and risk tolerance, before deployment.
- **Clinical AI Governance Hub:** Centralized management of impact assessments, vendor documentation, compliance evidence, and regulatory reporting, including the specific disclosures required by the Colorado AI Act.
- **AI Performance Monitoring Studio:** Continuous post-deployment monitoring for model drift, performance degradation, and emerging bias, the operational capability that transforms governance from paperwork into active patient safety protection.

Why Asher Informatics

Our founding team brings decades of medical device development experience. Our CTO, Dr. John F. Kalafut, served as Chief Science Officer at GE Healthcare and holds 44+ patents in medical imaging and computational analytics. We understand the nuances of FDA regulatory pathways, the technical realities of AI model behavior, and the operational constraints of healthcare delivery, because we've lived them.

Asher Informatics serves as an **independent validator and monitor**, not a conflicted vendor. We don't build or sell AI diagnostic tools. We help healthcare organizations govern, monitor, and maintain the AI tools they acquire from others. This independence is essential for the credibility and defensibility of any governance program, particularly when the affirmative defense provisions of emerging legislation depend on demonstrable, arms-length risk management.

Ready to close the governance gap? Contact Asher Informatics to learn how the AshMatics AI Governance Studio can help your organization meet emerging compliance requirements while protecting patients and reducing institutional risk. Visit www.asherinformatics.com or email info@asherinformatics.com.

References

1. Lin Z, et al. “Benefit-Risk Reporting for FDA-Cleared Artificial Intelligence−Enabled Medical Devices.” *JAMA Health Forum*. 2025. PMC12475944.

2. Vela D, et al. “Temporal quality degradation in AI models.” *Scientific Reports* (Nature Publishing Group). 2022;12:11654. PMC9270447.

3. U.S. Food and Drug Administration. “Artificial Intelligence and Machine Learning (AI/ML)-Enabled Medical Devices.” FDA.gov. Updated 2025.

4. Bipartisan Policy Center. “FDA Oversight: Understanding the Regulation of Health AI Tools.” December 2025.

5. Wongvibulsin S, et al. “FDA Approval of Artificial Intelligence and Machine Learning Devices in Radiology: A Systematic Review.” *JAMA Network Open*. 2025. PMC12595527.

6. “The Illusion of Safety: A Report to the FDA on AI Healthcare Product Approvals.” *PMC*. 2025. PMC12140231.

7. Stacke K, et al. “Empirical data drift detection experiments on real-world medical imaging data.” *Nature Communications*. 2024;15:1887.

8. Subasri V, et al. “Detecting and Remediating Harmful Data Shifts for the Responsible Deployment of Clinical AI Models.” *JAMA Network Open*. 2025. PMC12138723.

9. Colorado General Assembly. Senate Bill 24-205: Consumer Protections for Artificial Intelligence. Enacted May 17, 2024. Enforcement effective June 30, 2026.

10. Foley & Lardner LLP. “The Colorado AI Act: Implications for Health Care Providers.” February 2025.

11. National Institute of Standards and Technology (NIST). “Artificial Intelligence Risk Management Framework (AI RMF 1.0).” January 2023.

15. American Hospital Association. “AHA Letter to FDA on AI-Enabled Medical Devices.” December 2025.

16. Rajaram R, et al. “Evaluating Transparency in AI/ML Model Characteristics for FDA-Reviewed Medical Devices.” *npj Digital Medicine*. November 2025.

17. Bayram F, Ahmed B, Kassler A. “From Concept Drift to Model Degradation: An Overview on Performance-Aware Drift Detectors.” *Knowledge-Based Systems*. 2022;245:108632.

12. Mintz LLP. “Colorado AI Systems Regulation: What Health Care Deployers and Developers Need to Know.” June 2024.

13. U.S. Food and Drug Administration. “Clinical Decision Support Software: Guidance for Industry and Food and Drug Administration Staff.” Originally issued September 2022; updated January 6, 2026. U.S. Food and Drug Administration.

14. The White House. “Executive Order on Eliminating State Law Obstruction of National Artificial Intelligence Policy. December 11, 2025.The White House, Presidential Actions.

About Asher Informatics PBC

Asher Informatics PBC is a Public Benefit Corporation focused on healthcare AI governance and oversight solutions. Our software provides healthcare organizations with tools to govern the full lifecycle of clinical AI, from policy creation and vendor evaluation through monitoring and regulatory compliance. Our mission is to democratize access to these critical capabilities, particularly for the rural hospitals, community health centers, and resource-constrained organizations that serve America’s most underserved communities. Learn more at www.asherinformatics.com.

