

Perspective

The AI Hospital Formulary: A Practical Governance Framework for Prescribing, Monitoring, and Deprescribing Artificial Intelligence in Hospitals

Francisco Epelde ^{1,2} 

¹ Internal Medicine Department, Parc Taulí Hospital Universitari, Institut d'Investigació i Innovació Parc Taulí (I3PT-CERCA), Universitat Autònoma de Barcelona, 08208 Sabadell, Spain; fepelde@gmail.com

² APTIMA Centre Clinic, Mutua de Terrassa, 08221 Terrassa, Spain

Abstract

Background: Artificial intelligence (AI) is increasingly entering hospital practice through diagnostic, predictive, workflow, operational, and generative applications. Hospitals often govern these systems as procurement or information-technology projects rather than as clinical–organizational interventions requiring indication, evaluation, monitoring, accountability, and withdrawal. **Objective:** To refine the concept of an “AI Hospital Formulary” as an operational, proportional, and accountable framework for the safe, equitable, and sustainable adoption of hospital AI. **Design and Methods:** This is a perspective article using a structured, non-systematic narrative synthesis and conceptual framework development. Targeted literature and policy sources were identified through purposive searches and citation chaining through 20 July 2026. The synthesis compares the formulary with existing oversight approaches, maps its lifecycle gates to regulatory and risk management duties, and applies the framework to a worked example based on published evaluations of the Epic Sepsis Model. The EQUATOR reporting-guideline selection tool was consulted, and SANRA was used to strengthen the narrative synthesis component. **Framework:** The revised framework combines a hospital-wide AI register, a standardized formulary monograph, six lifecycle gates, proportional review pathways, governance-of-governance safeguards, cloud and data-sovereignty controls, continuous monitoring of technical and behavioral feedback loops, and explicit renewal or deprescribing criteria. The worked example shows how version-specific evidence can lead to local validation, controlled implementation, restriction, suspension, or renewal rather than automatic adoption. **Conclusions:** Hospitals should not merely purchase, install, and update AI systems. They should prescribe, monitor, audit, renew, restrict, and, when necessary, deprescribe them. The AI Hospital Formulary is proposed as a complementary institutional layer that converts external standards and existing governance approaches into documented portfolio decisions at the hospital level.



Academic Editor: Mojtaba Vaismoradi

Received: 6 June 2026

Revised: 20 July 2026

Accepted: 21 July 2026

Published: 22 July 2026

Copyright: © 2026 by the author.

Licensee MDPI, Basel, Switzerland.

This article is an open access article distributed under the terms and conditions of the [Creative Commons Attribution \(CC BY\)](https://creativecommons.org/licenses/by/4.0/) license.

Keywords: artificial intelligence; hospitals; governance; algorithmovigilance; patient safety; clinical decision support; digital health; AI regulation; cybersecurity; data sovereignty; implementation science; ethics

1. Introduction

Artificial intelligence has become one of the dominant narratives in contemporary hospital innovation. Its potential applications range from medical image interpretation,

early warning scores, triage support, operating room scheduling, bed management, discharge planning, clinical documentation assistance, and patient communication to large language model-based summarization and decision support [1–3]. Yet the translation of AI from technical performance to measurable clinical and organizational benefit remains difficult [4]. The diagnostic AI literature illustrates this gap: systematic reviews have documented limitations in study design, reporting standards, risk of bias, and claims about performance compared with clinicians [5,6].

Existing regulations and guidance increasingly emphasize risk management, transparency, post-market monitoring, human oversight, lifecycle evaluation, and trustworthy AI governance [7–14]. Existing hospital governance models, model-documentation tools, patient-safety approaches, and reporting standards also provide important components for institutional oversight [15–24]. However, these sources do not fully resolve the hospital-specific portfolio questions of which system and version should be authorized, for which indication and users, under what local restrictions, for how long, and with what documented exit plan. The proposed formulary is therefore intended to complement—not replace—existing regulatory, quality, safety, and governance mechanisms.

The hospital is an especially important site for AI governance. Hospitals concentrate high-risk care, complex workflows, multidisciplinary decision-making, scarce resources, vulnerable patients, and dense digital infrastructures. In this environment, AI failure may be technical, clinical, ethical, operational, cybersecurity-related, or social. A sepsis model may perform well in one hospital but poorly in another. A bed-management algorithm may improve throughput while disadvantaging certain patient groups. A generative AI tool may reduce documentation burden but introduce hallucinated content, data leakage, or dependency on a cloud service. A predictive model may remain technically functional but become clinically irrelevant after changes in clinical practice, coding, population mix, information systems, or clinician response to its outputs [25–32].

Hospital-level governance frameworks for predictive tools and AI applications have begun to address these concerns [15,16,23,24]. Reporting standards such as TRIPOD+AI, CONSORT-AI, SPIRIT-AI, DECIDE-AI, and PROBAST+AI improve the design, reporting, and critical appraisal of AI studies [17–21]. Recent work on generative AI further emphasizes privacy, provenance, safety evaluation, and the need for governance that follows systems after integration into clinical workflows [9,33–37]. Hospitals nevertheless need an operational mechanism that brings these contributions together in an enforceable institutional authorization and renewal process. This article proposes such a mechanism: the AI Hospital Formulary.

The central thesis is simple: hospital AI systems should not merely be purchased, installed, and updated; they should be prescribed, monitored, audited, renewed, restricted, and, when necessary, deprescribed.

2. Research Design and Methods

2.1. Study Design

This article is a Perspective, as defined by the journal, and uses a structured, non-systematic narrative synthesis to develop a conceptual and operational governance framework. It is not a systematic review, scoping review, consensus study, empirical validation study, or legal analysis. Its purpose is normative and translational: to integrate relevant bodies of knowledge and propose a hospital-level process that can subsequently be tested in practice.

2.2. Source Identification and Narrative Synthesis

During manuscript development and revision, sources were identified through targeted, iterative searches of PubMed and the official websites of regulatory, standards, and public health bodies through 20 July 2026. Search concepts included artificial intelligence, hospital governance, predictive models, clinical decision support, lifecycle oversight, algorithmovigilance, model cards, generative AI, large language models, cloud security, data sovereignty, and sepsis prediction. Reference list review and citation chaining were used to identify additional foundational and recent publications. Sources were selected purposively for their direct relevance, authority, recency, and coverage of complementary governance domains. No protocol was registered, and no exhaustive search, duplicate screening, quantitative synthesis, or formal risk-of-bias assessment was undertaken; these limitations are consistent with the Perspective design.

2.3. Framework Development and Reporting Approach

Concepts from clinical pharmacology, medical device regulation, patient safety, implementation science, algorithmovigilance, AI ethics, cybersecurity, procurement, and hospital governance were mapped to the functions of a hospital formulary. They were iteratively organized into a formulary monograph, six lifecycle gates, committee responsibilities, governance-of-governance safeguards, a regulatory crosswalk, an algorithmovigilance dashboard, and a worked example of a sepsis model. The EQUATOR Network reporting-guideline selection tool was consulted [38]. No dedicated EQUATOR checklist was identified for a Perspective or conceptual framework article; therefore, SANRA was used as a quality and structural aid for the narrative synthesis component, without claiming that the article is a systematic review [39].

3. Why Hospitals Need an AI Formulary

Hospital formularies emerged to promote the rational, safe, evidence-based, and cost-conscious use of medicines. They do not simply list available drugs. They define which agents are approved, for which indications, under what restrictions, and with what monitoring requirements. They also provide a governance process for evaluating new therapies, reviewing evidence, managing costs, identifying safety concerns, and removing obsolete or harmful treatments.

Hospitals need an analogous structure for AI. The number of AI systems entering hospitals is increasing, but governance is often fragmented across information technology departments, innovation units, clinical departments, data science teams, procurement offices, legal services, compliance teams, quality and safety units, and individual enthusiasts. The fragmentation makes it difficult to maintain a complete inventory, apply consistent standards, assign ownership, and compare risks across tools.

First, AI adoption may become vendor-driven rather than needs-driven. Tools may be acquired because they are available, fashionable, bundled into existing software, or promoted through enterprise cloud contracts, rather than because they address a well-defined clinical or operational problem. Second, evidence may be evaluated inconsistently. Third, accountability may be unclear. Fourth, monitoring may be insufficient. Fifth, systems may persist after they are no longer useful. These risks are amplified when hospitals cannot determine where data are processed, which sub-processors have access, whether prompts or outputs are retained for model improvement, how a cloud service changes over time, or how the hospital can recover data and workflows if a vendor relationship ends.

An AI Hospital Formulary would address these risks by creating a visible, structured, and accountable institutional process for AI stewardship. It would define the approved AI portfolio, classify tools by risk and purpose, specify approved conditions of use, re-

quire proportionate local validation before deployment, mandate post-implementation monitoring, and establish criteria for renewal, restriction, suspension, or withdrawal. The formulary should be integrated with—not isolated from—procurement, medical device governance, cybersecurity, privacy, clinical safety, quality improvement, and enterprise architecture.

Importantly, the formulary should not be limited to regulated AI medical devices. Hospitals increasingly use AI in domains that may fall outside classic medical device regulation but still affect care quality, safety, equity, professional work, resource allocation, or institutional decisions. Examples include staffing algorithms, waiting-list prioritization, discharge prediction, automated coding, patient messaging, ambient documentation, translation, and administrative triage. The scope should therefore include any AI system that can materially influence patient care, professional work, hospital operations, or access to services.

4. From Drug Formularies to AI Formularies

The analogy between medicines and AI is not perfect, but it is useful. Both require a defined indication, evidence appraisal, risk–benefit assessment, monitoring, and governance. Both can be overused, misused, or used outside approved conditions. Both may perform differently in populations not adequately represented in development studies. Both may create harm through interactions with other interventions or through inappropriate reliance by professionals.

AI systems differ from drugs in several important ways. First, AI performance is context-dependent. A model trained in one health system may perform differently in another because of differences in population, clinical practice, measurement, coding, infrastructure, or workflow. Second, AI systems may be dynamic even when the model weights are static because the surrounding data environment and clinical pathways change. Sudden systemic shifts, such as the COVID-19 pandemic, illustrate how rapidly data-generating processes and predictive assumptions can change [40]. Third, AI systems are socio-technical: their impact depends on how clinicians interpret, trust, ignore, override, or overuse their outputs [41]. Fourth, outputs can create feedback loops by changing clinician behavior and, consequently, the future data used to evaluate the system. Fifth, many AI systems are opaque or dependent on vendor-controlled updates, making subtle failure difficult to detect.

These differences do not weaken the formulary analogy; they strengthen the need for it. A hospital formulary for AI should not copy drug governance mechanically. It should translate formulary principles—indication, restrictions, ownership, monitoring, renewal, and withdrawal—into the AI lifecycle while accounting for software versioning, workflow dependence, cybersecurity, and feedback effects.

4.1. Relationship to Existing Governance Approaches

The AI Hospital Formulary is not presented as the first framework for healthcare AI governance. Bedoya et al. describe institutional oversight and local deployment of predictive models, including a real portfolio of locally governed models [15]. Reddy et al. provide a broad governance model for healthcare AI [16]. Model cards standardize documentation of model purpose, performance, limitations, and subgroup considerations [22]. Other proposals show how current patient-safety infrastructure can be extended to AI and how clinical organizations can monitor AI-related hazards [23,24]. Regulatory and risk management frameworks define provider and deployer responsibilities at a higher level [7–14,33].

The distinctive contribution of the formulary is the combination of these elements in a hospital-wide authorization and portfolio management process modeled on a familiar clinical institution. It covers clinical, operational, administrative, and patient-facing AI; requires an approved indication and restrictions; links evidence review to procurement and workflow authorization; names clinical, technical, monitoring, and executive owners; assigns an expiry date to approval; and makes managed withdrawal an expected lifecycle decision. A model card can be incorporated into the evidence dossier, a local oversight program can perform validation, and patient-safety systems can receive incident reports; the formulary supplies the institutional decision record that determines whether and how the tool may remain in use. Table 1 summarizes how the proposed formulary complements these selected approaches.

Table 1. Relationship of the AI Hospital Formulary to Selected Existing Approaches.

Approach	Primary Contribution	How the Formulary Complements It
Local predictive-model oversight (e.g., Bedoya et al.)	Institutional review, validation, and oversight of predictive models in local practice.	Extends the same discipline to a hospital-wide portfolio including operational, administrative, patient-facing, and generative AI, with explicit indication, expiry, renewal, and deprescribing.
General healthcare AI governance models (e.g., Reddy et al.)	Principles, roles, and governance domains for responsible healthcare AI.	Converts governance domains into a standardized authorization record, six decision gates, restrictions, named owners, and an exit decision.
Model cards	Structured documentation of intended use, performance, limitations, and subgroup behavior.	Uses model cards as evidence inputs but adds local validation, workflow authorization, surveillance, procurement conditions, and institutional approval status.
Patient-safety infrastructure	Incident reporting, safety review, hazard analysis, and organizational learning.	Connects incidents and near misses to predefined thresholds for investigation, suspension, restriction, renewal, or withdrawal.
EU AI Act, FDA/IMDRF, WHO, and NIST frameworks	Legal duties, regulatory expectations, and risk management functions.	Maps external duties to hospital tasks and evidence artifacts while recognizing that a formulary does not itself establish legal compliance.

4.2. The Formulary Monograph

A drug formulary asks: What is the medicine? What is it for? What evidence supports it? Who should receive it? What are the contraindications? What monitoring is required? What are the alternatives? What is the cost? When should it be stopped? An AI formulary asks parallel questions about the system, its version, decision or workflow influenced, approved users and populations, evidence, local performance, data and cloud dependencies, human oversight, costs, and withdrawal criteria. The proposed monograph domains are summarized in Table 2.

Table 2. Proposed AI Hospital Formulary Monograph.

Domain	Required Content
Identification	Name, version, vendor/developer, model type, regulatory status, update mechanism, and owner department.
Intended use	Clinical, operational, administrative, or patient-facing purpose; decision influenced; approved setting; prohibited or off-formulary uses.

Table 2. *Cont.*

Domain	Required Content
Target population	Inclusion and exclusion criteria; age groups; specialties; inpatient/outpatient scope; vulnerable populations.
Evidence base	Development data, external validation, prospective evaluation, clinical impact evidence, limitations, and version-specific evidence.
Local validation	Performance on local data; calibration; subgroup analysis; alert/workload simulation; usability and workflow fit.
Human oversight	Required reviewer; authority to override; escalation pathway; documentation expectations; competence and AI literacy requirements.
Workflow integration	Where the output appears, timing, alert logic, user interface, failure mode, downtime procedure, and alternative workflow.
Data dependencies	Input variables, source systems, missing-data handling, data quality requirements, lineage, and change controls.
Equity assessment	Performance and impact by clinically relevant demographic, language, disability, and socioeconomic groups.
Safety risks	False positives, false negatives, automation bias, alert fatigue, delayed care, overdiagnosis, deskilling, and feedback loops.
Cloud and data sovereignty	Data location and jurisdiction, cross-border transfers, sub-processors, encryption and key control, retention, vendor reuse, service continuity, portability, and exit plan.
Cybersecurity and privacy	Threat model, access controls, audit trails, incident response, vulnerability management, privacy impact assessment, and vendor obligations.
Ethical justification	Expected benefit, affected groups, non-AI alternatives, autonomy and transparency implications, and proportionality.
Governance and contestability	Named owners, committee decision, conflicts of interest, patient/public input, appeal route, and reconsideration date.
Economic assessment	Purchase, integration, cloud, maintenance, monitoring, training, staff time, downstream utilization, switching, and opportunity costs.
Monitoring plan	Metrics, frequency, responsible team, baseline, thresholds, incident route, and actions after a threshold breach.
Deprescribing criteria	Drift, harm, inequity, low value or adoption, excessive burden, superior alternative, vendor failure, cybersecurity concern, or regulatory change.

5. Proposed Lifecycle of the AI Hospital Formulary

The AI Hospital Formulary should be more than a catalog. It should function as a lifecycle governance system with six decision gates. Review should be proportional: a low-risk tool that summarizes non-clinical meeting notes may follow an abbreviated pathway, whereas an autonomous or high-consequence clinical tool should require full evidence review, silent validation, controlled implementation, and intensive surveillance.

5.1. Gate 1: Request and Triage

Any proposed AI system should enter through a standardized request process. The request should identify the problem to be solved, intended users, the affected population, expected benefit, non-AI alternatives, consequences of error, the data and cloud pathway, and the proposed owner. Risk classification should consider the severity of potential harm, autonomy, scale, reversibility, vulnerability, clinical dependence, data sensitivity, and whether the system may meet a legally defined high-risk category.

5.2. Gate 2: Evidence Review

The committee should evaluate the external evidence and the exact version proposed. Has the tool been externally validated? Was it tested prospectively? Was it compared with usual care? Were reporting standards followed? Were relevant subgroups analyzed? Does the evidence reflect the hospital's population and workflow? For generative AI, the evaluation should also address factuality, instability, prompt sensitivity, provenance, privacy, and whether outputs can change without transparent version control [33–36].

5.3. Gate 3: Local Validation

No consequential AI system should be deployed solely because it performed well elsewhere. Local validation should test discrimination, calibration, subgroup performance, usability, alert burden, workflow fit, and the reliability of input data. Silent evaluation may be appropriate before active use. Validation must be version-specific: a later vendor model should not inherit authorization automatically from evidence about an earlier model.

5.4. Gate 4: Controlled Implementation

Implementation should begin with a controlled release rather than immediate hospital-wide deployment. Pilot units should have trained users, clear escalation pathways, an alternative workflow, and predefined technical, clinical, equity, workload, and adoption metrics. Implementation science frameworks can distinguish adoption, acceptability, fidelity, feasibility, penetration, sustainability, and cost [42–44].

5.5. Gate 5: Active Surveillance

Post-deployment monitoring should cover performance, calibration, data quality, utilization, override behavior, alert burden, equity, outcomes, cybersecurity events, and unintended consequences. Surveillance should combine automated dashboards with user feedback, incident reporting, case review, and periodic audit.

5.6. Gate 6: Renewal, Restriction, or Deprescription

AI authorization should expire unless renewed. At defined intervals—or sooner after a trigger—the committee should decide whether the exact system and version remain approved, require restriction or retraining, should be replaced, or should be deprescribed. The decision should include transition planning so that withdrawal does not itself create clinical risk. The six lifecycle gates and their minimum deliverables are summarized in Table 3.

Table 3. Lifecycle Gates for the AI Hospital Formulary.

Gate	Main Question	Minimum Deliverable	Possible Decision
1. Request and triage	Is there a legitimate problem, an appropriate AI use case, and a proportionate pathway?	Use-case proposal, alternatives, ownership, data/cloud map, and risk classification.	Reject, revise, fast-track, or proceed.
2. Evidence review	Is the version-specific external evidence sufficient for local consideration?	Evidence dossier, model documentation, and critical appraisal.	Reject, restrict, request evidence, or validate locally.
3. Local validation	Does the system perform safely and equitably in this context?	Local technical, subgroup, usability, workflow, and alert-burden report.	Reject, modify, pilot, or proceed.
4. Controlled implementation	Can trained users employ it safely in real workflows?	Pilot evaluation with predefined metrics and an alternative workflow.	Stop, adapt, scale, or continue the pilot.
5. Active surveillance	Are value, safety, fairness, security, and appropriate use being maintained?	Dashboard plus incident, feedback, and case-review record.	Continue, investigate, suspend, restrict, or update.
6. Renewal or deprescription	Should this exact system and version remain authorized?	Periodic renewal report and transition plan.	Renew, restrict, retrain, replace, or retire.

5.7. Operational Alignment with Regulation and Risk Management Frameworks

The formulary is a governance mechanism, not a substitute for legal classification, conformity assessment, regulatory reporting, or professional judgment. Provider, deployer, manufacturer, and healthcare-organization obligations depend on intended purpose, product status, jurisdiction, and contractual roles. Nevertheless, the formulary can create the local evidence trail and assignments needed to operationalize many duties. Under the EU AI Act, for example, Gate 1 can document whether a system may fall within high-risk categories; Gates 2 and 3 can receive risk management, data-governance, technical-documentation, transparency, human-oversight, accuracy, robustness, and cybersecurity evidence; Gate 4 can enforce use according to instructions and AI literacy requirements; Gate 5 can support deployer monitoring, log retention, and escalation; and Gate 6 can suspend or withdraw local use. Provider post-market monitoring and serious-incident reporting remain distinct legal duties [10]. An illustrative crosswalk is provided in Table 4.

Table 4. Illustrative Mapping from External Requirements to Formulary Tasks.

External Requirement or Function	Formulary Operationalization	Evidence Retained by the Hospital
EU AI Act risk classification and high-risk determination (including Article 6 where applicable) [10]	Gate 1 screens intended purpose, autonomy, affected decisions, regulated-product status, and potential high-risk classification; legal/regulatory review confirms the classification.	Risk-classification rationale, intended use, prohibited uses, and legal review.
EU AI Act risk management, data governance, documentation, logs, transparency, human oversight, accuracy, robustness, and cybersecurity (Articles 9–15 for high-risk systems) [10]	Gates 2–4 require the provider dossier, local data assessment, instructions for use, oversight design, technical and cybersecurity review, and pilot controls.	Evidence dossier, data lineage, local validation report, training record, oversight protocol, and go-live decision.
EU AI Act deployer duties, monitoring, input data relevance, log retention, suspension/escalation, and impact assessment where applicable (Articles 26–27) [10]	Gates 4–5 define authorized users, local data checks, monitoring ownership, log retention, incident escalation, and suspension authority.	Usage logs, monitoring dashboards, incident files, impact assessment, and suspension decisions.
EU AI Act provider post-market monitoring and serious-incident reporting (Articles 72–73) [10]	Vendor contracts require notification and cooperation; Gate 5 receives local safety signals and transmits them through regulatory channels as required.	Vendor notices, local signal assessments, communications, and regulatory reports.
FDA/IMDRF lifecycle expectations and real-world performance monitoring [11–13]	Version control, change assessment, local validation, and performance/drift surveillance are embedded across Gates 2–6.	Version history, validation protocol, change log, drift reports, and renewal decision.
NIST AI RMF and Generative AI Profile: Govern, Map, Measure, Manage [14,33]	Committee charter and ownership govern; use-case analysis maps; validation and surveillance measure; restrictions, mitigation, and withdrawal manage risk.	Risk register, test results, mitigations, residual-risk acceptance, and review schedule.

6. The AI Formulary Committee

A formulary requires a committee with authority. Hospitals should create an AI Formulary Committee or expand an existing digital health governance board. Membership should be multidisciplinary: clinicians, nurses, pharmacists, informaticians, data scientists, quality and safety experts, ethicists, legal and compliance officers, privacy and cybersecurity specialists, patient representatives, procurement professionals, and hospital managers. The committee should not be purely technical because AI governance is a clinical, organizational, ethical, legal, and operational responsibility.

The committee's mandate should include approval of AI systems, maintenance of the AI registry, risk classification, evidence appraisal, local validation requirements, monitoring standards, incident review, and deprescribing decisions. It should define prohibited, restricted, approved, and experimental uses. For generative AI, policy should address prompts, privacy, output verification, documentation, disclosure, model updates, and acceptable use.

The committee should also manage "shadow AI." Many professionals already use public tools for summarization, translation, drafting, coding, teaching, literature review, and administrative tasks. Prohibiting all use may be unrealistic; ignoring it is unsafe. A visible classification of approved, restricted, prohibited, and experimental tools can move hospitals from informal, invisible use toward transparent and governable practice.

Governance of the Governance

The committee itself requires a formal charter. The charter should define its scope, delegated authority, quorum, voting rules, documentation standards, review timelines, conflicts-of-interest declarations, recusal, confidentiality, emergency suspension authority, and periodic independent evaluation. Decisions should include a written rationale, conditions of use, minority views where material, the evidence considered, and a date for reconsideration. Commercial sponsorship or departmental ownership should not exempt a proposal from review.

Applicants should have a route to request reconsideration when relevant evidence, proposed safeguards, or circumstances change. A separate executive, clinical safety, or ethics body should hear unresolved disputes or allegations that the process was not followed. An appeal should not automatically delay an urgent safety suspension. The institution should publish an annual, public-facing summary of the AI portfolio, categories of decisions, material incidents, and governance improvements without disclosing protected information.

Patient and public participation should be meaningful rather than symbolic. Hospitals can use compensated, rotating patient/public panels; accessible summaries of high-impact proposals; targeted consultation for patient-facing systems and tools affecting triage or resource allocation; and documented explanations of how input changed—or did not change—the decision. Participation should occur at agenda setting, review, monitoring, and renewal, not only through one committee seat.

7. Monitoring and Algorithmovigilance

AI monitoring should be broader than model performance. A hospital system can fail even when accuracy appears stable. Clinicians may ignore alerts, over-trust outputs, misunderstand scores, or use the tool outside its intended scope. Conversely, a statistically strong model may create operational harm by increasing workload, unnecessary testing, or inequitable prioritization. Algorithmovigilance therefore combines technical surveillance with process, outcome, equity, cybersecurity, and human-factors monitoring [24,31,32].

A practical monitoring plan should preserve event-level traceability: the system and version, timestamp, relevant inputs or input-quality flags, output, whether and how the output was displayed, the user response or override, escalation, downstream actions, and clinically appropriate outcome labels. This record makes feedback loops observable. For example, if repeated alerts lead clinicians to order tests earlier, the apparent lead time of the model and the future training data may change. If users selectively ignore alerts in a particular unit or subgroup, overall discrimination may conceal an unsafe workflow effect.

Not every outcome is available in real time. Hospitals should distinguish near-real-time leading indicators—missing inputs, distribution shifts, service latency, alert volume,

views, acknowledgments, overrides, and workflow delays—from lagging indicators such as complications, readmissions, or mortality. Monitoring can use rolling windows and comparison with a prespecified baseline; statistical process-control charts; distributional checks such as population stability or two-sample tests, where appropriate; calibration-in-the-large and calibration slope; threshold-specific sensitivity, specificity, positive predictive value, and number needed to evaluate; and subgroup-stratified process and outcome measures. A statistical threshold should trigger review, not automatic retraining or automatic shutdown, because changes may reflect legitimate epidemiologic or workflow shifts.

Monitoring should include at least five categories: technical performance; process and use; clinical or operational outcomes; equity; and human factors. Cybersecurity and privacy signals should be integrated rather than reviewed separately. For generative systems, additional measures may include unsupported-content rates, clinically material omissions, prompt-injection or data-leakage events, human verification rates, and changes in output after vendor model updates [33–36].

The analogy with pharmacovigilance is valuable. Pre-implementation studies cannot capture all future contexts, workflow adaptations, or population shifts. Hospitals should create a recognizable reporting channel for AI-related incidents and near misses, with non-punitive review in the broader patient-safety tradition [23,24,45,46]. “Algorithmic morbidity and mortality rounds” could review cases involving harm, delay, inequity, confusion, or missed opportunity and link lessons to training, restrictions, monitoring, and renewal [47]. Table 5 presents the minimum dashboard and illustrative trigger logic.

Table 5. Minimum Algorithmovigilance Dashboard and Trigger Logic.

Domain	Illustrative Measures	Review Frequency	Illustrative Trigger and Response
Data and service integrity	Missingness, coding/source changes, distribution shift, latency, downtime, API/model-version change.	Near-real time to monthly, according to risk.	Unexpected source or version change, persistent missingness, or service failure → verify lineage, invoke the downtime process, and consider suspension.
Technical performance	Calibration, sensitivity, specificity, PPV, false-negative rate, alert burden, number needed to evaluate.	Monthly to quarterly after labels mature.	Prespecified degradation or confidence intervals crossing the threshold → case review, recalibration assessment, restriction, or new validation.
Use and behavior	Eligibility, display/view rate, acknowledgment, override, response time, unit/user variation, use outside indication.	Weekly to monthly.	Rapid change, extreme between-unit variation, or off-label use → workflow investigation, training, interface change, or restricted access.
Clinical/operational outcomes	Time to treatment, adverse events, downstream testing, length of stay, throughput, workload.	Monthly to semiannual.	No expected benefit or evidence of harm/cost displacement → controlled evaluation, restriction, or renewal refusal.
Equity	Performance, access, overrides, actions, and outcomes across relevant subgroups.	Quarterly and after material change.	Unexplained disparity with clinical significance → equity review, mitigation, narrower indication, or suspension.
Human factors and safety reports	Alert fatigue, trust, comprehension, workarounds, incidents, near misses, complaints.	Continuous reporting; periodic synthesis.	Serious incident or recurring pattern → immediate safety review, vendor notification, and possible suspension.

8. Ethical Governance and Accountability

The AI Hospital Formulary is also an ethical governance mechanism. Hospitals adopt AI in environments characterized by uncertainty, vulnerability, asymmetries of knowledge, time pressure, scarce resources, and unequal access to care. Governance cannot be reduced to model performance, regulatory status, or procurement compliance; it must address beneficence, non-maleficence, justice, autonomy, transparency, privacy, accountability, and public trust [48–51].

Ethical review should begin before procurement. The committee should ask whether the problem is appropriate for AI at all. Some problems are better addressed through staffing, pathway redesign, interoperability, standardization, or communication. AI should not automate dysfunction, conceal scarcity, or transfer institutional burdens to clinicians and patients. A request should explain why AI is needed, who may benefit or be harmed, and which non-AI alternatives were considered.

Fairness should be assessed before and after deployment. Statistical subgroup evaluation is necessary but insufficient: similar performance can still lead to unfair consequences if outputs alter access to beds, specialist review, diagnostic testing, discharge planning, or follow-up. Conversely, unequal metrics require clinical interpretation of baseline risk, harms, mitigation, and whether the intervention improves care for underserved populations. The formulary therefore requires both technical fairness analysis and clinical equity judgment.

Transparency should be proportionate to the decision. Every user should understand the intended use, limitations, inputs, the appropriate response, and the circumstances requiring override. Patients should receive meaningful information when AI materially affects triage, diagnosis, treatment prioritization, discharge, or the allocation of scarce resources. Transparency should support understanding, contestability, and trust rather than operate as a generic disclaimer.

Accountability requires named ownership. Each system should have a clinical owner, technical owner, monitoring owner, and executive sponsor, together with an escalation pathway. The professional authority to override should be explicit. Failure should be treated as a safety and governance signal, not simply as user error.

Data Sovereignty, Cybersecurity, and Cloud Infrastructure

Cloud-based AI can distribute data and control across hospitals, vendors, hyper-scale cloud providers, and sub-processors. A procurement review should document the location and jurisdiction of storage and processing; cross-border transfers; tenant isolation; encryption in transit and at rest; who controls encryption keys; identity, access, and privileged-administrator controls; logging; backup and recovery; retention and deletion; vulnerability disclosure and patching; business continuity; and notification timelines for security events. The hospital should know whether prompts, inputs, outputs, or telemetry are retained or used to train or improve vendor models and whether that practice can be contractually disabled.

Data sovereignty also concerns effective institutional control, not only physical location. Contracts should define permitted purposes, sub-processor changes, audit rights, government or third-party access requests, service-level obligations, model and API change notification, portability, interoperability, and the return or verified deletion of data upon termination. Hospitals should test an exit plan before dependency becomes irreversible. Where clinical continuity depends on a remote service, the formulary monograph should identify a safe downtime workflow and the maximum tolerable interruption.

For generative AI, public or consumer services should not receive identifiable patient information unless an approved legal, privacy, security, and contractual basis exists. Ap-

proved systems should specify whether prompts and outputs enter the clinical record, how human verification is documented, and how prompt injection, retrieval poisoning, unauthorized data exfiltration, or vendor-side model changes are detected and managed [9,33–36].

Ethical governance must include the possibility of saying no. Some systems should not be implemented even if technically impressive. Others should be restricted to particular settings or users or suspended after evidence of drift, inequity, security risk, or harm. A hospital that cannot safely withdraw an AI system has not fully governed it.

9. AI Deprescribing

Deprescribing is central to the framework because hospitals often accumulate digital tools without systematic retirement. This leads to alert fatigue, redundant systems, hidden costs, outdated models, and unclear accountability. An AI system may deserve removal even without a single dramatic safety event: low adoption, marginal benefit, excessive maintenance burden, unresolved inequity, cybersecurity risk, or replacement by a better tool may justify withdrawal.

The formulary does not make deprescribing easy. Withdrawal may be technically, commercially, clinically, and politically difficult, particularly after users and workflows become dependent on a tool. The formulary makes deprescribing planned, evidence-based, and accountable by defining triggers, authority, evidence, communication, and transition steps before deployment.

Every monograph should specify conditions that trigger review, suspension, or retirement: performance below threshold, calibration drift, unexplained subgroup disparities, repeated incidents, major workflow or data source changes, unannounced model updates, vendor discontinuation, cybersecurity vulnerabilities, loss of regulatory support, or failure to demonstrate expected value. Sudden changes in epidemiology or care delivery may also invalidate assumptions and require temporary suspension or revalidation [40].

Deprescribing should consider human dependency. Sudden removal may itself create risk. Withdrawal may require communication, alternative workflows, retraining, data export, contract termination, and temporary safeguards. Responsible discontinuation is therefore a sign of institutional maturity, not a failure of innovation.

10. Economic and Strategic Value

The AI Hospital Formulary also has economic value. Purchase price is only one component of the total cost. Hospitals must account for integration, validation, training, monitoring, governance, cybersecurity assessment, legal review, workflow redesign, downstream testing, cloud compute and data-egress charges, vendor management, model updating, service outages, switching and exit costs, and opportunity cost. A nominally inexpensive system may be costly if it generates unnecessary work or creates vendor lock-in.

Conversely, a formulary can prioritize investments that align with institutional strategy. Hospitals should favor tools that solve important problems, have credible evidence, integrate into workflows, improve outcomes or efficiency, reduce burden, and can be monitored sustainably. The formulary thus becomes both a safety instrument and a strategic portfolio management tool.

The formulary can strengthen procurement. Vendors can be required to provide an evidence dossier, a model card or equivalent documentation, cybersecurity evidence, data-flow and sub-processor information, subgroup performance, update and change-control policies, audit rights, monitoring cooperation, portability, incident notification, and termination support. Contracts should distinguish hospital responsibilities from vendor responsibilities and require version-specific notice before material changes.

11. Illustrative Worked Example: Epic Sepsis Model

The following example applies the formulary to a widely implemented proprietary sepsis prediction model using published evidence. It is an illustrative monograph, not a local empirical validation, endorsement, legal classification, or current authorization decision for any hospital. A real decision would require the exact contracted version, local data, workflow design, current vendor documentation, and hospital-specific thresholds.

A published external evaluation of the original Epic Sepsis Model reported limited performance in one health system: an area under the receiver operating characteristic curve of 0.63, sensitivity of 33%, specificity of 83%, and a positive predictive value of 12% at a commonly used threshold, with a substantial alert burden [52]. A multicenter prospective validation published in 2026 evaluated version 2 in 227,091 encounters across four health systems. Encounter-level discrimination ranged from 0.82 to 0.92, but the positive predictive value remained 0.13 to 0.26, the alert burden was high, and thresholds and performance varied across institutions [53]. These findings demonstrate why an authorization must identify the version and why improved external evidence still does not remove the need for local validation and workflow evaluation. The resulting illustrative monograph and gate decisions are summarized in Table 6.

Table 6. Illustrative Formulary Monograph and Gate Decisions for a Proprietary Sepsis Prediction Model.

Monograph Field	Illustrative Entry and Decision
Identification and version	Epic Sepsis Model; exact deployed version, build, threshold logic, input mapping, and update date must be recorded. Evidence for version 1 must not be assumed to apply to version 2 or later versions [52,53].
Intended use	Early identification of hospitalized adults at elevated risk of sepsis to prompt timely clinical reassessment—not autonomous diagnosis or automatic treatment.
Risk and triage	High-consequence clinical decision support because false negatives may delay recognition and false positives may create alert fatigue, testing, antibiotics, or resource use. Full pathway required.
External evidence	Version 1 evaluation showed poor discrimination and low PPV in one institution [52]. Version 2 showed improved but institutionally variable discrimination and a persistently low PPV/high alert burden across four systems [53].
Local validation	Run silently on local adult inpatient data using the exact production inputs. Assess calibration, sensitivity, PPV, number needed to evaluate, lead time, subgroup performance, unit-level variation, missingness, and simulated alert volume at candidate thresholds.
Human oversight and workflow	The alert prompts clinician reassessment; it does not mandate antibiotics. Define the responsible role, response window, documentation, override/escalation, duplicate-alert suppression, and downtime workflow.
Equity and safety	Compare performance, alert receipt, response, and downstream actions across age, sex, race/ethnicity where legally and methodologically appropriate, language, disability, service, and other locally relevant groups. Review false-negative and false-positive cases.
Vendor, cloud, and change control	Confirm whether scoring is local or remote, the data transmitted, availability, audit logs, threshold control, update notices, support, and the ability to continue or exit after contract change.

Table 6. *Cont.*

Monograph Field	Illustrative Entry and Decision
Monitoring	Weekly leading indicators during the pilot; monthly/quarterly performance as labels mature; track alerts, acknowledgment, response time, overrides, testing, antibiotics, outcomes, alert burden, subgroup effects, and version/input changes.
Gate 1 decision	Proceed to evidence review only if the hospital identifies a clinical need, a responsible owner, non-AI alternatives, and a feasible response pathway.
Gate 2 decision	Do not approve solely on vendor claims. The evidence supports version-specific local validation; mixed evidence and low PPV justify caution rather than automatic rejection or adoption.
Gate 3 decision	Proceed only if local silent validation meets prespecified safety, calibration, equity, and workload criteria. Choose the threshold locally rather than importing one from another institution.
Gate 4 decision	Limited pilot in selected units with trained staff, a comparator/baseline plan, an alert-silencing strategy, safety huddles, and stop criteria.
Gate 5 decision	Continue only while benefit, workload, calibration, equity, system availability, and appropriate clinician response remain acceptable. Investigate performance or behavior changes promptly.
Gate 6 decision	Renew the exact version for a fixed period, restrict the indication or units, recalibrate/revalidate after change, suspend after safety signals, or retire if value is not demonstrated or a better alternative exists.

This walkthrough illustrates proportionality and feasibility. The committee does not reproduce a regulator’s conformity assessment or demand a randomized trial for every local change. It asks a sequence of practical questions, records evidence and uncertainty, and limits authorization to a defined version, indication, workflow, and time period. The example also shows that “better” performance in a later version is not equivalent to universal transportability: large between-site differences and the alert burden remain governance issues [53].

12. Discussion

The AI Hospital Formulary translates governance principles into a hospital authorization and portfolio management process. Its contribution is not a new technical validation method or a replacement for existing oversight. Rather, it provides a familiar institutional container in which local predictive-model governance, model documentation, patient-safety infrastructure, cybersecurity, procurement, regulatory duties, ethics, and implementation science can produce one traceable decision about whether a system may be used, under which restrictions, and for how long.

This approach can avoid two extremes: uncritical enthusiasm, in which tools are adopted because they are innovative or commercially attractive, and regulatory paralysis, in which uncertainty prevents potentially useful innovation. Proportionality is essential. A non-clinical meeting-summary tool can receive an abbreviated review focused on confidentiality, acceptable use, and vendor terms. A sepsis alert, autonomous triage system, or patient-facing treatment recommendation should receive the complete pathway.

The framework also strengthens accountability. Without a formulary, responsibility can be diffused among vendors, clinicians, information technology teams, innovation units, and executives. A formulary assigns named ownership, creates a decision record, and links monitoring to the authority for restriction or suspension. Governance-of-governance safe-

guards reduce the risk that the committee itself becomes opaque, commercially captured, or inaccessible to patients and professionals.

Implementation challenges remain. Smaller hospitals may lack the expertise to operate a full committee; regional or network-level services may therefore support shared evidence review, procurement, and validation while retaining local workflow accountability. Vendors may resist transparency. Clinicians may perceive review as bureaucracy unless timelines and requirements are proportional. Monitoring may be constrained by weak data infrastructure, delayed outcomes, and incomplete demographic data. Cloud concentration and vendor lock-in can limit institutional leverage.

This article has important limitations. It is a conceptual Perspective based on a purposive, non-systematic narrative synthesis; it does not provide exhaustive evidence identification, formal risk-of-bias assessment, consensus validation, or prospective testing of the complete framework. The sepsis example uses published data and does not constitute a new empirical evaluation or local implementation study. The regulatory mapping is illustrative and cannot replace jurisdiction-specific legal analysis. The framework may also require adaptation for national health systems, hospital networks, and smaller institutions. Future research should pilot the formulary, measure feasibility and decision time, assess effects on procurement quality and incident detection, and compare implementation and deprescribing outcomes with existing fragmented governance.

Nevertheless, hospitals already maintain governance mechanisms for medicines, devices, infection control, transfusion, radiation safety, clinical pathways, and patient safety. AI now requires a similarly visible form of stewardship that builds on, rather than duplicates, those structures.

13. Conclusions

AI will enter hospitals as a growing ecosystem of models, assistants, alerts, predictions, interfaces, cloud services, and automated workflows. Managing this ecosystem through ordinary procurement and isolated projects is insufficient.

The AI Hospital Formulary offers a practical institutional layer for safe, equitable, secure, and sustainable adoption. Each system should have a clear indication, version, evidence base, local validation proportionate to risk, monitoring plan, ethical justification, accountability structure, cloud and data-sovereignty assessment, approval expiry, and deprescribing criteria. The formulary also creates institutional memory of what is in use, why it was approved, how it is performing, and when it should be reconsidered.

Hospitals should not merely install AI. They should prescribe it, monitor it, audit it, renew it, restrict it, and deprescribe it when necessary.

Funding: This research received no external funding, and the author declares no financial relationships relevant to this manuscript.

Institutional Review Board Statement: Not applicable.

Informed Consent Statement: Not applicable.

Data Availability Statement: No new data were created or analyzed in this study. Data sharing is not applicable to this article.

Conflicts of Interest: The author declares no conflicts of interest.

References

1. Topol, E.J. High-performance medicine: The convergence of human and artificial intelligence. *Nat. Med.* **2019**, *25*, 44–56. [[CrossRef](#)] [[PubMed](#)]
2. Rajkomar, A.; Dean, J.; Kohane, I. Machine learning in medicine. *N. Engl. J. Med.* **2019**, *380*, 1347–1358. [[CrossRef](#)] [[PubMed](#)]

3. Yu, K.H.; Beam, A.L.; Kohane, I.S. Artificial intelligence in healthcare. *Nat. Biomed. Eng.* **2018**, *2*, 719–731. [CrossRef] [PubMed]
4. Kelly, C.J.; Karthikesalingam, A.; Suleyman, M.; Corrado, G.; King, D. Key challenges for delivering clinical impact with artificial intelligence. *BMC Med.* **2019**, *17*, 195. [CrossRef] [PubMed]
5. Nagendran, M.; Chen, Y.; Lovejoy, C.A.; Gordon, A.C.; Komorowski, M.; Harvey, H.; Topol, E.J.; Ioannidis, J.P.; Collins, G.S.; Maruthappu, M. Artificial intelligence versus clinicians: Systematic review of design, reporting standards, and claims of deep learning studies. *BMJ* **2020**, *368*, m689. [CrossRef] [PubMed]
6. Liu, X.; Faes, L.; Kale, A.U.; Wagner, S.K.; Fu, D.J.; Bruynseels, A.; Mahendiran, T.; Moraes, G.; Shamdas, M.; Kern, C.; et al. A comparison of deep learning performance against health-care professionals in detecting diseases from medical imaging: A systematic review and meta-analysis. *Lancet Digit. Health* **2019**, *1*, e271–e297. [CrossRef] [PubMed]
7. World Health Organization. *Ethics and Governance of Artificial Intelligence for Health: WHO Guidance*; World Health Organization: Geneva, Switzerland, 2021. Available online: <https://www.who.int/publications/i/item/9789240029200> (accessed on 6 June 2026).
8. World Health Organization. *Regulatory Considerations on Artificial Intelligence for Health*; World Health Organization: Geneva, Switzerland, 2023. Available online: <https://www.who.int/publications/i/item/9789240078871> (accessed on 6 June 2026).
9. World Health Organization. *Ethics and Governance of Artificial Intelligence for Health: Guidance on Large Multi-Modal Models*; World Health Organization: Geneva, Switzerland, 2025. Available online: <https://www.who.int/publications/i/item/9789240084759> (accessed on 6 June 2026).
10. European Parliament, Council of the European Union. Regulation (EU) 2024/1689 of the European Parliament and of the Council of 13 June 2024 Laying Down Harmonised Rules on Artificial Intelligence. 2024. Available online: <http://data.europa.eu/eli/reg/2024/1689/oj> (accessed on 20 July 2026).
11. US Food and Drug Administration. *Artificial Intelligence/Machine Learning-Based Software as a Medical Device Action Plan*; US Food and Drug Administration: Silver Spring, MD, USA, 2021. Available online: <https://www.fda.gov/media/145022/download> (accessed on 6 June 2026).
12. International Medical Device Regulators Forum. *Software as a Medical Device: Clinical Evaluation*. IMDRF/SaMD WG/N41FINAL:2017. 2017. Available online: <https://www.imdrf.org/documents/software-medical-device-samd-clinical-evaluation> (accessed on 6 June 2026).
13. International Medical Device Regulators Forum. *Good Machine Learning Practice for Medical Device Development: Guiding principles*. IMDRF/AIML WG/N88 FINAL:2025. 2025. Available online: <https://www.imdrf.org/documents/good-machine-learning-practice-medical-device-development-guiding-principles> (accessed on 6 June 2026).
14. National Institute of Standards and Technology. *Artificial Intelligence Risk Management Framework, AI RMF 1.0*; National Institute of Standards and Technology: Gaithersburg, MD, USA, 2023. [CrossRef]
15. Bedoya, A.D.; Economou-Zavlanos, N.J.; Goldstein, B.A.; Young, A.; Jelovsek, J.E.; O'Brien, C.; Parrish, A.B.; Elengold, S.; Lytle, K.; Balu, S.; et al. A framework for the oversight and local deployment of safe and high-quality prediction models. *J. Am. Med. Inform. Assoc.* **2022**, *29*, 1631–1636. [CrossRef] [PubMed]
16. Reddy, S.; Allan, S.; Coghlan, S.; Cooper, P. A governance model for the application of AI in health care. *J. Am. Med. Inform. Assoc.* **2020**, *27*, 491–497. [CrossRef] [PubMed]
17. Collins, G.S.; Moons, K.G.M.; Dhiman, P.; Riley, R.D.; Beam, A.L.; Van Calster, B.; Ghassemi, M.; Liu, X.; Reitsma, J.B.; Van Smeden, M.; et al. TRIPOD+AI statement: Updated guidance for reporting clinical prediction models that use regression or machine learning methods. *BMJ* **2024**, *385*, e078378. [CrossRef] [PubMed]
18. Liu, X.; Cruz Rivera, S.; Moher, D.; Calvert, M.J.; Denniston, A.K. SPIRIT-AI and CONSORT-AI Working Group. Reporting guidelines for clinical trial reports for interventions involving artificial intelligence: The CONSORT-AI extension. *Nat. Med.* **2020**, *26*, 1364–1374. [CrossRef] [PubMed]
19. Cruz Rivera, S.; Liu, X.; Chan, A.W.; Denniston, A.K.; Calvert, M.J. SPIRIT-AI and CONSORT-AI Working Group. Guidelines for clinical trial protocols for interventions involving artificial intelligence: The SPIRIT-AI extension. *Nat. Med.* **2020**, *26*, 1351–1363. [CrossRef] [PubMed]
20. Vasey, B.; Nagendran, M.; Campbell, B.; Clifton, D.A.; Collins, G.S.; Denaxas, S.; Denniston, A.K.; Faes, L.; Geerts, B.; Ibrahim, M.; et al. Reporting guideline for the early-stage clinical evaluation of decision support systems driven by artificial intelligence: DECIDE-AI. *BMJ* **2022**, *377*, e070904. [CrossRef] [PubMed]
21. Moons, K.G.M.; Damen, J.A.A.; Kaul, T.; Hooft, L.; Andaur Navarro, C.; Dhiman, P.; Beam, A.L.; Van Calster, B.; Celi, L.A.; Denaxas, S.; et al. PROBAST + AI: An updated quality, risk of bias, and applicability assessment tool for prediction models using regression or artificial intelligence methods. *BMJ* **2025**, *388*, e082505. [CrossRef] [PubMed]
22. Mitchell, M.; Wu, S.; Zaldivar, A.; Barnes, P.; Vasserman, L.; Hutchinson, B.; Spitzer, E.; Raji, I.D.; Gebru, T. Model cards for model reporting. In *Proceedings of the Conference on Fairness, Accountability, and Transparency*; Association for Computing Machinery: New York, NY, USA, 2019; pp. 220–229. [CrossRef]

23. Fleisher, L.A.; Economou-Zavlanos, N.J. Artificial intelligence can be regulated using current patient safety procedures and infrastructure in hospitals. *JAMA Health Forum* **2024**, *5*, e241369. [CrossRef] [PubMed]
24. Ratwani, R.M.; Bates, D.W.; Classen, D.C. Patient safety and artificial intelligence in clinical care. *JAMA Health Forum* **2024**, *5*, e235514. [CrossRef] [PubMed]
25. Challen, R.; Denny, J.; Pitt, M.; Gompels, L.; Edwards, T.; Tsaneva-Atanasova, K. Artificial intelligence, bias and clinical safety. *BMJ Qual. Saf.* **2019**, *28*, 231–237. [CrossRef] [PubMed]
26. Obermeyer, Z.; Powers, B.; Vogeli, C.; Mullainathan, S. Dissecting racial bias in an algorithm used to manage the health of populations. *Science* **2019**, *366*, 447–453. [CrossRef] [PubMed]
27. Gianfrancesco, M.A.; Tamang, S.; Yazdany, J.; Schmajuk, G. Potential biases in machine learning algorithms using electronic health record data. *JAMA Intern. Med.* **2018**, *178*, 1544–1547. [CrossRef] [PubMed]
28. Finlayson, S.G.; Bowers, J.D.; Ito, J.; Zittrain, J.L.; Beam, A.L.; Kohane, I.S. Adversarial attacks on medical machine learning. *Science* **2019**, *363*, 1287–1289. [CrossRef] [PubMed]
29. Subbaswamy, A.; Saria, S. From development to deployment: Dataset shift, causality, and shift-stable models in health AI. *Biostatistics* **2020**, *21*, 345–352. [CrossRef] [PubMed]
30. Davis, S.E.; Greevy, R.A., Jr.; Lasko, T.A.; Walsh, C.G.; Matheny, M.E. Detection of calibration drift in clinical prediction models to inform model updating. *J. Biomed. Inform.* **2020**, *112*, 103611. [CrossRef] [PubMed]
31. Embi, P.J. Algorithmovigilance—Advancing methods to analyze and monitor artificial intelligence-driven health care for effectiveness and equity. *JAMA Netw. Open* **2021**, *4*, e214622. [CrossRef] [PubMed]
32. Balendran, A.; Benchoufi, M.; Evgeniou, T.; Ravaud, P. Algorithmovigilance, lessons from pharmacovigilance. *npj Digit. Med.* **2024**, *7*, 270. [CrossRef] [PubMed]
33. National Institute of Standards and Technology. *Artificial Intelligence Risk Management Framework: Generative Artificial Intelligence Profile*; NIST AI 600-1; National Institute of Standards and Technology: Gaithersburg, MD, USA, 2024. [CrossRef]
34. Ong, J.C.L.; Chang, S.Y.H.; William, W.; Butte, A.J.; Shah, N.H.; Chew, L.S.T.; Liu, N.; Doshi-Velez, F.; Lu, W.; Savulescu, J.; et al. Ethical and regulatory challenges of large language models in medicine. *Lancet Digit. Health* **2024**, *6*, e428–e432. [CrossRef] [PubMed]
35. Freyer, O.; Wiest, I.C.; Kather, J.N.; Gilbert, S. A future role for health applications of large language models depends on regulators enforcing safety standards. *Lancet Digit. Health* **2024**, *6*, e662–e672. [CrossRef] [PubMed]
36. Huo, B.; Boyle, A.; Marfo, N.; Tangamornsuksan, W.; Steen, J.P.; McKechnie, T.; Lee, Y.; Mayol, J.; Antoniou, S.A.; Thirunavukarasu, A.J.; et al. Large language models for chatbot health advice studies: A systematic review. *JAMA Netw. Open* **2025**, *8*, e2457879. [CrossRef] [PubMed]
37. Everson, J.; Nong, P.; Richwine, C. Uptake of generative AI integrated with electronic health records in US hospitals. *JAMA Netw. Open* **2025**, *8*, e2549463. [CrossRef] [PubMed]
38. EQUATOR Network. Selecting the Appropriate Reporting Guideline for Your Article. Available online: <https://www.equator-network.org/toolkits/selecting-the-appropriate-reporting-guideline/> (accessed on 20 July 2026).
39. Baethge, C.; Goldbeck-Wood, S.; Mertens, S. SANRA—A scale for the quality assessment of narrative review articles. *Res. Integr. Peer Rev.* **2019**, *4*, 5. [CrossRef] [PubMed]
40. Misra, N.K.; Das, S.; Satpathy, S.; Addula, S.R.; Trivedi, M.C. COVID-19 pandemic: A worldwide critical review with the machine learning model-based prediction. *J. Inst. Eng. India Ser. B* **2025**, *106*, 339–349. [CrossRef]
41. Bates, D.W.; Kuperman, G.J.; Wang, S.; Gandhi, T.; Kittler, A.; Volk, L.; Spurr, C.; Khorasani, R.; Tanasijevic, M.; Middleton, B. Ten commandments for effective clinical decision support: Making the practice of evidence-based medicine a reality. *J. Am. Med. Inform. Assoc.* **2003**, *10*, 523–530. [CrossRef] [PubMed]
42. Greenhalgh, T.; Wherton, J.; Papoutsis, C.; Lynch, J.; Hughes, G.; A’Court, C.; Hinder, S.; Fahy, N.; Procter, R.; Shaw, S. Beyond adoption: A new framework for theorizing and evaluating nonadoption, abandonment, and challenges to the scale-up, spread, and sustainability of health and care technologies. *J. Med. Internet Res.* **2017**, *19*, e367. [CrossRef] [PubMed]
43. Proctor, E.; Silmere, H.; Raghavan, R.; Hovmand, P.; Aarons, G.; Bunker, A.; Griffey, R.; Hensley, M. Outcomes for implementation research: Conceptual distinctions, measurement challenges, and research agenda. *Adm. Policy Ment. Health* **2011**, *38*, 65–76. [CrossRef] [PubMed]
44. Damschroder, L.J.; Aron, D.C.; Keith, R.E.; Kirsh, S.R.; Alexander, J.A.; Lowery, J.C. Fostering implementation of health services research findings into practice: A consolidated framework for advancing implementation science. *Implement. Sci.* **2009**, *4*, 50. [CrossRef] [PubMed]
45. Institute of Medicine. *To Err Is Human: Building a Safer Health System*; National Academies Press: Washington, DC, USA, 2000. [CrossRef] [PubMed]
46. Institute of Medicine. *Crossing the Quality Chasm: A New Health System for the 21st Century*; National Academies Press: Washington, DC, USA, 2001. [CrossRef] [PubMed]

47. Orlander, J.D.; Barber, T.W.; Fincke, B.G. The morbidity and mortality conference: The delicate nature of learning from error. *Acad. Med.* **2002**, *77*, 1001–1006. [[CrossRef](#)] [[PubMed](#)]
48. Vayena, E.; Blasimme, A.; Cohen, I.G. Machine learning in medicine: Addressing ethical challenges. *PLoS Med.* **2018**, *15*, e1002689. [[CrossRef](#)] [[PubMed](#)]
49. Char, D.S.; Shah, N.H.; Magnus, D. Implementing machine learning in health care--addressing ethical challenges. *N. Engl. J. Med.* **2018**, *378*, 981–983. [[CrossRef](#)] [[PubMed](#)]
50. Rajkomar, A.; Hardt, M.; Howell, M.D.; Corrado, G.; Chin, M.H. Ensuring fairness in machine learning to advance health equity. *Ann. Intern. Med.* **2018**, *169*, 866–872. [[CrossRef](#)] [[PubMed](#)]
51. Amann, J.; Blasimme, A.; Vayena, E.; Frey, D.; Madai, V.I. Precise4Q Consortium. Explainability for artificial intelligence in healthcare: A multidisciplinary perspective. *BMC Med. Inform. Decis. Mak.* **2020**, *20*, 310. [[CrossRef](#)] [[PubMed](#)]
52. Wong, A.; Otles, E.; Donnelly, J.P.; Krumm, A.; McCullough, J.; DeTroyer-Cooley, O.; Pesttrue, J.; Phillips, M.; Konye, J.; Penzo, C.; et al. External validation of a widely implemented proprietary sepsis prediction model in hospitalized patients. *JAMA Intern. Med.* **2021**, *181*, 1065–1070. [[CrossRef](#)] [[PubMed](#)]
53. Wong, A.; Currey, D.; Schwinne, M.; Park-Egan, B.; Meyer, S.; Gutting, A.; Cao, J.; Khan, S.; Dantes, R.; Pan, T.; et al. Multicenter prospective validation of an updated proprietary sepsis prediction model. *JAMA Netw. Open* **2026**, *9*, e260181. [[CrossRef](#)] [[PubMed](#)]

Disclaimer/Publisher's Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.