Viewpoint

Institutional Approval Is Not Public Authorization: Legitimacy as a Governance Requirement for Clinical AI

Jialin Liu[1,2*], MD; Siru Liu[3*], PhD

[1]Information Center, West China Hospital of Sichuan University, Chengdu, Sichuan, China

[2]Department of Otolaryngology-Head and Neck Surgery, West China Hospital of Sichuan University, Chengdu, Sichuan, China

[3]Department of Biomedical Informatics, Vanderbilt University Medical Center, Nashville, TN, United States

[*]all authors contributed equally

Corresponding Author:
Siru Liu, PhD
Department of Biomedical Informatics
Vanderbilt University Medical Center
2525 West End Ave, Suite 1475
Nashville, TN, 37212
United States
Phone: 1 615 875 5216
Email: siru.liu@vumc.org

Abstract

Clinical AI governance has advanced rapidly in technical validation, regulatory compliance, risk management, human oversight, and postdeployment monitoring. Yet a more fundamental question often remains implicit: under what conditions are health institutions justified in exercising consequential authority through AI over patients, clinicians, and communities? In this Viewpoint, we argue that legitimacy should be recognized as an independent governance requirement, distinct from technical assurance, legal and ethical compliance, trust, acceptance, and institutional approval. Our aim is to distinguish institutional approval from public authorization and to propose a normative framework for authorizing consequential clinical AI deployments. We address health system leaders, clinical AI governance and procurement teams, frontline clinicians, and patient and community representatives. Drawing on theories of legitimacy and public value and published clinical AI governance research, we make the following contributions: (1) the concept of the unclaimed decision, defined as a consequential choice about purpose, threshold, population, or priority for which no actor clearly assumes responsibility for public justification; (2) a tripartite framework comprising procedural authorization, substantive justification, and relational renewal; and (3) practical governance tools to support public justification, contestation, and renewal. These tools include 3 minimum auditable conditions: visibility of deployment; affected-party input before governing conditions are fixed; and conditional reviewability, which requires a predefined trigger that can be invoked by a named actor outside the deploying team. A 5-question public authorization test examines purpose, distribution, exclusion, meaningful choice and contestability, and renewal. Governance intensity should increase with clinical consequence; distributive reach; difficulty of reversal; and the absence of genuine, informed choice. For smaller health care organizations, we propose ways to perform the same core governance functions through existing advisory bodies, targeted interviews, shared regional resources, and standardized documentation. We also outline a heuristic for identifying affected publics, procedures for specifying participation authority and resolving disagreement, and procurement provisions that preserve reviewability across vendor boundaries. These are normative proposals requiring evaluation. Public authorization requires continuing substantive justification; it cannot be inferred from compliance, popularity, or completion of an audit checklist.

(J Med Internet Res 2026;28:e109467) doi: 10.2196/109467

KEYWORDS

artificial intelligence; AI; clinical AI; AI governance; legitimacy; public authorization; public justification; social license; patient engagement; contestability; accountability

XSL•FO
RenderX

Introduction

Clinical AI governance increasingly emphasizes technical validation, regulatory compliance, risk management, and life cycle monitoring. However, published evaluations have identified important limitations in the evidence available to support oversight. A review of 1012 US Food and Drug Administration (FDA) device summaries through December 2024 reported a mean transparency score of 3.3 out of 17 [1]. A separate examination of 903 devices listed through August 2024 found that 505 (55.9%) reported clinical performance studies, of which 41 (8.1%) were prospective and 12 (2.4%) were randomized [2]. These findings concern publicly reported evidence rather than providing proof that unreported evaluations never occurred. The European Union AI Act and professional governance proposals provide additional legal, ethical, and organizational safeguards [3-7]. Nevertheless, technical assurance and regulatory compliance do not by themselves establish public authorization.

We use institutional approval to describe the regulatory, professional, and organizational warrants that permit an institution to deploy AI [4,6,7]. Affected publics remain incompletely represented in these processes [8,9]. Public authorization concerns whether the resulting exercise of authority is publicly justifiable, meaningfully open to affected-party influence, and subject to continuing contestation and revision [10]. A clinical AI system may therefore be safe, effective, legally permissible, and institutionally approved while still lacking public legitimacy if affected publics lack meaningful opportunities to shape its governing conditions or contest consequential uses. For example, AI referral triage may reorder access without directly diagnosing disease [11,12]. Its definition of urgency still requires justification beyond evidence of classification accuracy or organizational approval [13].

Experience with health data initiatives illustrates this distinction. Analyses of England's care data program identified deficiencies in public engagement, transparency, and safeguards despite an asserted public benefit rationale [14,15]. The Royal Free–DeepMind collaboration raised concerns about data use, transparency, and accountability [16,17]. These cases illustrate challenges of public justification; they do not establish the effectiveness of the framework proposed here. Social license emphasizes the continuing acceptability of institutional practices [14,18]; public authorization additionally asks whether the authority exercised can be justified to those subject to it, including dissenters [18,19].

Organizational and political theories offer distinct accounts of legitimacy [20,21]. Drawing on these accounts, we examine the public justification of deployment authority. Public value theory helps frame this inquiry: authorizing environment in health care combines regulatory mandates, institutional accountability, professional judgment, and affected-public participation [22]. Deployment choices may fall between these sources of authority, creating what we call the unclaimed decision: a consequential choice about purpose, threshold, population, or priority for which no actor clearly assumes responsibility for public justification.

This Viewpoint distinguishes institutional approval from public authorization and proposes a normative framework for governing consequential clinical AI deployments [19,23,24]. We advance four key messages: (1) legitimacy is an independent governance requirement that cannot be inferred from technical assurance, legal or ethical compliance, trust, acceptance, or institutional approval; (2) the unclaimed decision identifies gaps in responsibility for publicly justifying deployment choices; (3) public authorization requires procedural authorization, substantive justification, and relational renewal, supported by 3 minimum auditable conditions (visibility, affected-party input before governing conditions are fixed, and conditional reviewability) and complemented by the 5-question public authorization test; and (4) governance intensity should reflect clinical consequence, distributive reach, difficulty of reversal, and the absence of genuine, informed choice. The practical recommendations address institutional decision-making, procurement, clinical oversight, and patient and community participation. We also propose candidate process indicators to support empirical evaluation of public authorization practices.

The Legitimacy Concept: From Political Philosophy to Clinical AI

Legitimacy as a Distinct Normative Requirement

In political philosophy, legitimacy concerns the normative justification of authority: whether an institution has a defensible right to exercise power over those subject to its decisions [25,26]. It is distinct from legality, effectiveness, ethical compliance, and sociological acceptance. Authority may be lawful and effective yet remain illegitimate if its exercise lacks a publicly defensible justification [26]; conversely, public acceptance does not itself establish legitimacy [25].

This distinction is particularly important in clinical AI because health institutions increasingly exercise authority through computational systems as well as through clinicians and formal policies. AI systems stratify risk, prioritize patients, recommend treatments, flag diagnoses, and allocate clinical attention [27,28]. In doing so, they shape the distribution of treatment opportunities, risks, responsibilities, and other socially important goods [27-29]. Institutions must therefore be able to justify to affected publics why they are entitled to exercise clinical authority through a particular system, for a particular purpose, and under particular conditions [26,28,30]. Moreover, AI design and deployment embed value choices that technical validation cannot resolve: they concern how institutional power should be exercised and therefore require public justification [26,28,31].

The Structural Legitimacy Deficit in Clinical AI

Current clinical AI governance may treat legitimacy as a presumed consequence of institutional approval. A device may receive regulatory clearance [32], its local deployment may be classified as nonresearch quality improvement or receive a favorable institutional determination [33], and a health system committee may approve its integration [34]. Together, these processes may establish technical, legal, and organizational permissibility [28,32-34], but they do not by themselves

XSL•FO
RenderX

establish public authorization to govern through the technology [26,29].

The deficit is not an absence of governance. Patients, communities, and frontline clinicians may nonetheless have limited, late-stage, or advisory-only influence over a system's purposes, embedded trade-offs, target populations, or conditions for modification or withdrawal [8,35,36]. Those most affected by AI-mediated authority may therefore have little meaningful opportunity to shape, contest, or reauthorize its exercise. Existing oversight mechanisms remain necessary, but they do not by themselves establish why an institution is entitled to make particular deployment-level choices on behalf of those subject to them [26,29].

The Unclaimed Decision and the Limits of Existing Authorization

This problem is best understood as a jurisdictional gap. Regulatory mandates, institutional accountability, professional judgment, and affected-public participation each provide a partial warrant for clinical AI deployment, yet responsibility for publicly justifying deployment-level decisions about purpose, threshold, population, priority, and acceptable trade-offs may remain unclear [28,34,37]. The concept of the unclaimed decision introduced above captures this gap in responsibility. Accumulating existing forms of approval does not resolve the gap: regulatory, institutional, professional, and participatory processes answer different questions and provide complementary rather than substitutable warrants [26,28,29,34].

Individual consent remains important but cannot authorize every systemic consequence of clinical AI. Organizational uses may reshape access, workload, and resource allocation beyond the consenting encounter [28,34,38,39], particularly when AI-mediated processes are embedded in routine care and difficult to avoid. We distinguish 3 complementary mechanisms: voice, meaning influence over governing purposes and conditions; contestability, meaning the ability to challenge a consequential decision and obtain meaningful reconsideration; and exit, meaning a practical opportunity to avoid the AI-mediated pathway. A patient may challenge an AI-mediated decision without being able to influence deployment policy or decline an encounter-level AI tool while remaining subject to an AI-managed referral queue.

Why Legitimacy Becomes Unavoidable in Clinical AI

The legitimacy deficit is not merely conceptual. Research identifies concerns about trust in institutional AI use and the inclusion of affected communities in governance [35,40]. Insufficient public justification may compound these concerns, although its effects on resistance and inequities require empirical evaluation. We propose that stronger public authorization could help identify value conflicts and distributional concerns before deployment arrangements become embedded in routine care [8,35]. Legitimacy does not guarantee acceptance or better outcomes; it provides a basis for institutions to assess whether a proposed exercise of deployment authority is justifiable despite disagreement [25-27,29].

The legitimacy question becomes more consequential as AI is embedded in clinical workflows [28,34]. As validated tools become integrated into professional practice, clinicians may face expectations to use them, making nonuse as consequential as use [28,41], while AI-mediated care may become increasingly difficult for patients to avoid. The question is therefore not only whether AI is safe, effective, or trustworthy, but also on what grounds institutions subject patients to consequential algorithmic processes [26,29,34]. Patients may not have chosen these processes and may have limited ability to contest them [27,28].

Clinical AI governance must account for 2 constituencies simultaneously. Patients may be subject to AI-mediated prioritization, recommendations, communication, and access decisions, while frontline clinicians may experience changes in professional discretion, workload, accountability, and standards of care [34,37,42,43]. Public authorization therefore cannot be reduced to patient acceptance alone [35,40]. Legitimate governance must account for both those who receive AI-mediated care and those who remain professionally answerable for delivering it, particularly because clinical authority is embedded in fiduciary and relational obligations [44]. When AI redistributes judgment, attention, and responsibility, the authority exercised through it must remain publicly justifiable, contestable, and renewable [26,29,37,42,44].

A Tripartite Framework for Public Authorization

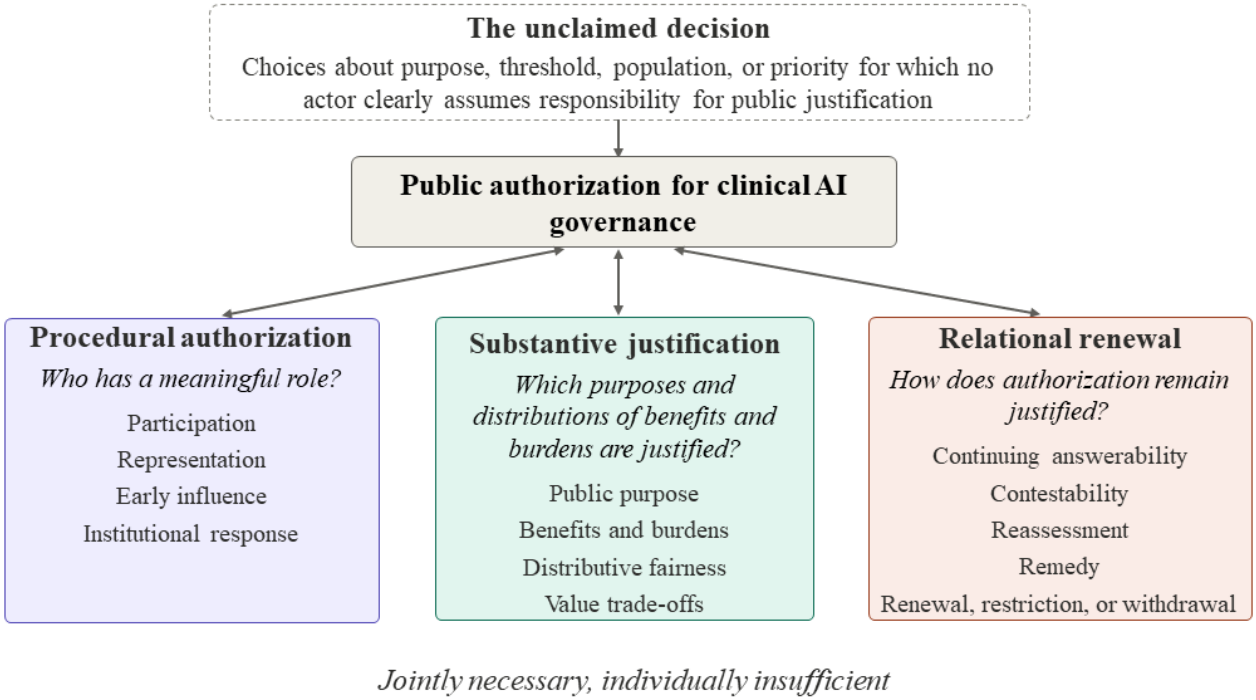
Conceptual Basis

Our framework connects political theories of legitimacy and public value theory [22,25,26,29] with governance concerns discussed in the clinical AI literature. These include limited affected-party influence over deployment decisions [8,35,36], insufficient public justification of purposes and distributive trade-offs [28,30,31], and the need for mechanisms to support contestation and reauthorization as systems and uses change [45,46]. We interpret these concerns as aspects of the unclaimed decision. This normative-conceptual synthesis is not a systematic review or a consensus exercise.

On this basis, we propose 3 interdependent dimensions of public authorization: procedural authorization, substantive justification, and relational renewal (Figure 1). Each addresses a distinct question: Who can meaningfully influence deployment decisions? Which purposes and distributions of benefit and burden can be publicly justified? How does authorization remain contestable and justified over time?

XSL•FO
RenderX

Figure 1. Tripartite framework for public authorization of clinical AI. The unclaimed decision arises when no actor clearly assumes responsibility for publicly justifying deployment choices about purpose, threshold, population, or priority. Public authorization addresses this gap through 3 interdependent dimensions: procedural authorization, substantive justification, and relational renewal. These dimensions are jointly necessary but individually insufficient and must remain aligned as systems, uses, and conditions change.



Procedural Authorization

Procedural authorization concerns how AI deployment decisions are made and whether those affected have meaningful influence over them. Affected constituencies, including disproportionately affected groups, should be involved before the purposes, trade-offs, and conditions of use are established [35]. Participation should involve deliberation or co-design where appropriate, with documented institutional responses showing how input influenced decisions [35]. Input should address choices that remain open to influence; product-level constraints may require engagement with developers or vendors.

In a national survey of US adults conducted in 2023, 62.7% rated as “very true” the statement that notification of AI use in their care is important [47]. Two UK citizens’ juries, each involving 18 participants over 5 days, favored accuracy over explainability in the health care scenarios considered [48]. These findings inform notification and deliberation but do not establish legitimacy or clinical benefits. Broader research identifies context-dependent expectations concerning public benefit, privacy, commercial involvement, and institutional trustworthiness [49-52]. Participation should therefore shape governing decisions rather than merely secure acceptance [35,51].

Mapping the Affected Public

We propose a 4-step process for mapping the affected public for each deployment. First, identify whose care, access, work, or responsibilities may change. Second, trace indirect effects, including queue displacement, caregiver burdens, consequences for data-contributing communities, and unequal exposure to error. Third, identify groups at risk of underrepresentation because of language, disability, digital access, geography, or social disadvantage. Fourth, publish the proposed boundary, the reasons for inclusion and exclusion, recruitment routes, and omissions; invite corrections; and revisit the map after material changes. This engagement-informed heuristic is not a validated sampling method [35,43,53].

For a hospital-wide risk model, relevant groups may include patients whose care is intensified or deprioritized, staff receiving alerts, and subgroups with unequal error burdens. Referral triage also affects people displaced within the queue and referring services. Ambient documentation affects patients whose encounters are recorded and clinicians whose documentation changes; secondary data uses may warrant wider inclusion. The boundary should reflect distributive reach rather than current use or ease of recruitment.

Participation Authority and Disagreement

Institutions should specify participants’ influence before engagement. For low-impact, readily avoidable uses, input may be advisory but should be accompanied by a reasoned response. For consequential uses, affected representatives should help set conditions, monitoring priorities, and review triggers. For high-impact, rights-sensitive, or effectively unavoidable uses, we propose formally delegated powers to require reconsideration or to pause deployment, subject to applicable law and institutional authority. These collective powers are distinct from an individual veto over a shared service.

A named accountable institutional body should resolve disagreement through documented deliberation, disclose conflicts of interest, publish reasons and dissent, and provide escalation to a body outside the deploying team. Delegated

XSL•FO
RenderX

powers should be explicitly defined, while the deploying institution retains accountability. Majority support cannot waive safety obligations or protections against discrimination. Unresolved concerns about fundamental interests should prompt modification, restriction, or suspension pending review [19,28,45].

Substantive Justification

Substantive justification concerns whether the purposes and distributive consequences of AI deployment can be publicly defended. It requires explicit scrutiny of what a system is designed to optimize, which trade-offs it privileges, and whose interests it prioritizes [28,31]. A legitimate deployment should pursue a clearly stated clinical or social purpose, generate benefits proportionate to its risks and burdens, and avoid creating or exacerbating inequities [27,28,31]. Efficiency and cost reduction may be relevant benefits, but they do not by themselves justify adverse effects on autonomy, access, workforce burden, or relational care, nor do they justify commercial extraction of public health data without adequate safeguards [28,31,51].

This dimension draws on the Belmont principles of respect for persons, beneficence, and justice [54], adapting them to clinical AI deployment beyond research settings. It requires health systems to assess and publicly report how AI-related benefits, burdens, errors, and opportunity costs are distributed. Health AI assurance laboratories can provide essential evidence on validity, safety, fairness, and real-world performance [7]. However, technical assurance cannot by itself establish the public justification of a deployment's purpose or distributive trade-offs. Technical metrics can inform but cannot resolve normative choices about tolerable errors, the time available for relational care, or the allocation of scarce attention [7,31].

For data-intensive or commercially developed systems, substantive justification should also address reciprocity: whether data-contributing communities receive proportionate benefits and whether commercial value extraction is consistent with the stated public purpose [28,51]. Where technically reasonable options remain normatively contestable, value pluralism itself becomes part of the governance problem [26,29,31].

Relational Renewal

Relational renewal concerns the continuing answerability of institutions that exercise authority through AI. Legitimacy cannot be secured once at procurement or deployment because systems, workflows, institutional incentives, and public expectations change over time [35,45,46]. Trust in clinical AI is relational rather than a fixed property of the system [40,49,55]. It should be treated as a possible consequence of legitimate governance, not as evidence of legitimacy itself [25,26,29]. Declining trust, rising complaints, repeated override patterns, or persistent subgroup concerns should therefore function as warning signals that prompt review, rather than as direct measures of legitimacy [45,46,56]. Relational renewal requires accessible mechanisms for patients and clinicians to identify AI involvement, challenge consequential outputs, and obtain human reconsideration, explanation, or remedy [28,45,56]. It also requires transparent communication, continued professional responsibility, and predefined review triggers when models, vendors, data sources, populations, purposes, workflows, or degrees of autonomy materially change [28,42,45,46]. Authorization should remain capable of renewal, restriction, suspension, or withdrawal when its original conditions no longer hold [45,46].

The 3 dimensions are jointly necessary but individually insufficient. Participation without substantive safeguards can create an appearance of legitimacy while leaving harmful deployment choices inadequately justified; technical fairness without meaningful participation can preserve paternalistic authority; and initial authorization without relational renewal can become obsolete. Public authorization therefore depends on the continuing alignment of fair participation, publicly defensible purposes, and accountable relationships.

The 5-Question Public Authorization Test

Purpose and Scope of the Test

The 5-question public authorization test (Table 1) translates the tripartite framework into 5 governance questions addressing purpose, distribution of benefits and burdens, exclusion, meaningful choice and contestability, and renewal. Table 1 links each question to minimum evidence and warning signs. These domains reflect established health AI governance priorities concerning public benefit, equity, accountability, human oversight, contestation and remedy, monitoring, and change management [3,28,45,57]. The test is a normative governance diagnostic, not a validated assessment instrument or a substitute for regulatory assessment.

It complements existing technical, ethical, legal, regulatory, and risk management assessments by examining the public justification of a specified deployment [26,28]. Its domains span procedural, substantive, and relational concerns. The test should be applied before deployment, at predefined review intervals, and after material changes in systems or deployment conditions [3,45,46,57]. Each answer should identify publicly defensible evidence, unresolved concerns, and any further justification, modification, or review needed.

XSL•FO
RenderX

Table 1. The 5-question public authorization test: questions, minimum evidence, and warning signs for clinical AI.

Question number	Public authorization questions[a]	Minimum evidence	Warning signs
1	Does the AI system serve a publicly defensible clinical or social purpose?	Meaningful clinical, patient-centered, safety, workflow, or population benefit compared with an appropriate alternative; justification for using AI; and public reason and value statement[b]	Benchmark performance or efficiency gains without evidence of meaningful benefit
2	Who benefits, and who bears the burdens?	Subgroup-specific benefits and harms and distribution of errors, costs, workload, risks, and opportunity costs	Average performance only; burdens shifted to patients, clinicians, caregivers, or disadvantaged groups
3	Does deployment create, reduce, or reproduce exclusion?	Assessment across relevant dimensions of disadvantage, including language, disability, geography, digital access, and socioeconomic context	Dependence on resources, literacy, language, or access that systematically excludes some groups
4	Do affected people have meaningful choice and access to contestation?	Proportionate disclosure; accessible non-AI option without undue cost, delay, or loss of quality or access; language and disability accommodations; reconsideration of choice; accessible explanation, appeal, meaningful human reconsideration, and remedy; and documented justification and safeguards where exit is infeasible	Invisible or effectively unavoidable AI use without adequate justification or safeguards; nominal opt-out with unreasonable disadvantages; inaccessible explanations or appeals; and no meaningful human reconsideration
5	Does authorization remain justified over time?	Periodic reviews and reviews after material changes; complaint and incident trends; corrective actions; and criteria for renewal, restriction, suspension, or withdrawal	Permanent approval; material changes without review; ignored complaints or incidents; and persistent subgroup concerns

[a]Procedural authorization is assessed primarily through the 3 minimum auditable conditions described separately. Question 4 examines meaningful choice, contestability, and human reconsideration; voice in deployment governance is addressed through affected-party input. Where a practical non-AI pathway is unavailable, its absence should be explicitly acknowledged and should inform the justification, safeguards, and intensity of review.

[b]A concise statement of the deployment's purpose, major value trade-offs, rationale for using AI, and limits of authorized use.

Interpreting the Test

The test addresses a distinct governance question: whether an institution is justified in exercising authority through a particular AI system, for a particular purpose and population, and under particular conditions [26,29]. It is noncompensatory and should not be reduced to an additive score: strength in 1 domain cannot offset an unresolved failure in another. Technical performance cannot compensate for authority that affected people cannot meaningfully challenge, nor can participation justify substantively unacceptable distributions of benefit and burden. Prior authorization cannot substitute for renewed justification when the system or its conditions materially change. An answer that reveals an authorization gap or leaves a material concern unresolved should therefore trigger further justification, modification, or review.

The test identifies where authorization remains insufficient and what evidence or governance response is needed before deployment proceeds or continues. The following recommendations translate these requirements into institutional governance mechanisms.

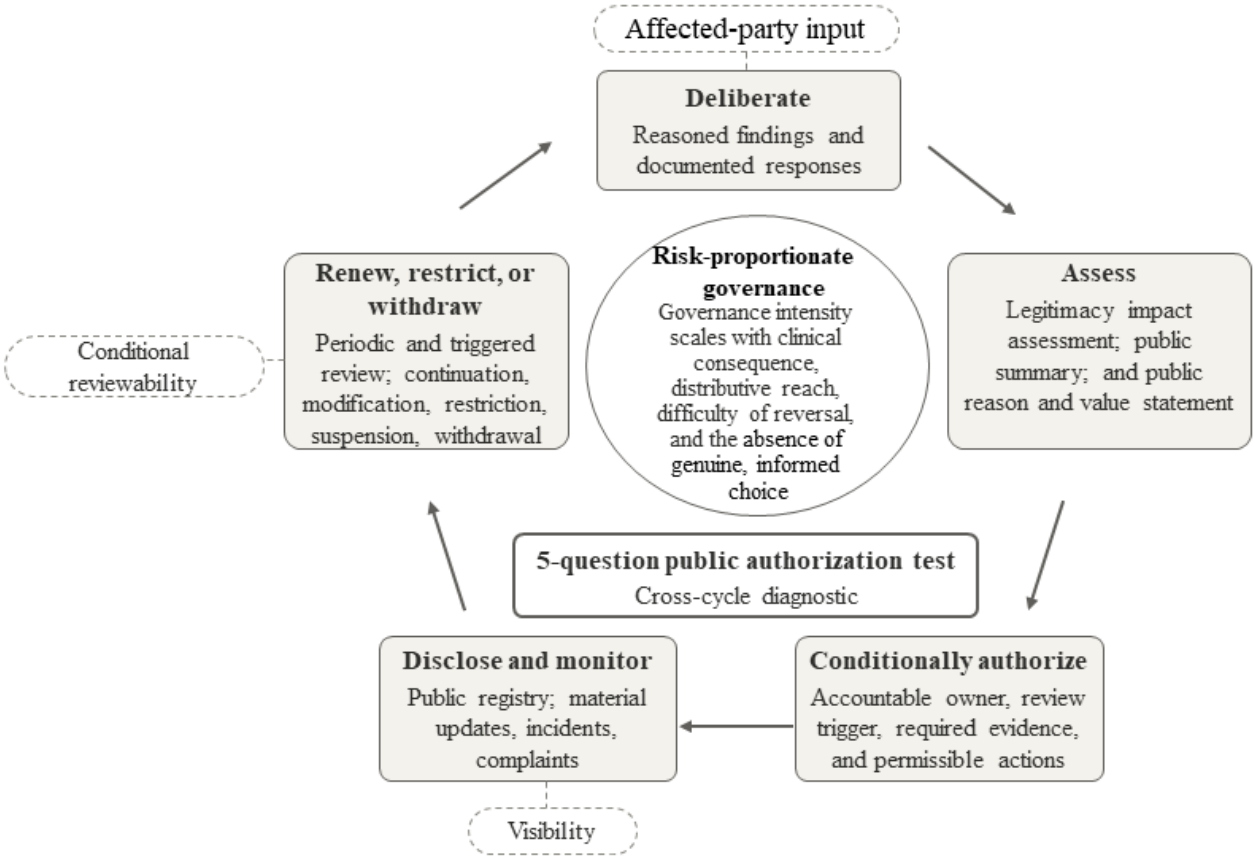
Operationalizing Legitimacy: Practical Recommendations

Overview

The tripartite framework supports a recurring, risk-proportionate governance cycle of deliberation, legitimacy impact assessment, conditional authorization, disclosure and monitoring, and renewal, restriction, or withdrawal (Figure 2). The 5-question public authorization test applies across the cycle, while the 3 minimum auditable conditions provide a procedural foundation. The cycle specifies governance functions; it does not require a new committee or a separate assessment document at every stage.

XSL•FO
RenderX

Figure 2. Operational pathway for public authorization of clinical AI. Public authorization operates as a recurring governance cycle comprising deliberation; legitimacy impact assessment; conditional authorization; disclosure and monitoring; and renewal, restriction, or withdrawal. The 5-question public authorization test serves as a cross-cycle diagnostic. Visibility, affected-party input, and conditional reviewability constitute the 3 minimum auditable conditions. The dashed tags indicate the stages at which these conditions are most directly operationalized, although they may extend across the entire cycle.



Minimum Auditable Conditions for Public Authorization

Translating the tripartite framework into practice requires mechanisms that support, document, and renew public authorization. We propose 3 minimum auditable conditions for consequential clinical AI deployments: visibility, affected-party input, and conditional reviewability. Visibility means that a deployment is discoverable by those it affects. Affected-party input must be structured and provided before the purposes and governing conditions are established. Conditional reviewability requires at least 1 predefined review trigger that can prompt formal reconsideration and be invoked by a named actor outside the deploying team. In operational terms, deployments should be locatable, claimable, and revisable.

These are minimum procedural conditions, not a comprehensive standard or proof of legitimacy. Together, they make AI-mediated authority identifiable, allow affected parties to influence choices before institutional embedding, and permit authorization to be reopened when circumstances change or its justification weakens. They are consistent with governance principles emphasizing transparency, stakeholder engagement, contestation, and redress [29,35,45].

Risk-Proportionate Public Deliberation

Deliberation should address the purposes and major trade-offs of proposed deployments [28,51]. It should occur early enough to influence procurement, workflow design, conditions of use, and the scope of authorized uses [35,45], and should produce reasoned findings that receive documented institutional responses. Public deliberation should complement, not replace, technical, regulatory, ethical, and clinical review.

We propose 4 determinants of governance intensity: clinical consequence; distributive reach across access, attention, opportunity, and burden; difficulty of reversal; and the absence of genuine, informed choice. Meaningful exit requires proportionate disclosure of AI involvement, an accessible non-AI pathway without unreasonable disadvantage in cost, waiting time, quality, or access, appropriate language and disability accommodations, and an opportunity to reconsider the choice. Where AI use is invisible or effectively unavoidable, stronger affected-party influence, accessible contestation, and more frequent review are warranted. Where a non-AI alternative is infeasible, institutions should explain why and justify continued AI use under clearly specified safeguards [28,38,45].

XSL•FO
RenderX

Proportionate Implementation in Smaller Health Care Organizations

Every health care organization using consequential AI should identify an accountable owner, make deployment visible, obtain timely affected-party input, publicly explain purposes and trade-offs, and maintain contestation and invocable review. These are minimum functions within the proposed framework; their organizational form can vary with risk and resources. They do not necessarily require dedicated panels, stand-alone assessments, bespoke registries, or in-house audit teams. Shared technical assessments can be reused, but local judgment is still needed to assess effects on workflows, identify excluded populations, and determine whether patients have a practical opportunity to opt out of AI-mediated care [9,34,45].

A community clinic or smaller behavioral health care organization could use its patient advisory body, sampled interviews, including interviews with underrepresented patients, and a clinician working group; document decisions using standardized assessment and disclosure templates; publish information on an existing service page; and designate an existing safety lead to receive concerns. Where a patient body is absent, trusted community partners can support recruitment. Regional networks or professional associations could share assessment expertise, contract templates, and independent escalation mechanisms. The deploying health care organization retains responsibility for its local decision.

Low-impact, readily reversible, and genuinely optional uses may receive abbreviated review. Consequential uses require documented affected-party input and conditional authorization; high-impact, broadly distributive, or effectively unavoidable uses warrant stronger representation and independent scrutiny, which can be obtained through shared services where necessary. To limit participation fatigue, institutions can reuse engagement where contexts are comparable and combine related reviews while preserving local input; participation should be accessible, compensated where feasible, and followed by an explanation of how input affected decisions. Periodic reassessment should continue, with additional review triggered by material changes, complaints, or emerging risks. Resource constraints should change the mechanism, not excuse failure to perform a necessary function [35,45,53].

Legitimacy Impact Assessment and Public Reason

Legitimacy impact assessment should ordinarily be integrated into, rather than duplicate, existing algorithmic, privacy, clinical safety, ethical, or fundamental rights impact assessments [3,45,57]. Beyond bias, fairness, and technical performance, evaluation should consider the meaningful involvement of affected publics, the public defensibility of the deployment's clinical or social purpose [25,26], the distribution of benefits, burdens, and opportunity costs, and the availability of mechanisms for contestation and remedy [3,28,31,45,57].

A public summary should document major concerns, institutional responses, and how affected-party input influenced the final decision. It should include a concise public reason and value statement specifying what the system is intended to optimize, the principal value trade-offs accepted, why AI is preferred to feasible alternatives, and which uses fall outside the scope of authorization. Relevant trade-offs include whose outcomes are prioritized, which errors are considered acceptable, how efficiency is balanced against relational care, and how burdens are distributed across affected publics [31,57].

Conditional Authorization and Periodic Renewal

For consequential clinical AI deployments, we recommend conditional authorization with predefined review intervals rather than presumptively indefinite approval. Authorization should specify the approved purpose, population, workflow, degree of autonomy, human oversight requirements, accountable institutional owner, and conditions for continued use.

Authorization should remain valid only while its specified conditions continue to hold. Formal reassessment should occur at predefined intervals and after material changes to the system, its intended use, population, vendor, or deployment context [3,45,46,57]. For AI-enabled devices regulated by the US FDA, relevant modification procedures should be aligned with any authorized predetermined change control plan [58]. Review should permit continuation, modification, restriction, suspension, or withdrawal of authorization.

At least 1 predefined review trigger should be invocable by a named actor outside the deploying team, such as a patient or community governance body, clinical safety office, ethics committee, or independent review body. Each authorization should specify the trigger, evidence required, and permissible actions. Review should not depend solely on self-initiation by the deploying team.

Disclosure and Monitoring

Publicly accessible deployment records should be maintained for consequential clinical AI through a registry or an existing service page, with the format adapted to the health care organization's scale and resources. These records should identify the system, vendor, intended purpose, deployment context, and accountable institutional owner, and summarize local validation, population and equity analyses, affected-party input, key findings from the legitimacy impact assessment, contestation and remedy pathways, authorization conditions, and the date or trigger for subsequent review [3,45,59,60].

These records should also report material updates and aggregated or deidentified information on incidents, complaints, corrective actions, and significant subgroup concerns [3,45,57], with proportionate safeguards for privacy, cybersecurity, and legitimate commercial confidentiality [3,57]. They should include the public reason and value statement, allowing affected publics to assess the stated purpose, accepted trade-offs, and institutional justification for deployment. Together, these records should support risk-proportionate public scrutiny, contestation, and review of AI-mediated institutional authority, without requiring a public referendum on every algorithm.

XSL•FO
RenderX

Extending Accountability Beyond Institutional Boundaries

Cross-Institutional Accountability and Remedy

Health systems should provide accessible pathways for patients and clinicians to challenge consequential AI-mediated decisions, request meaningful human reconsideration, and obtain documented responses and corrective action [3,28]. These pathways should include independent review beyond the deploying team for unresolved concerns or systemic problems.

Accountability must also extend across organizational and contractual boundaries. Legal counsel and procurement teams, working with clinical and governance leads, should negotiate contractual provisions that make conditional reviewability operational. Contracts should specify notification requirements for material changes affecting performance or authorized use; access to version histories, audit logs, and relevant validation and subgroup evidence; investigation and remediation responsibilities and timelines; and audits by qualified independent reviewers.

Health systems should secure contractual rights and operational means to pause, restrict, or disable use, revert to a prior version where feasible, and terminate the contract and transition to an alternative service when essential authorization conditions cannot be maintained. Access should be sufficient for oversight, with proportionate confidentiality safeguards, without necessarily requiring unrestricted source code disclosure. Urgent safety or security updates should follow an expedited notification and review process. These are proposed contractual safeguards, not universally available legal rights [3,45,57].

Deployment records should distinguish locally configurable choices from vendor-controlled constraints, including optimization targets, thresholds, updates, data dependencies, and modification limits. Outsourcing can deepen the problem of the unclaimed decision when an institution remains responsible for care but lacks the information, authority, or capacity to investigate and correct failures [37,41]. Shared contract templates and pooled procurement expertise may help smaller health care organizations secure these review conditions. Where essential review conditions cannot be secured, institutions should document the gap and restrict or decline the proposed use [45,57].

Specified Human Oversight and Relational Answerability

Human oversight should be specified rather than presumed. Requirements for human oversight are insufficient unless institutions define the conditions for meaningful review. These include the identity of the person responsible for reviewing AI-mediated decisions, when review occurs, what information and scope for independent judgment are available, what authority reviewers have to override recommendations or reverse decisions, and how review is documented [45,61]. A nominal human checkpoint may preserve procedural compliance while leaving algorithmic influence largely unchanged [3,45].

Accountability must also be relational rather than merely documentary. Patients and clinicians should be able to identify who remains answerable for a consequential AI-assisted decision, request meaningful human reconsideration, understand the reasons for the resulting decision, and obtain correction or remedy when a decision produces avoidable harm or unjustified disadvantage [28,44,45].

Limitations and Future Research

Limitations

The framework is normative and has not been empirically validated. Its dimensions, test, proposed governance arrangements, and process indicators require evaluation across applications, populations, and organizational settings. The examples illustrate proposals rather than demonstrated implementation effects.

No participation procedure can guarantee that all affected interests are represented or eliminate disagreement. Cultural and legal contexts, resource disparities, and institutional power may distort participation [28,35,53]. Mapping and publicly explaining representation decisions can expose these limitations but cannot resolve them mechanically. Advisory input should not be described as delegated decision authority.

Governance requirements may delay beneficial deployment, create participation fatigue, favor large organizations or established vendors, duplicate existing oversight, and widen inequalities between health care organizations [53,62]. The proportionate pathways proposed here are intended to mitigate those risks, but their costs and effectiveness remain uncertain.

Future Research

Measurement and construct validation. Studies should develop measures of perceived legitimacy distinct from trust, technology acceptance, satisfaction, perceived usefulness, and willingness to use AI [25,26,49,63]. Conventional acceptance frameworks address perceived usefulness, behavioral intention, and technology use [64-66]. Measures of perceived legitimacy should assess the perceived rightfulness of authority, the opportunity for procedural voice, the adequacy of public justification, distributive fairness, contestability, and access to remedy. Validation should examine reliability, construct validity, and measurement invariance across populations and jurisdictions. Perceived legitimacy cannot, by itself, establish normative legitimacy.

Participatory governance and implementation. Comparative studies should examine whether participation changes deployment purposes, conditions, or safeguards; whether legitimacy assessments identify concerns missed by conventional review; and whether procurement provisions and review mechanisms enable corrective action [3,35,45,51,60]. Comparisons across health care organization sizes should assess abbreviated, integrated, and shared governance arrangements, including costs, staff time, participation burden, deployment delays, and opportunity costs [53,62]. Evaluation should examine whether affected-public mapping and specified participation powers change whose concerns influence decisions

XSL•FO
RenderX

and whether shared resources preserve local influence and institutional responsibility.

Authorization processes and process indicators. Three candidate indicators warrant evaluation: authorization coverage, the proportion of consequential deployments with documented evidence for all 3 minimum auditable conditions; time to affected-party input, measured from formal initiation of deployment consideration to the first structured input from a materially affected group; and invocable review trigger rate, the proportion of consequential deployments with a predefined trigger invocable by a named actor outside the deploying team. The review trigger rate provides a component-specific measure that may identify weaknesses obscured by overall authorization coverage. These indicators should be assessed alongside evidence of whether decisions remained open to influence, meaningful choice, actual participation influence, and the timeliness and consequences of review. They should not be combined into a numerical legitimacy score.

Comparative and jurisdictional governance. Research should examine how regulation, professional standards, ownership, commercial incentives, financing, and Indigenous and community governance shape public authorization [3,28,45,60,67]. Indigenous data governance frameworks foreground collective benefit, authority to control, responsibility, and ethics [67]. Studies should include both patients and frontline clinicians.

Conclusions

Institutional approval and public authorization are not equivalent. The unclaimed decision identifies consequential deployment choices for which no actor clearly assumes responsibility for public justification. We propose a normative framework comprising procedural authorization, substantive justification, and relational renewal, supported by 3 minimum auditable conditions and the 5-question public authorization test. Governance intensity should reflect clinical consequence; distributive reach; difficulty of reversal; and the absence of genuine, informed choice. The proposed arrangements require empirical evaluation and do not constitute a certification standard or legitimacy score. Institutional approval permits deployment; public authorization requires continuing justification of the authority exercised through it.

Funding

The authors declare that no external financial support or grants were received from any public, commercial, or not-for-profit organization for the research, authorship, or publication of this article.

Data Availability

Data sharing is not applicable to this article, as no new data were created or analyzed in this study.

Generative AI Use

After the initial draft was completed, ChatGPT (version 6.1; OpenAI) was used to assist with language refinement. The authors independently evaluated all AI-generated revisions, made further modifications as needed, and take full responsibility for the content of the final manuscript.

Authors' Contributions

JL and SL conceived and designed the Viewpoint, performed the literature review, and drafted the manuscript. Both authors critically revised the manuscript for substantive intellectual content and approved the final version.

Conflicts of Interest

None declared.

References

1. Mehta V, Komanduri A, Bhadouriya RS, Mehta V, Johnson MD, Shrestha P, et al. Evaluating transparency in AI/ML model characteristics for FDA-reviewed medical devices. NPJ Digit Med. Nov 17, 2025;8(1):673. [FREE Full text] [doi: 10.1038/s41746-025-02052-9] [Medline: 41249460]
2. Windecker D, Baj G, Shiri I, Kazaj PM, Kaesmacher J, Gräni C, et al. Generalizability of FDA-approved AI-enabled medical devices for clinical use. JAMA Netw Open. Apr 01, 2025;8(4):e258052. [FREE Full text] [doi: 10.1001/jamanetworkopen.2025.8052] [Medline: 40305017]
3. Regulation (EU) 2024/1689 of the European Parliament and of the Council of 13 June 2024 laying down harmonised rules on artificial intelligence and amending Regulations (EC) No 300/2008, (EU) No 167/2013, (EU) No 168/2013, (EU) 2018/858, (EU) 2018/1139 and (EU) 2019/2144 and Directives 2014/90/EU, (EU) 2016/797 and (EU) 2020/1828 (Artificial Intelligence Act). European Union. URL: https://eur-lex.europa.eu/eli/reg/2024/1689/oj [accessed 2026-06-28]
4. Solomonides AE, Koski E, Atabaki SM, Weinberg S, McGreevey JD, Kannry JL, et al. Defining AMIA's artificial intelligence principles. J Am Med Inform Assoc. Mar 15, 2022;29(4):585-591. [FREE Full text] [doi: 10.1093/jamia/ocac006] [Medline: 35190824]

XSL•FO
RenderX

5. Wang C, Liu S, Yang H, Guo J, Wu Y, Liu J. Ethical considerations of using ChatGPT in health care. J Med Internet Res. Aug 11, 2023;25:e48009. [FREE Full text] [doi: 10.2196/48009] [Medline: 37566454]
6. Eaneff S, Obermeyer Z, Butte AJ. The case for algorithmic stewardship for artificial intelligence and machine learning technologies. JAMA. Oct 13, 2020;324(14):1397-1398. [doi: 10.1001/jama.2020.9371] [Medline: 32926087]
7. Shah NH, Halamka JD, Saria S, Pencina M, Tazbaz T, Tripathi M, et al. A nationwide network of health AI assurance laboratories. JAMA. Jan 16, 2024;331(3):245-249. [doi: 10.1001/jama.2023.26930] [Medline: 38117493]
8. Mello MM, Nag S, Tuttle CT, Char D. Engaging patients in AI governance. BMJ Digit Health Ai. Apr 10, 2026;2(1):e000099. [doi: 10.1136/bmjdh-2026-000099] [Medline: 42712282]
9. McKay F, Williams BJ, Prestwich G, Treanor D, Hallowell N. Public governance of medical artificial intelligence research in the UK: an integrated multi-scale model. Res Involv Engagem. May 21, 2022;8(1):21. [FREE Full text] [doi: 10.1186/s40900-022-00357-7] [Medline: 35598004]
10. Beetham D. The Legitimation of Power. Basingstoke, UK. Palgrave Macmillan; 2013.
11. Abdel-Hafez A, Jones M, Ebrahimabadi M, Ryan C, Graham S, Slee N, et al. Artificial intelligence in medical referrals triage based on clinical prioritization criteria. Front Digit Health. Oct 27, 2023;5:1192975. [FREE Full text] [doi: 10.3389/fdgth.2023.1192975] [Medline: 37964894]
12. Vergara PO, Oliveira JD, Mattiello R, Montelongo A, Roman R, Katz N, et al. Accuracy of artificial intelligence for gatekeeping in referrals to specialized care. JAMA Netw Open. Jun 02, 2025;8(6):e2513285. [FREE Full text] [doi: 10.1001/jamanetworkopen.2025.13285] [Medline: 40459894]
13. Habli I, Lawton T, Porter Z. Artificial intelligence in health care: accountability and safety. Bull World Health Organ. Apr 01, 2020;98(4):251-256. [FREE Full text] [doi: 10.2471/BLT.19.237487] [Medline: 32284648]
14. Carter P, Laurie GT, Dixon-Woods M. The social licence for research: why care.data ran into trouble. J Med Ethics. May 2015;41(5):404-409. [FREE Full text] [doi: 10.1136/medethics-2014-102374] [Medline: 25617016]
15. Review of data security, consent and opt-outs. United Kingdom Government. 2017. URL: https://www.gov.uk/government/publications/review-of-data-security-consent-and-opt-outs [accessed 2026-06-28]
16. Powles J, Hodson H. Google DeepMind and healthcare in an age of algorithms. Health Technol (Berl). 2017;7(4):351-367. [FREE Full text] [doi: 10.1007/s12553-017-0179-1] [Medline: 29308344]
17. Google DeepMind and class action lawsuit. Information Commissioner's Office. URL: https://ico.org.uk/for-the-public/ico-40/google-deepmind-and-class-action-lawsuit/ [accessed 2026-06-28]
18. Bharti N, O'Donovan C, Smallman M, Wilson J. Public trust, deliberative engagement and health data projects: beyond legal provisions. Engag Sci Technol Soc. Oct 05, 2021;7(1):125-133. [doi: 10.17351/ests2021.1197]
19. Daniels N, Sabin JE. Accountability for reasonableness: an update. BMJ. Oct 09, 2008;337:a1850. [doi: 10.1136/bmj.a1850] [Medline: 18845595]
20. Suchman MC. Managing legitimacy: strategic and institutional approaches. Acad Manag Rev. Jul 1995;20(3):571-610. [doi: 10.2307/258788]
21. Buchanan A, Keohane RO. The legitimacy of global governance institutions. Ethics Int Aff. Mar 25, 2011;20(4):405-437. [doi: 10.1111/j.1747-7093.2006.00043.x]
22. Moore MH. Creating Public Value: Strategic Management in Government. Cambridge, MA. Harvard University Press; 1995.
23. Schmidt VA. Democracy and legitimacy in the European Union revisited: input, output and ‘throughput’. Polit Stud. Aug 06, 2012;61(1):2-22. [doi: 10.1111/j.1467-9248.2012.00962.x]
24. Cohen T, Suzor NP. Contesting the public interest in AI governance. Internet Policy Rev. Sep 30, 2024;13(3). [doi: 10.14763/2024.3.1794]
25. Peter F. Political legitimacy. In: Zalta EN, Nodelman U, editors. The Stanford Encyclopedia of Philosophy. Stanford, CA. Metaphysics Research Lab, Stanford University; 2023.
26. Peter F. The Grounds of Political Legitimacy. Oxford, UK. Oxford University Press; 2023.
27. Vayena E, Blasimme A, Cohen IG. Machine learning in medicine: addressing ethical challenges. PLoS Med. Nov 6, 2018;15(11):e1002689. [doi: 10.1371/journal.pmed.1002689]
28. Ethics and governance of artificial intelligence for health: WHO guidance. World Health Organization. URL: https://www.who.int/publications/i/item/9789240029200 [accessed 2026-06-28]
29. Lazar S. Legitimacy, authority, and democratic duties of explanation. In: Sobel D, Wall S, editors. Oxford Studies in Political Philosophy. Oxford, UK. Oxford University Press; 2024.
30. Bartlett B. Towards accountable, legitimate and trustworthy AI in healthcare: enhancing AI ethics with effective data stewardship. New Bioeth. Dec 2024;30(4):285-309. [FREE Full text] [doi: 10.1080/20502877.2025.2482282] [Medline: 40135972]
31. Mittelstadt BD, Allo P, Taddeo M, Wachter S, Floridi L. The ethics of algorithms: mapping the debate. Big Data Soc. Dec 01, 2016;3(2):1-21. [doi: 10.1177/2053951716679679]
32. List of artificial intelligence-enabled medical devices. U.S. Food & Drug Administration. URL: https://www.fda.gov/medical-devices/software-medical-device-samd/artificial-intelligence-enabled-medical-devices [accessed 2026-06-28]

XSL•FO
RenderX

33. Quality improvement activities FAQs. U.S. Department of Health and Human Services. URL: https://www.hhs.gov/ohrp/regulations-and-policy/guidance/faq/quality-improvement-activities/index.html [accessed 2026-06-28]
34. Liao F, Adelaine S, Afshar M, Patterson BW. Governance of clinical AI applications to facilitate safe and equitable deployment in a large health system: key elements and early successes. Front Digit Health. Aug 24, 2022;4:931439. [FREE Full text] [doi: 10.3389/fdgth.2022.931439] [Medline: 36093386]
35. Abejirinde IO, Itanyi IU, Kornas K, Zaim R, Guha S, Chui V, et al. Principles and practices of community engagement in AI for population health: formative qualitative study of the AI for diabetes prediction and prevention project. J Particip Med. Oct 02, 2025;17:e69497. [FREE Full text] [doi: 10.2196/69497] [Medline: 41037755]
36. Ulizio CW, Dua D, Mukul NM, Areti S, Kostick-Quenet K, Rahimzadeh VN. Exploring the ethical and practical considerations of artificial intelligence in real-world health care settings: stakeholder focus group study. JMIR AI. Apr 02, 2026;5:e85163. [FREE Full text] [doi: 10.2196/85163] [Medline: 41927044]
37. Lang BH, Nyholm S, Blumenthal-Barby J. Responsibility gaps and black box healthcare AI: shared responsibilization as a solution. Digit Soc. Dec 2023;2(3):52. [doi: 10.1007/s44206-023-00073-z] [Medline: 38596344]
38. Rose SL, Shapiro D. An ethically supported framework for determining patient notification and informed consent practices when using artificial intelligence in health care. Chest. Sep 2024;166(3):572-578. [FREE Full text] [doi: 10.1016/j.chest.2024.04.014] [Medline: 38788895]
39. Sheehan M, Thompson R, Fistein J, Davies J, Dunn M, Parker M, et al. Authority and the future of consent in population-level biomedical research. Public Health Ethics. Oct 30, 2019;12(3):225-236. [FREE Full text] [doi: 10.1093/phe/phz015] [Medline: 32082417]
40. Nong P, Platt J. Patients' trust in health systems to use artificial intelligence. JAMA Netw Open. Feb 03, 2025;8(2):e2460628. [FREE Full text] [doi: 10.1001/jamanetworkopen.2024.60628] [Medline: 39951270]
41. Price WNII, Gerke S, Cohen IG. Liability for use of artificial intelligence in medicine. In: Research Handbook on Health, AI and the Law. Cheltenham, UK. Edward Elgar Publishing; Jul 16, 2024.
42. Schuitmaker L, Drogt J, Benders M, Jongsma K. Physicians' required competencies in AI-assisted clinical settings: a systematic review. Br Med Bull. Jan 16, 2025;153(1):ldae025. [FREE Full text] [doi: 10.1093/bmb/ldae025] [Medline: 39821209]
43. Vo V, Chen G, Aquino YS, Carter SM, Do QN, Woode ME. Multi-stakeholder preferences for the use of artificial intelligence in healthcare: a systematic review and thematic analysis. Soc Sci Med. Dec 2023;338:116357. [FREE Full text] [doi: 10.1016/j.socscimed.2023.116357] [Medline: 37949020]
44. Ludewigs S, Narchi J, Kiefer L, Winkler EC. Ethics of the fiduciary relationship between patient and physician: the case of informed consent. J Med Ethics. Dec 23, 2024;51(1):59-66. [FREE Full text] [doi: 10.1136/jme-2022-108539] [Medline: 36564172]
45. Lekadir K, Frangi AF, Porras AR, Glocker B, Cintas C, Langlotz CP, et al. FUTURE-AI: international consensus guideline for trustworthy and deployable artificial intelligence in healthcare. BMJ. Feb 05, 2025;388:e081554. [FREE Full text] [doi: 10.1136/bmj-2024-081554] [Medline: 39909534]
46. Keyes T, Callahan A, Pandya AS, Ambers N, Banda JM, Fuentes M, et al. Why and how to monitor deployed AI systems in health care. NEJM Catal Innov Care Deliv. Jun 2026;7(6):CAT250372. [doi: 10.1056/CAT.25.0372] [Medline: 42418612]
47. Platt J, Nong P, Carmona G, Kardia S. Public attitudes toward notification of use of artificial intelligence in health care. JAMA Netw Open. Dec 02, 2024;7(12):e2450102. [FREE Full text] [doi: 10.1001/jamanetworkopen.2024.50102] [Medline: 39661391]
48. van der Veer SN, Riste L, Cheraghi-Sohi S, Phipps DL, Tully MP, Bozentko K, et al. Trading off accuracy and explainability in AI decision-making: findings from 2 citizens' juries. J Am Med Inform Assoc. Sep 18, 2021;28(10):2128-2138. [FREE Full text] [doi: 10.1093/jamia/ocab127] [Medline: 34333646]
49. Bracic A, Spector-Bagdady K, Towle S, Zhang R, James CA, Price WNII. Factors for patient trust and acceptance of medical artificial intelligence. JAMA Netw Open. Mar 02, 2026;9(3):e260815. [FREE Full text] [doi: 10.1001/jamanetworkopen.2026.0815] [Medline: 41784960]
50. Beets B, Newman TP, Howell EL, Bao L, Yang S. Surveying public perceptions of artificial intelligence in health care in the United States: systematic review. J Med Internet Res. Apr 04, 2023;25:e40337. [FREE Full text] [doi: 10.2196/40337] [Medline: 37014676]
51. Muller SH, Kalkman S, van Thiel GJ, Mostert M, van Delden JJ. The social licence for data-intensive health research: towards co-creation, public value and trust. BMC Med Ethics. Aug 10, 2021;22(1):110. [FREE Full text] [doi: 10.1186/s12910-021-00677-5] [Medline: 34376204]
52. Busch F, Hoffmann L, Xu L, Zhang LJ, Hu B, García-Juárez I, COMFORT consortium, et al. Multinational attitudes toward AI in health care and diagnostics among hospital patients. JAMA Netw Open. Jun 02, 2025;8(6):e2514452. [FREE Full text] [doi: 10.1001/jamanetworkopen.2025.14452] [Medline: 40493367]
53. Ocloo J, Garfield S, Franklin BD, Dawson S. Exploring the theory, barriers and enablers for patient and public involvement across health, social care and patient safety: a systematic review of reviews. Health Res Policy Syst. Jan 20, 2021;19(1):8. [FREE Full text] [doi: 10.1186/s12961-020-00644-3] [Medline: 33472647]

XSL•FO
RenderX

54. Department of Health, Education, Welfare, National Commission for the Protection of Human Subjects of Biomedical and Behavioral Research. The Belmont Report. Ethical principles and guidelines for the protection of human subjects of research. J Am Coll Dent. 2014;81(3):4-13. [Medline: 25951677]
55. Goisauf M, Cano Abadía M, Akyüz K, Bobowicz M, Buyx A, Colussi I, et al. Trust, trustworthiness, and the future of medical AI: outcomes of an interdisciplinary expert workshop. J Med Internet Res. Jun 02, 2025;27:e71236. [FREE Full text] [doi: 10.2196/71236] [Medline: 40455564]
56. Mathias R, Vasey B, Chalkidou A, Riedemann L, Melvin T, Gilbert S. Safe AI-enabled digital health technologies need built-in open feedback. Nat Med. Feb 2025;31(2):370-375. [doi: 10.1038/s41591-024-03397-6] [Medline: 39905271]
57. Tabassi E. Artificial Intelligence Risk Management Framework (AI RMF 1.0). NIST Trustworthy and Responsible AI, National Institute of Standards and Technology. 2023. URL: https://www.nist.gov/publications/artificial-intelligence-risk-management-framework-ai-rmf-10 [accessed 2026-10-23]
58. Marketing submission recommendations for a predetermined change control plan for artificial intelligence-enabled device software functions: guidance for industry and Food and Drug Administration staff. U.S. Food & Drug Administration. Aug 2025. URL: https://www.fda.gov/regulatory-information/search-fda-guidance-documents/marketing-submission-recommendations-predetermined-change-control-plan-artificial-intelligence [accessed 2026-06-28]
59. Health data, technology, and interoperability: certification program updates, algorithm transparency, and information sharing. Federal Register. 2024. URL: https://www.federalregister.gov/documents/2024/01/09/2023-28857/health-data-technology-and-interoperability-certification-program-updates-algorithm-transparency-and [accessed 2026-10-23]
60. Labkoff S, Oladimeji B, Kannry J, Solomonides A, Leftwich R, Koski E, et al. Toward a responsible future: recommendations for AI-enabled clinical decision support. J Am Med Inform Assoc. Nov 01, 2024;31(11):2730-2739. [FREE Full text] [doi: 10.1093/jamia/ocae209] [Medline: 39325508]
61. Toro-Tobon D, Ponte OP, Montori VM, Brito JP. Clinician in the loop: a flawed solution for AI oversight. BMJ. May 05, 2026;393:e089213. [FREE Full text] [doi: 10.1136/bmj-2025-089213] [Medline: 42086290]
62. Bergholtz J, Wolf A, Crine V, Cleeve H, Santana MJ, Björkman I. Patient and public involvement in healthcare: a systematic mapping review of systematic reviews - identification of current research and possible directions for future research. BMJ Open. Sep 19, 2024;14(9):e083215. [FREE Full text] [doi: 10.1136/bmjopen-2023-083215] [Medline: 39304210]
63. Gordon SC, Huber GA. The empirical study of legitimate authority: normative guidance for positive analysis. In: Knight J, Schwartzberg M, editors. Political Legitimacy: NOMOS LXI. New York, NY. New York University Press; 2019.
64. Davis FD. Perceived usefulness, perceived ease of use, and user acceptance of information technology. MIS Q. Sep 1989;13(3):319-340. [doi: 10.2307/249008]
65. Venkatesh V, Morris MG, Davis GB, Davis FD. User acceptance of information technology: toward a unified view. MIS Q. 2003;27(3):425-478. [doi: 10.2307/30036540]
66. Dingel J, Kleine AK, Cecil J, Sigl AL, Lermer E, Gaube S. Predictors of health care practitioners' intention to use AI-enabled clinical decision support systems: meta-analysis based on the unified theory of acceptance and use of technology. J Med Internet Res. Aug 05, 2024;26:e57224. [FREE Full text] [doi: 10.2196/57224] [Medline: 39102675]
67. Carroll SR, Garba I, Figueroa-Rodríguez OL, Holbrook J, Lovett R, Materechera S, et al. The CARE principles for indigenous data governance. Data Sci J. 2020;19:43. [doi: 10.5334/dsj-2020-043]

Abbreviations

FDA: Food and Drug Administration

Edited by I Steenstra; submitted 12.Aug.2026; peer-reviewed by J-V Steiskal, SS Furusa; comments to author 08.Oct.2026; revised version received 18.Oct.2026; accepted 20.Oct.2026; published 02.Nov.2026

Please cite as:
Liu J, Liu S
Institutional Approval Is Not Public Authorization: Legitimacy as a Governance Requirement for Clinical AI
J Med Internet Res 2026;28:e109467
URL: https://www.jmir.org/2026/1/e109467
doi: 10.2196/109467
PMID:

©Jialin Liu, Siru Liu. Originally published in the Journal of Medical Internet Research (https://www.jmir.org), 02.Nov.2026. This is an open-access article distributed under the terms of the Creative Commons Attribution License (https://creativecommons.org/licenses/by/4.0/), which permits unrestricted use, distribution, and reproduction in any medium, provided the original work, first published in the Journal of Medical Internet Research (ISSN 1438-8871), is properly cited. The

XSL•FO
RenderX

complete bibliographic information, a link to the original publication on https://www.jmir.org/, as well as this copyright and license information must be included.

XSL•FO
RenderX