Original Article

Building an Explainability Contract between Clinical Algorithms and Practicing Pharmacists

Mark Anderson[1]*, Lisa Wong[2], Sarah Lee[1]

[1]Department of Explainable AI in Pharmacy, College of Pharmacy, University of Florida, Gainesville, United States. [2]Department of Algorithm Transparency and Pharmacy Practice, Faculty of Pharmacy, National University of Singapore, Singapore.

Abstract

Clinical algorithms increasingly support medication-risk identification, prescription prioritization, alert management, and other pharmacy activities. Yet explainability is commonly treated as a technical property of a model rather than an accountable relationship among developers, health-care organizations, practising pharmacists, and affected patients. An explanation may appear plausible while omitting uncertainty, local applicability, data limitations, workflow consequences, or mechanisms for professional challenge. This article proposes an explainability contract for pharmacist-facing clinical algorithms. The contract is a non-legal, sociotechnical governance construct that links each algorithmic role in the medication-decision lifecycle to defined parties, pharmacist rights, reciprocal duties, required disclosures, and procedures for challenge, correction, and escalation. Its principal components include intended-use boundaries, data provenance, patient-specific rationale, uncertainty and calibration, local-validity information, professional override, version traceability, and documented resolution of material concerns. The framework distinguishes model output from medication authorization, explanation plausibility from fidelity, technical performance from pharmaceutical usefulness, and implementation from demonstrated benefit. Contract performance would require staged technical, clinical, human-factors, organizational, medication-safety, and equity evaluation. Relevant outcomes include explanation fidelity, appropriate reliance, workflow accessibility, challenge responsiveness, correction completeness, subgroup performance, and recurrence of known failures. The proposed contract does not establish that explanation improves medication decisions, transfer liability to pharmacists, guarantee equitable outcomes, or constitute a validated clinical, legal, regulatory, or deployment standard. Its original contribution is to reposition explainability from optional information supplied by an algorithm into a reciprocal and contestable governance relationship requiring empirical validation.

Keywords: Digital pharmacy, Artificial intelligence, Pharmacy practice, Medication safety, Human–AI collaboration, Clinical decision support

Introduction

Explanations without Obligations

Explainable artificial intelligence is often presented as a response to the opacity of clinical algorithms. However, current health-care explainability methods cannot be assumed to reveal faithful patient-level reasoning or provide independent safety assurance [1]. Explanation requirements also vary among technical developers, clinicians, patients, ethicists, and organizational decision-makers [2]. Consequently, the existence of an explanation does not establish that the explanation is clinically useful, sufficiently complete, or appropriate for the person expected to act on it.

Clinical decision-support systems can improve information access and task organization, but their consequences depend on data quality, interface design, workflow integration, alert burden, professional interpretation, and governance [3].

These dependencies are especially consequential in pharmacy, where algorithmic outputs may affect prescription verification, interaction assessment, dose review, monitoring, prioritization, or communication with prescribers and patients. The pharmacist must connect a computational output to treatment history, comorbidity, organ function,

Address for correspondence: Mark Anderson, Department of Explainable AI in Pharmacy, College of Pharmacy, University of Florida, Gainesville, United States.
mark.anderson@ufl.edu
Received: 12 January 2025; **Accepted:** 08 April 2025

How to cite this article: Anderson M, Wong L, Lee S. Building an Explainability Contract between Clinical Algorithms and Practicing Pharmacists. Arch Pharm Pract. 2025;16(2):26-34. https://doi.org/10.51847/rgnlCe8U0W

therapeutic goals, medicine availability, adherence, and the practical consequences of acting or not acting.

Medication-alert applications illustrate the gap between explanation and obligation. Artificial intelligence has been investigated as a means of improving alert relevance, but evaluation methods and evidence of real-world value remain heterogeneous [4]. Machine learning has also been used to identify prescriptions with an elevated risk of medication error, demonstrating a plausible role in directing pharmacist attention [5]. Neither application establishes that an algorithm can explain why a particular output should govern a medication decision or who must respond when the explanation is incomplete, misleading, outdated, or inconsistent with pharmaceutical knowledge.

This article therefore proposes an explainability contract: a sociotechnical governance arrangement connecting algorithmic participation in a medication decision to required disclosures, pharmacist rights, reciprocal duties, and mechanisms for challenge, correction, and closure. "Contract" denotes structured accountability rather than automatic legal enforceability. The construct is intended to organize future evaluation, procurement, implementation, and professional governance. It does not presume that explanation is always beneficial or that one explanation format is appropriate across pharmacy tasks.

Why Generic Explainability Is Insufficient in Pharmacy

Explanation is not equivalent to unrestricted technical disclosure. Research across artificial intelligence and the social sciences characterizes useful explanations as selective, contrastive, audience-dependent, and connected to the question being asked [6]. Health-care explainability therefore requires explicit choices concerning the explanation recipient, intended purpose, scope, representation, and evaluation method [7]. A pharmacist deciding whether to clarify a dose needs different information from a developer debugging a model or a governance committee reviewing aggregate performance.

Generic explainability is also insufficient because a plausible post-hoc account may not faithfully describe the model mechanism that produced an output. For some high-stakes decisions, an interpretable model may be preferable to a black box accompanied by an explanatory approximation [8]. Health-care explanations can otherwise persuade users without reliably establishing model validity, causal reasoning, or clinical correctness [9]. This risk is material when an interface highlights apparently important variables but conceals missing data, poor calibration, an inapplicable population, or an unstable association.

For pharmacy practice, explanation adequacy should therefore be judged against the medication task. Generic explanation quality is jointly constrained by audience, explanatory purpose, and the risk of persuasive but non-faithful post-hoc accounts [6, 7, 9]. This article proposes six insufficiency tests: recipient fit, task fit, uncertainty disclosure, applicability disclosure, action relevance, and recourse. An explanation is contractually incomplete when it does not enable the pharmacist to understand what the output addresses, where it may fail, what action remains professionally authorized, and how a material concern can be challenged.

Alternative approaches remain possible. Better interface design, more interpretable models, improved training, or stronger validation may reduce particular explanation failures. None alone establishes reciprocal accountability across the medication-decision lifecycle. The contract is proposed to connect these approaches without treating any single technical mechanism as sufficient.

Table 1 organizes the evidence, constructs, relationships, uncertainties, and boundary conditions required for definitions, components, relationships, and intended uses for the article Building an Explainability Contract between Clinical Algorithms and Practising Pharmacists.

Table 1. Definitions, components, relationships, and intended uses for the article Building an Explainability Contract between Clinical Algorithms and Practising Pharmacists.

Component or scientific dimension	Problem addressed	Inputs or determinants	Proposed mechanism or relationship	Expected contribution	Evidence required	Failure or uncertainty risk	Boundary statement
Explainability contract	Explanations may be available without accountable action or correction [1].	Algorithmic role, decision consequence, parties, local governance	**Proposed:** connect disclosure, professional authority, challenge, and closure	Organize reciprocal accountability	Prospective governance and workflow evaluation	Documentation may exist without meaningful recourse	Not automatically a legally enforceable contract
Recipient fit	Different stakeholders require different explanations [2].	Role, knowledge, decision authority, information need	**Proposed:** tailor explanation access to the responsible recipient	Reduce irrelevant disclosure and hidden information	Comprehension and task-performance studies	Oversimplification or information overload	User satisfaction does not prove fidelity

Task fit	Clinical decision support is dependent on workflow and use context [3].	Medication task, urgency, reversibility, alternatives	**Proposed:** link explanation content to the medication decision	Improve pharmaceutical relevance	Task-specific human–AI studies	Technically correct information may remain unusable	Model performance is not medication-task performance
Pharmacy-specific output	Medication alerts and prescription-risk tools may direct pharmacist attention [4, 5].	Prescription data, patient context, prioritization rule	**Proposed:** explain why attention is directed and what remains unassessed	Support contextual verification	External, temporal, and prospective pharmacy evaluation	False reassurance, alert burden, omitted context	Prioritization is not authorization
Explanation fidelity	Post-hoc explanation may not represent relevant model behaviour [9].	Model architecture, explanation method, perturbation stability	**Proposed:** disclose explanation type and test fidelity	Distinguish rationale from persuasive display	Fidelity, robustness, and sensitivity testing	Plausible but misleading explanation	Explanation is not proof of correctness
Interpretability choice	High-stakes tasks may justify directly interpretable models [8].	Task complexity, required accuracy, available model classes	**Proposed:** justify black-box use and explanation strategy	Make design trade-offs reviewable	Comparative technical and clinical evaluation	Interpretability may be claimed without usability	Interpretable form does not guarantee valid use
Uncertainty and applicability	Explanation may omit confidence, missingness, or local mismatch	Calibration, data quality, population, temporal context	**Proposed:** pair rationale with uncertainty and applicability disclosures	Support cautious reliance	Calibration, shift, subgroup, and missing-data studies	Hidden uncertainty or false precision	Disclosure does not correct invalid predictions
Recourse	Users may receive outputs without a meaningful challenge route	Override authority, documentation, governance access	**Proposed:** guarantee challenge, review, disposition, and traceability	Convert explanation into a contestable process	Implementation, safety, and organizational studies	Challenge without response or closure	Recourse does not establish that every challenge is correct

Parties, Rights, and Duties within the Explainability Contract

Responsibility for clinical artificial intelligence cannot be assigned only at the point of use. Ethical issues arise during design, data selection, validation, implementation, and clinical operation [10]. Health-care AI governance similarly requires coordinated oversight across developers, organizations, clinicians, and other affected stakeholders [11]. The proposed contract therefore treats explainability as a distributed obligation rather than a feature delivered by a vendor and interpreted at the pharmacist's sole risk.

Trust does not remove this distribution of responsibility. Clinician trust in artificial intelligence is shaped by the user, task, system, institution, and consequences of reliance [12]. AI-supported decisions also create unresolved accountability questions when a professional remains formally responsible but cannot inspect development choices, data provenance, or model limitations [13]. In pharmacy, professional judgment should remain active, but it should not become a mechanism through which inaccessible technical or organizational failures are transferred to the practising pharmacist.

Figure 1 presents the original conceptual synthesis for explainability contract across the medication-decision lifecycle, showing how the article's principal components, evidence relationships, uncertainties, and decision boundaries are connected.

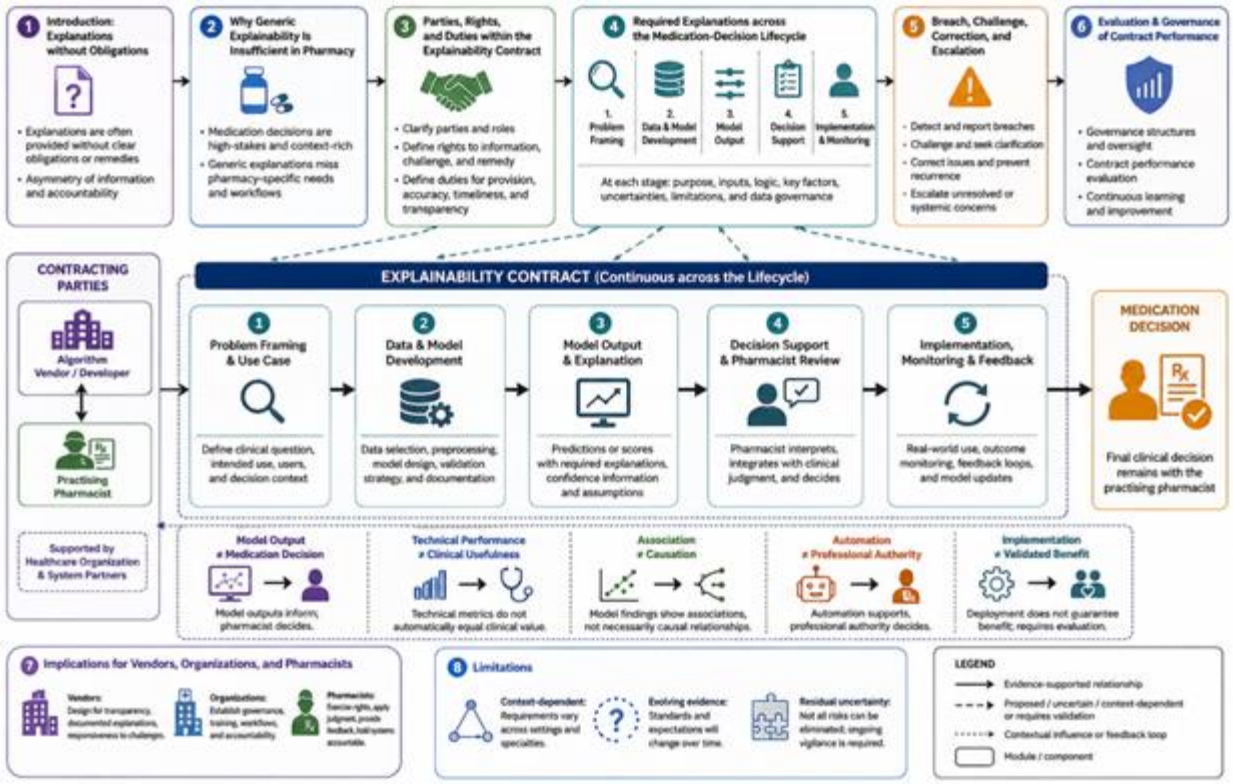

Figure 1. Explainability Contract across the Medication-Decision Lifecycle

The figure is an original conceptual synthesis developed for "Building an Explainability Contract between Clinical Algorithms and Practising Pharmacists." Arrows and grouping indicate proposed or evidence-supported relationships rather than measured effect sizes. The figure does not represent a validated predictive model, clinical recommendation, regulatory determination, or deployment-ready system.

The proposed contract recognizes seven parties: the vendor or developer; deploying organization; pharmacy leadership and medication-safety function; practising pharmacist; data, informatics, and security stewards; multidisciplinary governance body; and affected patient or caregiver. Their responsibilities overlap but are not interchangeable. A vendor can document intended use and known limitations but cannot determine whether an output is pharmaceutically appropriate for an individual patient. A pharmacist can challenge or override an output but cannot independently correct an inaccessible model. An organization can authorize implementation but should also maintain local validation, monitoring, escalation, and suspension mechanisms.

The pharmacist's proposed rights include access to intended and prohibited uses, relevant data provenance, material missingness, model and explanation version, patient-specific rationale, uncertainty, calibration, local applicability, known failure conditions, and available recourse. The pharmacist should also be able to override an output, document the reason, obtain review, and receive a traceable disposition without retaliation for appropriate professional challenge.

These rights create reciprocal duties. Vendors should communicate material limitations and changes; organizations should ensure local access, validation, governance, and monitoring; pharmacists should contextualize outputs and document consequential concerns; and governance bodies should investigate recurrent or high-severity failures. Patients retain an interest in understandable communication, equitable treatment, and accountability for consequential medication decisions. These are proposed allocations, not established legal rights or universally applicable professional standards.

Table 2 organizes the evidence, constructs, relationships, uncertainties, and boundary conditions required for evaluation criteria, evidence requirements, and failure modes for the article Building an Explainability Contract between Clinical Algorithms and Practising Pharmacists.

Table 2. Evaluation criteria, evidence requirements, and failure modes for the article Building an Explainability Contract between Clinical Algorithms and Practising Pharmacists.

Architecture layer or process stage	Core function	Information flow	Interaction with other components	Evaluation criterion	Potential failure mode	Human or organizational responsibility	Readiness boundary
Scope and intended use	Define what the algorithm may and may not support	Vendor to organization and pharmacist	Constrains every later explanation	Intended-use comprehension and use conformity	Use outside validated task or population	Vendor documents; organization enforces	Documentation alone does not establish local readiness
Data and provenance	Identify relevant inputs, sources, missingness, and transformations	Data steward and system to pharmacist	Conditions rationale and applicability	Traceability and material-data completeness	Hidden missingness or unreliable source data	Data steward investigates; pharmacist contextualizes	Provenance cannot establish causal validity
Explanation generation	Provide task-relevant rationale	Model or explanation layer to pharmacist	Must align with scope and patient context	Fidelity, stability, and pharmaceutical relevance [8, 9].	Persuasive but non-faithful explanation	Vendor tests; organization verifies access	Explanation quality is not model correctness
Professional interpretation	Integrate output with medication and patient context	Algorithm to pharmacist, then clinical team	Connects technical output to medication action	Comprehension and appropriate reliance [12].	Automation bias, rejection bias, or misplaced trust	Pharmacist interprets; organization supports training	Pharmacist review does not cure upstream defects
Challenge and override	Permit disagreement and temporary containment	Pharmacist to governance and technical owners	Activates review and correction	Accessibility, documentation, and response completeness	Override without review or challenge without closure	Pharmacy leadership protects and reviews challenge	Challenge volume alone is not a quality measure
Accountability allocation	Prevent responsibility from being displaced to the final user	Across vendor, organization, governance, and pharmacist	Defines ownership of investigation and correction	Clear responsibility for each failure class [10, 11].	Responsibility gaps or defensive transfer of blame	Governance body assigns and audits ownership	Allocation is proposed, not a legal determination
Correction and versioning	Resolve confirmed defects and preserve traceability	Governance and vendor back to users	Updates scope, explanations, and monitoring	Version integrity and documented disposition	Silent update, incomplete correction, or recurrence	Vendor corrects; organization revalidates	Technical correction is not demonstrated clinical benefit
Patient and equity interest	Examine differential effects and communication needs	Organization and pharmacist to patient and governance	Cross-cuts use, challenge, and monitoring	Subgroup evaluation and accessible communication	Unequal error, access, explanation, or recourse	Organization and governance retain oversight	Transparency alone cannot eliminate inequity

Required Explanations across the Medication-Decision Lifecycle

The contract requires explanation to change with the medication-decision stage. Before use, stakeholders need information that is understandable, actionable, and compatible with the intended workflow [14]. At data capture, the relevant questions concern provenance, completeness, timing, and whether medication history or clinical context is missing. At prescription generation, the explanation should clarify which order, patient characteristic, interaction, or deviation created the signal.

At algorithmic assessment, a numerical score should not be presented without information needed to interpret it. Predictive discrimination does not establish that estimated probabilities are reliable; calibration is therefore a material part of risk communication [15]. Local usefulness may also deteriorate when populations, prescribing behaviour, data systems, formularies, or clinical practices differ from development conditions [16]. The proposed contract consequently requires access to uncertainty, calibration, applicable population, relevant exclusions, and known conditions of instability.

At pharmacist verification, the explanation should support a pharmaceutical question rather than merely describe model computation. It should indicate what the system detected, what information materially influenced the output, what was not considered, and which alternative conditions might change the recommendation. The pharmacist should remain able to compare the output with treatment history, dose, route, laboratory results, interactions, adherence, patient goals, and feasible alternatives. The explanation must not imply that automation has displaced professional authority.

At medication action and communication, the record should distinguish the algorithmic output from the pharmacist's interpretation and the resulting clinical action. Where an output affects counselling, prescriber communication, escalation, or withholding of a medicine, material uncertainty

should be communicated at a level appropriate to the recipient. This requirement does not mean that technical model details must be presented indiscriminately to patients.

Finally, explanation obligations continue after use. Clinical machine learning can produce unintended consequences that are not captured by conventional performance measures [17]. Monitoring should therefore examine overrides, recurrent challenge categories, false reassurance, alert burden, subgroup patterns, workflow delay, version changes, and unresolved corrections. The proposed lifecycle package comprises intended use, provenance, patient-specific rationale, uncertainty, applicability, action boundaries, version identity, challenge status, and correction history. It is an organizing framework requiring prospective validation, not a universally established minimum dataset or proof that explanation improves medication safety.

Breach, Challenge, Correction, and Escalation

An explanation contract is meaningful only if a practising pharmacist can contest an output without being required to prove that the algorithm is defective. Experimental evidence indicates that health professionals may follow erroneous artificial-intelligence advice, demonstrating that access to decision support does not itself ensure appropriate reliance [18]. A challenge should therefore be triggered by a material inconsistency with patient information, pharmaceutical knowledge, expected model use, or the explanation supplied.

The proposed breach taxonomy includes omission, fidelity, context, uncertainty, temporal-validity, access, traceability, recourse, equity, and security breaches. Equity cannot be addressed solely through technical fairness adjustments because unequal data, institutions, access, and downstream decisions may remain uncorrected [19]. Clinical algorithms may also be vulnerable to corrupted or adversarial inputs, although the prevalence and significance of such failures must be established for individual pharmacy systems [20].

The proposed response pathway is: detection; pharmacist challenge and documentation; immediate containment or independent verification; assessment of severity, recurrence, scope, and reversibility; clinical and technical review; correction and versioning; local revalidation; and documented closure. Responsible clinical machine learning requires continuing monitoring and maintenance after deployment [21]. Accordingly, a breach pathway must address erroneous advice, manipulable inputs, and post-deployment safety monitoring rather than treating explanation delivery as closure [18, 20, 21].

Containment may involve disregarding an output, obtaining an independent medication review, restricting a function, or temporarily suspending use. These responses remain conditional on urgency, potential harm, available alternatives, organizational authority, and applicable professional requirements. The pathway is proposed and does not establish that a particular escalation route improves clinical outcomes.

Evaluation and Governance of Contract Performance

Contract performance should not be reduced to predictive accuracy or user satisfaction. Evaluation must examine transparency, reproducibility, ethics, clinical relevance, and effectiveness [22]. For the proposed contract, this requires separate assessment of the algorithm, the explanation, pharmacist interaction, organizational response, and downstream medication process.

Minimum model information should include intended use, data provenance, development context, validation methods, performance limitations, and reproducibility-relevant details [23]. Early clinical evaluation should additionally examine human–AI interaction, workflow effects, error patterns, safety concerns, and differences between intended and actual use [24]. Prediction-model reporting should make development, validation, performance, and applicability sufficiently transparent for independent appraisal [25]. Compliance with reporting guidance, however, is not evidence of local validity or readiness.

Proposed contract-performance domains are explanation fidelity, pharmaceutical relevance, uncertainty comprehension, calibration, appropriate reliance, challenge accessibility, correction completeness, version traceability, workload, subgroup performance, and recurrence of recognized breaches. Evaluation should combine technical testing, pharmacist simulation, usability research, prospective observation, qualitative inquiry, and organizational process analysis. Thresholds should be established empirically for defined tasks and settings rather than inferred from the conceptual architecture.

Governance should review both failures and silent non-use. Low challenge frequency may reflect reliable operation, but it may also reflect inaccessible reporting, excessive workload, weak professional protection, or normalization of algorithmic authority. Contract evaluation must therefore interpret metrics within their organizational context.

Implications for Vendors, Organizations, and Pharmacists

Implementation requires more than adding explanatory text to an interface. Health technologies may be rejected, abandoned, or fail to scale when interacting technological, organizational, professional, and contextual complexities are neglected [26]. Implementation also depends on intervention characteristics, the inner and outer setting, participating individuals, and the implementation process [27]. Real-world clinical-AI integration further requires operational ownership, workflow redesign, stakeholder participation, and continuing adaptation [28].

For vendors, the contract implies responsibility for intended-use documentation, explanation-method disclosure, version control, known-failure communication, and support for

correction. Organizations require authority to validate locally, monitor use, protect professional challenge, restrict functions, and demand vendor response. Pharmacy leaders must connect governance to medication-safety processes and ensure that pharmacists can access explanations during actual work.

Pharmacist-facing evidence indicates that perceived usefulness can depend on how medication-order outputs are presented and whether the decision unit corresponds to professional practice [29]. Pharmacists should interpret rather than merely accept outputs, document consequential disagreement, and escalate recurrent or potentially systemic failures. Sustainable implementation therefore depends on technology complexity, organizational context, and integration into clinical work [26-28].

Figure 2 presents the original conceptual synthesis for breach, challenge, correction, and escalation pathway, showing how the article's principal components, evidence relationships, uncertainties, and decision boundaries are connected.

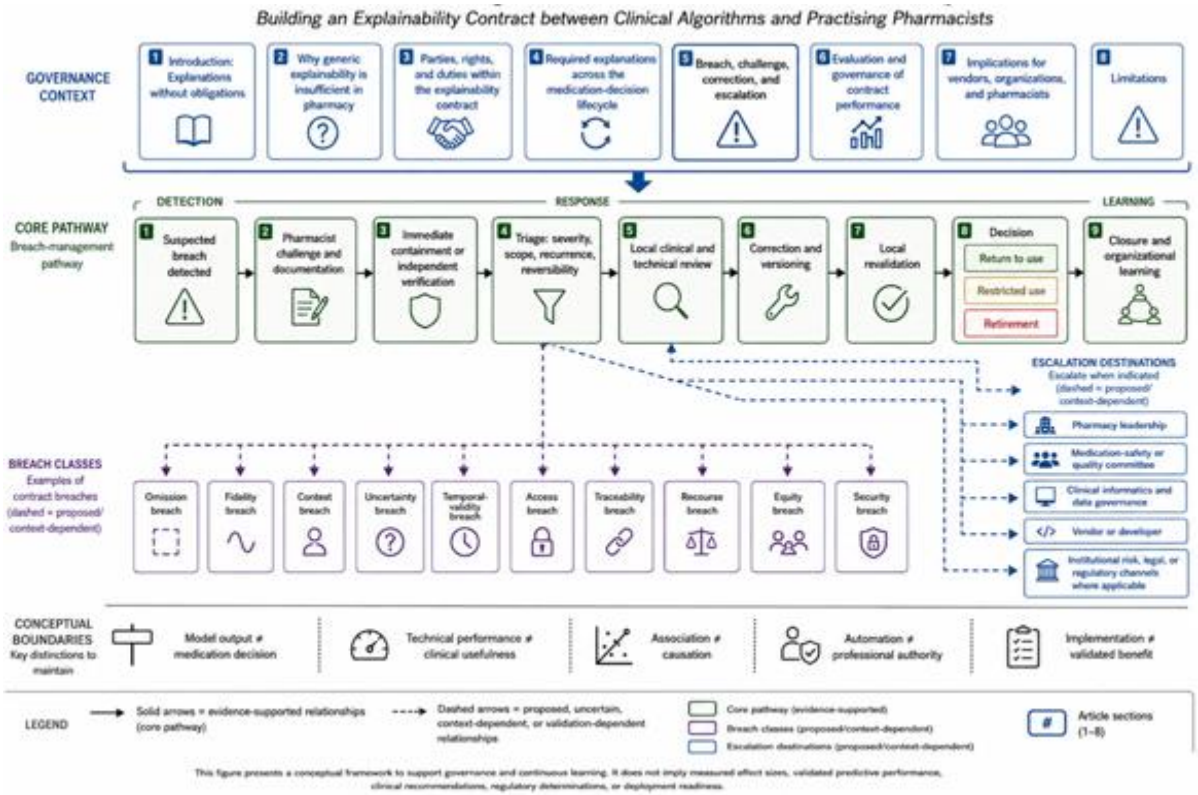

Figure 2. Breach, Challenge, Correction, and Escalation Pathway

The figure is an original conceptual synthesis developed for "Building an Explainability Contract between Clinical Algorithms and Practising Pharmacists." Arrows and grouping indicate proposed or evidence-supported relationships rather than measured effect sizes. The figure does not represent a validated predictive model, clinical recommendation,

Table 3 organizes the evidence, constructs, relationships, uncertainties, and boundary conditions required for governance responsibilities, implementation conditions, and monitoring requirements for the article Building an Explainability Contract between Clinical Algorithms and Practising Pharmacists.

Table 3. Governance responsibilities, implementation conditions, and monitoring requirements for the article Building an Explainability Contract between Clinical Algorithms and Practising Pharmacists.

Governance domain	Responsible actor	Required authority	Documentation requirement	Monitoring indicator	Escalation trigger	Corrective action	Accountability boundary
Intended use	Vendor and deploying organization	Define and restrict permitted tasks	Model purpose, exclusions, version, validation context	Use outside defined task	Recurrent off-label or context-incompatible use	Restrict function, revise documentation, retrain users	Intended-use documentation does not establish local benefit

Local implementation	Organization and pharmacy leadership	Approve workflow and allocate resources	Local validation, workflow map, assigned owners	Non-use, workarounds, delay, alert burden	Material divergence from intended workflow	Redesign, retrain, restrict, or withdraw	Adoption is not evidence of safety or sustainability
Professional challenge	Practising pharmacist and pharmacy leadership	Override and request review	Reason for challenge and immediate action	Challenge accessibility and unresolved cases	Potential harm, recurrence, or inability to obtain review	Independent verification and formal escalation	Pharmacists are not solely responsible for upstream defects
Technical investigation	Vendor, informatics, and data stewards	Inspect data, model, explanation, and versions	Investigation record and affected versions	Drift, missingness, instability, security signal	Reproducible or systemic technical concern	Correct, version, test, and communicate	Technical correction does not establish clinical effectiveness
Medication-safety oversight	Medication-safety or quality committee	Review cross-case risk and containment	Severity, scope, recurrence, and disposition	Recurrent false reassurance or harmful workflow consequence	Multi-patient or high-severity concern	Restrict use, initiate broader review, require revalidation	Governance review does not replace clinical judgment
Equity governance	Organization and multidisciplinary governance body	Examine differential performance and recourse	Subgroup methods, limitations, and response	Unequal error, challenge, override, or access patterns	Material disparity or race-based assumption	Investigate data, model, workflow, and institutional causes	Transparency alone cannot remove structural inequity
Closure and learning	Organization with vendor participation	Determine return, restriction, or retirement	Correction, revalidation, communication, final disposition	Recurrence of a previously recognized breach	Incomplete correction or repeated failure	Reopen investigation or retire the function	Closure requires evidence; it is not administrative completion alone

Limitations

The proposed contract cannot correct inequity merely by making algorithmic reasoning more visible. Proxy targets can encode unequal access or expenditure and produce discriminatory allocation even when protected attributes are not explicitly used [30]. Aggregate model performance may also conceal higher error among underserved groups [31]. Language models may reproduce race-based clinical assumptions in generated health information, although this evidence should not be generalized to every clinical algorithm [32]. Fairness assessment must therefore remain connected to health-equity goals, clinical context, and downstream consequences [33].

The contract is conceptual, jurisdiction-dependent, and unvalidated. It may increase documentation or workload, encourage defensive practice, or produce superficial compliance. Explanation may remain technically unfaithful, and challenge mechanisms may be inaccessible in under-resourced settings. The framework does not determine legal liability, define universal explanation content, replace professional judgment, or establish clinical or regulatory readiness.

Conclusion

Clinical algorithms enter pharmacy practice through concrete medication decisions, but their explanations are often supplied without reciprocal obligations. The proposed explainability contract reframes explanation as a relationship among vendors, organizations, pharmacists, governance bodies, and affected patients. It connects intended use, provenance, rationale, uncertainty, applicability, professional challenge, correction, and traceability across the medication-decision lifecycle.

Its central professional boundary is that algorithmic output does not independently authorize medication action. Its central organizational boundary is that responsibility cannot be displaced onto the practising pharmacist when relevant technical or governance information remains inaccessible. Its central scientific boundary is that explanation availability does not demonstrate fidelity, usefulness, safety, equity, or improved outcomes.

Future validation must examine technical performance, pharmacist comprehension, appropriate reliance, workflow consequences, medication-safety signals, correction processes, organizational feasibility, and differential effects across populations and settings. The contract should therefore be treated as a testable governance proposal rather than a clinical standard. Its value lies in making explanation duties, professional rights, failure pathways, and unresolved accountability visible enough to be examined, challenged, and revised.

ACKNOWLEDGMENTS: None
CONFLICT OF INTEREST: None
FINANCIAL SUPPORT: None
ETHICS STATEMENT: None

References

1. Ghassemi M, Oakden-Rayner L, Beam AL. The false hope of current approaches to explainable artificial intelligence in health care. Lancet Digit Health. 2021;3(11). doi:10.1016/S2589-7500(21)00208-9

2. Amann J, Blasimme A, Vayena E, Frey D, Madai VI; Precise4Q Consortium. Explainability for artificial intelligence in healthcare: A multidisciplinary perspective. BMC Med Inform Decis Mak. 2020;20(1):310. doi:10.1186/s12911-020-01332-6
3. Sutton RT, Pincock D, Baumgart DC, Sadowski DC, Fedorak RN, Kroeker KI. An overview of clinical decision support systems: Benefits, risks, and strategies for success. NPJ Digit Med. 2020;3:17. doi:10.1038/s41746-020-0221-y
4. Graafsma J, Murphy RM, van de Garde EMW, Karapinar-Çarkit F, Derijks HJ, Hoge RHL, et al. The use of artificial intelligence to optimize medication alerts generated by clinical decision support systems: A scoping review. J Am Med Inform Assoc. 2024;31(6):1411-22. doi:10.1093/jamia/ocae076
5. Corny J, Rajkumar A, Martin O, Dode X, Lajonchère JP, Billuart O, et al. A machine learning-based clinical decision support system to identify prescriptions with a high risk of medication error. J Am Med Inform Assoc. 2020;27(11):1695-704. doi:10.1093/jamia/ocaa154
6. Miller T. Explanation in artificial intelligence: Insights from the social sciences. Artif Intell. 2019;267:1-38. doi:10.1016/j.artint.2018.07.007
7. Markus AF, Kors JA, Rijnbeek PR. The role of explainability in creating trustworthy artificial intelligence for health care: A comprehensive survey of the terminology, design choices, and evaluation strategies. J Biomed Inform. 2021;113:103655. doi:10.1016/j.jbi.2020.103655
8. Rudin C. Stop explaining black box machine learning models for high stakes decisions and use interpretable models instead. Nat Mach Intell. 2019;1(5):206-15. doi:10.1038/s42256-019-0048-x
9. Babic B, Gerke S, Evgeniou T, Cohen IG. Beware explanations from AI in health care. Science. 2021;373(6552):284-6. doi:10.1126/science.abg1834
10. Char DS, Shah NH, Magnus D. Implementing machine learning in health care—addressing ethical challenges. N Engl J Med. 2018;378(11):981-3. doi:10.1056/NEJMp1714229
11. 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(3):491-7. doi:10.1093/jamia/ocz192
12. Asan O, Bayrak AE, Choudhury A. Artificial intelligence and human trust in healthcare: Focus on clinicians. J Med Internet Res. 2020;22(6). doi:10.2196/15154
13. Price WN 2nd, Gerke S, Cohen IG. Potential liability for physicians using artificial intelligence. JAMA. 2019;322(18):1765-6. doi:10.1001/jama.2019.15064
14. Benda NC, Das LT, Abramson EL, Blackburn K, Thoman A, Kaushal R, et al. "How did you get to this number?" Stakeholder needs for implementing predictive analytics: A pre-implementation qualitative study. J Am Med Inform Assoc. 2020;27(5):709-16. doi:10.1093/jamia/ocaa021
15. Van Calster B, McLernon DJ, van Smeden M, Wynants L, Steyerberg EW. Calibration: The Achilles heel of predictive analytics. BMC Med. 2019;17(1):230. doi:10.1186/s12916-019-1466-7
16. Subbaswamy A, Saria S. From development to deployment: Dataset shift, causality, and shift-stable models in health AI. Biostatistics. 2020;21(2):345-52. doi:10.1093/biostatistics/kxz041
17. Cabitza F, Rasoini R, Gensini GF. Unintended consequences of machine learning in medicine. JAMA. 2017;318(6):517-18. doi:10.1001/jama.2017.7797
18. Gaube S, Suresh H, Raue M, Merritt A, Berkowitz SJ, Lermer E, et al. Do as AI say: Susceptibility in deployment of clinical decision-aids. NPJ Digit Med. 2021;4(1):31. doi:10.1038/s41746-021-00385-9
19. McCradden MD, Joshi S, Mazwi M, Anderson JA. Ethical limitations of algorithmic fairness solutions in health care machine learning. Lancet Digit Health. 2020;2(5). doi:10.1016/S2589-7500(20)30065-0
20. Finlayson SG, Bowers JD, Ito J, Zittrain JL, Beam AL, Kohane IS. Adversarial attacks on medical machine learning. Science. 2019;363(6433):1287-9. doi:10.1126/science.aaw4399
21. Wiens J, Saria S, Sendak M, Ghassemi M, Liu VX, Doshi-Velez F, et al. Do no harm: A roadmap for responsible machine learning for health care. Nat Med. 2019;25(9):1337-40. doi:10.1038/s41591-019-0548-6
22. Vollmer S, Mateen BA, Bohner G, Király FJ, Ghani R, Jonsson P, et al. Machine learning and artificial intelligence research for patient benefit: 20 critical questions on transparency, replicability, ethics, and effectiveness. BMJ. 2020;368. doi:10.1136/bmj.l6927
23. Norgeot B, Quer G, Beaulieu-Jones BK, Torkamani A, Dias R, Gianfrancesco M, et al. Minimum information about clinical artificial intelligence modeling: The MI-CLAIM checklist. Nat Med. 2020;26(9):1320-4. doi:10.1038/s41591-020-1041-y
24. Vasey B, Nagendran M, Campbell B, Clifton DA, Collins GS, Denaxas S, et al. Reporting guideline for the early-stage clinical evaluation of decision support systems driven by artificial intelligence: DECIDE-AI. Nat Med. 2022;28(5):924-33. doi:10.1038/s41591-022-01772-9
25. Collins GS, Moons KGM, Dhiman P, Riley RD, Beam AL, Van Calster B, et al. TRIPOD+AI statement: Updated guidance for reporting clinical prediction models that use regression or machine learning methods. BMJ. 2024;385. doi:10.1136/bmj-2023-078378
26. Greenhalgh T, Wherton J, Papoutsi C, Lynch J, Hughes G, A'Court C, et al. 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(11). doi:10.2196/jmir.8775
27. Damschroder LJ, Reardon CM, Widerquist MAO, Lowery J. The updated Consolidated Framework for Implementation Research based on user feedback. Implement Sci. 2022;17(1):75. doi:10.1186/s13012-022-01245-0
28. Sendak MP, Ratliff W, Sarro D, Alderton E, Futoma J, Gao M, et al. Real-world integration of a sepsis deep learning technology into routine clinical care: Implementation study. JMIR Med Inform. 2020;8(7). doi:10.2196/15182
29. Hogue SC, Chen F, Brassard G, Lebel D, Bussières JF, Durand A, et al. Pharmacists' perceptions of a machine learning model for the identification of atypical medication orders. J Am Med Inform Assoc. 2021;28(8):1712-18. doi:10.1093/jamia/ocab071
30. Obermeyer Z, Powers B, Vogeli C, Mullainathan S. Dissecting racial bias in an algorithm used to manage the health of populations. Science. 2019;366(6464):447-53. doi:10.1126/science.aax2342
31. Seyyed-Kalantari L, Zhang H, McDermott MBA, Chen IY, Ghassemi M. Underdiagnosis bias of artificial intelligence algorithms applied to chest radiographs in under-served patient populations. Nat Med. 2021;27(12):2176-82. doi:10.1038/s41591-021-01595-0
32. Omiye JA, Lester JC, Spichak S, Rotemberg V, Daneshjou R. Large language models propagate race-based medicine. NPJ Digit Med. 2023;6(1):195. doi:10.1038/s41746-023-00939-z
33. Rajkomar A, Hardt M, Howell MD, Corrado G, Chin MH. Ensuring fairness in machine learning to advance health equity. Ann Intern Med. 2018;169(12):866-72. doi:10.7326/M18-1990