Original Article

The Pharmacy Data Commons Needed for Safe and Equitable Algorithm Development

Kristian Eriksen[1]*, Solveig Berg[1], Magnus Dahl[2]

[1]Department of Pharmacy Data Governance and Sharing, Faculty of Pharmacy, University of Oslo, Oslo, Norway. [2]Department of Equitable Algorithm Development, Faculty of Pharmacy, University of Bergen, Bergen, Norway.

Abstract

Pharmacy algorithms are increasingly developed from data dispersed across dispensing systems, electronic health records, claims, laboratories, pharmacovigilance repositories, and patient-facing technologies. Fragmentation is not only a technical inconvenience: inconsistent semantics, incomplete provenance, selective representation, unclear permissions, and weak accountability can propagate into algorithm development and medication-related decisions. This article proposes a pharmacy data commons as a governed, distributed capability for responsible data access, documentation, harmonization, analysis, evaluation, and benefit sharing. The commons is not defined as a single pooled database. Its proposed architecture links participating data assets, semantic and provenance services, computable use permissions, representation profiles, privacy-preserving analysis environments, institutional decision rights, community participation, and explicit boundaries between model output and professional medication authority. Access is treated as a conditional lifecycle rather than a one-time approval, with stewardship, representation review, safety and equity evaluation, output review, monitoring, incident learning, and retirement. Commons performance should therefore be assessed across data fitness, interoperability, provenance, access consistency, privacy, representation, human–AI workflow, medication-safety risk, community legitimacy, distributive benefit, and sustainability. Validation would require multisite technical audits, external and subgroup evaluation, human-factors studies, safety assessment, governance-process evaluation, privacy testing, and community-defined measures of legitimacy and benefit. The proposal remains non-empirical and context-dependent. It does not establish clinical benefit, regulatory acceptability, fairness, lawful use, or deployment readiness. Its original contribution is an integrated governance architecture that makes the conditions for safe and equitable pharmacy-algorithm development visible, testable, and contestable.

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

Introduction

Artificial-intelligence applications in pharmacy span medication review, dispensing support, risk prediction, information retrieval, and clinical-pharmacy services, but their tasks, settings, source data, development stages, and evaluation practices remain heterogeneous [1, 2]. This variability matters because pharmacy algorithms are rarely built from a self-contained record. Medication decisions may depend on prescriptions, dispensing histories, laboratory findings, diagnoses, adherence signals, adverse-event reports, socioeconomic context, and professional interpretation distributed across organizations.

Fragmentation includes technical incompatibility, but also divergent coding, local workflow conventions, missing contextual fields, and uncertain meaning. Digital medicine depends on interoperability that preserves both exchange and interpretation across systems [3]. A merged dataset can therefore remain clinically incoherent when identical fields arise from different processes or when apparently missing

Address for correspondence: Kristian Eriksen, Department of Pharmacy Data Governance and Sharing, Faculty of Pharmacy, University of Oslo, Oslo, Norway.
kristian.eriksen@farmasi.uio.no
Received: 12 May 2025; **Accepted:** 27 August 2025

How to cite this article: Eriksen K, Berg S, Dahl M. The Pharmacy Data Commons Needed for Safe and Equitable Algorithm Development. Arch Pharm Pract. 2025;16(3):54-62. https://doi.org/10.51847/Zuc4NBkGdt

information reflects access, documentation, or care pathways rather than absence of disease or treatment.

Routine-care data can also encode selection, measurement, documentation, and missingness biases [4]. In pharmacy, these mechanisms may distort exposure histories, adherence proxies, treatment opportunities, or the apparent frequency of medication-related harm. Alternative explanations must remain visible: observed differences may reflect disease burden, service availability, recording practices, or institutional policy rather than algorithmically relevant patient differences.

Clinical decision support is consequently a sociotechnical intervention. Its safety and usefulness depend on data quality, workflow integration, human factors, governance, and the capacity of professionals to interpret and act on outputs, not technical performance alone [5]. The problem is therefore not solved by adding a larger dataset or a more accurate model. A shared architecture is needed to govern how pharmacy data are described, accessed, transformed, represented, evaluated, returned to contributors, monitored, and separated from medication-decision authority.

Defining a Pharmacy Data Commons

A pharmacy data commons is proposed here as a governed, distributed capability through which participating organizations and communities make appropriately documented data available for approved algorithm-development and evaluation purposes. FAIR-oriented stewardship provides a useful foundation, but responsible reuse also requires contextual accountability, metadata practices, organizational capacity, and enforceable conditions [6, 7]. The commons is therefore neither an unrestricted open repository nor a synonym for data centralization.

Its defining unit is a participating data asset: a versioned dataset linked to provenance, semantic mappings, quality information, permitted uses, representation limitations, and accountable stewards. Computable data-use terms can support consistent matching between requests and restrictions, but they cannot replace accountable authorization [8]. **Figure 1** presents the original conceptual synthesis for technical and institutional architecture of a pharmacy data commons, showing how the article's principal components, evidence relationships, uncertainties, and decision boundaries are connected.

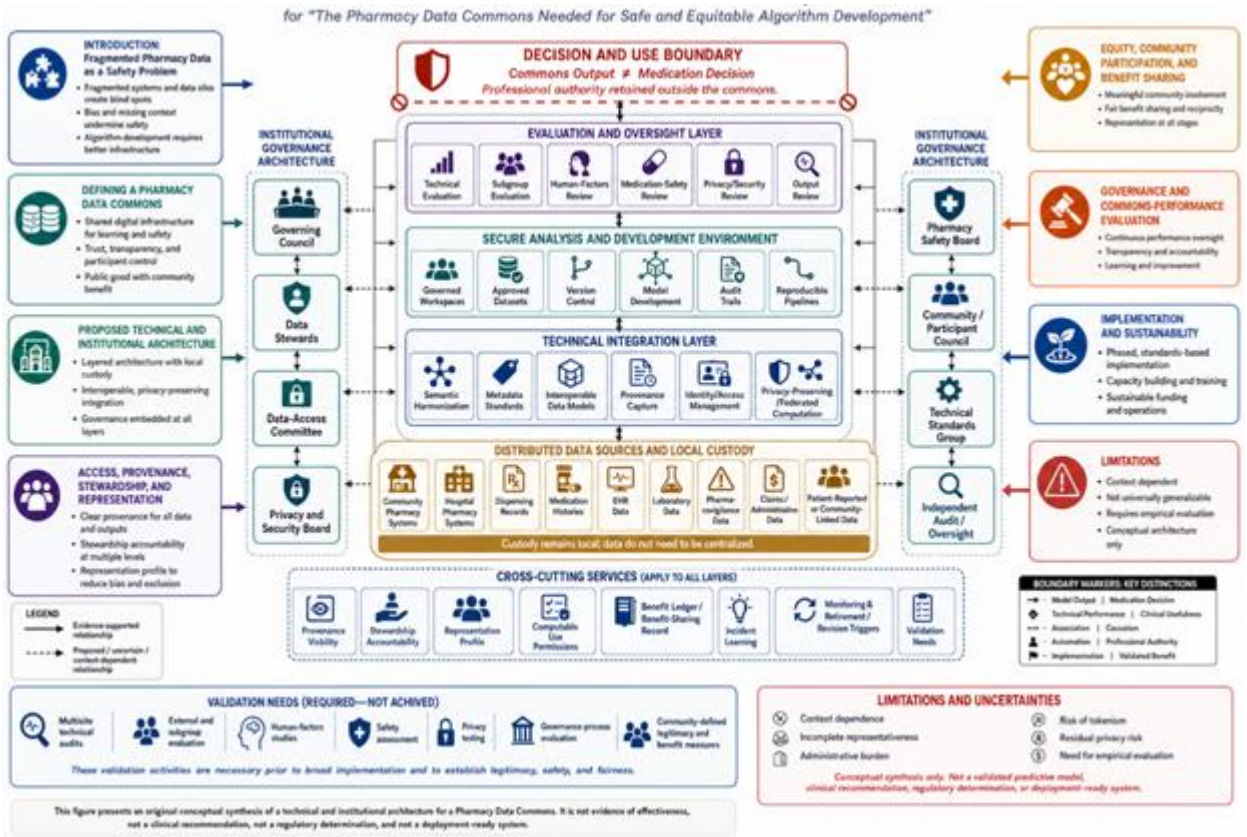

Figure 1. Technical and Institutional Architecture of a Pharmacy Data Commons

The figure is an original conceptual synthesis developed for "The Pharmacy Data Commons Needed for Safe and Equitable Algorithm Development." 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.

Commons membership also requires documentation capable of explaining why a dataset exists, how it was collected and processed, what it contains, which uses are intended or discouraged, and where uncertainty remains [9]. Such documentation does not prove fitness for use. It provides an inspectable basis for deciding whether data are sufficiently relevant, complete, interpretable, and representative for a specified task.

The commons may retain data under local custody and support federated analysis. Federation can reduce raw-data movement, but it does not eliminate heterogeneity, privacy risk, security obligations, governance, or external-validation requirements [10]. The proposed definition therefore combines shared rules and evidence artifacts with distributed institutional control. It separates access to data from permission to deploy an algorithm and separates algorithmic evidence from professional authority over medication decisions.

Table 1 organizes the evidence, constructs, relationships, uncertainties, and boundary conditions required for definitions, components, relationships, and intended uses for the article The Pharmacy Data Commons Needed for Safe and Equitable Algorithm Development.

Table 1. Definitions, components, relationships, and intended uses for the article The Pharmacy Data Commons Needed for Safe and Equitable Algorithm Development.

Component or scientific dimension	Problem addressed	Inputs or determinants	Proposed mechanism or relationship	Expected contribution	Evidence required	Failure or uncertainty risk	Boundary statement
Pharmacy data commons	Isolated, incompatible data	Distributed pharmacy and health data	Proposed shared rules plus local custody	Reusable, governed evidence infrastructure	Interoperability and governance evaluation [3].	Centralization, exclusion, or mission drift	Not an open database or clinical authority
Participating data asset	Data without usable context	Version, metadata, quality, and permissions	Proposed admission package	Inspectable fitness and limitations	Documentation-completeness assessment [9].	Outdated or misleading metadata	Admission does not prove suitability
Participant organization	Unclear rights and duties	Custodians, pharmacies, communities, and researchers	Proposed membership charter	Defined participation and exit	Institutional and legal review	Power imbalance or institutional capture	Membership does not transfer ownership
Data steward	Diffuse responsibility	Source knowledge, metadata, and remediation capacity	Proposed accountable stewardship	Maintained quality and provenance	Role-performance evaluation [6].	Under-resourcing or conflicts	Stewardship is not sole authorization
Common semantic layer	Non-comparable meanings	Standards, terminologies, and local mappings	Shared models with documented exceptions	Reproducible interpretation	Mapping and transformation tests [3].	False semantic equivalence	Standardization does not erase context
Provenance package	Untraceable transformations	Origin, linkage, exclusions, and versions	Traceable lineage attached to outputs	Reproducibility and contestability	Provenance audit [9].	Broken or incomplete lineage	Traceability does not establish validity
Computable access contract	Inconsistent permission review	Use terms, purpose, requester, and conditions	Automated matching plus human decision	More consistent access review	Agreement, error, and appeal testing [8].	Automation of normative judgment	Professional authority remains human
Representation profile	Hidden population limitations	Data generation, inclusions, and missingness	Proposed subgroup and context disclosure	Earlier identification of exclusions	External and subgroup assessment	Reporting without correction	Description does not guarantee fairness
Secure analysis workspace	Uncontrolled copies and code	Approved data, code, users, and outputs	Proposed versioned controlled environment	Reproducible and auditable analysis	Security and reproducibility testing [10].	Leakage, lock-in, or false assurance	Security controls do not prove safety
Governance bodies	Fragmented accountability	Institutional and community mandates	Proposed distributed decision rights	Review, escalation, audit, and appeal	Governance-process evaluation	Duplication, delay, or tokenism	No body can certify benefit alone
Clinical decision boundary	Model output mistaken for action	Evidence artifact, workflow, and professional judgment	Explicit handoff outside the commons	Preserved medication authority	Workflow and safety validation	Automation bias or responsibility gaps	Commons release is not deployment approval

Proposed Technical and Institutional Architecture

The proposed architecture separates technical services from institutional authority while requiring them to interact. Privacy-preserving and federated computation may reduce direct exposure of raw records, but they introduce trade-offs involving utility, attack surfaces, coordination, and implementation complexity [11]. They are therefore controls within the commons, not substitutes for stewardship or accountability.

A distributed node holds source data under local responsibility and executes approved transformations or analyses. Multisite model development without pooling patient-level records is technically feasible under controlled conditions [12]. In the proposed commons, however, each node must also expose versioned metadata, transformation logic, permission conditions, quality findings, and representation limitations to the shared governance layer.

A canonical interoperability layer can support repeatable creation of analysis-ready data marts from standardized records [13]. The architecture pairs this layer with local semantic-exception registers because the same medication, encounter, or laboratory concept may be recorded differently across organizations. Harmonization should remain reversible: an analyst must be able to trace each derived variable back to its source and transformation.

The approved algorithm workspace contains versioned data views, code, model configurations, outputs, and evaluation records. Federated infrastructure can coordinate learning across heterogeneous institutions, but successful coordination does not remove the need for task-specific external validation [14]. Proposed transition gates therefore separate data access, technical development, clinical-task evaluation, workflow evaluation, safety and equity review, and any later deployment process.

Decentralized coordination can reduce dependence on a central data custodian [15]. It does not decentralize accountability automatically. The institutional spine is therefore proposed to include a governing council, data-access committee, data-stewardship office, technical standards group, privacy and security board, pharmacy safety board, community or participant council, and independent audit function. Their mandates should be non-substitutable: technical teams cannot approve ethical acceptability, community participation cannot certify model performance, and a governing council cannot infer medication-safety benefit from administrative compliance.

Access, Provenance, Stewardship, and Representation

Access begins with a defined purpose, eligible requester, specified data elements, analytic plan, output controls, duration, and benefit pathway. Computable oversight can improve the consistency and scalability of matching requests to explicit use restrictions [16]. The proposed mechanism combines automated screening with accountable authorization, documented reasons, appeal, revocation, and an audit trail; automation supports review but does not decide whether a contested use is legitimate.

Provenance and data fitness should be evaluated together. Electronic health-record quality assessment is multidimensional, including completeness, correctness, concordance, plausibility, and currency [17]. The commons should not collapse these dimensions into an unqualified score. Instead, stewards should state which defects matter for the proposed task, how transformations changed the data, what uncertainty remains, and whether remediation, restriction, or rejection is warranted.

Representation is likewise more than a demographic table. Health-dataset reporting should describe populations, data-generating processes, inclusions, exclusions, missingness, and known limitations relevant to bias [18]. A proposed representation profile would connect these disclosures to the intended medication-related decision, target definition, subgroup risks, and expected users. The profile can reveal risk but cannot establish representativeness, fairness, or equitable benefit.

Finally, access and stewardship must operate across institutional and jurisdictional boundaries. International data-sharing experience shows that technical standards must align with policy, identity, access, and governance mechanisms [19]. The pharmacy commons adapts that principle but does not assume a single legal solution. Each participating organization retains duties under applicable law and professional regulation, while shared rules define minimum documentation, review, monitoring, and escalation. The resulting architecture is a proposed basis for validation, not evidence that safe or equitable algorithm development has already been achieved.

Table 2 organizes the evidence, constructs, relationships, uncertainties, and boundary conditions required for evaluation criteria, evidence requirements, and failure modes for the article The Pharmacy Data Commons Needed for Safe and Equitable Algorithm Development.

Table 2. Evaluation criteria, evidence requirements, and failure modes for the article The Pharmacy Data Commons Needed for Safe and Equitable Algorithm Development.

Architecture layer or process stage	Core function	Information flow	Interaction with other components	Evaluation criterion	Potential failure mode	Human or organizational responsibility	Readiness boundary

Data onboarding	Admit a documented asset	Source to catalog	Stewardship, provenance, and representation	Proposed completeness and currency check	Incomplete context	Data steward	Onboarding is not task approval
Semantic harmonization	Preserve comparable meaning	Local fields to common model	Data marts and analysis	Mapping accuracy and exception disclosure [3].	Silent semantic loss	Technical standards group	Interoperability is not equivalence
Provenance capture	Record lineage	Sources and transformations to audit trail	Quality, access, and outputs	Proposed traceability and version integrity [9].	Broken lineage	Steward and custodian	Provenance is not causal validity
Access review	Match purpose and restrictions	Request to permission decision	Governance and workspace	Decision consistency, auditability, and appeal [8].	Inappropriate approval or denial	Data-access committee	Permission is not clinical authorization
Data-fitness review	Test purpose-specific usability	Quality findings to project decision	Stewardship and redesign	Proposed completeness, plausibility, concordance, and currency	Composite score conceals defects	Data-quality team	No universal quality threshold
Representation review	Expose population limitations	Profile to equity gate	Target selection and evaluation	Proposed coverage and subgroup relevance	Exclusion or proxy mismatch	Equity panel and community council	Representation is not guaranteed fairness
Privacy and security	Limit exposure and misuse	Approved data to controlled workspace	Access, monitoring, and output review	Threat testing and residual-risk review [10].	Leakage or re-identification	Security and privacy board	Privacy technology is not lawful-use proof
Technical evaluation	Assess task performance	Model results to evaluation record	External and subgroup testing	Proposed discrimination, calibration, and uncertainty assessment	Overfitting or site dependence	Technical evaluation team	Performance alone cannot pass
Human–AI evaluation	Assess workflow use	Output to professional interpretation	Safety and implementation review	Proposed role clarity, actionability, and workload assessment	Automation bias or non-use	Pharmacy workflow team	Usability is not clinical benefit
Medication-safety review	Identify harm pathways	Errors, overrides, and near misses to board	Human–AI and governance gates	Proposed hazard and mitigation evidence	New or displaced error	Pharmacy safety board	No prescriptive use without validation
Community legitimacy	Assess participation quality	Deliberation to governance record	Access, representation, and benefit sharing	Proposed influence, burden, and transparency assessment	Tokenism or exclusion	Community council	Consultation is not shared authority
Benefit realization	Return value to contributors	Findings and resources to communities	Monitoring and sustainability	Proposed documented benefit pathway	Extractive participation	Governing and community councils	Promised benefit is not realized benefit
Monitoring and retirement	Detect change and stop unsuitable use	Drift and incidents to escalation	All architecture layers	Proposed triggers, response time, and retirement capacity	Persistence despite deterioration	Governing council	Continued operation is not readiness

Equity, Community Participation, and Benefit Sharing

Equity cannot be reduced to balanced sample counts or a single fairness statistic. A clinically convenient proxy may represent need differently across groups, while fairness criteria depend on the decision, target, error distribution, and health-equity objective [20, 21]. The proposed commons therefore requires an early equity gate examining who is represented, how targets and proxies were constructed, which errors matter, and who may receive or be denied benefit.

Legitimacy also depends on the conditions under which people permit health-data reuse. Public support is conditional on privacy protection, transparency, trust, credible governance, and recognizable public benefit [22]. A proposed community or participant council should consequently influence problem selection, access conditions, representation review, dissemination, and monitoring rather than provide retrospective consultation. Community engagement throughout the artificial-intelligence lifecycle can help align technical work with locally defined needs and social value [23]. Benefit sharing may include accessible findings, improved data quality, workforce capacity, services, or resources, but promised benefits must be documented and evaluated. Participation alone does not demonstrate equitable benefit distribution.

The proposed benefit-sharing mechanism contains three linked obligations. First, applicants should specify prospectively who is expected to benefit and which burdens may fall on contributing populations or organizations. Second, the governing and community councils should record whether proposed benefits are relevant, proportionate, and realistically deliverable. Third, realized benefits, unmet commitments, and unintended burdens should enter the monitoring record. These mechanisms are proposed rather than validated; their representativeness, influence, administrative burden, and distributive consequences require empirical assessment.

Governance and Commons-Performance Evaluation

Commons governance should treat privacy, bias, transparency, accountability, trust, and responsibility as interacting concerns [24]. Principles become operational only when institutions possess defined authority, procedures, incentives, escalation routes, and enforcement capacity [25]. The proposed governance structure therefore allocates non-substitutable responsibilities among data stewards, access reviewers, technical teams, pharmacy-safety reviewers, community representatives, and independent auditors. Technical teams cannot authorize contested data uses, community participation cannot certify model validity, and administrative compliance cannot establish medication-safety benefit.

Commons performance should be evaluated separately from model performance. Early clinical-artificial-intelligence evaluation guidance supports assessment of human–AI interaction, workflow integration, errors, safety events, and contextual change [26]. A responsible lifecycle also requires stakeholder involvement, external validation, monitoring, incident response, and retirement [27]. Proposed evaluation gates should therefore examine data fitness, provenance, access consistency, privacy, representation, technical validity, workflow actionability, medication-safety hazards, equity, legitimacy, and sustainability.

No single metric should determine passage through these gates. High technical performance cannot compensate for non-representative data, weak provenance, unmanageable workflow consequences, or an unresolved safety hazard. Conversely, complete governance documentation cannot compensate for poor external validity. Gate decisions should record the evidence examined, residual uncertainty, dissent, required corrective action, and conditions for reconsideration. Passing a commons evaluation gate permits only the next specified stage; it does not authorize clinical deployment.

Implementation and Sustainability

Implementation is shaped by interactions among technology, adopters, organizations, value propositions, wider systems, and adaptation over time [28]. The proposed commons should assess these conditions across outer setting, inner setting, individuals, intervention characteristics, and implementation processes [29]. **Figure 2** presents the original conceptual synthesis for data access, provenance, stewardship, representation, and benefit-sharing cycle, showing how the article's principal components, evidence relationships, uncertainties, and decision boundaries are connected.

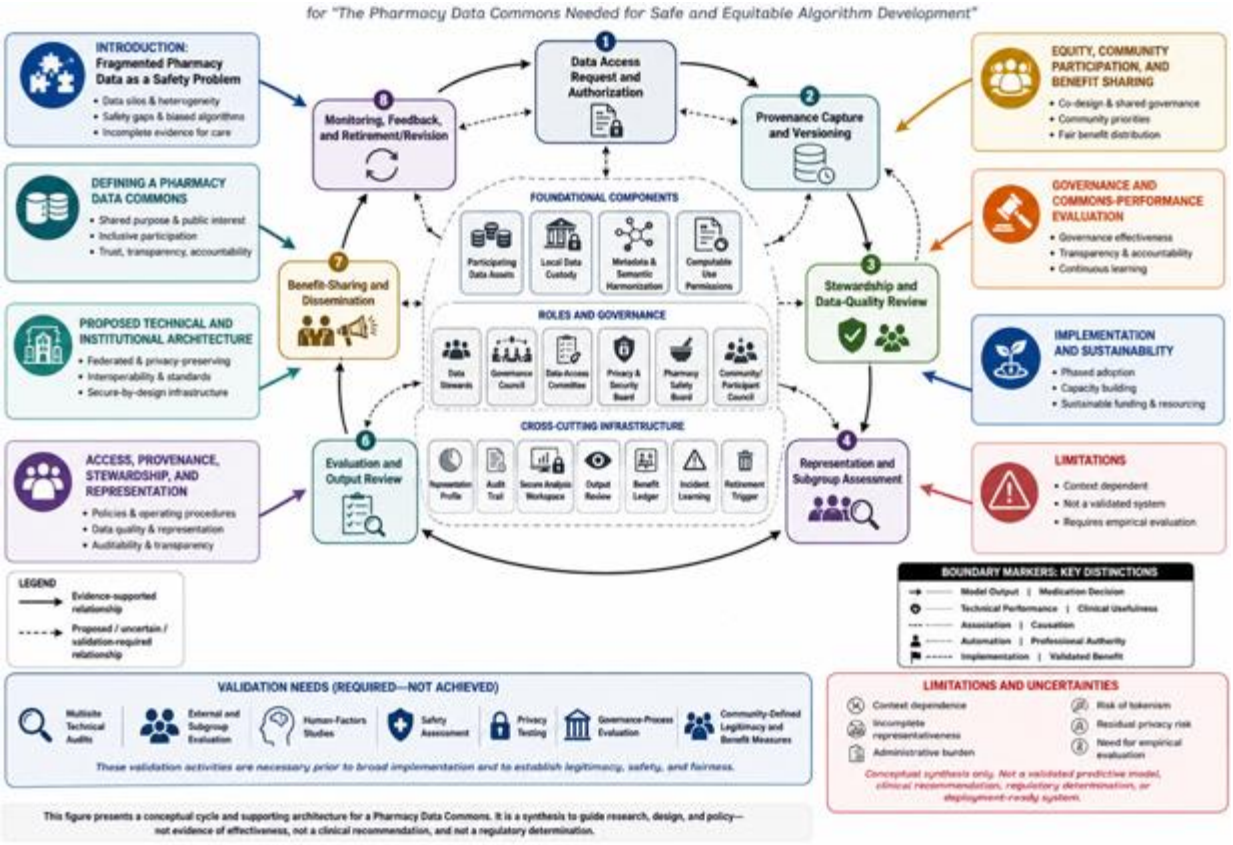

Figure 2. Data Access, Provenance, Stewardship, Representation, and Benefit-Sharing Cycle

The figure is an original conceptual synthesis developed for "The Pharmacy Data Commons Needed for Safe and Equitable Algorithm Development." 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.

Implementation must also support human–AI collaboration. Adoption depends on role clarity, workflow fit, calibrated trust, communication, and the capacity to act on outputs [30]. Sustainability requires continuing resources, adaptation, monitoring, and the ability to retire practices that no longer provide acceptable value [31]. The proposed operating cycle should therefore connect access authorization to provenance capture, representation assessment, secure analysis, multidimensional evaluation, output review, benefit sharing, monitoring, and retirement. Failed checks should return the project to correction or redesign rather than be treated as administrative delays.

Table 3 organizes the evidence, constructs, relationships, uncertainties, and boundary conditions required for governance responsibilities, implementation conditions, and monitoring requirements for the article The Pharmacy Data Commons Needed for Safe and Equitable Algorithm Development.

Table 3. Governance responsibilities, implementation conditions, and monitoring requirements for the article The Pharmacy Data Commons Needed for Safe and Equitable Algorithm Development.

Governance domain	Responsible actor	Required authority	Documentation requirement	Monitoring indicator	Escalation trigger	Corrective action	Accountability boundary
Membership and charter	Governing council	Admit, condition, suspend, or remove participants	Charter, conflicts, mandates, and decision rights	Participation and unresolved disputes	Repeated non-compliance or institutional capture	Apply conditions, suspend participation, or revise governance	Council membership does not certify safety [24, 25].

Dataset onboarding	Data steward and source custodian	Accept, restrict, remediate, version, or reject assets	Metadata, provenance, permissions, quality, and representation profile	Missing, inconsistent, or outdated documentation	Material undocumented change	Correct, restrict, re-version, or withdraw the asset	Onboarding does not establish task fitness [6, 9].
Access authorization	Data-access committee	Approve, deny, condition, revoke, or suspend use	Purpose, requested elements, permissions, reasons, appeal, and audit trail	Agreement, reversals, exceptions, and unauthorized use	Permission conflict, misuse, or scope expansion	Pause access and reassess conditions	Computable screening does not replace accountable judgment [16].
Representation and equity	Equity panel and community council	Require redesign, restriction, additional data, or additional evaluation	Population profile, targets, proxies, missingness, and subgroup plan	Exclusions, differential errors, and benefit distribution	Unaddressed proxy or subgroup risk	Redesign the target, collect evidence, or restrict use	Representation review cannot guarantee fairness [18, 21].
Medication safety	Pharmacy safety board	Require testing, mitigation, suspension, or retirement	Hazard analysis, workflow assumptions, overrides, incidents, and residual risks	New error pathways, unsafe workarounds, or non-actionable alerts	Credible medication-harm pathway	Redesign workflow, add controls, or stop progression	Technical performance is not clinical usefulness [26, 27].
Community participation and benefit	Community or participant council	Influence priorities, access conditions, dissemination, and benefit plans	Membership, deliberations, institutional response, and benefit ledger	Influence, burden, trust, unmet commitments, and delivered benefit	Tokenism, exclusion, or unfulfilled commitments	Reopen decisions and revise participation or benefit arrangements	Consultation alone is not shared authority [22, 23].
Implementation and workforce	Participating organizations and pharmacy leaders	Allocate resources, roles, training, and local workflow controls	Readiness plan, workflow map, role allocation, and training record	Workload, uptake, trust calibration, and unresolved responsibilities	Unsafe workarounds or responsibility gaps	Retrain, redesign workflow, restrict use, or suspend activity	Adoption is not validated benefit [28, 30].
Monitoring and retirement	Governing council with stewardship, technical, safety, and community advice	Reassess, restrict, suspend, or retire data assets, models, and approved uses	Change, drift, incident, remediation, and retirement records	Data, practice, model, governance, legitimacy, and resource change	Persistent deterioration or unsupported continued use	Recalibrate, correct, restrict, suspend, or retire	Continued operation does not establish sustainability [27, 31].

Limitations

Dataset shift may alter relationships across sites and time, so development evidence may not remain transportable [32]. Cross-institution performance can vary when models learn site-specific data-generating patterns [33]. Privacy risk persists beyond de-identification and depends on linkage, access, security, consent, and jurisdiction [34]. Fairness metrics encode normative choices and may conflict; they remain useful diagnostic tools but cannot independently establish equitable outcomes [35].

The proposed architecture may also create administrative burden, institutional capture, tokenistic participation, indicator gaming, or unequal capacity demands. Its legal applicability, operating costs, governance effectiveness, effects on medication safety, and distributional consequences require empirical evaluation. The architecture organizes validation questions but does not answer them.

Conclusion

A pharmacy data commons should be understood as a governed capability rather than merely a larger repository. The proposed architecture links distributed data custody with shared semantics, provenance, permissions, representation assessment, secure analysis, institutional accountability, community participation, benefit sharing, evaluation gates, monitoring, and retirement. Its central boundary is that access to data and production of an algorithmic output do not confer authority to make or automate a medication decision.

The contribution is an integrated framework for making assumptions, responsibilities, uncertainties, and failure conditions inspectable. It offers a basis for testing whether pharmacy data can be reused more consistently, safely, legitimately, and equitably across organizations. Those outcomes remain unproven. Validation must examine technical operation, professional workflow, medication-safety consequences, subgroup effects, privacy, organizational feasibility, community influence, benefit distribution, and sustainability in the intended context.

The architecture is therefore a research and governance proposition rather than a clinical recommendation, regulatory determination, universal standard, or deployment-ready system. Its value depends on whether future evaluation demonstrates that its technical and institutional components operate coherently without transferring hidden burdens, weakening professional responsibility, or reproducing existing inequities.

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

References

1. Hatzimanolis J, Riley B, El-Den S, Aslani P, Zhou J, Chaar BB. Applications of artificial intelligence in current pharmacy practice: A scoping review. Res Social Adm Pharm. 2025;21(3):134-41. doi:10.1016/j.sapharm.2024.12.007
2. Ranchon F, Chanoine S, Lambert-Lacroix S, Bosson JL, Moreau-Gaudry A, Bedouch P. Development of artificial intelligence powered apps and tools for clinical pharmacy services: A systematic review. Int J Med Inform. 2023;172:104983. doi:10.1016/j.ijmedinf.2022.104983
3. Lehne M, Sass J, Essenwanger A, Schepers J, Thun S. Why digital medicine depends on interoperability. NPJ Digit Med. 2019;2:79. doi:10.1038/s41746-019-0158-1
4. Gianfrancesco MA, Tamang S, Yazdany J, Schmajuk G. Potential biases in machine learning algorithms using electronic health record data. JAMA Intern Med. 2018;178(11):1544-7. doi:10.1001/jamainternmed.2018.3763
5. 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
6. Boeckhout M, Zielhuis GA, Bredenoord AL. The FAIR guiding principles for data stewardship: Fair enough? Eur J Hum Genet. 2018;26(7):931-6. doi:10.1038/s41431-018-0160-0
7. Wise J, de Barron AG, Splendiani A, Balali-Mood B, Vasant D, Little E, et al. Implementation and relevance of FAIR data principles in biopharmaceutical R&D. Drug Discov Today. 2019;24(4):933-8. doi:10.1016/j.drudis.2019.01.008
8. Lawson J, Cabili MN, Kerry G, Boughtwood T, Thorogood A, Alper P, et al. The Data Use Ontology to streamline responsible access to human biomedical datasets. Cell Genom. 2021;1(2):100028. doi:10.1016/j.xgen.2021.100028
9. Gebru T, Morgenstern J, Vecchione B, Vaughan JW, Wallach H, Daumé H III, et al. Datasheets for datasets. Commun ACM. 2021;64(12):86-92. doi:10.1145/3458723
10. Rieke N, Hancox J, Li W, Milletarì F, Roth HR, Albarqouni S, et al. The future of digital health with federated learning. NPJ Digit Med. 2020;3:119. doi:10.1038/s41746-020-00323-1
11. Kaissis GA, Makowski MR, Rückert D, Braren RF. Secure, privacy-preserving and federated machine learning in medical imaging. Nat Mach Intell. 2020;2(6):305-11. doi:10.1038/s42256-020-0186-1
12. Sheller MJ, Edwards B, Reina GA, Martin J, Pati S, Kotrotsou A, et al. Federated learning in medicine: Facilitating multi-institutional collaborations without sharing patient data. Sci Rep. 2020;10(1):12598. doi:10.1038/s41598-020-69250-1
13. Lenert LA, Ilatovskiy AV, Agnew J, Rudisill P, Jacobs J, Weatherston D, et al. Automated production of research data marts from a canonical fast healthcare interoperability resource data repository: Applications to COVID-19 research. J Am Med Inform Assoc. 2021;28(8):1605-11. doi:10.1093/jamia/ocab108
14. Dayan I, Roth HR, Zhong A, Harouni A, Gentili A, Abidin AZ, et al. Federated learning for predicting clinical outcomes in patients with COVID-19. Nat Med. 2021;27(10):1735-43. doi:10.1038/s41591-021-01506-3
15. Warnat-Herresthal S, Schultze H, Shastry KL, Manamohan S, Mukherjee S, Garg V, et al. Swarm learning for decentralized and confidential clinical machine learning. Nature. 2021;594(7862):265-70. doi:10.1038/s41586-021-03583-3
16. Cabili MN, Lawson J, Saltzman A, Rushton G, O'Rourke P, Wilbanks J, et al. Empirical validation of an automated approach to data use oversight. Cell Genom. 2021;1(2):100031. doi:10.1016/j.xgen.2021.100031
17. Lewis AE, Weiskopf N, Abrams ZB, Foraker R, Lai AM, Payne PRO, et al. Electronic health record data quality assessment and tools: A systematic review. J Am Med Inform Assoc. 2023;30(10):1730-40. doi:10.1093/jamia/ocad120
18. Alderman JE, Palmer J, Laws E, McCradden MD, Ordish J, Ghassemi M, et al. Tackling algorithmic bias and promoting transparency in health datasets: The STANDING Together consensus recommendations. Lancet Digit Health. 2025;7(1). doi:10.1016/S2589-7500(24)00224-3
19. Rehm HL, Page AJH, Smith L, Adams JB, Alterovitz G, Babb LJ, et al. GA4GH: International policies and standards for data sharing across genomic research and healthcare. Cell Genom. 2021;1(2):100029. doi:10.1016/j.xgen.2021.100029
20. 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
21. 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
22. Kalkman S, van Delden JJM, Banerjee A, Tyl B, Mostert M, van Thiel GJMW. Patients' and public views and attitudes towards the sharing of health data for research: A narrative review of the empirical evidence. J Med Ethics. 2022;48(1):3-13. doi:10.1136/medethics-2019-105651
23. Bazzano AN, Mantsios A, Mattei N, Kosorok MR, Culotta A. AI can be a powerful social innovation for public health if community engagement is at the core. J Med Internet Res. 2025;27. doi:10.2196/68198
24. Morley J, Machado CCV, Burr C, Cowls J, Joshi I, Taddeo M, et al. The ethics of AI in health care: A mapping review. Soc Sci Med. 2020;260:113172. doi:10.1016/j.socscimed.2020.113172
25. Mittelstadt B. Principles alone cannot guarantee ethical AI. Nat Mach Intell. 2019;1(11):501-7. doi:10.1038/s42256-019-0114-4
26. Vasey B, Nagendran M, Campbell B, Clifton DA, Collins GS, Denaxas S, et al. DECIDE-AI: New reporting guidelines to bridge the development-to-implementation gap in clinical artificial intelligence. Nat Med. 2022;28(5):924-33. doi:10.1038/s41591-022-01772-9
27. 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
28. 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 scale-up, spread, and sustainability of health and care technologies. J Med Internet Res. 2017;19(11). doi:10.2196/jmir.8775
29. 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
30. Henry KE, Kornfield R, Sridharan A, Linton RC, Groh C, Wang T, et al. Human-machine teaming is key to AI adoption: Clinicians' experiences with a deployed machine learning system. NPJ Digit Med. 2022;5(1):97. doi:10.1038/s41746-022-00597-7
31. Lennox L, Maher L, Reed J. Navigating the sustainability landscape: A systematic review of sustainability approaches in healthcare. Implement Sci. 2018;13(1):27. doi:10.1186/s13012-017-0707-4
32. 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
33. Zech JR, Badgeley MA, Liu M, Costa AB, Titano JJ, Oermann EK. Variable generalization performance of a deep learning model to detect pneumonia in chest radiographs: A cross-sectional study. PLoS Med. 2018;15(11). doi:10.1371/journal.pmed.1002683
34. Price WN 2nd, Cohen IG. Privacy in the age of medical big data. Nat Med. 2019;25(1):37-43. doi:10.1038/s41591-018-0272-7
35. 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