Original Article

Continual Learning without Silent Change in Medication-Support Systems

Daniel Kim[1]*, Jihoon Park[1], Min-seo Kang[2], Hyun-woo Jung[3]

[1]Department of Continual Learning and Medication Safety, College of Pharmacy, Kyung Hee University, Seoul, South Korea. [2]Department of AI Stability in Pharmacy Systems, College of Pharmacy, Chung-Ang University, Seoul, South Korea. [3]Department of Non-Silent AI Updates, College of Pharmacy, Kangwon National University, Chuncheon, South Korea.

Abstract

Medication-support systems may change after implementation because patient populations, clinical practices, data pipelines, coding conventions, medicine knowledge, model parameters, interfaces, and user responses evolve. Such change becomes unsafe to govern when it is not visible, version-linked, investigated, or subjected to evidence proportionate to its possible consequences. This article develops a proposed continual-learning control architecture for medication-support systems. It distinguishes continual learning from automatic production updating and defines silent change as a potentially consequential alteration in data, semantics, model behaviour, workflow position, human use, or decision effects that occurs without a contemporaneous and accountable assessment of whether requalification is required. The proposed architecture separates a locked operational plane from an isolated candidate-learning plane. Its principal components are an intended-use and change contract, multidimensional monitoring, change-event classification, an immutable version and evidence ledger, risk-proportionate requalification, staged release, rollback, suspension, and retirement. A no-silent-promotion boundary prevents a candidate version from replacing the operational version solely because retraining or automated updating has occurred. Evaluation must extend beyond discrimination to calibration, medication-task performance, workflow effects, human reliance, medication-safety signals, subgroup equity, traceability, and organizational readiness. Validation would require retrospective, external, prospective, human-factor, implementation, and post-release evidence matched to the system's intended use and potential medication harm. The architecture does not establish clinical effectiveness, regulatory conformity, universal thresholds, or deployment readiness. Its original contribution is an integrated governance structure through which continual learning may remain technically possible while operational change remains visible, reviewable, reversible, and professionally accountable.

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

Introduction

Artificial-intelligence systems are increasingly proposed for medication alerts, prescription prioritization, risk estimation, pharmacovigilance, reconciliation, and other forms of medication-related decision support. Their technical performance, however, does not independently establish clinical value. Translation requires evidence that the output is useful within the intended workflow, population, decision, and professional context [1]. This distinction is especially important for systems expected to learn after deployment because the object being evaluated may no longer remain stable.

Responsible clinical machine learning therefore requires controls that extend across development, implementation, use, monitoring, and maintenance [2]. Initial validation addresses only a defined model version, dataset, configuration, population, and use condition. It cannot prospectively establish the safety or usefulness of every later version. Nevertheless, continual-learning language sometimes implies that adaptation is intrinsically beneficial or that technical updating may proceed as an ordinary background process.

Clinical practice itself can change the data received by an algorithm. New testing patterns, altered prescribing, documentation behaviour, intervention effects, or changes in who receives care may modify the relationships on which a

Address for correspondence: Daniel Kim, Department of Continual Learning and Medication Safety, College of Pharmacy, Kyung Hee University, Seoul, South Korea.
daniel.kim@khu.ac.kr
Received: 17 November 2025; **Accepted:** 28 February 2026

How to cite this article: Kim D, Park J, Kang M, Jung HW. Continual Learning without Silent Change in Medication-Support Systems. Arch Pharm Pract. 2026;17(1):57-67. https://doi.org/10.51847/65Q1PxmAOE

model depends [3]. A system may consequently change even when its source code remains fixed. Conversely, retraining may alter its outputs despite stable headline performance. Neither condition is adequately described by viewing change only as a software release.

Sustainable clinical prediction consequently requires explicit maintenance strategies, including surveillance, recalibration, revision, and retirement [4]. Yet maintenance frameworks rarely integrate these activities with medication-specific risk, pharmacist workflow, accountable release authority, and an enforceable distinction between learning and production deployment. Postdeployment algorithmovigilance further indicates that monitoring must address effectiveness and equity rather than technical operation alone [5].

This article defines silent change as any potentially consequential alteration in data, semantics, model behaviour, clinical knowledge, system configuration, workflow position, human interaction, or decision effects that occurs without a visible, version-linked, and accountable determination of whether requalification is required. It proposes a controlled continual-learning lifecycle in which adaptation may occur within an isolated candidate environment, but operational promotion requires declared change, traceable evidence, staged evaluation, professional authorization, and rollback readiness. The contribution is conceptual and non-empirical. It does not claim that the proposed architecture has been validated, that continual learning is preferable to a fixed system, or that satisfying its gates establishes clinical safety or regulatory acceptance.

Sources and Consequences of Silent Change

Silent change extends beyond conventional data drift. Medical datasets can change through population composition, disease prevalence, acquisition technology, coding, measurement, clinical practice, and technical infrastructure [6]. In medication-support systems, corresponding sources include formulary revisions, new medicines, changed interaction knowledge, laboratory-unit changes, substitution policies, prescribing pathways, documentation practices, and altered access to care. A change may therefore originate upstream of the model, within the model, at the interface, or in the surrounding organization.

Temporal change may alter both feature distributions and the relationships between observed variables and clinical risk [7]. A laboratory pattern that previously identified a medication-related concern may become less informative after a new treatment pathway is introduced. This does not necessarily mean that the model has malfunctioned. Alternative explanations include altered case mix, improved care, changed outcome ascertainment, or a feedback effect created by the system itself.

Model degradation is similarly not uniform. Drift type, timing, label availability, and retraining strategy can influence whether updating improves, preserves, or worsens performance [8]. Retraining on recent data may reproduce temporary disruption, amplify documentation bias, or obscure rare but severe medication events. Accordingly, the presence of statistical change does not prove clinical harm, and the absence of aggregate performance loss does not prove stability.

External generalization failures demonstrate that models may depend on institution-specific markers or data-generation processes rather than transportable clinical relationships [9]. Within one organization, equivalent problems can arise after an electronic prescribing migration, new laboratory interface, alert redesign, or change in pharmacist review practices. A stable average metric may conceal changed subgroup errors, workload displacement, alert burden, or altered reliance.

Causal reasoning can help distinguish mutable correlates from relationships expected to remain more stable under specified shifts [10]. It cannot, however, eliminate uncertainty about future clinical practice or guarantee transportability. The proposed architecture therefore treats change as a chain:

source of change → changed system property → affected decision layer → observable consequence → required control response.

Six proposed change objects organize this chain: data and population; semantics and preprocessing; model or knowledge content; interface and workflow; human interpretation and action; and safety, equity, or organizational consequence. These categories are intended to prevent a model-centred account from overlooking medication-support harms that emerge through workflow or professional use.

Table 1 organizes the evidence, constructs, relationships, uncertainties, and boundary conditions required for definitions, components, relationships, and intended uses for the article Continual Learning without Silent Change in Medication-Support Systems.

Table 1. Definitions, components, relationships, and intended uses for the article Continual Learning without Silent Change in Medication-Support Systems.

Component or Scientific Dimension	Problem Addressed	Inputs or Determinants	Proposed Mechanism or Relationship	Expected Contribution	Evidence Required	Failure or Uncertainty Risk	Boundary Statement
Silent Change—Proposed Construct	Consequential alteration may occur without	Data, semantics, model, interface, workflow, human-	A change becomes silent when it is not linked to a visible	Expands attention beyond retraining	Temporal, technical, workflow, safety,	Change may be detectable but not	The construct does not establish

	explicit recognition	use, safety, or equity change	assessment and accountable response	and software versions	and subgroup evidence [6]	causally attributable	frequency or severity
Data and Population Change	Development data may cease to represent operational inputs	Case mix, prevalence, missingness, acquisition, access to care	Changed input distributions may alter calibration or errors	Supports monitoring of input validity	Longitudinal and external evaluation [6]	Aggregate stability may conceal subgroup change	Distribution change is not itself proof of harm
Semantic and Preprocessing Change	Identical values may acquire different meaning	Coding systems, units, mappings, extraction logic	Changed meaning alters the effective feature even when the field name persists	Supports semantic versioning	Data dictionaries, interface tests, clinical review	Undocumented mapping changes may evade model monitoring	Technical interoperability does not ensure semantic equivalence
Model or Knowledge Change	Updating may alter outputs and decision thresholds	Retraining data, parameters, rules, knowledge bases	Candidate modification creates a new decision-producing object	Makes updating subject to declared version control	Reproducibility, calibration, task and safety evidence [8]	Retraining may reproduce transient or biased patterns	A new version is not qualified merely because training completed
Workflow and Human-Use Change	System effects depend on where and how outputs are used	Alert placement, timing, user group, workload, training	Changed presentation may alter attention, reliance, overrides, and action	Connects technical change to pharmacy practice	Observation, usability, workload, and human-factor evidence	Human adaptation may occur without model change	Human presence does not itself constitute effective oversight
Transportability	Local signals may not remain valid elsewhere or later	Institutional practices, equipment, coding, population	Models may rely on local proxies rather than stable relationships	Requires external and temporal requalification	Cross-site and temporal testing [9]	Similar performance may mask different error mechanisms	One-site validity does not establish generalizability
Shift-Stability	Correlations may change when clinical processes change	Causal structure, interventions, policy, workflow	More stable relationships may be preferred when justified	Encourages explicit assumptions about mutable mechanisms	Causal analysis and prospective testing [10]	Causal structure may be misspecified	Causal reasoning does not guarantee future stability
Medication-Support Consequence	Model change may affect professional work before patient outcomes are observable	Prioritization, alerts, recommendations, escalation	Technical outputs influence medication decisions through workflow and human interpretation	Separates model performance from medication-safety effect	Task, workflow, safety, equity, and implementation evidence	Delayed or rare harm may not be visible in routine metrics	No single metric establishes medication safety

These proposed categories organize plausible change pathways; they do not rank changes, establish universal thresholds, or demonstrate that every detected alteration requires model updating.

Proposed Controlled Continual-Learning Lifecycle

Real-world clinical integration depends on workflow fit, professional roles, trust, infrastructure, and operational adaptation rather than model installation alone [11]. A continual-learning medication-support system must therefore be governed as a sociotechnical intervention whose data, model, users, organization, and consequences can all change.

Figure 1 presents the original conceptual synthesis for controlled continual-learning lifecycle for medication-support systems, showing how the article's principal components, evidence relationships, uncertainties, and decision boundaries are connected.

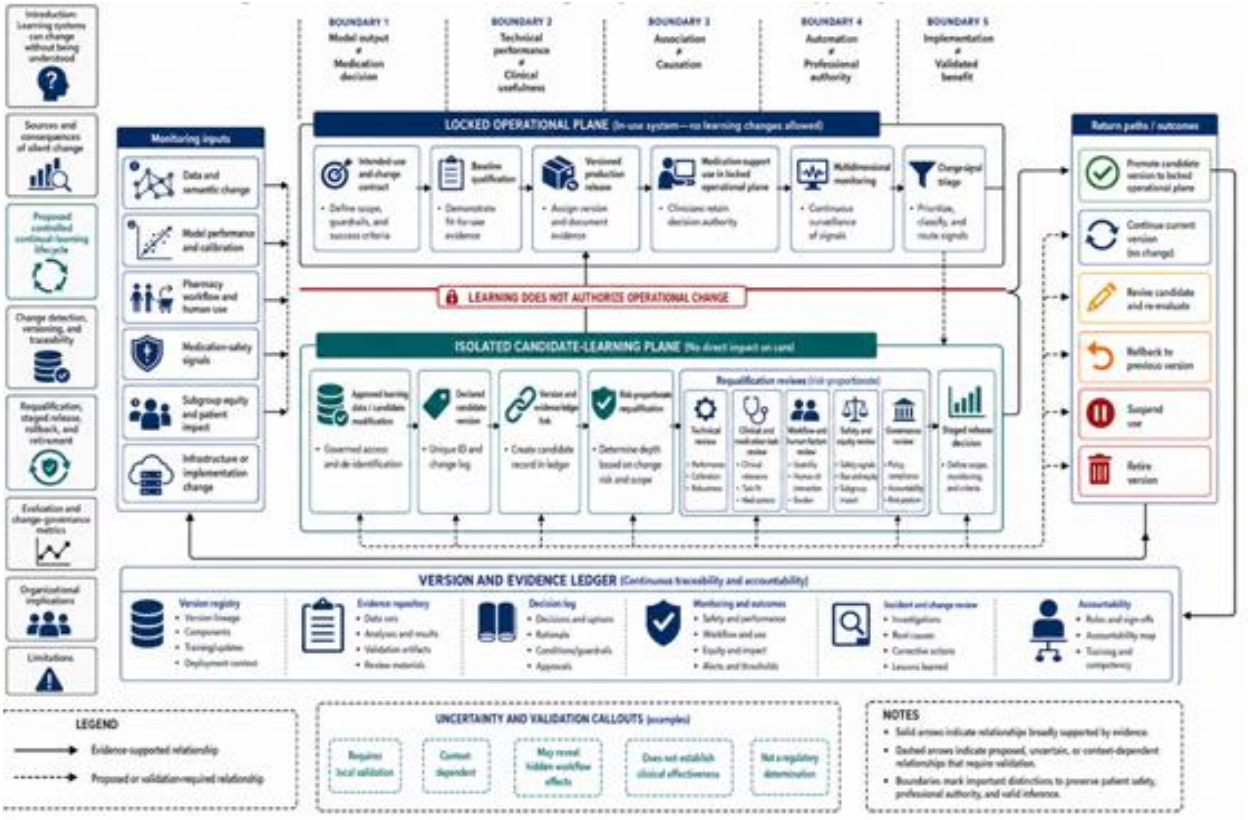

Figure 1. Controlled Continual-Learning Lifecycle for Medication-Support Systems

The proposed architecture contains two planes. The locked operational plane contains the qualified version currently available to users. It operates only within its declared intended use, input contract, workflow position, user population, decision role, and escalation conditions. The isolated candidate-learning plane may receive approved data and generate a modified candidate, but it cannot directly replace the operational version.

This separation addresses a central conceptual error: continual learning is not equivalent to continual deployment. A system may learn automatically while authorization to alter medication-support behaviour remains an accountable organizational decision. The architecture therefore proposes a no-silent-promotion rule: a candidate version cannot enter routine use without a declared version change, evidence linked to that version, risk-proportionate requalification, staged-release conditions, named authorization, and a tested rollback route.

The lifecycle begins with an intended-use and change contract defining the medication task, users, inputs, excluded uses, escalation rules, evidence expectations, and changes that require review. Patient-benefit evaluation requires transparency, reproducibility, ethics, and effectiveness to be considered together rather than as interchangeable properties [12]. Baseline qualification consequently includes technical performance, medication-task performance, workflow interaction, foreseeable safety consequences, subgroup evaluation, and organizational readiness.

After release, multidimensional surveillance observes data, semantics, calibration, task behaviour, pharmacist interaction, medication-safety signals, equity, infrastructure, and intended-use adherence. Pharmacovigilance offers relevant principles for signal detection, investigation, documentation, and continuing reassessment, although algorithm-related events are not equivalent to adverse drug reactions [13]. Monitoring signals initiate investigation rather than automatic retraining.

If change is plausible, an approved candidate may be produced in the isolated plane. Its provenance should identify the datasets, inclusion rules, collection conditions, transformations, limitations, code, parameters, thresholds, and knowledge sources used. Structured dataset documentation supports this requirement [14]. Semantic interoperability is also necessary because stable field names do not guarantee stable clinical meaning across systems or time [15].

The candidate then enters risk-proportionate requalification. The proposed evidence burden increases when the change affects intended use, decision thresholds, medication-harm severity, subgroup performance, workflow authority, or reversibility. Possible outcomes are further investigation, candidate revision, staged release, promotion, rollback,

suspension, or retirement. Every decision is entered into an immutable version and evidence ledger.

Table 2 organizes the evidence, constructs, relationships, uncertainties, and boundary conditions required for evaluation criteria, evidence requirements, and failure modes for the article Continual Learning without Silent Change in Medication-Support Systems.

Table 2. Evaluation criteria, evidence requirements, and failure modes for the article Continual Learning without Silent Change in Medication-Support Systems.

Architecture Layer or Process Stage	Core Function	Information Flow	Interaction With other Components	Evaluation Criterion	Potential Failure Mode	Human or Organizational Responsibility	Readiness Boundary
Intended-Use And Change Contract—Proposed	Define task, users, inputs, exclusions, decision role, and review-triggering change	Clinical, pharmacy, technical, safety, and organizational requirements enter the contract	Constrains monitoring, requalification, and release	Intended use is specific, testable, versioned, and reviewable	Scope remains implicit or expands without review	Pharmacy clinical owner and accountable release authority	A system cannot be qualified for undefined or silently expanded use
Locked Operational Plane—Proposed	Preserve the currently authorized version	Qualified version produces traceable outputs	Receives monitoring but not unapproved candidate updates	Output is reproducible and linked to a declared version	Candidate parameters or knowledge enter production without authorization	Technical operator under clinical governance	Operational continuity is not evidence of continued clinical validity
Isolated Candidate-Learning Plane—Proposed	Permit controlled adaptation without direct production replacement	Approved data create a declared candidate version	Sends evidence to requalification, not directly to users	Candidate is reproducible, documented, and separated from production	Data leakage, undocumented retraining, or accidental promotion	Model owner and data steward	Successful training does not establish release eligibility
Multidimensional Monitoring	Detect possible changes in data, model, workflow, safety, and equity	Operational evidence enters a version-linked surveillance process	May trigger investigation or continued observation	Indicators are sensitive to relevant change and interpreted in context	Overreliance on one aggregate model metric	Safety, pharmacy, informatics, and equity functions	A signal is not proof of harm or authorization to update
Provenance and Evidence Ledger—Proposed	Preserve an auditable history of versions, evidence, and decisions	Dataset, code, model, interface, intended-use, incident, and approval records are linked	Supports investigation, rollback, and accountability	Records permit reconstruction of what operated, when, and why	Incomplete lineage prevents attribution of an incident	Data steward, model owner, and governance secretariat	Documentation improves traceability but does not prove causal responsibility
Dataset Documentation	Record origin, composition, collection, processing, and limitations	Dataset evidence accompanies every candidate	Informs applicability and requalification	Provenance and limitations are explicit [14]	Recent data are assumed representative without assessment	Data steward with clinical review	Documented data are not necessarily unbiased or clinically sufficient
Semantic Interoperability	Preserve clinical meaning across data exchange and time	Coding, units, terminology, and mappings connect source systems to the candidate	Conditions both monitoring and model validity	Equivalent concepts remain semantically consistent [15]	Stable technical interfaces conceal changed meaning	Informatics and pharmacy-domain stewards	Interoperability is necessary but insufficient for safety
Requalification Gate—Proposed	Match evidence burden to possible consequence	Candidate evidence is compared with the contract and current version	Determines eligibility for staged release	Technical, task, workflow, safety, equity, and governance questions are answered	A narrow performance gain overrides unresolved safety evidence	Multidisciplinary review body	Gate passage applies only to the specified version, use, and release stage
Staged Release and Rollback—Proposed	Limit exposure and preserve reversibility	Candidate is introduced under predefined observation and escalation conditions	Feeds intensified monitoring and release decisions	Exposure is controlled and rollback is executable	Harm signal persists because authority or rollback route is unclear	Named release and incident authorities	Staged release does not establish final effectiveness or permanent authorization

The integrated lifecycle and its gates are proposed. The cited evidence supports individual monitoring, documentation, implementation, and evaluation requirements but does not validate the complete architecture.

Change Detection, Versioning, and Traceability

Change detection should answer whether an operational assumption may no longer hold, not merely whether a numerical distribution differs. Medication-support monitoring therefore requires a layered signal strategy. Data-layer signals include missingness, prevalence, coding, units, measurement processes, and subgroup composition. Model-layer signals include calibration, discrimination, uncertainty, threshold behaviour, and error distribution. Workflow-layer signals include alert volume, pharmacist review timing, overrides, escalation, reliance, and work displacement. Consequence-layer signals include medication incidents, near misses, subgroup disparities, unintended treatment delays, and changes in who receives professional review.

Calibration is particularly important because predicted risks may become systematically too high or too low even when rank discrimination remains acceptable. Calibration-drift detection can support decisions about when closer investigation or updating should be considered [16]. It should not operate as an automatic retraining trigger. Calibration change may result from altered prevalence, outcome ascertainment, clinical intervention, or data quality, each requiring a different response.

Drift detectors also vary in what they observe, the labels they require, and the changes they can detect. Empirical comparisons show that detection performance depends on the method and shift examined [17]. A label-free distribution alarm may provide early warning but cannot establish changed clinical accuracy. Conversely, outcome-based monitoring may be more directly interpretable yet arrive too late for rare or delayed medication harms. The architecture therefore proposes triangulation rather than dependence on one detector.

A change event should be opened when a monitored signal crosses a locally prespecified investigation condition, when an incident suggests changed behaviour, or when a declared modification affects a controlled component. The record should identify the initiating signal, affected version, data period, intended use, users, possible medication consequences, uncertainty, and immediate containment action. The signal is then classified as administrative, semantic, statistical, model-related, workflow-related, or potentially safety-, equity-, or intended-use-changing. These categories are proposed triage aids rather than validated risk levels.

Versioning must extend beyond the model file. Minimum clinical-AI reporting emphasizes documentation of model development, data, evaluation, and reproducibility [18]. For continual-learning systems, the version object should additionally contain or reference:

- Data-source and extraction version.
- Terminology, unit, and mapping version.
- Preprocessing and feature-definition version.
- Model code, parameters, and threshold version.
- Clinical knowledge and formulary version.
- Interface and workflow configuration.
- Intended-use and excluded-use declaration.
- User population and organizational location.
- Qualification evidence and decision record.
- Release, rollback, suspension, and retirement status.

Transparent reporting of model development, validation, and updating is also required by updated prediction-model guidance [19]. In operational use, these reporting principles must become executable traceability rather than retrospective manuscript description. Each output should therefore be attributable to the exact configuration that produced it. Each incident should be linkable to the contemporaneous data, model, threshold, interface, workflow, and authorization state.

The proposed version and evidence ledger is the architecture's persistent traceability mechanism. It records candidate creation, detected signals, evidence submissions, unresolved uncertainties, review decisions, staged-release conditions, overrides, incidents, rollback tests, and retirement actions. The ledger may be technically distributed across existing systems, but its logical links must remain reconstructable.

Traceability has limits. A complete record does not prove that a particular change caused an observed medication event. Multiple components may change together, outcome labels may be delayed, and user behaviour may adapt. The ledger instead supports disciplined investigation by making competing explanations visible. Accordingly, change detection is an investigative trigger; versioning establishes identity; and traceability preserves the evidence chain. None independently establishes clinical harm, successful remediation, or readiness for promotion.

Requalification, Staged Release, Rollback, and Retirement

Detecting change does not establish that a candidate update is preferable to the current operational version. Evidence on harmful distribution shifts shows that detection and remediation are separate problems: a method may recognize change without identifying its cause, and an attempted correction may introduce new errors or subgroup effects [20]. Requalification should therefore ask whether the changed system remains suitable for a defined medication-support function rather than whether retraining has technically succeeded.

The proposed architecture uses risk-proportionate requalification. A low-consequence documentation correction may require verification that operational behaviour is unchanged. A semantic remapping requires interface, unit, terminology, and clinical-meaning checks. Changes to model parameters, decision thresholds, intended users, or supported medication tasks require broader technical, clinical, workflow, safety, and equity evidence. The proposed tiers organize evaluation effort but are not validated risk categories or regulatory classifications.

Early live clinical evaluation should examine workflow fit, human interaction, safety, and implementation conditions in addition to model performance [21]. For medication-support systems, this includes whether alerts arrive at an actionable time, whether prioritization changes pharmacist workload, whether uncertainty is visible, whether escalation remains possible, and whether users can identify the operative version. Prospective evaluation should also report the system version, inputs, outputs, human interaction, errors, and relationship between model output and clinical action [22].

A candidate that passes requalification should enter a staged release rather than unrestricted substitution. Staging may limit location, user group, task, exposure, or duration while preserving comparison with the current version. Predefined escalation conditions should address technical instability, changed calibration, medication incidents, unexpected workload, subgroup disparity, or use outside the declared scope. Published drift research remains methodologically heterogeneous, so no universal detector or release threshold can be assumed appropriate across systems [23].

Rollback restores the most recent qualified configuration when a release condition is breached or uncertainty becomes unacceptable. It requires preserved software, data mappings, knowledge content, interface configuration, and operating instructions. Suspension differs from rollback because it may remove the function entirely while investigation continues. Retirement is appropriate when the intended use is no longer clinically meaningful, required data are unavailable, maintenance cannot be supported, or residual uncertainty exceeds the organization's capacity to govern the system.

Passing requalification establishes only eligibility for the specified next stage. It does not establish permanent validity, superior clinical outcomes, or universal readiness.

Evaluation and Change-Governance Metrics

Evaluation should distinguish model performance, medication-task performance, workflow consequence, safety, equity, implementation, and governance. Reviews of clinical-AI comparisons have identified weaknesses in study design, reporting, external validation, and claims of superiority [24]. A continual-learning system should therefore not be promoted because one retrospective metric improves while other decision-relevant properties remain unknown.

A multidimensional evaluation set may include data integrity, calibration, discrimination, clinically consequential error types, uncertainty, medication-alert burden, prescription-prioritization performance, pharmacist review time, override behaviour, subgroup performance, incident signals, traceability completeness, and rollback readiness. Trustworthy healthcare AI guidance similarly emphasizes fairness, universality, traceability, usability, robustness, and explainability across the lifecycle [25]. These dimensions should remain separate because strength in one does not compensate automatically for failure in another.

Medication-alert research shows that artificial intelligence may help prioritize or filter alerts, but available evidence is heterogeneous and frequently limited in external, prospective, workflow, or implementation evaluation [26]. Reduced alert volume must not therefore be equated with reduced medication harm. A system could suppress low-value alerts while also suppressing rare high-consequence warnings.

Machine-learning prescription prioritization may help direct pharmacist attention toward prescriptions with a greater estimated risk of error [27]. Its value remains conditional on local data, prescribing processes, error definitions, available pharmacist capacity, and the downstream action taken. Evaluation should consequently include both missed high-risk prescriptions and unnecessary review burden.

The proposed change-governance scorecard should record evidence by domain rather than collapse performance into a single readiness score. Release conditions must be prespecified for the local intended use, with non-compensable safety or accountability failures identified explicitly. No universal numerical threshold is proposed.

Organizational Implications

Continual-learning control is an organizational capability rather than a model feature. Implementation studies show that infrastructure, professional acceptance, collaboration, resources, workflow, and institutional context can facilitate or obstruct clinical-AI use [28]. An organization unable to identify the operative version, investigate a signal, suspend use, or restore a qualified configuration cannot reliably govern continual change.

Figure 2 presents the original conceptual synthesis for change detection, requalification, staged release, rollback, and retirement, showing how the article's principal components, evidence relationships, uncertainties, and decision boundaries are connected.

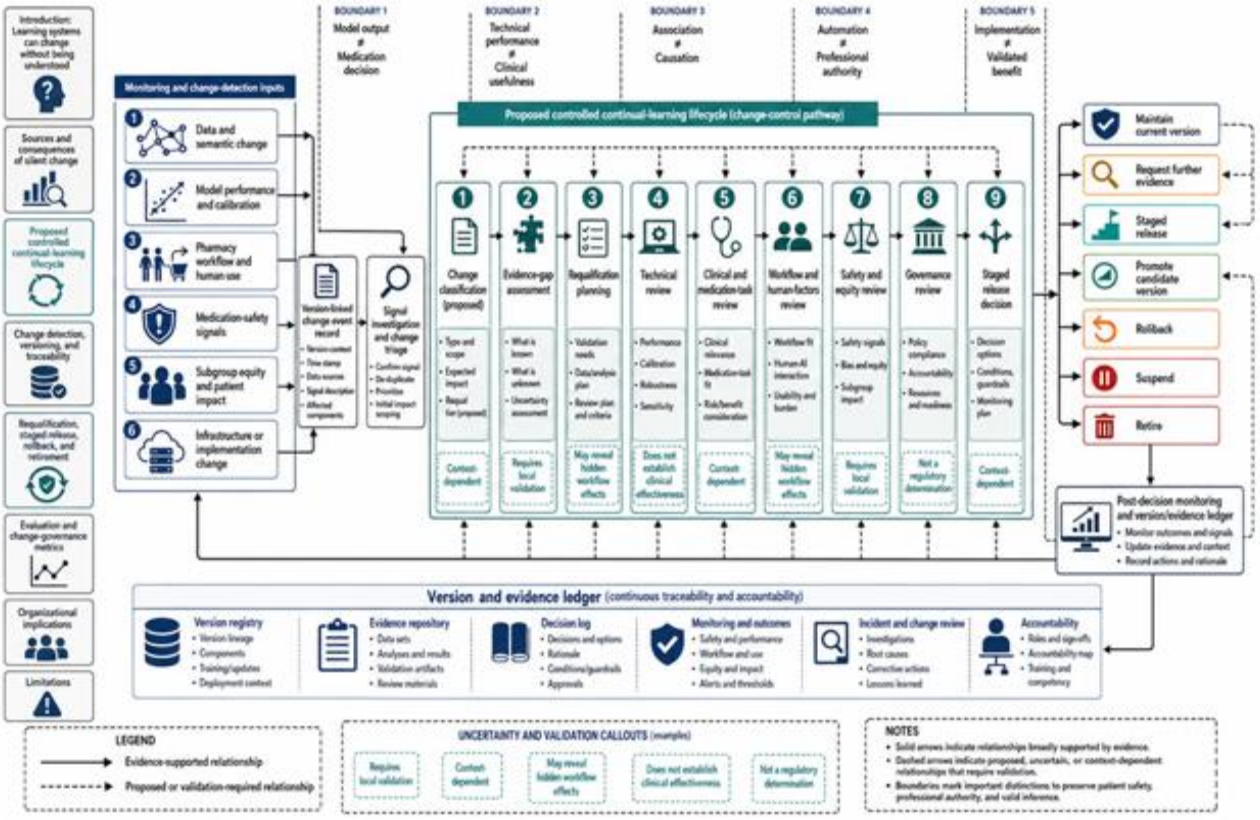

Figure 2. Change Detection, Requalification, Staged Release, Rollback, and Retirement

Human involvement is not automatically protective. Clinicians can be influenced by incorrect decision-aid recommendations, particularly when the output appears authoritative [29]. Meaningful oversight therefore requires time, competence, access to relevant evidence, authority to challenge the system, and a viable alternative action.

Governance must specify who can authorize, restrict, suspend, roll back, or retire a version. Organizational analyses of clinical-AI implementation emphasize that these decisions cannot be left solely to developers or informal local practice [30]. Pharmacy leadership, medication-safety personnel, informatics teams, data stewards, human-factors specialists, equity representatives, vendors, and institutional authorities may contribute different evidence, but accountability must remain identifiable.

Automation in medicine-safety surveillance similarly requires validation, expert interpretation, and continuing review [31]. Monitoring functions should therefore connect directly to a decision pathway rather than generate unowned dashboards or alerts.

Table 3 organizes the evidence, constructs, relationships, uncertainties, and boundary conditions required for governance responsibilities, implementation conditions, and monitoring requirements for the article Continual Learning without Silent Change in Medication-Support Systems.

Table 3. Governance responsibilities, implementation conditions, and monitoring requirements for the article Continual Learning without Silent Change in Medication-Support Systems.

Governance Domain	Responsible Actor	Required Authority	Documentation Requirement	Monitoring Indicator	Escalation Trigger	Corrective Action	Accountability Boundary
Clinical Intended Use—Proposed	Pharmacy clinical owner	Define and restrict medication-support use	Versioned task, users, exclusions, and escalation rules	Use outside declared scope	Scope expansion or changed clinical role	Restrict use and initiate requalification	Model output does not replace professional authority
Medication Safety	Medication-safety officer	Initiate investigation or suspension	Incidents, near misses, overrides, and reviews	Changed event pattern or credible severe signal	Possible preventable medication harm	Containment, investigation, rollback, or suspension	Signal does not prove causation

Technical Operation	Informatics and model owner	Maintain, isolate, and restore versions	Code, parameters, thresholds, interfaces, and release records	Failure, instability, or unauthorized change	Loss of reproducibility or traceability	Restore qualified version and preserve evidence	Technical authority does not confer clinical approval
Data And Semantics	Data and terminology steward	Halt use of invalid inputs or mappings	Source, extraction, unit, coding, and transformation versions	Missingness, coding, unit, or mapping change	Material semantic uncertainty	Correct pipeline and repeat qualification	Data availability does not establish clinical validity
Human Factors	Pharmacy workflow and human-factors lead	Require workflow reassessment	User group, interface, timing, training, and observed use	Workload, reliance, overrides, or delayed action	Changed use behaviour or unsafe interaction	Redesign, retrain, restrict, or suspend	Human presence alone is not meaningful oversight [29]
Release Governance—Proposed	Multidisciplinary release authority	Approve staged release, promotion, rollback, or retirement	Evidence package, dissent, conditions, and decision rationale	Unresolved evidence gaps or breached conditions	Failure of a non-compensable release condition	Hold, limit, rollback, suspend, or retire	Approval applies only to the declared version and use
Equity And Patient Relevance	Equity lead and patient representative	Require subgroup investigation and challenge assumptions	Subgroup definitions, access effects, and unresolved uncertainty	Changed subgroup error or allocation	Credible disparity or inadequate evidence	Restrict exposure and obtain further evidence	Statistical fairness does not establish health equity
Vendor Relationship—Proposed	Contract and governance owner	Require change disclosure, evidence access, and rollback support	Vendor updates, dependencies, incidents, and known limitations	Undeclared vendor-side change	Inability to reconstruct or control a version	Suspend update or system use	Outsourcing does not transfer institutional accountability
Independent Review—Proposed	Quality, safety, or audit function	Challenge release and incident decisions	Review findings and conflicts of interest	Repeated exceptions or unresolved dissent	Governance deviation or recurring failure	Escalate and mandate corrective review	Independence must be protected from operational pressure

These role allocations are proposed and must be adapted to local professional, legal, contractual, and regulatory arrangements.

Limitations

The architecture cannot eliminate structural bias or guarantee equitable consequences. An algorithm may produce inequity because its objective or proxy encodes unequal access or resource allocation despite apparently acceptable prediction [32]. Electronic health-record data may also reproduce differences in documentation, treatment, access, and observation [33].

Fairness evaluation must therefore be tied to clinical purpose and health-equity consequences rather than one mathematical criterion [34]. Technical mitigation cannot resolve every ethical question about which outcomes should be optimized, whose interests should govern trade-offs, or what disparities are acceptable [35].

Further limitations include delayed outcome labels, rare but severe medication events, uncertain causal attribution, vendor opacity, interacting system changes, human adaptation, and differences among advisory, prioritization, generative, and autonomous functions. The proposed classes, tiers, gates, and responsibilities have not been empirically validated. Their usefulness, feasibility, burden, and unintended effects require evaluation across diverse pharmacy settings and medication-risk contexts.

Conclusion

Continual learning in medication-support systems should not entail silent operational change. Data, semantics, models, interfaces, workflows, professional behaviour, and clinical context may evolve independently or together, making a fixed model-version view of maintenance inadequate.

This article proposes a controlled continual-learning architecture that separates candidate learning from production authorization. Its principal controls are a locked operational plane, an isolated learning plane, an intended-use and change contract, multidimensional monitoring, explicit versioning, an evidence ledger, risk-proportionate requalification, staged release, rollback, suspension, retirement, and named organizational authority. The no-silent-promotion rule requires every consequential production change to remain visible, evidence-linked, reviewable, and reversible.

The architecture also preserves critical boundaries. Technical performance is not equivalent to clinical usefulness; medication-task performance is not equivalent to medication safety; human presence is not equivalent to meaningful oversight; implementation is not evidence of benefit; and statistical fairness is not synonymous with health equity.

Continual learning may be appropriate for some systems and unnecessary or unsafe for others.

The proposed contribution is neither a validated standard nor a deployment manual. Its scientific value depends on whether future studies can test its propositions, compare governance strategies, quantify burden and benefit, and determine when controlled adaptation improves medication-support performance without creating new clinical, professional, organizational, or equity risks.

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

REFERENCES

1. Kelly CJ, Karthikesalingam A, Suleyman M, Corrado G, King D. Key challenges for delivering clinical impact with artificial intelligence. BMC Med. 2019;17(1):195. doi:10.1186/s12916-019-1426-2
2. 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
3. Finlayson SG, Subbaswamy A, Singh K, Bowers J, Kupke A, Zittrain J, et al. The clinician and dataset shift in artificial intelligence. N Engl J Med. 2021;385(3):283-6. doi:10.1056/NEJMc2104626
4. Davis SE, Embí PJ, Matheny ME. Sustainable deployment of clinical prediction tools—a 360° approach to model maintenance. J Am Med Inform Assoc. 2024;31(5):1195-8. doi:10.1093/jamia/ocae036
5. Embí PJ. Algorithmovigilance—advancing methods to analyze and monitor artificial intelligence-driven health care for effectiveness and equity. JAMA Netw Open. 2021;4(4). doi:10.1001/jamanetworkopen.2021.4622
6. Sahiner B, Chen W, Samala RK, Petrick N. Data drift in medical machine learning: Implications and potential remedies. Br J Radiol. 2023;96(1150):20220878. doi:10.1259/bjr.20220878
7. Duckworth C, Chmiel FP, Burns DK, Zlatev ZD, White NM, Daniels TWV, et al. Using explainable machine learning to characterise data drift and detect emergent health risks for emergency department admissions during COVID-19. Sci Rep. 2021;11(1):23017. doi:10.1038/s41598-021-02481-y
8. Rahmani K, Thapa R, Tsou P, Chetty SC, Barnes G, Lam C, et al. Assessing the effects of data drift on the performance of machine learning models used in clinical sepsis prediction. Int J Med Inform. 2023;173:104930. doi:10.1016/j.ijmedinf.2022.104930
9. 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
10. 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
11. 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
12. 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
13. Balendran A, Benchoufi M, Evgeniou T, Ravaud P. Algorithmovigilance, lessons from pharmacovigilance. NPJ Digit Med. 2024;7(1):270. doi:10.1038/s41746-024-01237-y
14. 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
15. 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
16. Davis SE, Greevy RA Jr, Lasko TA, Walsh CG, Matheny ME. Detection of calibration drift in clinical prediction models to inform model updating. J Biomed Inform. 2020;112:103611. doi:10.1016/j.jbi.2020.103611
17. Kore A, Abbasi Bavil E, Subasri V, Abdalla M, Fine B, Dolatabadi E, et al. Empirical data drift detection experiments on real-world medical imaging data. Nat Commun. 2024;15(1):1887. doi:10.1038/s41467-024-46142-w
18. 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
19. 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
20. Subasri V, Krishnan A, Kore A, Dhalla A, Pandya D, Wang B, et al. Detecting and remediating harmful data shifts for the responsible deployment of clinical AI models. JAMA Netw Open. 2025;8(6). doi:10.1001/jamanetworkopen.2025.13685
21. 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
22. Liu X, Cruz Rivera S, Moher D, Calvert MJ, Denniston AK, SPIRIT-AI and CONSORT-AI Working Group. Reporting guidelines for clinical trial reports for interventions involving artificial intelligence: The CONSORT-AI extension. Nat Med. 2020;26(9):1364-74. doi:10.1038/s41591-020-1034-x
23. Lacson R, Eskian M, Licaros A, Kapoor N, Khorasani R. Machine learning model drift: A systematic review. J Am Coll Radiol. 2022;19(10):1162-9. doi:10.1016/j.jacr.2022.05.030
24. Nagendran M, Chen Y, Lovejoy CA, Gordon AC, Komorowski M, Harvey H, et al. Artificial intelligence versus clinicians: Systematic review of design, reporting standards, and claims of deep learning studies. BMJ. 2020;368. doi:10.1136/bmj.m689
25. 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. 2025;388. doi:10.1136/bmj-2024-081554
26. 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
27. Corny J, Rajkumar A, Martin O, Dode X, Lajonchère JP, Billuart O. 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):1688-94. doi:10.1093/jamia/ocaa154
28. Strohm L, Hehakaya C, Ranschaert ER, Boon WPC, Moors EHM. Implementation of artificial intelligence applications in radiology: Hindering and facilitating factors. Eur Radiol. 2020;30(10):5525-32. doi:10.1007/s00330-020-06946-y
29. 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
30. Daye D, Wiggins WF, Lungren MP, Alkasab TK, Kottler N, Allen B Jr, et al. Implementation of clinical artificial intelligence in radiology: Who decides and how? Radiology. 2022;305(3):555-63. doi:10.1148/radiol.212151
31. Bate A, Hobbiger SF. Artificial intelligence, real-world automation and the safety of medicines. Drug Saf. 2021;44(2):125-32. doi:10.1007/s40264-020-01001-7
32. 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
33. 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
34. 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
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-