

Adaptive AI Change Control and Revalidation: A Risk-Based Lifecycle Framework for Pharmaceutical GMP, Medical Devices, and AI-Enabled Quality Systems

Dr. Pushpinder Paul

Director, QCE & Independent Researcher- Quality and Regulatory Compliance, Quality Compliance & Excellence (QCE), H.P., India.
Email: [pushparyan27\[at\]gmail.com](mailto:pushparyan27[at]gmail.com)

Abstract: Artificial intelligence (AI) is increasingly being integrated into pharmaceutical manufacturing, quality management systems (QMS), laboratory operations, process monitoring, medical-device software, and regulatory compliance activities. Unlike conventional deterministic software, adaptive artificial intelligence and machine-learning (ML) systems may change their outputs as data distributions evolve, models are retrained, parameters are updated, or operational environments change. These characteristics challenge traditional change-control and validation practices, which commonly assume that an approved system remains functionally stable until a formally authorised modification is introduced. In regulated environments, insufficient control of adaptive behavior may compromise product quality, patient safety, data integrity, traceability, and continued regulatory compliance. Conversely, indiscriminate full revalidation following every minor modification may create unnecessary operational burdens without proportionate improvements in assurance. This paper proposes an integrated, risk-based framework for adaptive AI change control and revalidation, combining conventional pharmaceutical and medical-device quality principles with model lifecycle management, continuous performance monitoring, change-impact assessment, and predefined acceptance criteria. The framework comprises eight interconnected elements: AI system classification, change detection, risk-based impact assessment, change authorization, tiered revalidation, controlled deployment, post-deployment monitoring, and lifecycle review. It distinguishes between software changes, data changes, model changes, changes in operational context, and emergent performance degradation. A proposed three-tier revalidation model—limited verification, targeted revalidation, and comprehensive revalidation—is presented alongside illustrative decision criteria, a change-control workflow, and practical applications in pharmaceutical manufacturing and AI-enabled medical devices. The framework is mapped conceptually to relevant quality risk management guidance, computerised-system controls, AI management standards, and medical-device regulatory approaches. It is a proposed framework rather than a prospectively validated method; empirical evaluation is required to establish its effectiveness. The approach aims to improve the proportionality, traceability, and scientific justification of AI lifecycle decisions while maintaining human accountability and patient protection.

Keywords: Adaptive artificial intelligence; machine learning; change control; revalidation; Software as a Medical Device

1. Introduction

1.1 Background

Artificial intelligence is transforming the way regulated organisations generate, interpret, and act on quality-related information. In pharmaceutical manufacturing, AI applications include predictive maintenance, process analytical technology support, deviation classification, out-of-specification investigation support, environmental monitoring, visual inspection, and predictive process control. In medical-device development and healthcare, AI-enabled software can support image interpretation, clinical decision support, diagnostic classification, and other functions that may influence clinical decisions. [1,2,6,8,9,11]

The performance of these systems depends not only on the original software implementation but also on model architecture, training data, input characteristics, operating conditions, and the relationships learned during development. A model that performs satisfactorily under its original validation conditions may become less reliable when the population, manufacturing process, sensor characteristics, data acquisition method, or clinical environment changes. [5,6,8,9]

Traditional change control provides a structured mechanism for documenting, evaluating, approving, implementing, and reviewing modifications. Conventional validation establishes documented evidence that a system or process meets specified requirements for its intended use. However, adaptive AI introduces an additional challenge: system behaviour can change even when the underlying software code remains unchanged. [1,2,7,9,11]

For example, an AI model used to predict equipment failure may exhibit deteriorating performance after a sensor replacement. A visual-inspection model may experience reduced sensitivity when lighting conditions or product appearance changes. An AI-enabled medical device may require renewed evaluation when retraining changes its classification thresholds or performance in a clinically relevant subgroup. [1,6,8,10]

These situations demonstrate that change control cannot be limited to software version updates. It must also address changes in data, model behavior, system configuration, operational context, and the assumptions underlying the original validation. [1,2,6,8]

Volume 15 Issue 9, September 2026

Fully Refereed | Open Access | Double Blind Peer Reviewed Journal

www.ijsr.net

1.2 Problem Statement

The central problem is the absence of a universally applicable operational method that connects adaptive AI behavior to proportionate, documented revalidation decisions across different regulated applications.

Three competing objectives must be balanced:

- 1) **Patient and product protection:** Prevent changes from introducing unacceptable safety, quality, or performance risks.
- 2) **Regulatory assurance:** Maintain traceability, validated state, data integrity, and compliance with applicable requirements.
- 3) **Operational efficiency:** Avoid unnecessary revalidation when objective evidence demonstrates that a change has no material effect on critical functions.

An approach based solely on fixed periodic validation may fail to detect deterioration between scheduled reviews. Conversely, automatic revalidation after every data refresh or minor software modification may consume resources without addressing the actual risk.

A lifecycle-oriented framework is therefore required to connect change detection, risk assessment, validation evidence, deployment decisions, and post-deployment monitoring.

1.3 Aim and objectives

The aim of this paper is to propose a risk-based framework for adaptive AI change control and revalidation in regulated pharmaceutical, medical-device, and quality-management environments.

The specific objectives are to:

- Identify the principal change mechanisms affecting adaptive AI systems.
- Integrate AI lifecycle risks into established quality risk management and change-control processes.
- Define a tiered revalidation approach based on intended use and potential impact.
- Establish decision criteria for continued use, restricted deployment, rollback, or comprehensive revalidation.
- Propose a traceable documentation structure and monitoring strategy.
- Identify a research methodology for future empirical evaluation of the framework.

1.4 Scope

This paper addresses AI systems used in regulated activities where their outputs may affect product quality, patient safety, regulatory decisions, or the integrity of quality records. It includes adaptive ML models and AI-enabled software, as well as AI tools used in pharmaceutical manufacturing and QMS processes.

The framework is intended to complement, not replace, applicable regulatory requirements, product-specific obligations, established validation procedures, or regulatory authorisation conditions.

2. Literature and Regulatory Context

2.1 Quality risk management

ICH Q9(R1) establishes principles for quality risk management in pharmaceutical development, manufacturing, distribution, and associated regulatory activities. It emphasizes that the level of effort, formality, and documentation should be commensurate with risk, and that quality risk assessments should be scientifically justified and linked to patient protection [1].

These principles are directly relevant to adaptive AI because changes differ in their potential consequences. A change affecting a non-critical administrative classification tool should not automatically receive the same validation intensity as a change to an AI model controlling a critical manufacturing parameter or supporting a high-risk clinical decision. [1,6,10]

A suitable AI change-control process should therefore consider the severity of potential harm, likelihood of failure, ability to detect failure, uncertainty associated with the change, and effectiveness of existing controls. Conventional failure mode and effects analysis (FMEA) can support this assessment, although its scores should not be treated as precise measurements of actual risk. [1,6,10]

2.2 Computerised systems in pharmaceutical GMP

EU GMP Annex 11 addresses computerised systems used in regulated pharmaceutical operations. Its control principles include lifecycle management, risk management, validation, change control, security, data integrity, and periodic evaluation [2].

These principles provide an established foundation for AI governance. However, an adaptive model may change operational performance without a conventional software release. Consequently, a risk-based AI procedure should extend the change inventory to include model retraining, data pipeline changes, model configuration changes, performance drift, and modifications to automated decision thresholds. [2,7,9,11]

A model's continuously updated state should not be assumed to remain validated merely because its software platform remains unchanged. [2,9,11]

2.3 Medical-device software and AI lifecycle management

The US Food and Drug Administration (FDA) issued final guidance in August 2025 on marketing submission recommendations for a Predetermined Change Control Plan (PCCP) for AI-enabled device software functions. The guidance describes planned modifications, the methodology for developing, validating, and implementing them, and the assessment of their impacts. An appropriately reviewed PCCP can provide a pathway for implementing specified modifications without a separate marketing submission for every covered change, subject to the applicable conditions [3]. [3,8,15]

This approach is important because it supports controlled iteration while retaining regulatory oversight. It does not establish that any adaptive update may be deployed automatically, nor does it remove the need to comply with applicable submission and change-reporting requirements. [3,8,15]

The FDA's January 2025 guidance on AI-enabled device software lifecycle management remains identified as draft guidance on the agency's guidance listing consulted for this paper [4]. Its recommendations should therefore be distinguished from binding legal requirements and final guidance. [4]

2.4 AI management systems and trustworthy AI

ISO/IEC 42001:2023 establishes requirements for an AI management system, including organizational processes for responsible AI development and use [5]. The NIST AI Risk Management Framework (AI RMF 1.0) provides a voluntary structure organized around four functions: Govern, Map, Measure, and Manage [6]. [5,6]

These resources reinforce the importance of defined accountability, risk identification, performance measurement, and continuing oversight. They are useful complements to sector-specific quality systems but do not independently establish pharmaceutical or medical-device regulatory compliance. [5,6,8]

2.5 Research gap

The sources reviewed provide relevant principles for quality risk management, computerized systems, AI governance, and controlled medical-device modifications. However, they do not constitute a single, universally applicable operational standard specifying when every type of adaptive AI change requires revalidation. [1–8,10–12]

This paper addresses that implementation gap through a proposed framework that links change categories, risk assessment, validation depth, deployment controls, and continuing monitoring. The framework is a synthesis of existing principles and a conceptual proposal; it is not presented as an empirically established standard or as proof that the literature has been exhaustively reviewed. [1–8,10–12]

3. Methodology

3.1 Research design

This study uses a conceptual framework-development approach supported by a targeted review of regulatory guidance and recognized AI and quality-management frameworks. [1,2,5,6,8]

The analysis follows four stages:

- 1) Identify regulatory and quality principles relevant to change control, validation, risk management, and AI lifecycle governance.
- 2) Categorize changes that can affect adaptive AI system behavior.

- 3) Develop a risk-based decision structure connecting each change category to validation and deployment controls.
- 4) Define implementation artifacts, illustrative applications, and measurable outcomes for future evaluation.

3.2 Analytical approach

The framework combines lifecycle analysis, qualitative risk assessment, and validation-impact reasoning. It distinguishes between evidence that a system has been correctly implemented and evidence that the resulting system remains suitable for its intended use. [1,6,8,9,11]

No original clinical trial, production dataset, or prospective industrial validation study was performed for this conceptual paper. Accordingly, the framework's expected benefits are hypotheses to be evaluated rather than measured findings.

3.3 Regulatory interpretation

Regulatory documents are interpreted according to their status and scope. Applicable legislation and binding requirements take precedence over voluntary standards and non-binding guidance. The precise obligations depend on the jurisdiction, product classification, intended use, system function, and regulatory pathway. [1–8,10–12]

4. Proposed Adaptive AI Change-Control Framework

4.1 Framework architecture

The proposed framework contains eight connected stages:

Stage 1- AI system classification: Identify intended use, users, critical outputs, affected processes, regulatory status, and potential consequences of failure. [1,2,5,6,8]

Stage 2-Change detection. Identify planned modifications and unplanned changes in data quality, input distributions, model performance, software dependencies, or operational conditions. [6,8,9]

Stage 3- Change-impact assessment: Evaluate effects on intended use, critical quality attributes, patient safety, product quality, data integrity, cybersecurity, bias, and existing risk controls. [1,6,8,10,12]

Stage 4- Change authorisation: Determine whether the change is permitted within the approved change envelope and establish the required independent review and regulatory assessment. [2,3,8,15]

Stage 5-Risk-based revalidation: Select the appropriate verification and validation activities using predefined criteria. [1,7,9,11,12]

Stage 6- Controlled deployment: Approve the model and its dependencies, use a staged deployment where appropriate, and retain rollback capability. [3,7,8,11]

Stage 7- Post-deployment monitoring: Evaluate performance, drift, safety signals, data quality, and compliance with established acceptance limits. [6,8,13]

Stage 8- Periodic lifecycle review: Review accumulated evidence, incidents, deviations, model updates, and continued fitness for intended use. [1,2,5,6,8]

The stages form a feedback loop rather than a one-time sequence. Monitoring may initiate a new change-control record, restrict model use, or trigger rollback.

4.2 Classification of AI changes

Changes should be classified by their origin and potential effect rather than by the size of the software modification alone. [1,6,8,11]

A change may belong to several categories. For example, retraining a model using data collected after a manufacturing equipment modification may involve simultaneous data, model, and operational changes. The combined effect should be assessed rather than evaluating each change in isolation. [1,6,8,10]

4.3 AI system criticality

A proposed criticality assessment should consider the following dimensions: [1,6,10]

- **Impact:** Potential severity of harm to patients, consumers, product quality, or regulated decisions.
- **Decision authority:** Whether AI provides information, recommends an action, or initiates an action automatically.
- **Detectability:** Whether an incorrect output is likely to be detected before causing harm.
- **Uncertainty:** The extent to which the change introduces unfamiliar data, model behavior, or operating conditions.
- **Exposure:** Frequency of use and number of affected products, batches, devices, or patients.
- **Reversibility:** The ability to restore a known, acceptable state promptly.

These dimensions should inform a documented risk assessment. They should not be combined into an unvalidated numerical formula and treated as an objective measure of safety. [1,6,10]

For high-criticality systems, unresolved uncertainty should lead to stronger evidence requirements, narrower deployment, or non-deployment until the relevant risk is controlled.

5. Risk-Based Revalidation Strategy

5.1 Principles of revalidation

Revalidation should establish that the modified system remains suitable for its intended use and that the change has not invalidated previous evidence. [1,2,7,9,11,12]

The scope should be determined by documented impact assessment and should include, as applicable: [1,7,9,11,12]

- Functional requirements and intended-use verification.
- Model performance against predefined acceptance criteria.
- Data integrity, lineage, and preprocessing controls.
- Robustness under relevant operating conditions.
- Bias and subgroup performance.
- Human–AI interaction and override mechanisms.
- Cybersecurity and software dependency controls.

- Interface and integration testing.
- Traceability between requirements, risks, tests, and results.
- Regulatory assessment and required approvals.

Revalidation is not synonymous with retraining. Retraining creates a candidate model; verification and validation establish whether that candidate can be used acceptably. [7,9,11]

5.2 Proposed Three-tier model

The following tiers are proposed for internal quality-system implementation. They are not official FDA, EMA, ISO, or ICH classifications. [1,6,8,10]

Tier 1: Limited verification

Appropriate when documented evidence indicates that the change is low impact and does not affect critical functionality, intended use, model behavior in a safety-relevant way, or essential controls.

Potential activities include configuration checks, regression testing of affected functions, interface testing, and verification of deployment artifacts.

Tier 2: Targeted revalidation

Appropriate when a change affects a defined subset of model behavior, data processing, operating conditions, or performance characteristics, but the overall intended use remains unchanged and risks can be adequately evaluated within a bounded scope.

Potential activities include testing on representative and independent datasets, subgroup analysis, robustness testing, comparison against the previously approved model, and reassessment of affected risk controls.

Tier 3: Comprehensive revalidation

Appropriate when a change materially affects intended use, model architecture, critical decision logic, key assumptions, risk controls, or performance in ways that cannot be adequately bounded by targeted testing. It may also be required when evidence is insufficient to establish that the existing validation remains applicable.

Activities may include requirements reassessment, updated hazard analysis, data and model qualification, comprehensive performance testing, integration testing, usability assessment, clinical or process-performance evaluation where applicable, and reassessment of regulatory submission obligations.

Comprehensive revalidation does not automatically mean that every activity must be repeated from the beginning. Previous evidence may be reused where its continuing applicability is scientifically justified and documented.

5.3 Revalidation Decision Matrix

The assigned tier is a starting point, not an automatic approval. A change that initially appears minor may require

escalation if testing reveals unexpected behavior or if multiple changes interact.

5.4 Acceptance Criteria [1,6,8,10,11]

Acceptance criteria should be defined before testing begins. They should reflect the intended use, known hazards, applicable regulatory obligations, and the performance demonstrated during initial validation.

Depending on the application, criteria may include:

- Sensitivity, specificity, precision, recall, or area under the receiver operating characteristic curve.
- Calibration error and predictive uncertainty.
- False-acceptance and false-rejection rates.
- Performance across clinically or operationally relevant subgroups.
- Process prediction error and control-limit excursions.
- Data completeness, latency, and availability.
- Rate of human overrides and unresolved alerts.
- Frequency and severity of AI-associated deviations.

A single aggregate metric should not conceal failure in a safety-critical subgroup or operating condition. Acceptance criteria must be selected according to the task; metrics that are suitable for one model may be unsuitable for another. [1,6,8,10]

6. Adaptive Behaviour and Continuous Monitoring

6.1 Distinguishing drift from a validated change

Adaptive AI systems can exhibit several types of drift: [6,8,13]

- **Data drift:** The distribution of incoming features changes.
- **Concept drift:** The relationship between inputs and outcomes changes.
- **Performance drift:** The observed predictive or operational performance deteriorates.
- **Operational drift:** Equipment, process, user, environmental, or workflow conditions change.
- **Label drift:** The prevalence or distribution of outcome labels changes. [6,8,13]

These phenomena are related but not interchangeable. A statistical shift in input data does not necessarily indicate unsafe performance, and apparently stable input data does not prove that the model remains effective.

Drift detection should therefore initiate investigation rather than automatically authorize retraining. Retraining without understanding the cause may reproduce existing errors, introduce new bias, or conceal an underlying process failure.

6.2 Monitoring architecture

A controlled monitoring system should capture, where relevant: [6,8,13]

- Input data quality and distribution.
- Model identity, version, configuration, and provenance.
- Prediction distributions and confidence measures.

- Performance against reference outcomes when reliable labels become available.
- Subgroup performance and bias indicators.
- Human overrides, complaints, deviations, and adverse-event signals.
- System availability, latency, cybersecurity, and integration failures.
- Drift alerts, investigation outcomes, and authorized interventions.

Monitoring thresholds should be justified and tested. A threshold breach should trigger a predefined response, such as intensified review, temporary restriction, use of a conventional fallback, rollback, or initiation of formal change control.

Where reliable outcome labels are delayed or unavailable, proxy measures may support surveillance but should not be treated as equivalent to demonstrated predictive accuracy. [5,6,8]

6.3 Controlled adaptation versus autonomous learning

A key governance decision is whether the AI system is permitted to update itself during operation. [3,4,8,15]

Controlled adaptation means that new data may be collected continuously, but candidate models are retrained, evaluated, approved, and released through a defined process.

Autonomous online adaptation means that model parameters or behavior may change during operational use without a separate conventional release event.

Autonomous adaptation presents additional challenges because the validated model state may change continuously. In a regulated environment, such adaptation should be permitted only when the applicable regulatory pathway and quality system allow it and when the adaptation mechanism, permitted boundaries, acceptance criteria, monitoring, intervention controls, and rollback strategy have been adequately established.

Where these conditions are not met, the safer approach is to separate continuous data collection from controlled model updates.

7. Governance, Responsibilities, and Documentation

7.1 Cross-functional governance

Adaptive AI change control should be governed by a cross-functional team whose composition reflects the application's risk. [1,5,6,12]

The model developer should not be the sole approver of a high-risk change. Independent review is particularly important when model updates affect product disposition, manufacturing control, diagnosis, or treatment decisions.

7.2 Minimum change-control record

A change-control record should contain: [1,2,5,7,9,11,12]

- Unique change-control identifier.
- AI system name and approved version.
- Intended use and criticality classification.
- Description and justification of the proposed change.
- Source and provenance of data or software modifications.
- Risk assessment and affected requirements.
- Predicted impact on performance and risk controls.
- Revalidation tier and documented rationale.
- Test protocols, datasets, acceptance criteria, and results.
- Deviations, anomalies, and investigation outcomes.
- Regulatory impact assessment.
- Approvals, deployment date, and rollback plan.
- Post-deployment monitoring plan.
- Final closure decision and references to related records.

For AI models, traceability should extend beyond the application version to the model artifact, training dataset version, preprocessing pipeline, relevant dependencies, configuration, and deployment environment.

7.3 Change history and reproducibility

The organization should be able to reconstruct which model produced a particular output and which approved data and configuration were used. This is essential for investigating quality events, responding to audits, and evaluating whether an earlier decision remains defensible. [5,6,8,11]

Where a third-party AI service is used, the organization should assess the supplier's controls, change notifications, service-version stability, data handling, and ability to provide evidence relevant to the regulated intended use.

A vendor's assurance statement does not by itself demonstrate that the configured AI system is suitable for the organization's specific application.

8. Implementation in Pharmaceutical Manufacturing

8.1 Predictive maintenance

Consider an AI model used to predict failures of a critical pharmaceutical manufacturing pump. [1,2,7,9]

A sensor replacement changes the distribution of vibration data. The model's predictions remain available, but its failure-warning performance has not been demonstrated under the new sensor conditions.

The proposed framework would require a documented change-impact assessment, confirmation of sensor compatibility, and comparison of model behavior against suitable reference data. If the sensor change affects critical failure detection, targeted revalidation may be required. If the model's failure-prediction capability cannot be established, the organization should use an appropriate approved alternative until acceptable performance is demonstrated.

The change should also be evaluated for effects on maintenance scheduling, equipment availability, and any associated product-quality risks.

8.2 AI-assisted visual inspection

An AI system used to detect defects in sterile product containers may be affected by changes in lighting, camera hardware, container appearance, or defect prevalence. [1,2,7,9]

A reduction in detection performance may increase the risk of defective units being accepted. The revalidation plan should therefore include representative defect classes, relevant operating conditions, and suitable reference assessments. Performance should be examined by defect type and other relevant subgroups rather than relying solely on aggregate accuracy.

Any change that affects the inspection strategy or the release of product must be assessed within the applicable validated process and quality-system controls.

8.3 AI-assisted deviation and CAPA classification

An AI tool may classify deviations by severity, identify recurring patterns, or recommend CAPA categories. If the model is updated using newly labeled records, it may change how similar events are classified. [1,2,5,6,7]

A targeted assessment should examine classification consistency, critical misclassification rates, and the effect on escalation and investigation workflows. The tool should not independently downgrade a potentially serious event unless such use is expressly justified, controlled, and permitted by the applicable procedures.

Human review should remain proportionate to the potential consequences of incorrect classification.

9. Implementation in AI-Enabled Medical Devices

9.1 Clinical performance and intended use

For AI-enabled medical devices, changes to training data, model weights, decision thresholds, input modalities, or target populations can affect clinical performance and the device's benefit–risk profile. [3,4,8,10,14]

The framework should require evaluation of whether the change remains within the established intended use and any applicable approved modification plan. Relevant testing may include performance against reference standards, clinically relevant subgroup evaluation, robustness to input variation, and usability or workflow assessment.

A technically improved model is not necessarily clinically safer. An increase in overall accuracy may coexist with reduced sensitivity for a clinically important subgroup.

9.2 Predetermined Change Control Plans

The FDA's final PCCP guidance provides a relevant regulatory example of planned and controlled AI modification [3]. A suitable plan identifies the modifications that are anticipated, explains the methods used to develop and validate them, and assesses their potential impacts. [3,8,15]

The proposed framework extends this logic operationally by connecting each permitted modification to monitoring triggers, risk-based validation evidence, deployment controls, and post-release review. [3,8,15]

However, the framework does not replace the PCCP or determine whether a particular change is covered by an approved plan. That determination must be made under the applicable regulatory requirements and the specific terms of the relevant authorization.

9.3 Post-deployment surveillance

For medical devices, post-deployment surveillance should capture relevant performance issues, complaints, incidents, and emerging risks. A model update should not be used to conceal or retrospectively erase the history of a performance failure. [3,4,8,10,14]

Where monitoring identifies a possible safety issue, the manufacturer should follow applicable incident-reporting, corrective-action, field-action, and regulatory notification requirements.

10. Proposed Standard Operating Procedure

The framework can be implemented through an AI lifecycle change-control SOP containing the following sequence. [1,2,3,7,9,11]

Step 1: Initiation: Record the proposed change or detected drift, affected system, justification, and responsible owner.

Step 2: Triage: Determine whether the event represents a planned change, an unplanned technical modification, a performance alert, a security event, or a potential quality incident.

Step 3: Containment: Where necessary, restrict use, suspend automated decisions, activate an approved fallback, or restore a previously approved model.

Step 4: Impact assessment: Evaluate the intended use, critical outputs, data lineage, model behaviour, risk controls, and applicable regulatory obligations.

Step 5: Approval of the test strategy: Assign the proposed revalidation tier, specify acceptance criteria, and approve the evidence plan before testing.

Step 6: Execution: Perform the authorized verification and validation activities. Record deviations and investigate unexpected results.

Step 7: Independent review: Confirm that requirements are met, residual risks are acceptable, documentation is complete, and required regulatory actions have been addressed.

Step 8: Controlled release: Approve deployment, confirm the correct model and configuration, and establish a documented rollback mechanism.

Step 9: Enhanced monitoring: Monitor the released change against predefined performance and risk thresholds.

Step 10: Closure and lifecycle review: Close the change only when required activities are complete and post-deployment responsibilities have been assigned. Reopen or initiate new change control if monitoring identifies unacceptable behavior.

The SOP should distinguish between change implementation and change closure. Where enhanced monitoring is necessary, the initial release decision and subsequent closure review may be separate documented events. [1,2,7,9,11]

11. Proposed Performance Indicators

The effectiveness of the framework should be assessed using predefined, measurable indicators.

These indicators should be interpreted together. A low incident rate alone does not prove effectiveness if incidents are under-reported, exposure is low, or the monitoring system cannot detect failures.

Similarly, a reduction in revalidation effort is not a successful outcome if it results from inadequate testing or inappropriate risk classification.

12. Discussion

12.1 Advantages of the proposed framework

The proposed framework integrates quality risk management with the technical characteristics of adaptive AI. Its principal advantage is the explicit link between the nature of a change, its potential consequences, the evidence required, and the decision to deploy. [1,5,6,8]

A tiered approach may reduce unnecessary repetition of unaffected tests while retaining comprehensive evaluation when intended use, critical functionality, or safety assumptions materially change.

The framework also recognizes that validation is not a permanent certificate of performance. It is a documented body of evidence whose continuing applicability depends on the system's configuration, operating context, and intended use. [1,2,6,8]

12.2 Limitations of fixed validation cycles

Periodic reviews remain important, but fixed intervals alone may not provide adequate control over rapidly changing systems. A critical performance decline can occur between scheduled reviews, while a low-impact software change may not justify comprehensive revalidation.

A hybrid approach is therefore proposed: event-triggered assessments for material changes and performance alerts, supported by periodic lifecycle reviews. [6,8,13]

12.3 Challenges in implementation

Implementation may be difficult where outcome labels are delayed, reference datasets are limited, suppliers do not disclose model details, or systems rely on third-party platforms that change without sufficient notice.

Organizations may also lack personnel with combined expertise in quality assurance, statistics, machine learning, cybersecurity, and regulatory affairs.

A further challenge is determining how much evidence is sufficient. Statistical significance alone does not establish clinical or manufacturing significance, and a model may satisfy an aggregate metric while failing in a critical subgroup.

These limitations require explicit acceptance criteria, suitable reference evidence, independent review, and escalation mechanisms rather than reliance on automated scores alone. [1,6,8,10]

12.4 Human oversight and accountability

AI systems can support decision-making but do not remove the organization's responsibility for regulated activities. The level of human oversight should depend on the system's intended use and potential consequences of failure.

For high-risk decisions, the quality system should define who can authorize deployment, restrict operation, accept residual risk, and initiate corrective action. An automated drift detector or model-monitoring platform should not be treated as a substitute for accountable review. [5,6,8,11]

12.5 Regulatory harmonization

ICH Q9(R1), EU GMP Annex 11, ISO/IEC 42001, NIST AI RMF, and FDA AI-device guidance address different but complementary aspects of risk, system lifecycle management, and governance [1–6].

The proposed framework provides a way to operationalize these concepts within a single internal process. It does not establish equivalence among the documents, and it does not imply that compliance with one framework satisfies all obligations under another. [1–8,10–12]

Organizations should maintain a jurisdiction-specific regulatory assessment and update it as relevant legislation, guidance, and standards evolve. [1–8,10–12]

13. Limitations and Future Research

This paper presents a conceptual framework derived from a targeted review of selected regulatory and technical resources. It does not report empirical validation, prospective deployment outcomes, or comparative performance data. [1–8,10–15]

The three-tier revalidation model requires evaluation across different AI applications and risk classes. Future studies should test whether the framework produces consistent decisions among independent reviewers, identifies material

changes reliably, and reduces unnecessary validation effort without increasing undetected failures. [6,8,13,14]

A multicentre implementation study could compare organizations using conventional change control with organizations implementing the proposed framework. Relevant outcomes could include change-assessment agreement, time to release, detection of material performance degradation, revalidation effort, and quality-event rates.

Further research should also address adaptive systems that update during operation, generative AI used in regulated documentation, foundation-model version changes, supplier-controlled model updates, and AI agents that execute actions within validated workflows. [4,5,6,8]

A practical next step would be to develop and prospectively evaluate a standardized AI change-impact assessment tool, together with a structured evidence package and an auditable mapping between identified risks and revalidation tests. [1,6,8,10,14]

14. Conclusion

Adaptive AI presents a significant challenge to conventional change-control and validation practices because system behavior may change through modifications to data, model parameters, software dependencies, configuration, or operational conditions. [1–8,10–12]

A sustainable regulatory approach must distinguish between changes that have negligible effects on validated functions and changes that may alter critical performance or introduce new risks. Revalidation should therefore be proportionate to the potential impact, supported by predefined acceptance criteria, traceable evidence, independent review, and controlled deployment. [1,2,3,6,8,10,11,12]

The framework proposed in this paper combines eight lifecycle stages with a three-tier revalidation strategy, continuous monitoring, documented escalation, and rollback controls. It is designed to support pharmaceutical GMP, medical-device quality management, and other regulated AI applications while preserving the principles of patient protection, product quality, data integrity, and accountability. [1–8,10–12]

The framework should be regarded as a conceptual proposal requiring empirical assessment rather than an already validated industry standard. Its practical value will depend on consistent implementation, appropriate system-specific acceptance criteria, reliable monitoring, and evidence that proportional revalidation maintains assurance while avoiding unnecessary testing.

Ultimately, the objective of adaptive AI change control is not to prevent all change, but to ensure that every material change is understood, justified, tested, authorized, traceable, and monitored throughout the system lifecycle. [1,2,3,6,8,10,11,12]

References

- [1] International Council for Harmonisation of Technical Requirements for Pharmaceuticals for Human Use (ICH). *ICH Harmonised Guideline: Quality Risk Management Q9(R1)*. Step 4, 18 January 2023. Geneva: ICH; 2023. Available at: <https://www.ema.europa.eu/en/scientific-guidelines/ich-q9-quality-risk-management>
- [2] European Commission. *EudraLex, Volume 4: EU Guidelines for Good Manufacturing Practice for Medicinal Products for Human and Veterinary Use, Annex 11: Computerised Systems.* Brussels: European Commission; 2011. Available at: https://health.ec.europa.eu/document/download/8d305550-dd22-4dad-8463-2ddb4a1345f1_en?filename=annex11_01-2011_en.pdf
- [3] US Food and Drug Administration (FDA). *Marketing Submission Recommendations for a Predetermined Change Control Plan for Artificial Intelligence-Enabled Device Software Functions: Guidance for Industry and Food and Drug Administration Staff.* August 2025. Available at: <https://www.fda.gov/regulatory-information/search-fda-guidance-documents/marketing-submission-recommendations-predetermined-change-control-plan-artificial-intelligence>
- [4] US Food and Drug Administration (FDA). *Artificial Intelligence-Enabled Device Software Functions: Lifecycle Management and Marketing Submission Recommendations: Draft Guidance for Industry and Food and Drug Administration Staff.* January 2025. Available at: <https://www.fda.gov/regulatory-information/search-fda-guidance-documents/artificial-intelligence-enabled-device-software-functions-lifecycle-management-and-marketing>
- [5] International Organization for Standardization (ISO) and International Electrotechnical Commission (IEC). *ISO/IEC 42001:2023 Information Technology—Artificial Intelligence—Management System.* Geneva: ISO; 2023. Available at: <https://www.iso.org/standard/81230.html>
- [6] Tabassi E. *Artificial Intelligence Risk Management Framework (AI RMF 1.0)*. NIST AI 100-1. Gaithersburg, MD: National Institute of Standards and Technology; 2023. <https://doi.org/10.6028/NIST.AI.100-1>
- [7] US Food and Drug Administration (FDA). *Computer Software Assurance for Production and Quality Management System Software: Guidance for Industry and Food and Drug Administration Staff.* February 2026. Available at: <https://www.fda.gov/regulatory-information/search-fda-guidance-documents/computer-software-assurance-production-and-quality-management-system-software>
- [8] International Medical Device Regulators Forum (IMDRF). Good Machine Learning Practice for Medical Device Development: Guiding Principles. IMDRF/AIML WG/N88 FINAL:2025. 29 January 2025. Available at: <https://www.imdrf.org/documents/good-machine-learning-practice-medical-device-development-guiding-principles>
- [9] US Food and Drug Administration (FDA). General Principles of Software Validation: Final Guidance for Industry and FDA Staff. 11 January 2002. Available at: <https://www.fda.gov/regulatory-information/search-fda-guidance-documents/general-principles-software-validation>
- [10] International Organization for Standardization (ISO). ISO 14971:2019 Medical Devices—Application of Risk Management to Medical Devices. Geneva: ISO; 2019. Available at: <https://www.iso.org/standard/72704.html>
- [11] International Electrotechnical Commission (IEC). IEC 62304:2006+AMD1:2015 CSV Medical Device Software—Software Life Cycle Processes. Geneva: IEC; 2015. Available at: <https://webstore.iec.ch/>
- [12] International Organization for Standardization (ISO). ISO 13485:2016 Medical Devices—Quality Management Systems—Requirements for Regulatory Purposes. Geneva: ISO; 2016. Available at: <https://www.iso.org/standard/59752.html>
- [13] US Food and Drug Administration (FDA), Health Canada, and Medicines and Healthcare products Regulatory Agency (MHRA). Good Machine Learning Practice for Medical Device Development: Guiding Principles. 27 October 2021. Available at: <https://www.fda.gov/medical-devices/software-medical-device-samd/good-machine-learning-practice-medical-device-development-guiding-principles>
- [14] International Medical Device Regulators Forum (IMDRF). Software as a Medical Device (SaMD): Clinical Evaluation. IMDRF/SaMD WG/N41 FINAL:2017. 21 September 2017. Available at: <https://www.imdrf.org/documents/software-medical-device-samd-clinical-evaluation>
- [15] US Food and Drug Administration (FDA), Health Canada, and Medicines and Healthcare products Regulatory Agency (MHRA). Predetermined Change Control Plans for Machine Learning-Enabled Medical Devices: Guiding Principles. October 2023. Available at: <https://www.fda.gov/medical-devices/software-medical-device-samd/predetermined-change-control-plans-machine-learning-enabled-medical-devices-guiding-principles>

Appendix A: Framework Tables

1) Change Category

Illustrative example	Principal concern	Principal concern
Data change	New training dataset or revised preprocessing	Representativeness, quality, bias, leakage
Model change	Retraining, architecture update, weight modification	Performance, generalisation, reproducibility
Configuration change	Decision threshold or hyperparameter adjustment	Sensitivity, specificity, false decisions
Software change	Library, interface, API, or runtime update	Functional compatibility and cybersecurity
Infrastructure change	Cloud platform, processor, deployment environment	Reproducibility, latency, availability
Intended-use change	New patient group, product type, or clinical task	New hazards and regulatory implications
Operational drift	Sensor degradation or changing input distribution	Loss of performance in actual use
Governance change	Revised human oversight or access permissions	Accountability and control effectiveness

2) Change scenario

Proposed initial tier	Required decision considerations	Required decision considerations
Cosmetic dashboard modification with no effect on calculations	Tier 1	Confirm unchanged calculations, permissions, and outputs
Routine dependency patch	Tier 1 or 2	Assess compatibility, security, and functional effects
Retraining using new representative data	Tier 2 or 3	Assess dataset quality, performance, drift, and risk
Change to a critical decision threshold	Tier 2 or 3	Assess false-positive and false-negative consequences
New intended use or patient population	Tier 3	Assess new hazards, evidence needs, and regulatory pathway
Material decline in critical model performance	Restrict use; assess Tier 2 or 3	Investigate cause and restore acceptable performance
Unauthorised model-weight change	Quarantine or rollback	Establish provenance, investigate, and revalidate before use
Change affecting an AI-controlled critical process parameter	Tier 2 or 3	Assess process capability, product quality, and control strategy

3) Function

Principal responsibility	Responsibility
Quality Assurance	Change-control governance, independent review, and release approval
AI/ML engineering	Model development, technical assessment, and reproducibility
Data governance	Data provenance, quality, representativeness, and access controls
Process owner or clinical expert	Intended-use assessment and operational relevance
Validation specialist	Test strategy, evidence review, and validation status
Regulatory Affairs	Regulatory impact assessment and submission obligations
IT and cybersecurity	Infrastructure, access control, security, and recovery
Management	Resources, accountability, and acceptance of residual risk

4) Indicator

Proposed measurement	Purpose	Purpose
Change-control completeness	Percentage of changes with complete required records	Documentation quality
Risk-assessment timeliness	Time from initiation to approved assessment	Process responsiveness
Revalidation effectiveness	Percentage of released changes meeting predefined acceptance criteria	Validation quality
Post-release incident rate	Relevant incidents per defined exposure period	Safety and quality outcomes
Drift detection performance	Detection delay and false-alert frequency	Monitoring effectiveness
Rollback effectiveness	Percentage of required rollbacks completed within approved response time	Recovery capability
Unauthorised model-change rate	Number of detected changes outside approved controls	Governance effectiveness
Revalidation burden	Effort and elapsed time by risk tier	Operational efficiency
Repeat deviation rate	Recurrence of related AI change-control failures	Corrective-action effectiveness