

AI-Assisted Deviation Investigations in Pharma

Deviation management in pharmaceutical quality systems

In pharmaceutical manufacturing and laboratory operations, a **deviation** is best understood operationally as a **departure from an approved or established requirement**—most commonly a written procedure, instruction, specification, or defined control. In U.S. finished-pharmaceutical CGMP, this concept is explicit: production and process control procedures must be followed and documented contemporaneously, and **any deviation from the written procedures must be recorded and justified.** ¹

Deviation management exists because deviations are one of the most direct signals that a process may have drifted, a control may have failed, or execution may not have matched the validated/approved intent. Modern pharmaceutical quality system frameworks explicitly treat deviations as key inputs into CAPA and continuous improvement. **ICH Q10** describes a CAPA system that implements corrective and preventive actions resulting from investigations of, among other sources, **deviations**; it also calls for a structured approach to investigations with effort and documentation commensurate with risk (aligned to ICH Q9). ² The EU GMP “Pharmaceutical Quality System” chapter similarly states that product and process monitoring results should be considered in batch release and in the **investigation of deviations**, with preventive action to avoid potential deviations. ³

Common deviation “types” in practice usually fall into a few operational buckets (knowing that exact categorizations vary by company and SOP):

- **Procedure/SOP adherence deviations** (steps not performed, performed out of sequence, use of unapproved procedure). ¹
- **Process parameter excursions** (documented parameter out of range, alarms, process interruptions, yield anomalies). ⁴
- **Material and component issues** (wrong status material used, expired/rejected component used). ⁵
- **Equipment/utilities/environmental events** (equipment failures, calibration/qualification gaps, environmental monitoring excursions, water system issues). ⁶
- **Documentation and data issues** (missing entries, late entries, corrections needing justification). ⁷

Planned vs unplanned deviations

Many regulated organizations distinguish **unplanned deviations** (unexpected departures) from **planned deviations** (a deliberate, one-time departure from an existing SOP/method/batch record due to an unforeseen situation) that is still expected to be **fully documented and justified**. A WHO deviation handling guidance used in vaccine/biologics contexts describes “planned deviation” in this way and notes that one-time planned deviations are typically documented and justified, while **permanent changes** should move into change control. ⁸

Minor, major, and critical classifications

Severity/criticality classification is widely used, but it is not globally standardized. The WHO deviation handling guidance provides examples of **minor**, **major**, and **critical** deviations using risk-focused criteria (impact to quality attribute/critical parameter and probability of patient impact), and it explicitly warns that classification should be based on objective, justified criteria and that examples may be categorized differently with proper justification. ⁸ This “risk-based, justified classification” posture is consistent with **ICH Q9(R1)**, which emphasizes that the level of effort, formality, and documentation should be commensurate with the level of risk. ⁹

Deviation lifecycle and how regulated systems expect it to work

A “full” deviation lifecycle is a controlled progression from detection through investigation, CAPA linkage, and closure, with lifecycle evidence retained and retrievable.

Identification, initiation, and immediate containment

At the front end, deviation systems aim to:

- capture an objective description of what happened (time, location, who discovered it, and what evidence exists), and
- implement **immediate actions (“corrections”)** to contain the issue and protect product/patient while the deeper investigation proceeds. ⁸

The same WHO deviation handling guidance provides an operational structure: minor deviations typically require description, correction, and documented efficacy/conclusion; major/critical deviations add batch disposition (if applicable), root cause investigation, CAPA, and effectiveness assessment of corrective action. ⁸

Impact assessment and scope determination

Regulators repeatedly treat **scope** as a core competency: the investigation must determine what lots/products/processes may be affected, not merely what happened in one record. U.S. CGMP explicitly requires that any unexplained discrepancy or failure to meet specifications be thoroughly investigated and that the investigation **extend to other batches** and potentially other drug products associated with the discrepancy, with a written record including conclusions and follow-up. ¹⁰ The WHO deviation guidance similarly encourages “horizontal” analysis—evaluating possible impacts on other lots or similar manufacturing processes in addition to “vertical” analysis for root cause. ⁸

Investigation, root cause analysis, CAPA linkage, closure, and effectiveness follow-up

At closure, regulators expect more than narrative: they expect a defensible chain from evidence → cause(s) → corrective/preventive action(s) → verification of effectiveness. **ICH Q10** is explicit that a structured investigation should aim to determine root cause and that formality/documentation should align with risk. ¹¹ The WHO deviation guidance likewise ties root cause investigation to objective evidence for corrective and preventive actions and explicitly calls for documented efficacy/effectiveness checks. ¹²

Record retention and retrievability are not optional features. U.S. CGMP requires that records be **readily available for authorized inspection during the retention period**, and retained as originals or true copies.

13

QA governance and decision rights in deviation systems

Quality assurance's role in deviations is not just "review" in a general sense; it is a set of defined authorities and oversight responsibilities that must be clear in procedures and demonstrable in records.

In U.S. CGMP, the quality control unit (often practically the "quality unit") must have the **responsibility and authority** to approve or reject materials and drug products and also the **authority to review production records** to assure no errors have occurred or, if they have, that they have been fully investigated. 14 This links QA/QCU directly to investigation adequacy—not only to record completion.

Operationally, QA involvement commonly concentrates in:

- **Triage and classification governance:** ensuring consistent severity classification and ensuring classification aligns with risk rationale rather than individual bias. The WHO deviation guidance explicitly warns against "natural bias" and pushes objective criteria. 8
- **Quality review and approval of immediate actions:** for many deviations, especially those affecting lots, QA approval of corrections prior to release is expected; the WHO deviation guidance describes QA approval expectations for corrections and lot-related decisions. 8
- **Oversight of investigation quality:** ensuring investigations are thorough, unbiased, evidence-based, and appropriately scoped. FDA's investigation expectations are visible in both regulation (scope extension requirements) and enforcement (frequent citations of inadequate investigations). 15
- **CAPA linkage and effectiveness verification:** ensuring corrective actions are tied to identified causes and that effectiveness is verified and documented (a theme repeated in CAPA/quality-system expectations). 16
- **Trend review and management review inputs:** deviations are explicitly part of management review inputs and monitoring systems in ICH Q10, and EU GMP links monitoring outcomes to deviation investigations and preventive action. 17

This is where AI often becomes attractive: QA is asked to maintain consistency, scope discipline, and trend insight across very large volumes of quality events. The challenge is achieving those benefits without shifting accountability away from humans or undermining record integrity expectations.

Investigation tools and evidence expectations

Deviation investigations are expected to be systematic, evidence-driven, and commensurate with risk.

Common investigation techniques in pharma practice

The WHO deviation guidance explicitly identifies **5 Whys** and the **Ishikawa (fishbone) diagram** as among the simplest and most commonly used root cause tools in deviation investigations, and describes 5 Whys as iterative questioning to uncover underlying/systemic causes. 8

Beyond these, pharma teams frequently use tools that are also recognized within quality risk management methods. **ICH Q9(R1)** lists and discusses methods including: - **Cause-and-effect (Ishikawa/fishbone) diagrams**,¹⁸ - **FMEA/FMECA**, and¹⁹ - **Fault Tree Analysis (FTA)**, which ICH Q9 describes as a structured approach to mapping causal chains and notes can be used to investigate complaints or **deviations** to understand root cause and ensure intended improvements fully resolve the issue (not “fix one thing, cause another”).²⁰

ICH Q9(R1) also underscores that modern “digitalization and emerging technologies” and “advanced data analysis methods” can reduce risk when fit for intended use, but can also introduce new risks that must be controlled—an important bridge to AI-based investigation support.⁹

Vertical and horizontal analysis

A frequent investigation failure mode is focusing only on the immediate event without evaluating impact elsewhere. The WHO deviation guidance explicitly pairs “vertical” analysis (root cause identification) with “horizontal” analysis (possible impact on other lots/products/processes).⁸ U.S. CGMP similarly requires investigations to extend to other batches and potentially other products associated with a discrepancy.¹⁰

What regulators mean by “scientifically sound” investigations

FDA’s OOS investigations guidance (while focused on laboratory OOS results) is widely used as a benchmark for investigation rigor because it states that investigations should be **thorough, timely, unbiased, well-documented, and scientifically sound**.²¹ It also emphasizes that an investigation is necessary even if a batch is rejected and that the investigation is needed to determine whether there are implications for other batches/products, echoing the broader CGMP scope principle.²²

This matters for deviations because many deviation investigations fail for the same reasons OOS investigations fail: weak hypotheses, unsupported “human error” conclusions, missing scope extension, and CAPA that does not address true causal factors.

Systemic weaknesses regulators repeatedly cite in deviation and investigation programs

Regulators’ enforcement language makes clear that weak investigations are not a documentation nuisance—they are a direct threat to reliable manufacturing and patient safety.

A recurring FDA warning letter message is that **inadequate investigations lead to unidentified root causes, ineffective CAPA, and recurring problems that compromise the ability to manufacture safe and effective products**. FDA has used this framing explicitly in warning letters such as *Glycerinas Industriales* (January 2024) and has repeated similar language in later letters.²³

Across warning letters, several patterns stand out as especially relevant to “AI-assisted deviation investigations” because these are exactly the gaps organizations often hope AI will help reduce:

Failure to extend investigation scope beyond the immediate record

FDA frequently criticizes firms for not extending investigations to other batches/products impacted by an event. In the *Excelvion Fareva* warning letter, FDA states the firm did not extend investigations to other batches to determine full scope and impact and did not implement effective CAPA to prevent recurrence.

²⁴ FDA also highlights scope extension failures in other domains such as environmental monitoring excursions, where a firm failed to extend investigations to other batches produced under potentially impacted conditions. ²⁵ These enforcement themes align directly with the explicit CGMP requirement in 21 CFR 211.192 that investigations extend to other batches and potentially other products. ²⁶

Investigations lacking sufficient detail or scientific rationale for root cause

FDA warning letters frequently cite investigations that “lack sufficient details” or do not establish root cause. For example, in *Tower Laboratories* FDA described investigations that lacked sufficient details and lacked root causes for failures. ²⁷ In other warning letters, FDA explicitly cites investigations into recurring problems that lacked adequate root-cause determination and/or sufficient scientific rationale. ²⁸ This is consistent with FDA’s OOS guidance warning that results should not be attributed to analytical error without completing an investigation that clearly establishes a laboratory root cause—an example of the broader “do not conclude without evidence” principle. ²¹

CAPA that does not prevent recurrence or lacks effectiveness verification

When regulators identify repeated events or repeated failures, they often conclude that CAPA is not effective or that root causes are not truly addressed. FDA warning letters frequently request that firms improve CAPA effectiveness and quality unit oversight when investigation systems are deficient. ²⁹

“Program-level” concern: calls for independent assessment of the investigation system

A strong signal that regulators see systemic weakness (not a single deviation failure) is a demand for a comprehensive, independent reassessment of the entire investigation system—covering deviations, discrepancies, complaints, OOS, and failures—with improvements in investigation competency, scope determination, root-cause evaluation, CAPA effectiveness, and quality unit oversight. FDA has requested this type of remediation in multiple letters. ³⁰

These patterns matter because they define where AI can help most (triage, scope discipline, recurrence detection, completeness checks) and where it can be dangerous (substituting for scientific judgment or becoming a “black box” root-cause engine).

Where AI can support deviation investigations and what it realistically requires

AI can support deviation investigations, but the most defensible uses are those that:

- improve **signal detection, consistency, or review efficiency,**
- remain **assistive** rather than decisive, and
- preserve clear human accountability for conclusions and approvals.

This framing aligns with regulators' emerging AI principles. FDA's and EMA's jointly developed "Good AI Practice" principles emphasize concepts like risk-based approaches, data governance/documentation, performance assessment, lifecycle management, and human-centric design. ³¹ FDA's AI credibility draft guidance further reinforces a **risk-based credibility assessment** tied to a defined "context of use," a concept directly relevant to deciding whether an AI model is merely drafting text versus influencing a regulated decision. ³²

Below are high-value AI use cases mapped to deviation management needs, with realism and boundaries.

Clustering similar deviations and detecting recurrence patterns

Intended value: Improve horizontal analysis by identifying similar events (same equipment, step, material, failure mode) and highlighting repeated "minor" deviations that may be trending toward systemic issues. This directly supports the scope discipline regulators demand and can reduce the risk of "single-record tunnel vision." ³³

Human review needed: High. Similarity matching should be treated as a lead list; humans decide whether the events are truly linked and whether scope expansion is required. ³⁴

Data requirements: A well-structured deviation dataset (date/time, site, line, product, batch, step, equipment IDs, deviation category, disposition) plus NLP over narrative fields and attachments where appropriate. ICH Q9's description of advanced data analysis needing fit-for-intended-use design and control is relevant here. ⁹

Key limitations: Inconsistent taxonomy and narrative-heavy records reduce clustering quality; poor data standardization produces false linkages or misses true ones.

AI-assisted triage and severity classification support

Intended value: Recommend an initial classification (minor/major/critical) and investigation pathway based on risk and similarity to historical classified events, improving consistency and reducing "natural bias." ⁸

Human review needed: Very high. Severity classification is a regulated decision input (investigation depth, release impact); AI should propose, not decide. WHO explicitly frames classification as objective and justified, not arbitrary. ⁸

Data requirements: Historical labeled deviations and a stable classification framework; linkage to risk assessments (ICH Q9) increases defensibility. ⁹

Key limitations: AI will learn existing bias if prior classifications are inconsistent or overly lenient/harsh; "model drift" can change recommendations over time.

Investigation "gap detection" and completeness checking

Intended value: Flag missing core investigation elements (timeline missing, batch impact not assessed, CAPA not linked, effectiveness checks not defined) before closure, reducing the frequency of regulator-observed "insufficient detail" findings. ³⁵

Human review needed: Moderate to high. AI can detect omissions against a checklist, but humans must determine adequacy and scientific validity. ³⁶

Data requirements: A defined "required elements" template per deviation type/severity and a controlled mapping to which elements are mandatory for closure. WHO provides a practical structure that can be translated into such templates. ⁸

Key limitations: AI may over-flag legitimate exceptions; over-strict checklists can create low-value bureaucracy.

Summarizing historical records and building an evidence timeline

Intended value: Reduce time spent reading long narratives and attachments; generate consistent summaries to support management review, trending, and cross-site comparisons. ³⁷

Human review needed: High for any summary that enters the controlled record. FDA’s investigation expectations for being well-documented and scientifically sound imply summaries must be verified against source evidence. ³⁸

Data requirements: Access to the controlled record set (deviation record, attachments, batch record excerpts, lab results). Record availability for inspection is a CGMP requirement, which also implies the summary and underlying sources must be retrievable. ³⁹

Key limitations: Generative AI can hallucinate or omit critical nuance; this risk rises when the model is not tightly grounded in the controlled record set. NIST’s generative AI risk profile emphasizes that generative systems introduce novel risks that require governance and testing. ⁴⁰

Suggesting potential root-cause categories and “next best questions”

Intended value: Support investigators by proposing candidate cause categories (equipment, measurement system, process execution, materials, environment, human factors) and recommending follow-up questions or data to collect—especially helpful for less experienced investigators.

Human review needed: Extremely high. Root cause is ultimately a scientific conclusion, and regulators penalize unsupported conclusions. FDA has criticized investigations that lack scientific rationale for root cause determination. ²⁸

Data requirements: A well-curated, high-quality training corpus of past investigations with verified root causes and CAPA outcomes; structured metadata; linkage to risk assessments and failure modes. ICH Q9 explicitly ties effective risk management to appropriate root cause analysis and causal factors, including human-related factors. ⁴¹

Key limitations: AI can amplify investigator bias if presented as “the likely cause.” The safest design is to present multiple hypotheses with confidence qualifiers, and to require explicit human rationale for acceptance/rejection.

Trend forecasting and proactive risk signals across sites/products/steps

Intended value: Identify emerging signals (increasing recurrence, drift in specific steps, suppliers, or equipment families) earlier than periodic manual trending. This supports ICH Q10 monitoring and CAPA systems and EU GMP’s preventive-action orientation. ¹⁷

Human review needed: High. Forecasts should drive prioritization and preventive action planning, not definitive conclusions without verification. ⁹

Data requirements: Harmonized taxonomies across sites, stable identifiers, sufficient event volume, and consistent data quality; otherwise forecasts are noise.

Key limitations: Shifts in documentation behavior can masquerade as real changes; model drift and changing processes can break historical comparability.

Compliance and risk controls for AI-assisted deviation investigations

AI introduces specific risks that QA must control so the investigation system remains credible and inspection-ready.

The core risk: turning AI suggestions into “the conclusion”

Many AI use cases are attractive exactly where regulators are strict: scope determination, root cause justification, and CAPA effectiveness. That is also where AI is most dangerous if it becomes a substitute for evidence-based reasoning. FDA warning letters show that unsupported root causes, poor scope extension, and ineffective CAPA are central enforcement themes. ⁴²

A defensible posture is: **AI outputs are leads, drafts, or prioritization signals—not investigation conclusions.** This aligns with FDA’s AI credibility framework concept (“context of use” and risk-based credibility) and NIST’s AI RMF focus on governance, validity, and documentation. ⁴³

Bias amplification and missing contextual nuance

AI can “learn” prior misclassifications, habitual root-cause shortcuts, or incomplete narratives. The WHO deviation guidance explicitly notes natural bias in classification and the need for objective criteria—exactly the kind of vulnerability AI can reinforce if training data reflects bias. ⁸

Validation and change control expectations grow with impact

If AI is embedded in a regulated computerized system (eQMS), the system must still meet electronic record/e-signature control requirements where applicable, including audit trails, access control, and validation expectations. 21 CFR Part 11 requires controls for closed systems such as operational checks, authority checks, and (importantly) secure audit trails for record changes. ⁴⁴ EU GMP Annex 11 likewise requires audit trails to be available, intelligible, and **regularly reviewed**, with reasons documented for GMP-relevant data changes/deletions. ⁴⁵

The compliance challenge with AI is that model updates, prompt changes, or retrieval-corpus changes can change outputs. FDA’s Good AI Practice principles explicitly emphasize lifecycle management and data governance/documentation, which implies organizations must manage evolving behavior under change control and maintain evidence of performance and limits. ⁴⁶

Documentation and inspection readiness for AI-enabled workflows

If AI is used in deviation workflows, inspection readiness typically requires being able to show: - the AI’s **intended use** and boundaries (drafting vs decision support vs automated action), ³²
- how outputs are reviewed and approved by humans (maintaining accountable decision-making), ⁴⁷
- what data the AI accessed and how confidentiality is protected, and
- how outputs and the underlying records remain retrievable for inspection (record retention and availability). ³⁹

A practical QA control concept that helps here is to treat AI-assisted deviation outputs as “**pre-decisional work products**” unless and until they are reviewed, corrected, and approved under controlled procedures.

AI tools that can support deviation investigations in real QA operations

Below are three realistic tool families that directly align to deviation investigation needs (triage, summarization, recurrence detection, and trend analytics). The emphasis is on tools that can plausibly operate inside or closely alongside regulated quality systems.

Sparta Systems ⁴⁸ and TrackWise AI

Best use cases: Deviation similarity detection, recurring issue detection, auto-summarization of quality events, and trend/pattern discovery across large volumes of quality records. The TrackWise AI product description explicitly highlights auto-categorization, auto-summarization, insights (correlations/trends/patterns), and a “Root Cause Advisor” that correlates similar anomalies using historical data. ⁴⁹

Strengths: Strong alignment to deviation workloads (classification consistency, faster synopses, and cross-record signal detection); explicit vendor positioning that AI augments decision-making and requires human approval, with audit trails reflecting human accountability. ⁴⁹

Limitations: Vendor claims of broad applicability still require buyer-side governance: QA must define what becomes part of the controlled record vs an internal aid, and must manage the risk that recommendations influence decisions without adequate evidence. FDA enforcement shows that investigation adequacy and scope discipline are frequent failure points; AI must not become a substitute for those fundamentals. ⁵⁰

Best for: Workflow support + text analysis + recurring-issue analytics (especially when the QMS is the system of record). ⁵¹

Veeva Systems ⁵² and Veeva AI Agents in the Vault ecosystem

Best use cases: In-platform assistance inside regulated content/workflows for investigations and CAPA-adjacent documentation, particularly if an organization already uses Vault-based quality systems.

Veeva describes Veeva AI Agents as agentic AI built into the Vault Platform with application-specific prompts/safeguards and direct access to application data/documents/workflows; it also states that safety and quality agents are planned for release in 2026. ⁵³

Strengths: The strongest practical advantage is **being embedded in the system that already enforces roles, permissions, workflows, and record structures**, which can reduce “shadow record” risk and support auditable use if configured correctly. ⁵⁴

Limitations: Agentic tooling can quickly cross from drafting support into decision influence; if outputs become relied upon for severity classification, root cause acceptance, or closure decisions, credibility evidence and lifecycle change control become more demanding. FDA’s AI credibility draft guidance and

Good AI Practice principles point directly to the need for defined context of use, risk-based credibility, and lifecycle management. ⁵⁵

Best for: Workflow support + controlled drafting assistance in established, validated systems. ⁵⁶

MasterControl ⁵⁷ and GxPAssist AI

Best use cases: Investigation drafting support, document summarization (including change summaries between versions), and productivity support for QA documentation-heavy work that surrounds deviations (SOP updates, investigation narratives, evidence summaries).

MasterControl’s announcement for its AI document summarizer emphasizes “assistive,” “low-risk,” “human review and modifications before acceptance,” and references a suite including Document Summarizer, Document Translator, and Exam Generator. ⁵⁸

Strengths: Clear emphasis on human-in-the-loop review; strong fit for reducing manual administrative load in documentation and review cycles that often slow deviation closure. ⁵⁸

Limitations: These tools are not inherently “root cause engines”; their value depends on governance and grounding (ensuring that summaries/drafts reflect the controlled record content and do not introduce hallucinated facts). NIST’s generative AI profile highlights the risk profile of generative outputs; QA must treat AI drafts as drafts. ⁵⁹

Best for: Drafting and summarization support around investigations and CAPA, especially in document-heavy environments. ⁶⁰

Comparison table for deviation-focused QA work

Tool family	Best for	Strengths	Limitations	Fit: text vs analytics vs workflow	Regulated-environment fit
TrackWise AI (Sparta Systems)	Recurrence detection, similarity matching, auto-summaries, trends	Built around QMS quality-event data; explicit auto-categorization/ auto-summarization/ insights/root-cause advisory concepts; human-approval emphasis	Requires strong taxonomy and governance; risk of over-trusting recommendations	Strong across text + analytics + workflow	Strong if used as assistive with clear human accountability and validated governance

Tool family	Best for	Strengths	Limitations	Fit: text vs analytics vs workflow	Regulated-environment fit
Veeva AI Agents (Veeva Systems)	In-system assistance in Vault quality workflows	Embedded access to controlled data/workflows with prompts/safeguards; planned quality agents in 2026	Agentic scope creep into decision influence; lifecycle change control burden	Strong in workflow + drafting	Strong when bounded to assistive use in validated workflows
GxPAssist AI (MasterControl)	Drafting and summarization around investigation documentation	Human review emphasized; accelerates document summaries and change summaries	Not a substitute for scientific investigation; hallucination risk if not grounded	Strong in text + drafting	Strong for draft-only use with controlled approval steps

Practical conclusion for QA teams

AI can meaningfully improve deviation investigation performance, but only if it is implemented to strengthen (not substitute for) the investigation fundamentals regulators already enforce.

Where AI can help deviations most today: - **Triage acceleration and consistency:** suggesting routing, checklists, and initial categories while humans decide final severity and required investigation depth. ⁶¹

- **Horizontal analysis support:** clustering similar deviations, identifying recurring patterns, and automatically proposing “other batches/products to assess” as prompts—directly targeting a frequent FDA enforcement weakness (scope). ⁶²

- **Drafting and summarization:** producing structured summaries of complex investigations to reduce cycle time, while retaining human verification and approval. ⁶³

- **Gap detection before closure:** checklist-based screening that warns when investigations lack required elements or when CAPA/effectiveness steps are missing—reducing “insufficient detail” closures. ⁶⁴

Where heavy human review remains essential: - **Root cause conclusions and scientific justification.** FDA enforcement repeatedly targets unsupported root causes and investigations lacking adequate rationale. AI may suggest hypotheses, but humans must test them against evidence. ⁶⁵

- **Scope decisions with product-quality implications.** Regulators require investigations to extend to other batches/products where relevant; AI can prompt, but QA must decide and document rationale. ⁶⁶

- **CAPA effectiveness decisions.** Effectiveness verification is frequently a regulator focus; AI can track and remind, but cannot replace accountable approval that actions worked. ⁶⁷

Controls that keep AI from undermining compliance: - Define **context of use**: “AI drafts; humans decide,” and document that boundary. ³²

- Keep investigation outputs attributable and auditable: if AI operates within the eQMS record lifecycle, ensure Part 11 / Annex 11-aligned controls (audit trails, access control, review expectations) remain effective. ⁶⁸

- Manage AI change like any other quality-impacting change: governance, testing, monitoring, and controlled updates consistent with Good AI Practice principles and risk management. ⁶⁹

A balanced conclusion is that AI **can** materially strengthen deviation investigations—especially by improving consistency, trend visibility, and scope discipline—if the organization uses it as assisted intelligence (not automated judgment), and if QA retains ownership of evidence, rationale, and the final decisions that regulators hold the firm accountable for. ⁷⁰

¹ ⁷ 21 CFR 211.100 -- Written procedures; deviations.

https://www.ecfr.gov/current/title-21/chapter-I/subchapter-C/part-211/subpart-F/section-211.100?utm_source=chatgpt.com

² ¹¹ ¹⁶ ¹⁷ ³⁷ Pharmaceutical Quality System Q10

https://database.ich.org/sites/default/files/Q10%20Guideline.pdf?utm_source=chatgpt.com

³ GMP Guide Chapter 1 Q10 implementation final

https://health.ec.europa.eu/system/files/2016-11/vol4-chap1_2013-01_en_0.pdf

⁴ ¹⁰ ¹⁵ ²⁶ ³³ ³⁴ ⁶² ⁶⁶ eCFR :: 21 CFR 211.192 -- Production record review.

<https://www.ecfr.gov/current/title-21/chapter-I/subchapter-C/part-211/subpart-J/section-211.192>

⁵ ⁸ ¹² ⁵⁷ ⁶¹ Risk Analysis as part of Deviation management

https://dcmn.org/wp-content/uploads/2016/03/who_guidance_deviation_and_risk_mgt_2013.pdf

⁶ ²⁵ Empower Clinic Services, LLC dba Empower Pharma

https://www.fda.gov/inspections-compliance-enforcement-and-criminal-investigations/warning-letters/empower-clinic-services-llc-dba-empower-pharma-700962-04022025?utm_source=chatgpt.com

⁹ ¹⁸ ¹⁹ ²⁰ ⁴¹ database.ich.org

https://database.ich.org/sites/default/files/ICH_Q9%28R1%29_Guideline_Step4_2022_1219.pdf

¹³ ³⁹ 21 CFR 211.180 -- General requirements.

https://www.ecfr.gov/current/title-21/chapter-I/subchapter-C/part-211/subpart-J/section-211.180?utm_source=chatgpt.com

¹⁴ eCFR :: 21 CFR 211.22 -- Responsibilities of quality control unit.

<https://www.ecfr.gov/current/title-21/chapter-I/subchapter-C/part-211/subpart-B/section-211.22>

²¹ ²² ³⁶ ³⁸ ⁶³ Investigating Out-of-Specification (OOS) Test Results for Pharmaceutical Production

<https://www.fda.gov/media/158416/download>

²³ ⁴² ⁷⁰ Glicerinas Industriales, S.A. de C.V. - 665472 - 01/16/2024

https://www.fda.gov/inspections-compliance-enforcement-and-criminal-investigations/warning-letters/glicerinas-industriales-sa-de-cv-665472-01162024?utm_source=chatgpt.com

²⁴ ⁵⁰ Excelvision Fareva MARCS-CMS 703245 — May 07, 2025

https://www.fda.gov/inspections-compliance-enforcement-and-criminal-investigations/warning-letters/excelvision-fareva-703245-05072025?utm_source=chatgpt.com

27 35 48 64 **Tower Laboratories Ltd. - 717456 - 12/23/2025**

https://www.fda.gov/inspections-compliance-enforcement-and-criminal-investigations/warning-letters/tower-laboratories-ltd-717456-12232025?utm_source=chatgpt.com

28 52 65 **Shiva Analyticals Private Limited - 707857 - 07/23/2025**

https://www.fda.gov/inspections-compliance-enforcement-and-criminal-investigations/warning-letters/shiva-analyticals-private-limited-707857-07232025?utm_source=chatgpt.com

29 67 **Unipack LLC - 716621 - 12/19/2025**

https://www.fda.gov/inspections-compliance-enforcement-and-criminal-investigations/warning-letters/unipack-llc-716621-12192025?utm_source=chatgpt.com

30 **Absolutely Natural - 715690 - 12/30/2025**

https://www.fda.gov/inspections-compliance-enforcement-and-criminal-investigations/warning-letters/absolutely-natural-715690-12302025?utm_source=chatgpt.com

31 46 69 **Guiding Principles of Good AI Practice in Drug Development**

https://www.fda.gov/media/189581/download?utm_source=chatgpt.com

32 43 55 **Considerations for the Use of Artificial Intelligence**

https://www.fda.gov/regulatory-information/search-fda-guidance-documents/considerations-use-artificial-intelligence-support-regulatory-decision-making-drug-and-biological?utm_source=chatgpt.com

40 59 **Artificial Intelligence Risk Management Framework**

https://nvlpubs.nist.gov/nistpubs/ai/NIST.AI.600-1.pdf?utm_source=chatgpt.com

44 68 **21 CFR 11.10 -- Controls for closed systems.**

https://www.ecfr.gov/current/title-21/chapter-I/subchapter-A/part-11/subpart-B/section-11.10?utm_source=chatgpt.com

45 **Annex 11: Computerised Systems**

https://health.ec.europa.eu/system/files/2016-11/annex11_01-2011_en_0.pdf?utm_source=chatgpt.com

47 49 51 **TrackWise AI for Life Sciences Quality Management - Honeywell**

<https://www.spartasystems.com/qualitywise-ai/>

53 54 56 **Veeva AI Agents to Be Released Across All Veeva Applications | Veeva**

<https://www.veeva.com/resources/veeva-ai-agents-to-be-released-across-all-veeva-applications/>

58 60 **MasterControl Announces General Availability of AI-Powered Document Summarizer**

<https://www.mastercontrol.com/news/mastercontrol-announces-general-availability-of-ai-powered-document-summarizer/>