AI for Pharmaceutical Risk Management: Applying Artificial Intelligence Within ICH Q9 Frameworks

A Practical GMP Perspective on AI-Assisted Quality Risk Management, FMEA, CAPA, and Predictive Risk Analytics

Pharmaceutical quality risk management is not a paperwork exercise. It is the structured process by which companies identify what could go wrong, estimate how serious it could be, decide whether controls are adequate, and review whether the risk profile has changed over time.

ICH Q9(R1) provides the core framework for this work. It describes quality risk management as a systematic process for the assessment, control, communication, and review of risks to drug product quality across the lifecycle. It also emphasizes two primary principles: risk evaluation should be based on scientific knowledge and linked to patient protection, and the level of effort, formality, and documentation should be proportional to the level of risk (ICH, 2023).

Artificial intelligence can support this framework by helping teams identify hazards earlier, analyze trends across large datasets, retrieve similar historical events, detect recurring risk patterns, and support more consistent risk scoring. But AI should not become the risk owner. In GMP, risk acceptance, risk control, deviation decisions, CAPA decisions, batch disposition, and regulatory conclusions must remain under qualified human oversight.

The realistic opportunity is not "AI replaces ICH Q9." The realistic opportunity is AI makes ICH Q9 implementation more evidence-based, more consistent, and more proactive.

Overview of ICH Q9: The Foundation for Pharmaceutical Quality Risk Management

ICH Q9(R1) is the internationally recognized guideline for quality risk management in pharmaceutical development, manufacturing, distribution, inspection, and regulatory review. It provides a systematic approach to identify, analyze, evaluate, control, communicate, and review risks to product quality (ICH, 2023).

The general ICH Q9 process includes:

ICH Q9 Element	Practical GMP Meaning
Risk assessment	Identify hazards, analyze risk, and evaluate risk against criteria
Risk control	Decide whether risk is acceptable or must be reduced
Risk communication	Share risk information among stakeholders and decision-makers
Risk review	Reassess risks as new knowledge, data, or events emerge
Risk management tools	Use structured methods such as FMEA, FTA, HACCP, HAZOP, PHA, and risk ranking

ICH Q9(R1) lists hazard identification, risk analysis, and risk evaluation as parts of risk assessment. It frames the three fundamental risk assessment questions as: what might go wrong, what is the likelihood it will go wrong, and what are the consequences (ICH, 2023).

AI can support each of these questions, but the answers must remain grounded in process knowledge, product knowledge, historical data, and patient risk.

Why Pharmaceutical Risk Management Often Struggles

Quality risk management is widely used in pharma, but it is not always applied well. Many risk assessments become static templates, subjective scoring exercises, or justification tools for decisions that were already made.

Weakness	GMP Impact
Generic hazard lists	Important site-specific failure modes may be missed
Subjective scoring	Different teams score the same risk differently
Poor historical data use	Prior deviations, CAPAs, complaints, and failures are not considered
Overreliance on FMEA numbers	Risk Priority Number becomes more important than scientific rationale
Weak detectability logic	Teams assume detection controls are stronger than they are
Risk assessments not updated	New deviations or changes do not trigger risk review
Missing cross-functional input	QA, validation, engineering, QC, manufacturing, and RA perspectives are incomplete
Poor linkage to CAPA	Risk assessment identifies issues, but controls are not implemented
Copy-paste assessments	Risk documents repeat old language without current evidence
Excessive formality for low risk	Teams waste resources on low-value assessments

ICH Q9(R1) directly recognizes the problem of subjectivity and states that subjectivity can affect the effectiveness of risk management activities and decisions; it also says subjectivity should be managed and minimized (ICH, 2023).

AI can help reduce some subjectivity by retrieving relevant evidence and highlighting historical patterns. But AI can also introduce new subjectivity if models are poorly trained, poorly validated, or overtrusted.

Where AI Fits Within the ICH Q9 Framework

AI can be mapped directly to the ICH Q9 process.

ICH Q9 Step	AI Support Possibility	Human Responsibility
Hazard identification	Retrieve similar events, detect recurring deviations, scan SOPs and process data	Confirm actual hazards and relevance
Risk analysis	Estimate likelihood using historical trends, failure rates, complaints, alarms, OOT results	Verify data quality and scientific rationale
Risk evaluation	Compare risk patterns against defined criteria and prior assessments	Decide acceptability based on patient/product risk
Risk control	Suggest possible controls from similar CAPAs or best practices	Select and approve controls
Risk communication	Generate summaries, dashboards, and escalation prompts	Communicate approved conclusions
Risk review	Monitor new events, model drift, CAPA effectiveness, process trends	Reassess and approve risk changes

The safest model is:

```
AI identifies evidence and patterns
        ↓
SMEs evaluate technical meaning
        ↓
QA challenges GMP adequacy
```

```
        ↓
Cross-functional team agrees risk conclusion
        ↓
Authorized decision-maker accepts or controls risk
        ↓
Risk profile is reviewed over time
```

AI should make the risk assessment more informed, not less accountable.

AI-Assisted Hazard Identification

Hazard identification asks: what might go wrong?

This is one of the strongest AI use cases because hazard identification often fails when teams rely only on memory, generic templates, or limited departmental experience.

AI can support hazard identification by searching:

- Deviations
- CAPAs
- Complaints
- OOS/OOT investigations
- Batch records
- APR/PQR reports
- Change controls
- Maintenance records
- Calibration failures
- Environmental monitoring excursions
- Audit observations
- Supplier issues
- Regulatory commitments
- SOP revisions
- Training failures

Risk Assessment Topic	AI-Retrieved Hazard Examples
New filling line component	Prior component jams, fill-weight variability, line stoppages, EM interventions
Extended hold time	Prior bioburden trends, stability data, batch deviations, temperature excursions
Cleaning process change	Prior residue failures, swab OOTs, detergent compatibility, operator errors
New supplier	Prior material deviations, complaint trends, incoming testing failures, change controls
LIMS calculation update	Prior calculation errors, audit trail findings, invalid reports, method changes
AI tool implementation	Hallucinated outputs, wrong recommendations, data integrity gaps, model drift

This supports the ICH Q9 expectation that hazard identification can use historical data, theoretical analysis, informed opinions, and stakeholder concerns (ICH, 2023).

The AI should not create a final hazard list alone. It should provide a stronger starting point for cross-functional review.

AI-Assisted Risk Analysis

Risk analysis asks: how likely is the risk, how severe is the harm, and how detectable is the failure?

ICH Q9(R1) describes risk analysis as estimating the risk associated with identified hazards, including the qualitative or quantitative process of linking likelihood of occurrence and severity of harms; in some tools, detectability also factors into risk estimation (ICH, 2023).

AI can support risk analysis by evaluating historical patterns such as:

Risk Factor	AI Data Source
Occurrence	Deviation frequency, complaint rate, equipment failure rate, OOT trend
Severity	Product impact history, batch rejection, patient risk, regulatory impact
Detectability	Alarm performance, in-process controls, review timing, audit trail review
Recurrence	Repeat deviations, failed CAPA effectiveness, recurring supplier issues
Control weakness	Missed inspections, overdue PMs, training gaps, procedure ambiguity
Trend direction	Increasing frequency, worsening process capability, rising complaint category

For example, instead of a team guessing that occurrence is “low,” AI may show that the same failure mode occurred five times in 18 months across three products, with two events requiring deviation investigations. That evidence does not automatically determine the score, but it improves the scoring discussion.

AI-Assisted Risk Evaluation

Risk evaluation asks: is this risk acceptable compared with defined criteria?

ICH Q9(R1) states that risk evaluation compares identified and analyzed risk against given risk criteria and considers the strength of evidence for the fundamental risk questions. It also emphasizes that robustness of the dataset matters because it determines the quality of the output, and that assumptions and uncertainty should be revealed (ICH, 2023).

AI can help by showing:

- Evidence supporting the risk score
- Similar historical risk assessments
- Previous risk acceptance decisions
- Related CAPAs or controls
- Data gaps and uncertainty
- Conflicting evidence
- Changes since the last risk review

AI Output	Human Review Question
“Three similar events occurred in the last year.”	Are they truly comparable?
“No complaints linked to this defect were found.”	Are complaint codes reliable?
“Detection depends on manual review.”	Is manual review timely and effective?
“Prior CAPA did not prevent recurrence.”	Should occurrence or detectability score increase?
“Dataset incomplete before 2023.”	How does uncertainty affect risk conclusion?

This is powerful because it forces the team to discuss uncertainty instead of hiding it behind a score.

AI-Assisted Risk Scoring

AI can support risk scoring, but this is also where caution is needed.

Risk scores in FMEA or similar tools are often based on severity, occurrence, and detection. AI can suggest scores based on historical data, but humans must approve final values.

Risk Dimension	AI Input	Human Decision
Severity	Prior product impact, patient risk category, batch rejection history	QA/medical/SME confirms severity
Occurrence	Frequency of similar deviations, complaints, failures	Team confirms comparability
Detection	Alarm history, IPC controls, audit trail review, sampling plan	SME confirms effectiveness
Risk priority	Score calculation and ranking	Team determines action threshold
Uncertainty	Missing data, weak metadata, conflicting records	Team decides whether conservative scoring is needed

AI should never silently auto-score GMP risks without rationale. If AI suggests a score, it should show the evidence behind that suggestion.

FMEA Example: AI-Supported Risk Assessment for a Filling Line

Process Step / Failure Mode	Potential Effect and Existing Control	AI-Detected Evidence	Scores	Proposed Action
Stopper feeding - stopper jam causes intervention	Increased contamination risk and rejects. Existing control: operator monitoring and intervention SOP.	6 similar minor stoppages in 3 months; two linked to the same component lot.	S=8, O=5, D=4, RPN=160	Engineering review, supplier lot review, and intervention trend monitoring.
Fill volume control - pump drift	Underfill/overfill risk. Existing control: in-process weight checks.	Fill weight variability increasing but still within limit.	S=7, O=4, D=3, RPN=84	Review pump maintenance and CPV trend.
Line clearance - prior label remains on line	Mislabeling risk. Existing control: QA line clearance.	One prior near-miss in the same packaging room.	S=9, O=2, D=3, RPN=54	Add visual aid and targeted retraining.
EM monitoring - viable excursion missed	Contamination risk. Existing control: routine EM sampling.	No direct signal, but recent room pressure alarms occurred.	S=10, O=2, D=5, RPN=100	Review pressure alarm trend and EM recovery data.

This example shows the correct AI role: it provides evidence and pattern detection. The team still owns the scoring and actions.

Predictive Risk Analytics

Predictive risk analytics uses historical and real-time data to identify risks that may emerge before a failure occurs.

Predictive Signal	Possible Risk
Increasing minor deviations in same process step	Future major deviation
Rising equipment alarm frequency	Equipment failure or process interruption
Slight worsening in process capability	Loss of process control
Repeated near-miss documentation errors	Data integrity or training issue
Complaint category increasing in one market	Product, packaging, or distribution issue

CAPA recurrence after closure	Ineffective corrective action
Supplier OOT results increasing	Material variability risk
Calibration drift increasing	Measurement reliability risk

This is where AI can make risk management more proactive. ICH Q9(R1) states that effective and proactive quality risk management can provide a means to identify and control potential quality issues during development, manufacturing, and distribution (ICH, 2023).

Predictive analytics is not a replacement for defined controls. It is an additional layer of risk visibility.

Trend-Based Risk Identification

Many risks are not visible in a single event. They emerge as trends.

AI can identify trend-based risks across systems:

Trend Source	AI Use
Deviations	Detect recurring root causes, repeated process steps, weak CAPAs
Complaints	Identify early quality signals and product experience patterns
CAPAs	Detect recurrence after CAPA closure
Maintenance	Link equipment failures with deviations or yield loss
Calibration	Detect drift patterns and measurement system risks
Training	Link repeat errors with training gaps or poor comprehension
Batch records	Detect repeated interventions, late entries, yield variability
CPV/APR/PQR	Identify drift, variability, or capability deterioration
Supplier quality	Detect material variability or recurring defects

ICH Q10 connects quality risk management with pharmaceutical quality system elements such as process performance and product quality monitoring, CAPA, change management, and management review (ICH, 2008). AI can strengthen those links by making signals searchable and trendable across systems.

Integration With Deviations and CAPAs

Deviation Integration

AI can help assess whether a deviation changes a known risk profile:

Deviation Finding	Risk Management Question
Same deviation occurred again	Should occurrence score increase?
Deviation escaped detection until final review	Should detectability score worsen?
Product impact greater than expected	Should severity be reassessed?
New failure mode identified	Should hazard list be updated?
Similar events in other products	Should risk assessment scope expand?

FDA's 21 CFR 211.192 requires unexplained discrepancies and batch failures to be thoroughly investigated, including extension to other batches of the same drug product and other drug products that may have been associated with the failure (FDA, 21 CFR 211.192). That requirement aligns strongly with AI-assisted similar-event retrieval and risk review.

CAPA Integration

AI can help evaluate whether CAPAs reduce risk as intended.

CAPA Type	AI Effectiveness Support
SOP revision	Monitor whether related deviations decrease
Training	Analyze recurrence by role, shift, or department
Equipment modification	Track alarms, failures, rejects, and deviations
Supplier corrective action	Monitor incoming defects and complaint patterns
Additional verification	Assess whether detection improved
Process change	Monitor CPV and APR/PQR trends

FDA's quality systems guidance describes CAPA as a CGMP concept focused on investigating, understanding, correcting discrepancies, and attempting to prevent recurrence (FDA, 2006). AI can support that by detecting recurrence more consistently.

Case Study 1: AI Identifies a Hidden Risk in Cleaning Validation

A site performs a risk assessment for adding a new product to a shared equipment train. The initial FMEA focuses on solubility, potency, toxicity, and cleaning method capability.

AI retrieves historical cleaning deviations, swab failures, and change controls. It identifies that a similar product with the same excipient had repeated visual residue issues after drying, even though analytical residue results were acceptable.

The team updates the hazard list to include visual residue persistence and drying behavior. Additional cleaning verification and operator inspection training are added.

Lesson: AI helps hazard identification by retrieving relevant historical knowledge that may not be obvious from the product formulation alone.

Case Study 2: AI Detects Increasing Risk in a Stable Process

A compression process has no OOS results and no batch rejections. Traditional risk assessment remains unchanged.

AI reviews CPV data and finds that tablet hardness variability has increased over the last 20 batches. It also identifies a rise in minor adjustments during compression and a recent change in tooling maintenance frequency.

The risk assessment is reopened. Occurrence score is increased for compression variability, and engineering initiates a tooling review.

Lesson: AI supports risk review by identifying drift before specifications fail.

Case Study 3: AI Challenges a Low Occurrence Score

A team scores occurrence as "rare" for a packaging line mix-up risk. AI searches deviation history and identifies four similar near-misses over two years, all coded differently.

The team revises occurrence from low to medium and adds a barcode verification enhancement.

Lesson: AI can reduce subjectivity when risk scores are based on incomplete memory.

Case Study 4: AI Creates a False Correlation

An AI dashboard suggests a relationship between a supplier change and increased complaints. QA reviews the data and finds that complaint volume increased because distribution volume doubled in the same market. Complaint rate per unit distributed did not increase.

The AI signal is documented as reviewed with no confirmed quality risk.

Lesson: AI can identify signals, but humans must normalize and interpret the data scientifically.

Risk Analysis: AI in ICH Q9 Workflows

AI Use Case	Potential Risk	GMP Impact	Required Control
AI suggests hazards	Misses critical hazard	Incomplete risk assessment	SME review and approved hazard checklist
AI retrieves historical events	Retrieves irrelevant events	Inflated or distorted scoring	Human relevance review
AI suggests risk scores	Overreliance on numerical output	Poor risk decision	Evidence-linked scoring and QA approval
AI predicts recurrence	False positive	Unnecessary controls or investigations	Threshold tuning and SME triage
AI misses trend	False sense of control	Delayed risk review	Traditional trending remains active
AI drafts FMEA	Generic or incorrect failure modes	Weak control strategy	Cross-functional review
AI supports CAPA effectiveness	Wrong recurrence linkage	Incorrect CAPA closure	QA confirmation
AI model changes	Different outputs over time	Loss of validated state	Change control and model versioning
AI uses poor data	Misleading conclusions	Incorrect risk acceptance	Data quality assessment

The safest principle is: AI may inform risk decisions, but it should not independently accept risk.

Validation Considerations for AI Risk Management Tools

If AI is used in a GMP risk management workflow, validation expectations depend on intended use.

Lower-risk uses include search similar risk assessments, retrieve prior deviations, suggest possible hazards, draft risk assessment summaries, and provide trend dashboards.

Higher-risk uses include assigning risk scores, recommending risk acceptance, routing CAPA requirements, triggering change control decisions, supporting batch disposition, and automatically updating risk registers.

FDA Part 11 applies to electronic records and signatures created, modified, maintained, archived, retrieved, or transmitted under FDA records requirements, and requires controls such as validation, access restriction, audit trails, authority checks, and record protection for closed systems (FDA, 21 CFR Part 11).

Validation should address:

Validation Area	Practical Question
Intended use	What GMP decisions can the AI influence?
Source data	Which systems feed the AI: QMS, LIMS, MES, CMMS, DMS, LMS?
Data quality	Are historical records accurate, complete, and consistently coded?
Output traceability	Can users see evidence behind AI suggestions?
Model control	Is the model/version/configuration controlled?
Performance testing	Can the AI retrieve known historical risks and events?

False negatives	Does it miss important hazards or trends?
False positives	Does it overload teams with irrelevant risk signals?
Human override	Can users reject AI suggestions with rationale?
Audit trail	Are AI outputs and human decisions retained where required?
Periodic review	Is model performance reviewed over time?

A peer-reviewed article on validation of AI-containing products across regulated healthcare industries notes that AI/ML introduces validation challenges across pharmaceuticals, medical devices, and diagnostics and highlights the need to align terminology and validation approaches across people and processes (Higgins & Johner, 2023).

Regulatory Concerns

Regulators are unlikely to object to companies using better tools for risk management. The concern will be whether AI is controlled, validated, scientifically justified, and appropriately supervised.

Likely inspection questions include:

Regulatory Question	Expected Company Response
What does the AI tool do?	Clear intended use statement
Does AI make final risk decisions?	No; qualified humans approve conclusions
How was the tool validated?	Risk-based validation evidence
What data does it use?	Controlled source list and data flow map
Are outputs traceable?	Source-linked evidence and audit trail
How are model changes controlled?	Change control and revalidation assessment
How are false outputs handled?	Deviation/system issue process if GMP impact occurs
How is performance reviewed?	Periodic review and metrics
How is subjectivity managed?	Cross-functional review and evidence-based rationale

EMA's AI reflection paper emphasizes that AI/ML risk depends on technology, data quality, context of use, and the degree of influence the AI exerts, and that developers and users should manage AI systems across their lifecycle using a risk-based approach (EMA, 2024).

That principle fits ICH Q9 perfectly: the more influence AI has over risk conclusions, the stronger the controls must be.

Practical Workflow: AI-Assisted ICH Q9 Risk Assessment

```
Risk question defined
        ↓
AI retrieves relevant historical data:
deviations, CAPAs, complaints, CPV, audits, changes
        ↓
AI suggests possible hazards and similar events
        ↓
Cross-functional team reviews relevance
        ↓
Team performs risk analysis:
severity, occurrence, detectability, uncertainty
        ↓
AI provides evidence-linked scoring support
        ↓
Team evaluates risk against criteria
        ↓
Risk controls selected and approved
        ↓
```

```
Risk communication documented
        ↓
AI monitors new data for risk review triggers
        ↓
Periodic human review confirms risk profile
```

This workflow keeps ICH Q9 intact while improving evidence retrieval and ongoing review.

Implementation Roadmap

Step 1: Start With Historical Retrieval

Use AI first to retrieve similar deviations, CAPAs, complaints, changes, and risk assessments. This is practical and lower risk.

Step 2: Define Approved Data Sources

Connect only controlled and trusted sources such as QMS, DMS, LIMS, MES, CMMS, LMS, APR/PQR repositories, and regulatory commitment trackers.

Step 3: Standardize Risk Taxonomy

Harmonize risk categories, failure modes, root cause codes, CAPA types, severity definitions, occurrence bands, and detectability criteria.

Step 4: Define Intended Use

Document whether AI is used for hazard suggestions, scoring support, trend monitoring, CAPA review, or risk register maintenance.

Step 5: Validate Based on Risk

Test the AI tool using historical cases where known hazards, failure modes, and similar events are already established.

Step 6: Require Source-Linked Output

AI should not simply state, “risk is high.” It should show the deviations, complaints, data trends, CAPAs, and assumptions behind its suggestion.

Step 7: Update Procedures

Risk management SOPs should describe acceptable AI use, human review expectations, documentation requirements, and escalation rules.

Step 8: Train Users

Train QA, validation, manufacturing, QC, engineering, RA, and process owners on both ICH Q9 and AI limitations.

Step 9: Monitor Performance

Track missed hazards, false signals, accepted/rejected AI suggestions, repeat deviations, CAPA recurrence, and risk assessment quality.

Step 10: Scale Gradually

After proving value in advisory use cases, expand into predictive risk dashboards and risk register monitoring.

FAQ: AI for Pharmaceutical Risk Management Under ICH Q9

Can AI perform an ICH Q9 risk assessment?

AI can support an ICH Q9 risk assessment by retrieving evidence, suggesting hazards, identifying similar events, and highlighting trends. It should not independently perform or approve the final risk assessment.

Can AI assign FMEA scores?

AI can suggest severity, occurrence, or detectability scores based on historical data, but qualified humans should approve final scores. Scores must be justified with scientific rationale and evidence.

What is the best first use case?

The best first use case is AI-assisted hazard identification and similar-event retrieval. This helps teams avoid missing known risks while keeping humans in control.

Can AI reduce subjectivity in risk assessments?

Yes, AI can reduce subjectivity by retrieving historical evidence and showing patterns. However, AI can also introduce bias if trained on poor historical data, so human review and data governance are essential.

Does AI risk management require validation?

If AI is used in GMP workflows or influences regulated risk decisions, it should be validated based on intended use and risk. Part 11 may apply if electronic records or signatures are involved.

Can AI help with risk review?

Yes. AI can monitor deviations, CAPAs, complaints, calibration failures, maintenance events, CPV data, and APR/PQR trends to identify when a risk assessment may need review.

What is the biggest risk?

The biggest risk is overreliance. If teams accept AI-generated scores or conclusions without understanding the evidence, the risk assessment may become less scientific, not more scientific.

Conclusion: AI Can Strengthen ICH Q9, but It Cannot Replace Quality Judgment

AI for pharmaceutical risk management is one of the most practical applications of artificial intelligence in GMP quality systems. ICH Q9 already provides the right structure: hazard identification, risk analysis, risk evaluation, risk control, risk communication, and risk review. AI can strengthen that structure by improving evidence retrieval, trend detection, predictive analytics, and ongoing risk review.

But AI does not remove GMP accountability. ICH Q9(R1) makes clear that quality risk management should be based on scientific knowledge, linked to patient protection, and proportional to the level of risk. It also emphasizes the need to manage subjectivity, define the risk question, reveal assumptions and uncertainty, and review risks over time (ICH, 2023).

The realistic future is not AI approving risk assessments. The realistic future is AI helping quality teams ask better questions, find better evidence, and review risks sooner.

For AIforQA.org, this is a strong cornerstone topic because it connects AI to one of the most important principles in modern pharmaceutical quality: better decisions come from better risk understanding.

References

- ICH. Q9(R1) Quality Risk Management. ICH Q9(R1) provides the foundational framework for quality risk management, including risk assessment, hazard identification, risk analysis, risk evaluation, risk control, risk communication, risk review, formality, risk-based decision-making, and minimizing subjectivity. https://database.ich.org/sites/default/files/ICH_Q9%28R1%29_Guideline_Step4_2023_0126.pdf
- ICH. Q10 Pharmaceutical Quality System. ICH Q10 describes quality risk management and knowledge management as enablers of the pharmaceutical quality system and connects QRM with process monitoring, CAPA, change management, and management review. https://database.ich.org/sites/default/files/Q10%20Guideline.pdf
- FDA. Guidance for Industry: Quality Systems Approach to Pharmaceutical CGMP Regulations. FDA's quality systems guidance discusses CAPA, risk management, management review, quality systems, and continuous improvement within pharmaceutical CGMP operations. https://www.fda.gov/media/71023/download
- FDA. 21 CFR Part 11 - Electronic Records; Electronic Signatures. Part 11 establishes requirements for trustworthy electronic records and electronic signatures, including validation, access controls, audit trails, authority checks, and record protection. https://www.ecfr.gov/current/title-21/chapter-I/subchapter-A/part-11
- FDA. 21 CFR 211.192 - Production Record Review. This regulation requires unexplained discrepancies and batch/component failures to be thoroughly investigated, including potentially associated batches or products, with written conclusions and follow-up. https://www.ecfr.gov/current/title-21/chapter-I/subchapter-C/part-211/subpart-J/section-211.192
- EMA. Reflection Paper on the Use of Artificial Intelligence in the Medicinal Product Lifecycle. EMA discusses AI/ML risk-based development, deployment, performance monitoring, context of use, data quality, degree of influence, and manufacturer responsibility for fit-for-purpose algorithms, models, datasets, and pipelines. https://www.ema.europa.eu/en/documents/scientific-guideline/reflection-paper-use-artificial-intelligence-ai-medicinal-product-lifecycle_en.pdf
- Higgins, D., & Johner, C. Validation of Artificial Intelligence Containing Products Across the Regulated Healthcare Industries. This article discusses validation challenges for AI/ML-containing products across regulated healthcare sectors, including pharmaceuticals, medical devices, and diagnostics. https://arxiv.org/abs/2302.07103