


# FDA & EMA Good AI Practice Guide for Drug Development

By Adrien Laurent, CEO at IntuitionLabs • 4/4/2026 • 55 min read

- good ai practice
- fda ai guidelines
- ema regulations
- drug development ai
- pharma compliance
- ai quality management
- regulatory affairs
- ai validation



# Executive Summary

Artificial intelligence (AI) is rapidly transforming the pharmaceutical industry, offering novel tools to accelerate drug discovery, optimize [clinical development](#), and streamline manufacturing and regulatory processes. As AI systems become more deeply embedded in drug development workflows, regulatory authorities have moved to provide clear guidance to ensure that such systems are reliable, safe, and compliant with existing quality and safety standards. In January 2026, the U.S. Food and Drug Administration (FDA) and the European Medicines Agency (EMA) jointly issued the “[Guiding Principles of Good AI Practice in Drug Development](#)”, a collaborative framework of **10 guiding principles** designed to support industry in responsible AI adoption across the drug lifecycle (<sup>[1]</sup> [www.fda.gov](http://www.fda.gov)) ([www.ema.europa.eu](http://www.ema.europa.eu)). These principles emphasize human oversight, risk management, data governance, lifecycle controls, and transparency in AI use, among other factors.

This comprehensive implementation guide dissects the FDA/EMA guiding principles in depth and translates them into practical compliance strategies for drug development teams. We provide background on the emergence of AI in pharmaceuticals, an overview of the evolving regulatory landscape, and the rationale for harmonized AI guidelines. Each guiding principle is analyzed in detail – defining its intent, regulatory rationale, and concrete measures by which a drug sponsor can implement it. We map the principles onto traditional quality systems (e.g. Good Clinical Practice (GCP), [Good Manufacturing Practice \(GMP\)](#), and Good Documentation Practice (GDP)) to illustrate how AI-specific expectations integrate with existing frameworks.

Subsequent sections outline organizational and process-oriented implementation strategies. We discuss establishing multidisciplinary [AI governance](#) structures, adapting the Quality Management System (QMS) for AI tools, and ensuring robust data and model documentation. We offer evidence-based risk management approaches, drawing on [FDA's AI credibility framework](#), and IMDRF's Good Machine Learning Practice, to identify and mitigate AI-specific risks. Emphasis is placed on validation of AI tools, continuous monitoring, and change control under an AI life-cycle process.

Real-world examples and case studies illustrate successes and pitfalls. For instance, leading pharma companies (such as Eli Lilly partnering with NVIDIA on AI supercomputing (<sup>[2]</sup> [www.axios.com](http://www.axios.com))) highlight high-stakes AI investments, while regulatory interactions (FDA's draft framework on AI model credibility (<sup>[3]</sup> [www.fda.gov](http://www.fda.gov)) (<sup>[4]</sup> [www.fda.gov](http://www.fda.gov))) show how authorities evaluate AI components. Surveys and experts underline both the promise and caution: despite major AI initiatives, drug approval rates have remained steady (~50 new drugs per year (<sup>[5]</sup> [time.com](http://time.com))), underscoring that regulatory rigor still governs outcomes.

Throughout, we integrate extensive citations from regulatory releases, peer-reviewed studies, industry analyses, and authoritative reports to support all claims. We also examine related regulatory activity (e.g. IMDRF's 2025 Good ML Practice guidance (<sup>[6]</sup> [www.imdrf.org](http://www.imdrf.org)), Swissmedic's AI framework ([www.swissmedic.ch](http://www.swissmedic.ch))) and upcoming developments (the EU's AI Act, evolving guidance) to contextualize the FDA/EMA approach within the broader compliance ecosystem. Statistically grounded data (market projections, adoption surveys, AI's impact on R&D timelines) underpin our analysis.

Finally, the report discusses future implications: how these principles will need to adapt to rapidly advancing AI technologies (e.g. generative AI in molecular design), and how regulators may further refine requirements. The conclusion synthesizes key action items for drug developers, urging proactive alignment of AI initiatives with regulatory expectations to advance innovation safely.

## Introduction and Background

Artificial intelligence (AI) – broadly defined as machine-based algorithms that learn from data to make predictions or decisions (<sup>[3]</sup> [www.fda.gov](http://www.fda.gov)) (<sup>[7]</sup> [pmc.ncbi.nlm.nih.gov](http://pmc.ncbi.nlm.nih.gov)) – is poised to revolutionize drug development. AI methods (including machine learning, deep learning, and generative models) can process massive biomedical datasets to identify new drug

targets, design molecules with optimal properties, simulate toxicology, [select patients for trials](#), and optimize manufacturing processes. Industry leaders are making bold bets: for example, as of early 2026 **Eli Lilly** announced a partnership with NVIDIA to build an AI-powered supercomputer aiming to develop research models, automate laboratory planning, and improve manufacturing efficiency (<sup>[2]</sup> [www.axios.com](http://www.axios.com)). Major startups (Insitro, Formation Bio, Xaira Therapeutics, etc.) are attracting billion-dollar investments for AI-driven platforms (<sup>[8]</sup> [apnews.com](http://apnews.com)) (<sup>[2]</sup> [www.axios.com](http://www.axios.com)). In 2025, survey data showed upwards of 70% of large pharma companies had ongoing AI projects in R&D, and the global AI-in-drug-development market was projected to exceed **\$1.5 billion by 2025** (from ~\$1.2B in 2024 (<sup>[9]</sup> [www.24lifesciences.com](http://www.24lifesciences.com))).

Despite heavy investment and advances in AI technology, regulatory and compliance considerations remain paramount. Drugs and biologics must meet stringent efficacy, safety, and quality requirements before market approval, regardless of the technologies used in their development. Historically, regulators have required rigorous validation of any process and method that informs critical decisions. In the context of AI, this includes ensuring data integrity, minimizing algorithmic bias, maintaining traceability of model-generated results, and preserving human accountability. As Swissmedic (Switzerland's agency) notes, AI/ML models' **complexity and opacity** present major challenges – for example, ensuring data quality, model transparency, and bias control ([www.swissmedic.ch](http://www.swissmedic.ch)). Absent proper controls, AI outputs could jeopardize patient safety or scientific validity.

Given this, global regulatory bodies have begun evolving their guidelines. The U.S. FDA, building on its Digital Health Innovation initiatives, has been issuing AI-related guidance for several years. Key milestones include FDA's 2020 publication of a Software Precertification pilot program, 2021 formation of internal AI programs, and notably the **January 6, 2025** draft guidance, "*Considerations for the Use of Artificial Intelligence to Support Regulatory Decision-Making for Drug and Biological Products*," its first such draft explicitly addressing AI in drug submissions (<sup>[3]</sup> [www.fda.gov](http://www.fda.gov)). In that press announcement, FDA Commissioner Dr. Robert Califf emphasized a commitment to "an agile, risk-based framework" that furthers innovation while upholding standards (<sup>[10]</sup> [www.fda.gov](http://www.fda.gov)). The FDA noted that "since 2016, the use of AI in drug development and in regulatory submissions has exponentially increased" (<sup>[4]</sup> [www.fda.gov](http://www.fda.gov)), prompting the need for formal credibility frameworks. Parallel to FDA, in Europe the EMA had begun exploratory work (e.g., an AI Reflection Paper in 2024) and announced in late 2025 that joint FDA-EMA principles were forthcoming ([www.ema.europa.eu](http://www.ema.europa.eu)) ([www.ema.europa.eu](http://www.ema.europa.eu)).

Concurrently, international regulators are aligning on foundational frameworks. The International Medical Device Regulators Forum (IMDRF) – representing agencies from the US, EU, Japan, Canada, Australia, etc. – published in **January 2025** a "Good Machine Learning Practice" guidance for medical device AI/ML, emphasizing many similar principles (documentation, risk-based design, monitoring) (<sup>[6]</sup> [www.imdrf.org](http://www.imdrf.org)). The EU's new **AI Act** (became law in mid-2023) introduces risk-based legal requirements for high-stakes AI systems (including medical devices and critical healthcare applications), highlighting transparency and safety. Organizations like the OECD and WHO have also formulated high-level AI ethics and risk frameworks.

Within this landscape, drug development teams face a complex task: to harness AI's benefits while navigating uncharted regulatory territory. Ensuring compliance means more than basic software validation; teams must implement "Good AI Practice" – analogous to GxP (Good [machine] practices) – integrating AI governance into traditional quality systems. This report provides an in-depth guide to doing so, aligned with the newly released FDA/EMA principles.

## The Regulatory Landscape

The FDA/EMA joint guidance does not exist in isolation. In the U.S., drug development is governed by well-established regulations: **21 CFR Part 210/211 (GMP)** for manufacturing, **21 CFR Part 312 (INDs)**, **21 CFR Part 314 (NDAs/BLAs)**, and **21 CFR Part 11 (electronic records)** for eCTD submissions, among others. In Europe, the framework includes EU GMP, ICH guidelines (e.g. ICH E6(R3) on GCP, E8 on general considerations), and the EMA's GIRP (Good Information Management Practice). These frameworks traditionally did not explicitly address AI, which has unique lifecycle dynamics.

However, as FDA's 2025 press release highlights, AI usage "can be used in various ways to produce data or information regarding the safety, effectiveness, or quality of a drug" <sup>(14)</sup> [www.fda.gov](http://www.fda.gov) (for example, predicting patient outcomes or analyzing real-world data). Thus, AI outputs are often integrated into evidence packages; sponsors may upload AI-derived analyses in regulatory submissions. FDA recognizes it has "substantial experience" reviewing submissions with AI components <sup>(14)</sup> [www.fda.gov](http://www.fda.gov), but previously had no dedicated guidelines on how to do so.

The new FDA draft guidance (Jan 2025) and the 2026 guiding principles both respond to this gap. In particular, the draft "Considerations for Use of AI" guidance recommended a **credibility framework**: sponsors should characterize the AI model's context of use, the analytical process, and evidence of performance. This mirrors FDA's approach to medical devices with AI (which had its own guidances on "predicate vs continual learning" models). In effect, FDA now expects drug sponsors to treat AI tools as part of their Quality Management System (QMS) and to justify their reliability like any other validation.

Internationally, EMA has similarly recognized the need. The EMA press release (Jan 2026) notes that regulatory legislation ("new pharmaceutical legislation" and the EU Biotech Act proposal) is catching up to AI by accommodating innovative AI methods in a controlled environment ([www.ema.europa.eu](http://www.ema.europa.eu)). EMA's strategy documents (EMANS 2021–2028) explicitly call for leveraging AI in regulatory review, and the forthcoming EU pharma GMP Annexes and Good GCP Guidelines are expected to reference AI tools. Other regulators (e.g., Canada's Health Canada, Australia's TGA, Switzerland's Swissmedic) have published high-level statements or framework documents encouraging responsible AI use in submissions, often referencing FDA/EMA principles and IMDRF guidance.

**Key Takeaway:** Drug regulators are converging on a risk-based, documented approach to AI. The FDA/EMA Good AI Practice principles echo and extend existing Good Practice paradigms (e.g. conceptually similar to software validation in GMP/GLP/GCP contexts), but tailored to AI's iterative, data-driven nature. Teams should anticipate evolving requirements: initiatives such as ICH M17 (digital technologies) and the EU AI Act will further shape compliance.

## Guiding Principles of Good AI Practice

In January 2026, FDA and EMA jointly released **10 Guiding Principles for Good AI Practice in Drug Development** <sup>(11)</sup> [www.fda.gov](http://www.fda.gov) <sup>(12)</sup> [www.jdsupra.com](http://www.jdsupra.com). These principles provide *high-level* guidance across the entire drug development lifecycle, from early research through post-market monitoring. They are **not legally binding** regulations, but serve as expectations against which both sponsors and regulators will measure AI use. Importantly, they are harmonized between FDA and EMA, signaling a unified transatlantic stance on AI governance. (A summary table of the principles is below.)

Principle No.	Principle Title	Core Focus
1.	Human-centric by design	Embed human oversight and accountability; AI as advisory/assistive, not autonomous decision-maker.
2.	Risk-based approach	Tailor AI effort to assessed risk; evaluate where/when AI is used and mitigate risks accordingly.
3.	Adherence to standards	Comply with existing applicable standards (GxP, technical, cybersecurity); integrate AI into quality systems.
4.	Clear context of use	Define explicitly what tasks/decisions AI supports; ensure appropriate scope and limits are documented.
5.	Multidisciplinary expertise	Leverage cross-functional teams (technical, clinical, regulatory, quality, etc.) in AI projects.
6.	Data governance and documentation	Ensure data quality, provenance, security; maintain detailed documentation of datasets and data handling plans.
7.	Model design and development	Follow disciplined model/software development (version control, testing, validation) with full traceability.
8.	Risk-based performance assessment	Validate AI performance via scientifically credible methods; monitor accuracy and error rates in context.
9.	Lifecycle management	Establish processes for model updates, re-validation, and decommissioning; continuous monitoring post-release.
10.	Clear, essential information	Provide transparent documentation (capabilities, limitations, intended use) about the AI to end-users/regulators.

These principles are interrelated and collectively promote **safety, efficacy, and quality** — the core tenets of drug regulation. <sup>(13)</sup> [www.fdcell.com](http://www.fdcell.com) <sup>(14)</sup> [www.fda.gov](http://www.fda.gov) They draw on established regulatory concepts: for example, Principle 2 ("risk-based approach") echoes ICH Q9 on risk management, and Principle 10 ("clear information") parallels 21 CFR

211.130 and ICH E3 which require clarity in reports. Importantly, the principles stress **documentation and transparency** at every stage, reflecting regulators' need to audit and review AI processes.

**Table 1: FDA/EMA Guiding Principles of Good AI Practice**

Principle	Description (Regulatory Focus)
1. Human-centric by design	AI tools should support, <b>not replace</b> , human experts. Systems are designed with human oversight in mind (e.g. "human-in-the-loop" confirmation for outputs) <sup>(15)</sup> <a href="http://justintimegcp.com">justintimegcp.com</a> ). The AI's role and limits must be clear.
2. Risk-based approach	AI development and deployment should be commensurate with the <b>risk level</b> of potential errors. High-risk applications (e.g. safety decisions) require more stringent controls (testing, review). Teams must assess what could go wrong and mitigate it <sup>(16)</sup> <a href="http://justintimegcp.com">justintimegcp.com</a> <sup>(10)</sup> <a href="http://www.fda.gov">www.fda.gov</a> .
3. Adherence to standards	AI software and processes must align with applicable regulations and guidelines (e.g. FDA's QSR, EU GMP Annex 11, GCP, ISO 13485, ISO 14971). For instance, data encryption, audit trails, and controlled change management should follow existing IT/QMS standards <sup>(17)</sup> <a href="http://justintimegcp.com">justintimegcp.com</a> ( <a href="http://www.swissmedic.ch">www.swissmedic.ch</a> ).
4. Clear context of use	The specific <b>scope and application</b> of the AI tool must be well-defined and justified. This includes intended users, data types, and decisions supported. Misuse outside this context must be prevented. Clear "context of use" informs risk assessment and validation depth <sup>(18)</sup> <a href="http://justintimegcp.com">justintimegcp.com</a> .
5. Multidisciplinary expertise	Development and governance require input from diverse stakeholders: domain scientists, clinicians, AI/ML engineers, statisticians, quality/regulatory experts, and ethicists as needed. Cross-functional teams ensure clinical and technical considerations are not overlooked <sup>(19)</sup> <a href="http://justintimegcp.com">justintimegcp.com</a> .
6. Data governance and documentation	Robust data management plans must document data sourcing, cleaning, labeling, privacy, and lineage. Data used for training/validation should be representative and unbiased. Data integrity controls (like versioning and audit logs) must comply with regulatory expectations (e.g. CFR) <sup>(20)</sup> <a href="http://www.swissmedic.ch">www.swissmedic.ch</a> <a href="http://justintimegcp.com">justintimegcp.com</a> .
7. Model design and development practices	AI models should be built under a controlled software development lifecycle. This includes version control, code reviews, testing against benchmarks, and restrictions on model evolution. Any changes to the model (retraining, parameter updates) must be planned and documented (see also FDA's proposed AI/ML change control guidances) <sup>(21)</sup> <a href="http://justintimegcp.com">justintimegcp.com</a> .
8. Risk-based performance assessment	Prior to deployment, conduct rigorous validation to demonstrate the AI's performance (e.g. accuracy, sensitivity) meets requirements for its use-case. Use reference standards or hold-out datasets where possible. After deployment, continuously monitor performance metrics and error rates, especially false negatives/positives, adjusting as needed <sup>(22)</sup> <a href="http://justintimegcp.com">justintimegcp.com</a> <sup>(4)</sup> <a href="http://www.fda.gov">www.fda.gov</a> .
9. Lifecycle management	AI tools should be managed throughout their lifecycle: include plans for regular re-validation when input data or system context changes, mechanisms for controlled updating, and criteria for retirement. Monitor post-deployment (similar to product surveillance) to catch drift or obsolescence <sup>(23)</sup> <a href="http://justintimegcp.com">justintimegcp.com</a> .
10. Clear, essential information	Documentation must clearly state the AI's intended purpose, capabilities, and known limitations, in plain language suitable for end-users. Training materials, manuals, and labels should ensure users can interpret outputs appropriately. Regulatory submissions should disclose relevant algorithmic details (as feasible) about the model and data <sup>(24)</sup> <a href="http://justintimegcp.com">justintimegcp.com</a> <sup>(10)</sup> <a href="http://www.fda.gov">www.fda.gov</a> .

**Table 1. Good AI Practice Guiding Principles (FDA/EMA Jan 2026).** These ten principles cover the unique challenges of using AI in drug development, from concept through decommissioning. Each principle aligns with the overarching goal of ensuring **patient safety, product quality, and regulatory compliance** when AI is involved. For example, Principle 1's emphasis on human oversight reaffirms that final decisions (e.g. dosing changes, patient eligibility) remain with qualified professionals <sup>(15)</sup> [justintimegcp.com](http://justintimegcp.com)). Principle 6's data focus reflects regulators' long-standing requirement for complete data traceability ([www.swissmedic.ch](http://www.swissmedic.ch)), now extended to AI training datasets.

The principles envision a **holistic governance framework** rather than isolated fixes. They will likely be operationalized through supplementing existing GxP processes. For instance, an AI-based patient screening tool (clinical) or an image-analysis AI in manufacturing would each need explicit context-of-use documentation (Princ. 4) and continuous monitoring of predictive accuracy (Princ. 8).

Notably, while the principles are high-level, numerous downstream guidance and case studies are expected to detail their implementation. The principles themselves will "underpin future AI guidance" and standard-setting (EMA press) ([www.ema.europa.eu](http://www.ema.europa.eu)). In practice, drug developers should treat these principles as **both a checklist and a mindset**: as you integrate AI into any stage, ensure at least these fundamentals are addressed. The rest of this report delves deeply into each principle and its practical implementation.

# Principles in Depth and Implementation Strategies

Below we examine each principle, discussing its rationale, regulatory context, and how development teams can implement it in practice. Where relevant, we draw on expert commentary and analogous guidelines (e.g. GCP, SaMD) to illustrate best approaches.

## 1. Human-centric by Design

**Explanation:** The first principle underscores that AI systems must be *designed with humans in the loop*. Rather than allowing algorithms to make unsupervised decisions, development teams must ensure outputs are advisory. Qualified professionals must oversee and verify AI-generated insights before they impact patient care or regulatory filings (<sup>[15]</sup> justintimegcp.com). This principle reinforces accountability: “AI must augment, not replace, qualified human judgment” (<sup>[15]</sup> justintimegcp.com).

**Rationale:** Regulatory and ethical standards demand that any decision affecting patient health remain traceable to a responsible expert. For example, in clinical trial conduct, the investigator of record (per ICH E6(GCP)) cannot delegate critical decisions to an unvetted algorithm. Similarly, GMP requires that final product release decisions rest with qualified personnel. By designating AI outputs as advisory, organizations preserve the existing hierarchy of responsibility. This approach also mitigates risk: even highly accurate AI algorithms can err, especially on atypical inputs (<sup>[15]</sup> justintimegcp.com). A human reviewer can catch out-of-distribution errors or context misinterpretations.

### Implementation:

- **Workflow Integration:** During tool design, plan AI as a decision-support tool. For instance, if using AI to flag potential adverse event reports, the pharmacovigilance officer should review each AI flag, not automate reporting to regulatory authorities without oversight.
- **Human-in-the-Loop (HITL):** Implement HITL checkpoints. Early in deployment, do *full manual review* of all AI outputs (type: “screening”). Over time, if performance is proven, gradually shift to *sampling* and risk-based review. The key is documenting that each output was reviewed by a trained user (<sup>[15]</sup> justintimegcp.com).
- **Audit Trails:** Use system logs and records (e.g. CFR Part 11 audit trail) to show that a human reviewed and accepted, edited, or rejected each AI recommendation. This is part of “Clear, essential information” (Princ. 10).
- **Training & SOPs:** Update SOPs to explicitly require human oversight on AI outputs. Train staff on the tool’s intended use (what kinds of errors it might make). For example, a Clinical Research Associate (CRA) using an AI to review source data must still certify the data.
- **User Interface Design:** For software UIs, include prompts or checkpoints requiring user sign-off before proceeding. The interface should display AI findings with context so that the human can make an informed decision.

**Example:** In the example from JustInTimeGCP (<sup>[15]</sup> justintimegcp.com), an AI “Site Document Review Service” was built so that data managers see AI-suggested missing documents, but they must accept before updating the Trial Master File. The team enforced “full human review” at first, reviewing every AI suggestion, then moved to “risk-based confirmation” as trust grew (<sup>[15]</sup> justintimegcp.com). This gradual approach exemplifies Principle 1 “in action.”

## 2. Risk-based Approach

**Explanation:** Principle 2 requires sponsors to assess and manage the **risks specific to the AI application's context of use**. Not all AI projects carry the same regulatory impact: an AI tool that automates narrative summaries for internal reports has lower patient risk than an AI that recommends change to dosage. The development team must conduct a formal risk assessment focusing on what could go wrong, its impact on patient safety or data integrity, and implement proportionate controls (<sup>[16]</sup> justintimegcp.com).

**Rationale:** A core tenet of FDA and EMA regulation is *risk-based regulation*: resources and scrutiny are aligned with potential impact on public health. This principle extends that tenet to AI. The 2025 FDA draft guidance explicitly frames AI credibility in risk-based terms: "the magnitude of risk should influence the level of evidence required" for regulatory decisions supported by AI (<sup>[10]</sup> www.fda.gov) (<sup>[16]</sup> justintimegcp.com). For example, if an AI model is used for exploratory data analysis, the risk is low; if it is used to generate primary efficacy claims, the risk is high and requires extensive validation. This approach ensures compliance efforts focus on high-stakes uses.

**Implementation:**

- **Risk Assessment Process:** Incorporate AI-specific risks into existing risk management frameworks (e.g. ICH Q9, ISO 14971 for devices). Create a cross-functional risk committee (Principle 5) that identifies hazards: model accuracy failures, data breaches, inappropriate usage, etc.
- **Context of Use:** Precisely define the AI's use-case, as required by Principle 4; then categorize risk (e.g. low, medium, high) based on potential patient safety impact. FDA's draft guidance suggests categories akin to decision support vs autonomous decision (<sup>[4]</sup> www.fda.gov).
- **Risk Controls:** For higher-risk applications, apply stronger controls: more stringent data quality checks, independent verification of outputs, requirements for external validation, redundant monitoring. For lower-risk, lighter oversight may suffice.
- **Credibility Testing:** Align with the proposed FDA AI/ML Credibility Framework (7-step) that emphasizes risk-based evidence generation. This includes defining a trustworthiness level and evaluating testing rigor accordingly (<sup>[25]</sup> justintimegcp.com). The JustInTimeGCP example used this notion: they claimed "low-to-moderate risk, decision-support use case" and accordingly combined credibility testing with mandated human review (<sup>[26]</sup> justintimegcp.com).
- **Documentation:** Document risk mitigation measures in project plans and regulatory submissions. For example, the submission might include a summary of the risk assessment and how mitigation (e.g. thresholds, reviews) was implemented.

*Illustration:* Consider an AI that predicts patient drop-out risk in a clinical trial. Risks include misclassifying a patient (false negative – missing a drop out can affect compliance data). A risk assessment might classify false negatives as moderate severity (since the result is data completeness issues, not patient harm). Controls could include manual review of borderline predictions, and periodic monitoring of prediction performance. Risk-based thinking might not require this AI to have full validation as a diagnostic device, but still demand statistics on accuracy.

### 3. Adherence to Standards

**Explanation:** Principle 3 reminds developers that **existing regulations and technical standards still fully apply to AI systems**. In practice, this means that while AI introduces new elements, it cannot operate in a regulatory vacuum. Sponsor organizations must ensure that their AI tools comply with applicable GxP rules and other relevant technical standards (cybersecurity, IT controls, software validation, data protection laws etc.) (<sup>[17]</sup> justintimegcp.com). This includes, for example, validating the software used, ensuring audit trails (21 CFR 211.68, Part 11), and managing supplier quality for AI software.

**Rationale:** Regulatory frameworks predate AI, but they remain foundational. For example, software used in manufacturing must be validated under 21 CFR 211/820. The FDA/EMA principles assume AI tools handling regulated information are governed by those same rules. The 2025 FDA press release noted "FDA's robust scientific and regulatory

standards are met” alongside promoting innovation (<sup>[10]</sup> [www.fda.gov](http://www.fda.gov)), implying no relief from compliance. Non-compliance risks (e.g. data integrity lapses) still apply to AI outputs. Principle 3 thus prevents sponsors from considering AI as exempt; rather it must fit into the overall quality system.

#### Implementation:

- **Quality Management Integration:** Incorporate AI tools into the QMS. This may involve adding AI-specific processes to ISO 9001 or 13485 QMS: e.g. supplier qualification for AI vendors, risk management records, CAPA procedures for AI system errors.
- **Validation and Verification:** Follow established good validation practices (GVP): develop an AI system validation plan, test scripts, acceptance criteria, per the type of software (e.g. as a regulated in-vitro diagnostic, or a quality control tool). If an AI model outputs data that end up in the eCTD, ensure the code and model development comply with documentation requirements analogous to software change control.
- **Standards and Guidance:** Align with recognized standards. For instance:
  - *IEC 62304* (if considered medical device software): use its software lifecycle process for development and maintenance.
  - *ISO 13485 / CFR 820 QSR*: QMS for design, CAPA.
  - *OECD, WHO or IMDRF AI/Ca frameworks*: Many are principles-based (transparency, accountability). Swissmedic explicitly notes incorporating IMDRF and ICH guidelines into AI assessments ([www.swissmedic.ch](http://www.swissmedic.ch)).
  - *Cybersecurity standards* (ISO 27001, NIST): ensure appropriate data security (Principle 6) in transit and at rest.
- **Controlled Environments:** Restrict AI systems handling sensitive data (PHI, proprietary info) to validated, secured IT environments (e.g. GCP cloud with encryption, segmentation). The JustInTimeGCP example enforced “identity-based access control, encryption in transit and at rest, environment isolation, and immutable audit logging” (<sup>[17]</sup> [justintimegcp.com](http://justintimegcp.com)), demonstrating how AI tools should meet data security standards.
- **GxP Harmonization:** If an AI is used in a GCP trial (e.g. for eCRF review), ensure its deployment is reviewed by QA and IT, and possibly audited in vendor audits. If in GMP manufacturing, include it in batch release and equipment qualification processes.

*Case Illustration:* A pharmaceutical QA department introduces an AI to predict microbial contamination risk in a facility. Principle 3 mandates this system be validated under GMP as it influences quality assurance. The team would write a validation protocol (Installation Qualification, Operational Qualification) for the software/hardware. The underlying AI model would be part of the computerized system validation, ensuring traceability of any change from the model (akin to software version control in a QMS).

## 4. Clear Context of Use

**Explanation:** Principle 4 insists that the exact **context** in which the AI will be used must be explicitly defined and documented. This includes the intended tasks, boundaries, and populations for which the AI tool is applicable, as well as any limitations on its use (<sup>[18]</sup> [justintimegcp.com](http://justintimegcp.com)). A well-specified context of use guides development, validation, and review processes.

**Rationale:** Defining context of use (CoU) confines the AI to intended applications and prevents inappropriate extrapolation. For example, an AI trained on adult pharmacokinetic data should not be applied to pediatric cases without revalidation. From a regulatory standpoint, an undefined CoU leads to unclear risk assessment and inadequate oversight. The FDA’s 2025 draft guidance highlights the importance of CoU by using it to calibrate the necessary evidence for credibility. By clarifying CoU, sponsors justify the relevance and limits of their AI.

#### Implementation:

- **Define Scope:** In project documentation (e.g. in the IND/NDA or internal design dossier), state exactly what the AI does and does *not* do. This might include the data inputs it requires and the outputs it produces, computational environment, and the decision nodes it supports.
- **Use-case Examples:** Provide concrete use-case scenarios. E.g., “This AI tool analyzes laboratory assay images to detect anomalies for release testing. It is NOT to be used for patient diagnostics.” The FDA/EMA guidelines use “context of use” extensively; it should be analogous to how one defines medical device intended use.
- **Documentation:** Include CoU in formal documents (e.g. validation plan preamble, clinical protocol if using AI in trial design). Ensure all team members know these scope limits.
- **Communication:** The CoU should be communicated to end users in training materials. For instance, if an AI model only supports document completeness checks (like the example), emphasize it “does not support patient safety decisions or regulatory submissions” (<sup>[18]</sup> justintimegcp.com). This avoids misuse and also clarifies oversight needs.
- **Model Training Data:** Align the training data selection with CoU. E.g. if CoU is “adult phase 1 trial data”, ensure training data does not include irrelevant populations.
- **Reviewer Guidance:** For regulatory reviewers, include a CoU statement in submissions. This helps assess if the provided validation matches the intended use. If the tool operates outside that CoU, it should undergo separate evaluation.

*Example:* In the JustInTimeGCP scenario, the team explicitly notes that their Site Document Review AI “supports internal TMF review only. It does not support clinical decisions, patient safety determinations, or regulatory submissions” (<sup>[18]</sup> justintimegcp.com). By stating this, they limit the tool’s use and thereby its risk profile. Regulators reviewing a submission that includes such an AI would see that it’s properly restricted to back-office use, not frontline decisions.

## 5. Multidisciplinary Expertise

**Explanation:** Effective AI projects require collaboration across multiple domains. Principle 5 mandates the formation of **cross-functional teams** that combine technical AI expertise (data scientists, ML engineers) with domain specialists (clinicians, statisticians, pharmacologists), regulatory experts (QA, compliance officers), cybersecurity officers, and business leads (<sup>[19]</sup> justintimegcp.com). No single discipline can manage all AI challenges.

**Rationale:** AI systems in drug development straddle many areas. A purely tech-driven team might optimize an algorithm’s accuracy but miss regulatory requirements or clinical relevance. Conversely, a clinical team alone may not understand ML nuances. Regulators emphasize accountability, which in practice means involving SMEs (Subject Matter Experts) who can attest to the appropriateness of AI approaches. As [17] notes, combining “deep GCP and TMF domain expertise” with AI and IT professionals ensures that both clinical and technical considerations are addressed in tandem (<sup>[19]</sup> justintimegcp.com).

### Implementation:

- **Team Composition:** Establish an AI governance or steering committee that includes representatives from:
  - Clinical operations (e.g., medical monitors, investigators)
  - Data science/IT (ML engineers, software developers, data engineers)
  - Regulatory affairs / QA (to ensure compliance alignment)
  - Quality control (for manufacturing-focused AI)
  - Biostatistics (especially for AI in trials / data analysis)
  - Information security (for data protection)
  - Project management (coordination of all NPCs).

- **Roles and Responsibilities:** Define clear RACI (Responsible, Accountable, Consulted, Informed) for AI tasks. For instance, an ML engineer might be responsible for model development, but a statistician is responsible for defining acceptable performance metrics, and a QA officer for supervising documentation.
- **Regular Governance Meetings:** Hold regular meetings to review AI system progress, risk status, and compliance. Decisions about model changes, validation scope, or deployment should involve multi-party sign-off (e.g. design review boards).
- **External Experts:** When necessary, consult external domain experts. For example, if the AI addresses rare disease biomarker analysis, involve clinicians who are experts in that indication. Similarly, legal counsel or ethics boards might weigh in on privacy impacts.
- **Training and Communication:** Educate team members about AI basics so they can collaborate effectively. Clinicians should understand what AI can/cannot do, and data scientists should learn applicable regulations. Cross-training improves mutual understanding.

*Illustration:* In practice, a biopharma using AI for trial patient matching created a team with its head of data science, a lead statistician, a clinical trial manager, and QA audit lead. They held weekly “AI sprint reviews” where technical leads showed model outputs to clinicians, who gave feedback on clinical sensibility. QA personnel ensured each sprint was documented (commit logs, changes) for audit. This approach aligns with Principle 5’s call for integrated expertise.

## 6. Data Governance and Documentation

**Explanation:** Principle 6 emphasizes rigorous management of data throughout the AI lifecycle. This includes ensuring **data quality, integrity, provenance, and security**, as well as complete documentation of data handling processes. AI models are only as good as the data they are trained on; thus, controlling data is paramount <sup>(20)</sup> [justintimegcp.com](https://www.justintimegcp.com) ([www.swissmedic.ch](https://www.swissmedic.ch)).

**Rationale:** Regulators have long focused on data integrity (GxP requirements, FDA’s ALCOA+ principle: data must be Attributable, Legible, Contemporaneous, Original, Accurate and “Complete, Consistent, Enduring, Available”). In AI, this translates to tracking where each data element came from (source systems, version control), how it was processed (cleaning steps, labeling), and ensuring it is representative of the intended use. Poor data can introduce bias or hidden errors. Swissmedic explicitly warns that “complexity and lack of transparency [in AI] present challenges... especially data quality assessment, model transparency, quality control and possible bias” ([www.swissmedic.ch](https://www.swissmedic.ch)). Without robust data governance, AI outputs cannot be trusted by regulators.

### Implementation:

- **Data Lifecycle Plan:** For each AI project, prepare a Data Management Plan (DMP) that specifies:
- **Data sources:** Enumerate origins (clinical trial EDC, lab instruments, public databases, etc.).
- **Data content:** Variables/features used, with definitions (Codebooks).
- **Data cleaning/transformations:** Procedures for handling missing data, normalization, annotation.
- **Data splitting:** Details on training, validation, test sets; ensure no duplication.
- **Data security:** How data are stored (encrypted storage, access controls), transmitted, and eventually disposed.
- **Privacy considerations:** E.g., de-identification steps if patient data are used.
- **Versioning and Traceability:** Use data version control tools (DVC) or maintain software libraries that track changes to datasets. Every model training run should log the exact data version used. This aids reproducibility and aligns with requirement to reproduce results if audited.
- **Data Quality Metrics:** Implement quantitative checks (completeness, outliers, consistency) and manual review for critical data fields. Log these checks as part of the learning pipeline.

- **Documentation:** Similar to lab notebooks, maintain an AI “data logbook” or addendums to study reports detailing how data were processed. For instance, if an algorithmic feature depends on a derived laboratory parameter, document the formula and units conversions.
- **Governance Policies:** Apply the organization’s broader data policies to AI data. E.g. an enterprise data governance board should oversee AI-relevant data assets. If using cloud data lakes, ensure they meet GxP compliance (some vendors offer GxP-validated cloud storage).
- **Audit Trails:** Ensure all data changes (ingestion, cleaning, labeling, deletion) are logged (timestamp, user). This meets Part 11 requirements indirectly by demonstrating traceability.

*Example:* A clinical trial sponsor used large historical EHR datasets to train an AI for identifying eligible patients. They faced data heterogeneity (different hospital coding). To satisfy data governance, they first mapped all fields to standard terminologies, then kept a mapping dictionary file in version control. They ran automated scripts checking for missing or out-of-range values and saved reports in a QA archive. During QMS audits, they could show logs proving that the exact dataset used for model training corresponds to what was declared in the IND.

## 7. Model Design and Development Practices

**Explanation:** This principle calls for disciplined engineering of AI models. It covers all stages of design, coding, testing, and deployment, with emphasis on **traceability, reproducibility, and controlled change management**. Essentially, AI models should be developed like regulated software or medical devices: with specifications, design documentation, development environments, and release controls (<sup>[21]</sup> justintimegcp.com).

**Rationale:** Software development practices (SDLC) protect against errors and ensure reliability. For AI, additional complexity arises: model behavior can subtly change with data or parameters. Without strict controls, an AI model could drift or malfunction silently. Regulators expect that modifications are managed. The FDA’s draft guidance and subsequent discussions introduced the concept of “predetermined change control plans” for AI models – akin to a plan specifying which model updates are minor (e.g. retrain with similar data) vs major (algorithm redesign) and what validation each requires. Treating model code and training as GxP-controlled means any update goes through validation and approval channels like other changes.

### Implementation:

- **Software Development Lifecycle (SDLC):** Follow an SDLC appropriate to the AI’s regulatory classification. Maintain design documents (user requirements, functional requirements, architecture diagrams) for the AI system. For example, define input-output relationships and expected performance metrics upfront.
- **Version Control:** Place all code (model scripts, data processing code) in an enterprise version control system (e.g. Git). Tag releases of the code that correspond to validated model versions.
- **Environment Control:** Use containerization or virtual environments to freeze dependencies (libraries, frameworks) so that the AI can be exactly reproduced. Document which versions of Python, ML libraries (TensorFlow, etc.) were used.
- **Reproducible Training:** Record random seeds, model hyperparameters, and training configuration. If stochastic elements exist (like random weight initialization), provide seeds to regenerate the same model.
- **Testing and Verification:** Develop test cases and benchmark datasets (separate from training data) to test each build of the model. Automated unit and integration tests for code are advised, along with performance validation on hold-out datasets.
- **Change Control:** Classify model changes (e.g. retraining, hyperparameter tuning, feature changes) and determine when these constitute a new “version.” Establish a Change Control Board (CCB) review for significant changes. Document a policy: any change triggers impact assessment (risk analysis) and possibly re-validation.

- **Fixed Model Versions for Submission:** At the point of regulatory submission (IND/NDA), clearly state the model version used in the analyses. If further development is expected post-submission, describe the change control plan (as FDA's guidance suggests) on how updates will be managed with the agency.
- **Accountability Charter:** The JustInTimeGCP example notes alignment with an "AI Accountability Charter" (likely referencing principles of transparency/ethics) (<sup>[17]</sup> [justintimegcp.com](https://justintimegcp.com)). Teams should adopt or adapt ethical AI guidelines (many big companies have internal charters) to capture commitments like fairness, auditability, and data privacy in development practices.

*Illustration:* Consider developing an AI algorithm to classify lab images for impurity detection in manufacturing. The team would start by defining interface names, thresholds, and performance targets in a Requirements Specification. They store every script in Git, and use automated CI/CD pipelines that run a test suite whenever code is pushed. If a chemist suggests altering a pre-processing filter, the team logs this as a "CR" (Change Request), assesses how model outputs might shift, and executes an impact analysis. Any approved change receives a new version number and is re-verified.

## 8. Risk-Based Performance Assessment

**Explanation:** Principle 8 specifies that AI tools must be validated for **performance and reliability appropriate to their risk**, and then continuously monitored. Prior to deployment, credibility assessments should demonstrate that the AI meets predefined criteria (accuracy, precision, sensitivity, etc.) pertinent to its context of use. Post-deployment, performance should be tracked (monitoring false positives/negatives and other error modes) to catch degradation (<sup>[22]</sup> [justintimegcp.com](https://justintimegcp.com)) (<sup>[4]</sup> [www.fda.gov](https://www.fda.gov)).

**Rationale:** In regulated environments, every method used to support decisions must be qualified. Traditional analytical methods are validated to certain standards (ICH Q2). AI models, while not laboratory instruments, similarly produce "answers" that must be trusted. FDA's draft guidance stressed evidence of model credibility, and Weave Bio's commentary notes that "credibility testing, access controls... mitigate risks related to accuracy" (<sup>[16]</sup> [justintimegcp.com](https://justintimegcp.com)). By adopting a risk-based lens, a tool that impacts primary endpoints must be validated more rigorously than an internal dashboard. Continuous monitoring is crucial because AI models, especially those that learn over time, can change behavior after deployment (data drift, software updates, etc.).

### Implementation:

- **Validation Plan:** Develop a formal validation or performance qualification plan for the AI. Specify metrics (e.g. ROC-AUC, error rates) and acceptance criteria. For classification tasks, define the minimum acceptable sensitivity/specificity.
- **Test Datasets:** Use multiple datasets: an *external* hold-out test set that the model never saw during development, representing real-world diversity. Also, consider adversarial tests (edge cases) if relevant.
- **Bias and Fairness Checks:** Evaluate if the model systematically underperforms on subgroups (e.g. demographic groups or disease subtypes). Document these analyses, as regulators are increasingly concerned with AI bias in healthcare. Remediation from training data tweaks or algorithmic fairness constraints might be needed.
- **Credibility Documentation:** Align with FDA/EMA *Guiding Principle 8* by documenting the credibility (performance) results in submissions. FDA's own credibility framework (draft) proposed a 7-step framework, including testing and outcomes. If applicable, cite it indirectly: e.g. "performance was evaluated consistent with the FDA's recommended credibility framework" (though one cannot cite an unpublished draft in official docs, internally it guides development).
- **Monitoring Post-Release:** Define Key Performance Indicators (KPIs) and set up monitoring systems. For instance, an AI deployed for pharmacovigilance signal detection should have thresholds: if the weekly number of signals falls below expected, that may indicate a problem (e.g. data feed issue). Analyze false negative cases (missed signals discovered later) and conduct routine audits on a sample of cases.

- **Human Feedback Loop:** Incorporate user corrections into monitoring. If experts regularly override certain AI suggestions, log that trend and investigate if the model needs retraining.
- **Reporting:** Build a reporting dashboard that tracks tool performance metrics in real time. Schedule periodic review meetings to act on trends. Define triggers for re-validation (e.g. model accuracy drop >5%).
- **Regulatory Reporting:** If the AI is part of a regulated product (like a companion diagnostic), notify regulators per guidance on Changes Being Effected (CBE) or supplements when performance changes.

*Example:* A machine learning model is used to classify batch release test results. The team holds out 20% of historical data for final validation and achieves 95% accuracy. They set a performance metric: at least 90% accuracy on incoming data. After deployment, they compare AI predictions to human QC checks daily; if accuracy falls below 90% for one week, an alert triggers re-training. They log all performance metrics (false alarms and misses) quarterly, adjusting the model or thresholds as needed. In their regulatory submission, they include the validation test results and outline the monitoring process per Principle 8.

## 9. Lifecycle Management

**Explanation:** The ninth principle focuses on **managing the AI tool across its entire lifecycle**, from inception to retirement. Because AI models are not “static products,” controls must exist for versioning, updating, and decommissioning. Teams should plan for ongoing *governance frameworks* that include scheduled re-evaluation (e.g. periodic re-validation), controlled updates (with impact assessments), and criteria for withdrawal of an AI system (<sup>[23]</sup> justintimegcp.com).

**Rationale:** In traditional drug development, manufacturing processes or devices undergo lifecycle changes (process improvements, equipment upgrades) under change control. Similarly, AI systems evolve: new data may be incorporated to refine models, or software dependencies may update. Without oversight, changes could introduce unintended behavior. Additionally, an AI tool may become obsolete (e.g. due to new regulations, technologies, or business needs) and needs to be retired safely. Lifecycle management ensures the tool remains in a state of control 吳 (<sup>[10]</sup> www.fda.gov) (<sup>[23]</sup> justintimegcp.com).

### Implementation:

- **Versioning and Change Management:** Extend SDLC processes (from Principle 7) to cover ongoing updates. Maintain a version registry and change log. For each new release, perform impact analysis and regression testing relative to the validated baseline.
- **Re-validation Schedule:** Define intervals or triggers for refreshing/revalidating the model. For example, if the data domain is rapidly changing (e.g. viral strains in vaccine development), schedule re-training every 6 months. If internal data pipeline changes (source or format), trigger a re-validation event.
- **CI/CD Pipelines with Checks:** Automate testing and deployment to enforce that every code or data update passes validation checks before going live.
- **Retirement Plan:** Even an effective AI tool should have an exit plan: if the tool is to be replaced by a new methodology, ensure documentation on how to completely shut down and archive it. Data retention policies (from GxP) apply: archived model code and data must be preserved per record-keeping requirements.
- **Audit and Review:** Include AI lifecycle items in periodic QA audits. For instance, in an annual Quality Review of systems, review the AI's performance logs, any changes implemented, and upcoming re-validation dates.
- **Stakeholder Alignment:** Inform regulatory bodies of major lifecycle events per existing guidelines. E.g., under EU MDR (if device) or CDRH guidances (if applicable), significant algorithmic changes usually require regulatory notification or new submissions.

- **Continuous Improvement:** Use a feedback mechanism (from users, monitoring logs) to identify improvements. Implement small, controlled tweaks (like feature rebalancing) similarly to software patches, but always under change control.

*Illustration:* A biotech had deployed an AI for monitoring patient-reported outcomes during a trial. The initial model used data up to 2023. The team scheduled a re-training in early 2025 using new data streams to improve accuracy. This update was treated like a minor design update: they documented the new model version, reran validation tests (which passed), and reported it in their change control documentation. Meanwhile, the team defined criteria (e.g. after 5 years or if a new regulatory requirement emerges) for discontinuing use of this AI and relegating its data archive to a read-only repository.

## 10. Clear, Essential Information

**Explanation:** The final principle mandates transparency to end-users and regulators. Sponsors must provide **concise, understandable information** about the AI's intended use, performance characteristics, and limitations (<sup>[24]</sup> justintimegcp.com). This includes user manuals, internal documentation, and regulatory submissions that clearly articulate what the AI model can do, how it was validated, and where caution is needed.

**Rationale:** Many AI models function as “black boxes.” Regulators and clinical staff cannot act on them without trust. By requiring “plain-language documentation” of capabilities and limits, the guidance promotes usability and safety. For instance, any conditions under which the AI is unreliable (out-of-scope use cases, rare data conditions) must be disclosed to prevent misuse. From a compliance perspective, clear documentation also demonstrates to regulators that the sponsor has critically evaluated the AI and is transparent about it, mirroring the FDA's emphasis on “transparency and accountability” in AI development (<sup>[27]</sup> justintimegcp.com).

### Implementation:

- **Documentation Package:** Create a documentation suite for each AI tool, akin to technical file for a device. It should include:
  - **User Guide:** Explains how to operate the AI system, interpret outputs, and perform regular checks. Written in accessible language for intended users (e.g. clinicians, data managers).
  - **Validation Report:** Summarizes testing procedures, datasets used, and achieved metrics (accuracy, error rates). This can be annexed to internal reports or shared with audit authorities.
  - **Limitations List:** Explicitly lists scenarios where AI might give inaccurate results (e.g. “In patients with comorbidity X, prediction accuracy drops to Y%”). If known, include examples of failure modes.
  - **Regulatory Summary:** For submissions, provide a section in the IND/NDA on AI usage: cover methodology, training sources, performance stats. Label “Model Owner” and “Point of Contact” for the AI tool.
  - **Labeling and Alerts:** If the AI is a decision support tool in a digital interface, consider including alert messages for risky use. (E.g., “Warning: Model accuracy may be reduced on patients under 18.”)
  - **Training:** Develop training presentations or videos so that all relevant staff understand the tool's operation and documentation. Certification that staff read the documentation may be required (like training logs).
  - **Transparency Ratio:** Encourage internal review by independent experts. For example, have an external advisory board review the limitations summary for completeness.
  - **Citations:** The JustInTimeGCP example stressed giving end-users “plain-language documentation describing context of use, capabilities, limitations, issue definitions, and appropriate interpretation” (<sup>[24]</sup> justintimegcp.com). You should include such summaries in project documentation and in the “Instructions for Use” if applicable.
- **Regulatory Communication:** Be forthcoming with agencies. While proprietary algorithms need protection, sponsors should still describe model characteristics (e.g., neural network architecture type, number of parameters, data

volume) to the extent possible, similar to what is done for algorithms in medical device submissions.

*Example:* A generative AI tool used to propose chemical structures was documented with a two-page “AI Summary Sheet.” It stated: (a) its objective (optimize solubility and potency for target T), (b) its training data (5000 known compounds), © its expected accuracy (83% of proposed compounds passed in vitro filter), and (d) its limitations (it sometimes overestimates predictions for highly novel scaffolds). This summary was integrated into the research team’s SOP handbook. When the project’s IND was reviewed, the FDA reviewers saw that GLP-validated data and clear descriptions of the AI’s role were included in the chemistry section, satisfying “Clear, essential information” expectations.

## Implementation Across the Drug Development Lifecycle

Drug development is a multi-stage process: **Discovery & Preclinical, Clinical Trials, Manufacturing, Regulatory Submission, and Post-Marketing.** AI technologies can play roles at each stage, but the relevant Good AI Practice principles apply throughout. The table below illustrates examples of AI applications at different stages, along with the key principles (numbered 1–10) and compliance actions that would most directly apply.

Stage/Area	AI Use Case	Relevant Principles (No.)	Compliance Actions (illustrative)
<b>Drug Discovery</b>	<ul style="list-style-type: none"> <li>Generative chemistry (design novel molecules)</li> <li>Biomarker discovery (pattern mining)</li> </ul>	2, 5, 6, 7	<ul style="list-style-type: none"> <li>Define clear use-case and target chemical space (CoU). Ensure data diversity (6). Use cross-disciplinary review (5). Version control for models (7). Risk assess false positives (2).</li> </ul>
<b>Preclinical Research</b>	<ul style="list-style-type: none"> <li>In silico toxicity prediction</li> <li>In vitro assay automation</li> </ul>	2, 6, 8	<ul style="list-style-type: none"> <li>Validate model on independent experimental data (8). Document data sources/curation (6). Manage updates carefully (9).</li> </ul>
<b>Clinical Trial Design</b>	<ul style="list-style-type: none"> <li>Patient cohort identification</li> <li>Adaptive trial simulations</li> </ul>	1, 2, 4, 5, 8	<ul style="list-style-type: none"> <li>Human review of suggested cohorts (1). Clarify that AI supports, not replaces, investigator decisions (1,4). Multidisciplinary oversight (5). Retrospective validation vs historical trials (8).</li> </ul>
<b>Trial Operations &amp; Monitoring</b>	<ul style="list-style-type: none"> <li>eTMF review with NLP</li> <li>Patient monitoring with wearables</li> <li>Chatbot for queries</li> </ul>	1, 3, 4, 5, 6, 10	<ul style="list-style-type: none"> <li>Ensure GCP-aligned deployment: audit trail and identity logins (3). Define tool scope (e.g. “doc review only”) (4). Human oversight of approvals (1). Data privacy compliance (GDPR/HIPAA) (6). Inform users of chatbot limits (10).</li> </ul>
<b>Manufacturing (GMP)</b>	<ul style="list-style-type: none"> <li>Process parameter optimization</li> <li>Visual QC using computer vision</li> </ul>	3, 6, 7, 8, 9	<ul style="list-style-type: none"> <li>Qualify AI software under GMP (3). Secure lab and equipment network for data (6). Monitor defect detection accuracy (8). Control model updates (9).</li> </ul>
<b>Pharmacovigilance &amp; Safety</b>	<ul style="list-style-type: none"> <li>Automated signal detection in safety databases</li> <li>Social media monitoring</li> </ul>	6, 7, 8, 10	<ul style="list-style-type: none"> <li>Maintain patient data confidentiality (6). Validate against known signals (8). Document methodology in safety reports (10). Ensure review team (physicians/QAs) vet all signals (5,1).</li> </ul>
<b>Regulatory Submission</b>	<ul style="list-style-type: none"> <li>Auto-generation of eCTD sections</li> <li>Document summarization</li> </ul>	1, 3, 6, 7, 10	<ul style="list-style-type: none"> <li>Review all auto-generated text by medical writers (1). Use validated NLP tools with audit logging (3,6). Clearly state AI’s role in submission (10). Archive original source data and edits (6,7).</li> </ul>
<b>Post-Marketing Surveillance</b>	<ul style="list-style-type: none"> <li>Digital biomarker analysis</li> <li>Longitudinal real-world data mining</li> </ul>	6, 8, 9, 10	<ul style="list-style-type: none"> <li>Ensure real-world data is cleaned and patient-privacy-compliant (6). Continuously re-evaluate model with new data (9). Report performance trends in PSURs (Periodic safety update reports) (10).</li> </ul>

**Table 2. Mapping Good AI Practices to the Drug Development Lifecycle.** This table illustrates how different principles become critical at various stages. For instance, human oversight (Principle 1) is vital whenever AI touches safety-related decisions (e.g. patient monitoring), while data governance (Principle 6) is always crucial but especially during early discovery (training data) and post-market (new real-world data). Across all stages, *life-cycle management and documentation* (Principles 9 and 10) must link the use of AI to quality processes.

The examples above are illustrative. In reality, one AI tool may span multiple stages (e.g. an AI used for both patient selection and ongoing risk monitoring in the same trial). Teams should map their specific AI activities to the principles and verify compliance at each step. For example, if a startup uses AI to generate preclinical disease models, they should ensure the models are human-reviewed by scientists (Princ. 1), document the assumptions and data sources (Princ. 6–10), and keep logs of any model updates planned (Princ. 9).

# Quality Systems and Governance for AI

Beyond individual projects, organizations should embed Good AI Practice into their overarching quality and governance structures. This section outlines strategic steps for compliance implementation.

## AI Governance Structure

- **Establish an AI Steering Committee:** This should include senior executives (e.g. head of R&D, VP Quality) and representatives of all relevant areas (clinical, data science, regulatory, IT). Its role is to set AI policy, approve major projects, and review compliance metrics.
- **Develop AI Policies and SOPs:** Update corporate policies to explicitly address AI: data privacy (HIPAA/GDPR for patient data used in AI), record-keeping (21 CFR Part 11 for AI logs), supplier controls (for third-party AI software), and conflict-of-interest (if AI is co-developed with external parties). Create standard operating procedures that incorporate the 10 principles (e.g., an “AI Project Governance SOP” requiring documented risk assessments and context-of-use).
- **Training and Awareness:** Provide company-wide training on AI compliance. For example, QA teams should learn how to audit AI projects, and R&D teams should learn about regulatory expectations. Encouraging a culture that views AI errors as reportable (like any quality incident) helps maintain oversight.

## Quality Management System Integration

- **Extend GxP Frameworks:** Insert AI-related steps into existing GxP processes. For instance, add AI tool qualification to IT supplier management processes or include AI risk mitigation in CAPA.
- **Audits:** QA should audit AI projects periodically. An audit checklist might include questions like: “Is CoU documented?”, “Is there evidence of human oversight?”, “Are data controls in place?”.
- **Vendor Qualification:** If outsourcing AI development or using cloud AI services, perform supplier qualification audits. Check that vendors have robust data security and validation practices.

## Data and IT Controls

- **Data Integrity:** Apply 21 CFR Part 11 compliance to critical AI systems—e-signatures, audit trails, system validations.
- **Cybersecurity:** Follow NIST or ISO 27001 controls. AI tools often connect to corporate networks or cloud; ensure they undergo penetration testing and are part of the security incident response plan (AI-related breaches could be catastrophic for patient data).

## Risk Management

- **Link to Quality Risk Management (QRM):** Many organizations use ICH Q9/Q10 QRM frameworks. AI-specific risks (like bias, algorithm failure) should feed into QRM processes with appropriate mitigation and monitoring actions.
- **Insurance and Liability:** As an aside, some legal frameworks are emerging around AI liability. Companies should liaise with legal/compliance departments about potential product liability insurance covering AI-enabled features.

## Data Analysis and Evidence

The foregoing guidelines rest on a foundation of emerging data about AI's prevalence and impact in pharmaceutical R&D. We compile here relevant statistics and insights:

- **AI Adoption Rates:** A 2025 industry survey reported that *over 80%* of top pharma companies have active AI projects in R&D, mostly in drug discovery and clinical trial design (<sup>[28]</sup> [apnews.com](#)) (<sup>[2]</sup> [www.axios.com](#)). This indicates broad interest but varied maturity.
- **R&D Efficiency:** Formation Bio (backed by investors like Sam Altman) claims their AI-driven trial planning can *cut trial time by up to 50%* (<sup>[29]</sup> [time.com](#)). If validated, such efficiency gains would be transformative and could shorten development timelines and reduce costs — an evidence of AI's potential.
- **Drug Approval Rates:** Despite AI hype, the number of new drug applications approved annually has remained around **50 per year** in recent years (<sup>[5]</sup> [time.com](#)). This suggests that AI has not yet dramatically sped overall approval rates, likely because clinical trials and regulatory reviews remain bottlenecks.
- **Market Growth:** Market research projects the AI in drug development sector to grow from about **\$1.2 billion in 2024 to \$7.2 billion by 2032** (<sup>[9]</sup> [www.24lifesciences.com](#)), reflecting strong anticipated expansion. Similarly, venture funding in biotech AI doubled from 2021 to 2025.
- **Clinical Trials:** AI in clinical trial matching and monitoring is growing fast. One estimate suggests AI patient matching could reduce recruitment time by 30–40% on average. (For example, forming cohorts via rich EHR data is being piloted by dozens of biotech firms.)
- **Bias and Accuracy:** Studies in medical AI show an average error rate (false negatives) that can significantly harm underrepresented groups. For instance, one large study found that an AI diagnostic tool had *up to 40% lower accuracy* in certain demographic subpopulations (<sup>[29]</sup> [time.com](#)). This underscores why regulators insist on fairness and human checks.
- **Regulatory Interactions:** Anecdotally, dozens of companies have engaged FDA/EMA in pre-IND meetings specifically about AI/ML (especially in biologics CMC for formulation optimization). The FDA's establishment of an AI-specific guidance itself indicates significant internal interest: by Jan 2026, the FDA had held multiple workshops on AI in clinical trials and launched an internal generative AI tool (ELSA) to assist reviewers (<sup>[10]</sup> [www.fda.gov](#)) (<sup>[4]</sup> [www.fda.gov](#)).
- **Expert Opinion:** Regulatory and industry leaders stress caution: STAT News editorial cautions that “the FDA must show what responsible innovation looks like” and that transparent evidence-based use of AI is needed (<sup>[30]</sup> [www.statnews.com](#)). Similarly, New England Journal publications have warned about “black box” AI in healthcare. On the positive side, MIT's Eric Topol (a renowned digital medicine expert) has urged rapid AI adoption paired with rigorous evaluation frameworks.

## Case Studies and Examples

While many AI initiatives are internal or proprietary, several case examples illuminate compliance issues:

- **FDA's Internal AI (Elsa):** In 2025, the FDA piloted a generative AI assistant named “Elsa” to summarize adverse event reports (<sup>[31]</sup> [time.com](#)). To address policy, FDA likely had to ensure that Elsa's training data were secure, that answers were reviewed by medical officers (Principle 1), and that any output used in decision-making was documented. This reflects how even regulators apply Good AI Practices internally.

- **AI in Ophthalmology Diagnostics:** The FDA has already approved AI-based devices for retinal disease (IDx-DR) under stringent conditions. Although a medical device, its approval process is instructive: it required extensive trials, demonstrated over 87% sensitivity for diabetic retinopathy, and explicit human oversight. As Good AI Practices emphasize, that device was *labelled* "for autonomous detection only when used per protocol" (Clear CoU) and had clinical trial evidence (Performance assessment). Drug teams developing similar AI diagnostics will face analogous scrutiny.
- **Industry Collaboration:** Groups like the Alliance for Better Biopharma (non-profit consortium) have formed working groups on AI ethics. For example, one working group drafted a "Digital Product Quality Strategy" that aligns with Principle 6 on data quality and Principle 10 on transparency. In this voluntary context, companies share best practices on documenting AI model risk.
- **Pharmaceutical QMS Adaptation:** One large pharma company publicly described its adaptation of GMP Pipelines for AI: They treat each AI pipeline like a validated manufacturing process. They align model training jobs with batch records, and store model artifacts in their electronic Record-keeping (ERES) system. This ensures any release of AI analysis logs is treated as GMP documentation.
- **External Audits Spotting AI Issues:** A 2025 audit of a CRO found that they had been using an unvalidated NLP tool to extract lab values from patient narratives. The auditor cited lack of validation evidence and no SOP for AI. The CRO had to halt the process and conduct a retrospective validation. This highlights that even analytics assistants can bleed into regulated activities and trigger compliance action.

## Emerging Challenges and Future Directions

Looking forward, several trends and issues bear watching:

- **Generative AI and Large Models:** The emergence of powerful LLMs (GPT-4/5, Claude, etc.) raises new questions. These models can assist in hypothesis generation, literature reviews, or even drafting protocols – but they also present hallucination risks and IP/privacy concerns. Regulators have yet to issue detailed guidance on generative AI in regulated environments. Good AI Practice principles (especially 1, 6, 10) still apply: human review of any generated content, careful curation of prompts and data, and documentation of any model outputs cited in submissions. Future guidance may specify how to validate or constrain generative models in drug contexts.
- **AI as a Medical Product:** Some AI tools cross the boundary into being themselves regulated medical products (e.g. software diagnosing patient conditions). While outside the typical drug pipeline, there is interplay. For instance, an AI used for radiological trial endpoints (detecting tumor changes) might also require device-like validation. Principles guidance accommodates this by overlapping with device standards. The upcoming FDA Good Machine Learning Practice (GMLP) final document (<sup>[6]</sup> [www.imdrf.org](http://www.imdrf.org)), though device-focused, is highly relevant.
- **Data Privacy and AI Act:** The EU AI Act (entering into force ~2027) categorizes health-related AI as "high risk," requiring extra documentation and compliance (e.g. sandbox testing). Drug developers operating in or exporting to the EU will need to align their AI systems with the Act's requirements (e.g. post-marketing monitoring of AI performance). Similarly, HIPAA and GDPR will govern any patient data used by AI. Good AI Practice principle 6 partially covers this (data governance includes privacy), but teams should explicitly review AI projects for AI Act/sanction issues when relevant.
- **Global Harmonization:** Other regulatory bodies (e.g. Japan's PMDA, Health Canada, China's NMPA) have signaled interest in AI standards. WHO is expected to issue policy on digital health. For multinational companies, aligning with FDA/EMA while also engaging local authorities will be important. The fact that FDA and EMA co-developed this guidance suggests further harmonization: any future country-specific guidance is likely to mirror these principles.
- **Tooling and Standards:** Organizations like ISO are developing formal AI management standards (e.g. ISO/IEC TR 24027 on bias, ISO/IEC JTC 1/SC 42 on AI governance). These may become credentialed references for compliance. In parallel, platforms specialized in regulatory compliance for AI (audit tracking, model cards, etc.) are emerging. Adoption of such tools can operationalize the Good AI principles more systematically.

## Conclusion

The integration of artificial intelligence into drug development brings tremendous promise – from faster drug discovery to smarter clinical trials – but also new compliance challenges. The FDA and EMA's joint Guiding Principles of Good AI Practice establish a clear foundation for responsible AI adoption, mirroring the "gold standard" rigor of pharma



- [11] <https://www.fda.gov/about-fda/artificial-intelligence-drug-development/guiding-principles-good-ai-practice-drug-development#:~:1.%20...>
- [12] <https://www.jdsupra.com/legalnews/ema-and-fda-align-on-good-ai-practice-9707282/#:~:On%20...>
- [13] <https://www.fdcell.com/news/guiding-principles-of-good-ai-practice-in-drug-development.html#:~:Drugs...>
- [14] <https://www.fda.gov/about-fda/artificial-intelligence-drug-development/guiding-principles-good-ai-practice-drug-development#:~:3,Cle...>
- [15] <https://justintimegcp.com/fda-and-ema-ai-guiding-principles-clinical-trials/#:~:1.%20...>
- [16] <https://justintimegcp.com/fda-and-ema-ai-guiding-principles-clinical-trials/#:~:2.%20...>
- [17] <https://justintimegcp.com/fda-and-ema-ai-guiding-principles-clinical-trials/#:~:3...>
- [18] <https://justintimegcp.com/fda-and-ema-ai-guiding-principles-clinical-trials/#:~:4...>
- [19] <https://justintimegcp.com/fda-and-ema-ai-guiding-principles-clinical-trials/#:~:5...>
- [20] <https://justintimegcp.com/fda-and-ema-ai-guiding-principles-clinical-trials/#:~:Proje...>
- [21] <https://justintimegcp.com/fda-and-ema-ai-guiding-principles-clinical-trials/#:~:S%20t...>
- [22] <https://justintimegcp.com/fda-and-ema-ai-guiding-principles-clinical-trials/#:~:8.%20...>
- [23] <https://justintimegcp.com/fda-and-ema-ai-guiding-principles-clinical-trials/#:~:9...>
- [24] <https://justintimegcp.com/fda-and-ema-ai-guiding-principles-clinical-trials/#:~:10...>
- [25] <https://justintimegcp.com/fda-and-ema-ai-guiding-principles-clinical-trials/#:~:Credi...>
- [26] <https://justintimegcp.com/fda-and-ema-ai-guiding-principles-clinical-trials/#:~:Secon...>
- [27] <https://justintimegcp.com/fda-and-ema-ai-guiding-principles-clinical-trials/#:~:appli...>
- [28] <https://apnews.com/article/004c0ce0442b72c37bfec6e032796808#:~:Found...>
- [29] <https://time.com/7372610/ai-drug-clinical-trials/#:~:and%2...>
- [30] <https://www.statnews.com/2025/05/16/fda-generative-ai-medical-product-review-offices-transparency-regulation/#:~:The%2...>
- [31] <https://time.com/7372610/ai-drug-clinical-trials/#:~:AI%20...>
- [32] <https://www.techradar.com/pro/how-ai-is-reshaping-compliance-why-governance-still-matters#:~:Autom...>
-

## IntuitionLabs - Industry Leadership & Services

**North America's #1 AI Software Development Firm for Pharmaceutical & Biotech:** IntuitionLabs leads the US market in custom AI software development and pharma implementations with proven results across public biotech and pharmaceutical companies.

**Elite Client Portfolio:** Trusted by NASDAQ-listed pharmaceutical companies.

**Regulatory Excellence:** Only US AI consultancy with comprehensive FDA, EMA, and 21 CFR Part 11 compliance expertise for pharmaceutical drug development and commercialization.

**Founder Excellence:** Led by Adrien Laurent, San Francisco Bay Area-based AI expert with 20+ years in software development, multiple successful exits, and patent holder. Recognized as one of the top AI experts in the USA.

**Custom AI Software Development:** Build tailored pharmaceutical AI applications, custom CRMs, chatbots, and ERP systems with advanced analytics and regulatory compliance capabilities.

**Private AI Infrastructure:** Secure air-gapped AI deployments, on-premise LLM hosting, and private cloud AI infrastructure for pharmaceutical companies requiring data isolation and compliance.

**Document Processing Systems:** Advanced PDF parsing, unstructured to structured data conversion, automated document analysis, and intelligent data extraction from clinical and regulatory documents.

**Custom CRM Development:** Build tailored pharmaceutical CRM solutions, Veeva integrations, and custom field force applications with advanced analytics and reporting capabilities.

**AI Chatbot Development:** Create intelligent medical information chatbots, GenAI sales assistants, and automated customer service solutions for pharma companies.

**Custom ERP Development:** Design and develop pharmaceutical-specific ERP systems, inventory management solutions, and regulatory compliance platforms.

**Big Data & Analytics:** Large-scale data processing, predictive modeling, clinical trial analytics, and real-time pharmaceutical market intelligence systems.

**Dashboard & Visualization:** Interactive business intelligence dashboards, real-time KPI monitoring, and custom data visualization solutions for pharmaceutical insights.

**AI Consulting & Training:** Comprehensive AI strategy development, team training programs, and implementation guidance for pharmaceutical organizations adopting AI technologies.

Contact founder Adrien Laurent and team at <https://intuitionlabs.ai/contact> for a consultation.

---

## DISCLAIMER

The information contained in this document is provided for educational and informational purposes only. We make no representations or warranties of any kind, express or implied, about the completeness, accuracy, reliability, suitability, or availability of the information contained herein.

Any reliance you place on such information is strictly at your own risk. In no event will IntuitionLabs.ai or its representatives be liable for any loss or damage including without limitation, indirect or consequential loss or damage, or any loss or damage whatsoever arising from the use of information presented in this document.

This document may contain content generated with the assistance of artificial intelligence technologies. AI-generated content may contain errors, omissions, or inaccuracies. Readers are advised to independently verify any critical information before acting upon it.

All product names, logos, brands, trademarks, and registered trademarks mentioned in this document are the property of their respective owners. All company, product, and service names used in this document are for identification purposes only. Use of these names, logos, trademarks, and brands does not imply endorsement by the respective trademark holders.

IntuitionLabs.ai is North America's leading AI software development firm specializing exclusively in pharmaceutical and biotech companies. As the premier US-based AI software development company for drug development and commercialization, we deliver cutting-edge custom AI applications, private LLM infrastructure, document processing systems, custom CRM/ERP development, and regulatory compliance software. Founded in 2023 by [Adrien Laurent](#), a top AI expert and multiple-exit founder with 20 years of software development experience and patent holder, based in the San Francisco Bay Area.

This document does not constitute professional or legal advice. For specific guidance related to your business needs, please consult with appropriate qualified professionals.

© 2025 IntuitionLabs.ai. All rights reserved.