

FDA-EMA Good AI Practice Guidelines in Drug Development

By Adrien Laurent, CEO at IntuitionLabs • 3/3/2026 • 45 min read

fda ema ai guidelines

good ai practice

ai in drug development

pharmaceutical ai compliance

ai governance

regulatory affairs

machine learning pharma

ai risk management



FDA-EMA Good AI Practice Guidelines in Drug Development

intuitionlabs.ai

Executive Summary

The adoption of artificial intelligence (AI) technologies across the pharmaceutical product lifecycle is accelerating rapidly, promising to **boost innovation and dramatically shorten drug development timelines** (^[1] [pharmprom.net](#)) (^[2] [www.fiercepharma.com](#)). However, AI's deployment raises novel challenges for **product quality, patient safety, and data integrity**. In January 2026, the U.S. Food and Drug Administration (FDA) and the European Medicines Agency (EMA) issued a joint set of **10 Guiding Principles of Good AI Practice in Drug Development** (^[3] [www.fda.gov](#)) ([www.ema.europa.eu](#)). These principles – ranging from “Human-centric by design” to “Clear, essential information” – provide a high-level framework to ensure AI tools are developed and used responsibly throughout all phases of drug R&D, manufacturing, and **pharmacovigilance** (^[4] [www.fda.gov](#)) (^[1] [pharmprom.net](#)). This report provides a comprehensive analysis of these Guidelines in context: the **historical impetus** for regulating AI in medicine, the **technical and ethical rationale** behind each principle, and the **implications for industry and regulators**. We examine how these principles align with existing regulations (GxP, ICH Q9, GDPR, etc.), feature case studies (e.g. AI-driven **patient recruitment in clinical trials** (^[5] [pmc.ncbi.nlm.nih.gov](#)), predictive manufacturing analytics (^[6] [www.sciencedirect.com](#))), and survey data (e.g. 85% of top pharma leaders now view AI as an “immediate priority” (^[2] [www.fiercepharma.com](#))). Finally, we discuss future directions, such as anticipated guidance and harmonisation in an evolving global AI regulatory landscape.

Introduction and Background

Artificial intelligence (AI) and machine learning (ML) techniques — including deep learning, natural language processing, and predictive analytics — are increasingly being applied in pharmaceutical research and development (^[1] [pharmprom.net](#)) (^[7] [www.fda.gov](#)). These tools can generate new scientific insights from large datasets, optimize experimental design, repurpose existing drugs, and improve safety monitoring. AI-driven models have already shown promise in **molecular design, patient stratification, manufacturing process control, and pharmacovigilance** (post-market safety signal detection) (^[1] [pharmprom.net](#)) ([www.ema.europa.eu](#)). For example, modern AI systems have been used to speed patient matching in clinical trials – one report showed an AI platform reduced patient-screening time by 34% (^[5] [pmc.ncbi.nlm.nih.gov](#)) – and to analyze real-world health records for early detection of adverse drug reactions ([www.ema.europa.eu](#)).

Industry surveys confirm this rapid uptake: a 2025 survey of senior executives from leading pharmaceutical companies found that **85% of the top 20 biopharma firms consider AI an “immediate priority”** (^[2] [www.fiercepharma.com](#)). Most have ramped up investment, with 80% creating dedicated **AI governance structures** focused on **ethics and data security** (^[8] [www.fiercepharma.com](#)). Encouraged by AI's potential, companies invest in diverse use cases — from medical writing and document translation to predictive modeling of clinical outcomes. Conversely, regulators and public stakeholders worry about **unintended consequences**: biased models, poor data practices, and opaque decision-making could jeopardize patient safety and erode trust ([www.ema.europa.eu](#)) (^[9] [www.axios.com](#)).

Regulators have increasingly recognized the need for guardrails on AI in healthcare. In late 2024, EMA released a “Reflection Paper” on AI in the medicinal product lifecycle, outlining opportunities (e.g. reducing animal testing through predictive toxicology models) as well as challenges (e.g. algorithmic bias, technical failures) ([www.ema.europa.eu](#)) ([www.ema.europa.eu](#)). Similarly, in January 2025, FDA issued its **first draft guidance** on AI/ML in drug development, introducing a **risk-based framework** for assessing model credibility and emphasizing the importance of clearly defining a model's “*context-of-use*” (^[10] [www.fda.gov](#)). Industry input (800+ external comments and 500 actual AI-enabled submissions since 2016 (^[11] [www.fda.gov](#))) highlighted that a one-size-fits-all approach is impractical. Instead, oversight must **scale with risk**: a data-mining tool in exploratory research should be evaluated less stringently than an AI system that influences dose selection or adverse event reporting (^[12] [www.linkedin.com](#)) (^[10] [www.fda.gov](#)).

Against this backdrop, the FDA and EMA jointly identified ten high-level principles to **harmonize international expectations** for AI in drug development (^[3] www.fda.gov) (www.ema.europa.eu). These principles are not prescriptive rules but broad commitments (“good AI practice”) that stakeholders can interpret within existing regulatory frameworks. They embody enduring values like patient welfare, scientific rigor, and accountability, as applied to AI [systems](#). As EMA Commissioner Várhelyi emphasized, these principles represent a “*first step of a renewed EU-US cooperation*” in biotechnology and aim to bring both innovation and safeguards to medicine development (www.ema.europa.eu).

This report chronicles the emergence of these principles, unpacks each one in depth, examines data and case studies on AI effectiveness and risks, compares regulatory approaches, and explores future regulatory paths. Throughout, **citations to primary sources** and recent studies support every key point, ensuring an evidence-based, professional analysis.

Historical Context and Regulatory Drivers

Evolution of AI in Drug Development

Historically, drug development processes have evolved from empirical chemistry to complex biotechnology. Over the last decade, digitization and data generation (genomics, proteomics, electronic health records, etc.) have created fertile ground for AI tools. Early AI applications were limited (e.g. QSAR modeling, image analysis), but recent advances in computational power and algorithms have vastly expanded possibilities. For instance, generative models can now design novel molecules in silico, as exemplified by biotech companies like Exscientia and Insitro collaborating with major pharma on drug discovery projects (^[13] apnews.com) (^[14] apnews.com). These partnerships underscore that cutting-edge quantitative biology (‘omics’) paired with ML can uncover subtle patterns in disease biology that manual analysis might miss (^[13] apnews.com).

Simultaneously, the pharmaceutical industry’s cost and time pressures (the average time to develop a new drug remains ~10-12 years) have made efficiency gains via AI highly attractive. As Insitro’s CEO remarks, ML can “unravel the underlying complexity of heterogeneous diseases” and test therapies on nuanced patient subsets that traditional methods might overlook (^[15] apnews.com). Companies like AstraZeneca and GSK have similarly announced broad AI initiatives across drug discovery and development stages (^[2] www.fiercepharma.com). Real-world analytics for safety and adherence, such as AI-driven pharmacovigilance (studied in recent literature (^[16] pmc.ncbi.nlm.nih.gov)), further illustrate how AI is weaving into regulatory affairs and post-marketing surveillance.

This rapid growth, however, has outpaced formal guidance. Regulators worldwide (beyond FDA/EMA) have released general ethics guidelines for AI. Notably, the World Health Organization’s 2021 report on AI in health explicitly called for putting “*ethics and human rights... at the heart of [AI’s] design, deployment and use*” (^[9] www.axios.com), echoing what would become Principle 1 of the FDA-EMA framework. In parallel, global initiatives like the EU’s AI Act (2023) and the U.S. executive orders on trustworthy AI set broad expectations for transparency, fairness, and human oversight. Yet those act at a macro level; specific guidance for **AI in medicine development** remained fragmented.

Precedent Food and Drug Guidelines

Before the joint principles, the FDA had begun articulating how AI should fit within existing regulatory processes. The January 2025 FDA news release (^[17] www.fda.gov) made clear this was the *first* agency guidance on AI in drugs. FDA Commissioner Califf highlighted a “*risk-based framework... that promotes innovation and ensures... robust scientific and regulatory standards*” (^[18] www.fda.gov). The release noted that FDA reviewers had already encountered AI components in over **500 submissions** since 2016 (^[11] www.fda.gov), confirming AI’s ubiquity. FDA’s draft guidance (open for comment) laid out concepts of **context-of-use** and **model credibility**, requiring sponsors to tailor validation activities to each model’s intended decision-making role (^[10] www.fda.gov). Similarly, FDA’s earlier “Good Machine Learning Practice” for

medical devices (2021) had advocated for rigorous software development lifecycle and risk management for ML algorithms (^[19] www.fda.gov).

On the European side, the EMA (and its network) convened expert working groups. The 2024 EMA Reflection Paper (www.ema.europa.eu) (www.ema.europa.eu) invited public consultation on AI usage, stressing both scientific promise and regulatory challenges. Its findings – that AI can aid everything from “replacing animal models” to post-marketing signal detection (www.ema.europa.eu) – highlighted the **breadth of AI applications**. EMA also emphasized fundamental constraints: AI in medicines must comply with EU law and ethics (“human-centric”, respect rights) (www.ema.europa.eu). At a joint FDA-EMA meeting in April 2024, both agencies pledged closer cooperation on novel technologies, setting the stage for the 2026 principles. (The EMA press release quoted Commissioner Várhelyi on working “together on the two sides of the Atlantic... ensuring the highest level of patient safety” (www.ema.europa.eu).

Other stakeholders have similarly advocated guidelines. For example, industry consortia and NGOs have produced AI governance frameworks (e.g. IBM’s AI Ethics). In pharma itself, quality and AI thought-leaders have written on best practices; a July 2024 ISPE article even laid out a detailed AI governance framework *within* the existing GxP QMS (^[20] ispe.org) (^[21] ispe.org). Such efforts converged on key themes (risk, transparency, interdisciplinary oversight), which the FDA-EMA collaboration distilled into the 10 Principles.

The FDA-EMA Guiding Principles of Good AI Practice

The **ten Guiding Principles** announced in January 2026 (^[4] www.fda.gov) (^[1] pharmprom.net) are the cornerstone of this report. They are stated as high-level directives for **everyone involved in medicine development** – sponsors, developers, regulators, etc. – to ensure AI methods are used “safely and effectively” across the drug lifecycle (^[22] formiventos.com). Below, each principle is examined in depth: its meaning, rationale, and key considerations for compliance, along with relevant literature and examples.

1. Human-Centric by Design

Principle statement: “The development and use of AI technologies must align with ethical and human-centric values.” (^[23] pharmprom.net) (www.ema.europa.eu).

Explanation: This principle underscores that AI tools should ultimately serve people (patients, clinicians, scientists) and respect their rights. “Human-centric” encompasses both ethical safeguards and practical oversight. It means that AI systems should be designed *from the outset* to involve qualified human expertise, ensure fairness, and protect patient welfare (^[9] www.axios.com) (^[24] www.linkedin.com). Ethical values (privacy, informed consent, equity) must guide development as a default.

For example, the WHO’s “AI in Healthcare” report insisted that “*only if ethics and human rights are put at the heart of [AI’s] design*” can society realize AI’s benefits (^[9] www.axios.com). Similarly, EMA’s reflection paper explicitly stated a human-centric approach, demanding compliance with ethics and fundamental rights (www.ema.europa.eu). Practically, this means no algorithm should make critical decisions without human enquiry. Designers must build in mechanisms like **human-in-the-loop (HITL)**: the system architecture should assume skilled professionals review and, if necessary, override AI outputs. Crucially, those humans must themselves have domain expertise. Quoting AI experts, one must ask: “*if your ‘human oversight’ is someone who couldn’t do the work themselves, you haven’t met the principle*” (^[25] www.linkedin.com).

Rationale: Medical practice is fundamentally trust-based — patients trust clinicians, and with AI in the mix, that trust must not erode. Clear accountability is vital. “Human-centric” also implies that AI should *augment* human judgment, not replace

it. It demands transparency: patients (and providers) should be informed when AI is used (e.g. in diagnostics or trial matching). Moreover, the principle echoes the healthcare credo “do no harm” applied to AI systems (^[9] www.axios.com). In effect, AI developers must anticipate ethical issues like bias, and incorporate bias-mitigation strategies and audit trails.

Implementation considerations:

- **Ethical review:** Integrate ethical risk assessment (e.g. as part of Institutional Review Board processes) whenever AI tools are developed. Ensure patient data usage complies with privacy laws (HIPAA, GDPR) and cultural norms.
- **Human oversight:** Define explicitly *who* will review AI outputs. This should be an expert capable of independently understanding the task. Regulators have emphasized that AI systems should not autonomously drive high-stakes decisions (e.g. drug dosing) without qualified human judgement (^[12] www.linkedin.com).
- **Fairness and inclusivity:** Check models for biases (e.g. by demographics) especially if the AI tool will impact diverse patient populations. For instance, an AI model trained only on European patient data may underperform in other ethnic groups.
- **Transparency to patients:** When AI systems affect patient care or safety reporting, clearly communicate their role in patient-friendly terms (see principle 10). This fosters trust and informed consent.

Example: Consider an AI algorithm that predicts which trial participants are likely to respond to a new oncology drug. A human-centric design would involve oncologists in both training and evaluating the model, and ensure that the model can explain its reasoning (e.g. highlighting key clinical factors). Patients selected by the model would still have clinicians confirm eligibility. In contrast, a purely automated “black-box” screening (with no expert review) would violate Principle 1 and Principle 4 (context of use).

2. Risk-Based Approach

Principle statement: “Validation, risk mitigation, and oversight should be proportionate to the context of use and the determined model risk.” (^[26] pharmprom.net) (^[27] www.fdcell.com).

Explanation: Not all AI systems pose equal risk. This principle mandates a **risk classification** of AI tools based on their potential impact on patient safety, product quality, and regulatory decisions. High-impact applications (e.g. AI used to adjust dosing or detect adverse events) require rigorous validation and controls, whereas low-stakes pilot or hypothesis-generating models need less stringent oversight (^[12] www.linkedin.com) (^[10] www.fda.gov). The guidance explicitly states that oversight “should be proportionate to the context of use and...model risk” (^[27] www.fdcell.com).

Regulatory guidance already adopts risk-based frameworks for medical devices (Class I–III, etc.) and clinical trials (minimal vs. more-than-minimal risk). Here, the idea is analogous. The FDA’s draft guidance on AI/ML in drugs defines “context of use” as how the model is used to answer a specific question and emphasizes framing model credibility requirements accordingly (^[10] www.fda.gov). In practice, an AI model that performs exploratory data mining in bioinformatics (low immediate risk) might undergo cursory checks. In contrast, an AI tool embedded in a clinical monitoring dashboard must be validated against real patient outcomes and include fail-safes.

Data/Evidence: Stakeholders have resonated with this approach. In LinkedIn commentary, AI practitioners noted that “not all AI risk is equal” and shared an experience from pharmacovigilance modeling: exploratory literature mining was “informing thinking” and errors were tolerable, but when AI was used to flag drug safety signals (which could influence patient safety), it required “stricter validation, clearer audit trails, [and] explicit human oversight” (^[12] www.linkedin.com). This anecdote underlines the common-sense that AI with potential to harm patients or mislead regulators demands far more control.

Implementation considerations:

- **Context-of-use definition:** Before developing or validating an AI model, clearly document *what decision* it supports. FDA guidance treats this like a machine learning analog of a “predetermined application”—for example, “predicting tumor response in Phase II trials”.
- **Model risk assessment:** Convene cross-functional teams (regulatory, quality, technical) to classify the model's risk. Consider factors such as severity of potential harm if wrong, complexity of the model, and maturity of the technology.
- **Validation scope:** For high-risk use-cases, perform extensive performance testing (e.g. retrospective/prospective studies, simulations under stress scenarios). For low-risk pilots, lighter-weight testing may suffice.
- **Tiered governance:** Align with established quality risk management (QRM) processes. For instance, the ISPE guidance on AI suggests integrating AI risk classification into ICH Q9 processes (^[20] [ispe.org](https://www.ispe.org)). Document decisions just as one would for any clinical system.

Example: An AI model that optimizes dosages of a precision medicine would be tagged “high risk,” triggering comprehensive validation including independent dataset testing and an on-site audit of code. By contrast, a web-based chat interface helping researchers retrieve literature citations (a lower-risk tool) might simply have routine code review and end-user feedback monitoring, with less formal validation.

3. Adherence to Standards (GxP)

Principle statement: “AI technologies must adhere to relevant legal, ethical, technical, and scientific standards, including Good Practices (GxP) and cybersecurity requirements.” (^[28] pharmprom.net) (^[29] www.fdcell.com).

Explanation: AI systems in drug development must be treated like any other regulated tool. This means compliance with existing standards for pharmaceuticals and medical devices. In practice, adherence spans multiple domains:

- **GxP (Good Practices):** If AI software is used in clinical trials, Good Clinical Practice (GCP) and 21 CFR Part 11 for electronic records apply. In manufacturing, Good Manufacturing Practice (GMP) and Good Automated Manufacturing Practice (GAMP®) guidelines govern. For laboratory data, Good Laboratory Practice (GLP) may apply. The principle explicitly calls out “Good Practices” (GxP) (^[29] www.fdcell.com). AI governance documents have noted that AI tools should be integrated into the existing Quality Management System rather than treated as exceptions (^[21] [ispe.org](https://www.ispe.org)).
- **Data Integrity:** Any data used by AI must meet integrity criteria (ALCOA+ – attributable, legible, contemporaneous, original, accurate). Regulators have repeatedly enforced 21 CFR Part 11 and PIC/S Annex on computerized systems. An AI pipeline must have secure, version-controlled data and code, with audit trails.
- **Cybersecurity:** AI systems can be vulnerable to data breaches or adversarial attacks (e.g. maliciously altered inputs). The FDA's broader push for a Software Bill of Materials (SBOM) and robust cybersecurity in medical products also extends here. Principle 3 specifically includes cybersecurity as a key standard (^[28] pharmprom.net).
- **International Standards:** Emerging standards such as ISO/IEC 42001 (AI management systems) and ISO 13485 (medical devices QMS) should inform implementation. For example, Mintanciyan et al. recommend integrating AI governance with pharmacopeial quality standards, implying use of GMP guidance (PIC/S) and even ISO standards as appropriate (^[21] [ispe.org](https://www.ispe.org)).
- **Ethical/Legal Standards:** Beyond technical, this includes legal (data privacy regulations, intellectual property rights) and ethical standards (codes of conduct, industry ethics panels).

Data/Evidence: A preprint study of pharmaceutical manufacturing highlights this synergy. It notes that successful AI deployment in GMP environments “requires a thorough understanding of Good Manufacturing Practice (GMP) requirements combined with practical predictive performance” (^[6] www.sciencedirect.com). In other words, AI models must not contravene production quality rules; they should be considered part of the validated control strategy. Similarly, recent FDA internal communications (though context is device-centric) stress that proprietary AI algorithms need to be documented under the same validation SOPs as other software.

Implementation considerations:

- **Quality Systems:** Explicitly incorporate AI tools into the organization's QMS. This includes change control for AI model updates, CAPA processes for AI-related issues, and periodic review. The ISPE article suggests implementing ICH Q9 processes for AI and using existing QMS frameworks (PIC/S, 21 CFR 820) to provide oversight (^[20] ispe.org) (^[21] ispe.org).
- **Documentation:** Maintain complete records of model development and operations per GxP documentation standards. This includes code repositories, data provenance logs, validation test reports, and user manuals.
- **Cybersecurity Measures:** Follow FDA/DPI guidance on software security. Perform penetration testing on AI platforms, ensure data encryption and access controls, and plan response measures for breaches.
- **Third-Party Tools:** If using vendor software or cloud AI platforms, verify they comply with GxP requirements. For example, cloud providers often have GxP-compliant services, but it is incumbent on sponsors to document this compliance.

Example: A machine learning model predicting drug stability might be developed in an R&D environment, but before applying it to GMP release testing, its entire workflow (data inputs, processing, model code, outputs) should be validated similarly to a conventional analytical instrument per GMP Annex 11. This ensures traceability and that any changes (e.g. model retraining) are controlled via formal change management.

4. Clear Context of Use

Principle statement: *“Every technology must have a well-defined context of use – clearly specifying the role and scope for why it is being used.”* (^[30] pharmprom.net).

Explanation: Context of Use (COU) is a keystone concept borrowed from regulatory science (e.g. FDA's Biomarker COU in drug trials). In the AI domain, COU means precisely stating **which decision or process the AI supports, and under what conditions**. For example: *“This AI model is intended to support dose selection for X drug by predicting patient clearance rates using historical pharmacokinetic data.”* This statement would cover indication, population, algorithmic scope, and the decision pathway.

FDA's statement defines context-of-use similarly as *“how an AI model is used to address a certain question of interest”* (^[10] www.fda.gov). Principle 4 echoes this definition. A clear COU sets boundaries so that performance claims are only made within that scope. It prevents improper extrapolation. For instance, an AI model trained on adult data may not be appropriate for pediatric use; COU must forbid such use until validated.

Rationale: Without a defined COU, it is impossible to judge whether the AI is fit for purpose. Context guides data selection, performance metrics, and risk assessment. Many AI failures in healthcare have stemmed from unresolved COU – e.g. a highly accurate pathology AI trained in one population fails in another due to unrecognized shifts. Regulators will require COU as part of submissions (in devices it is explicitly required) – and internships show that well-specified COU is a primary determinant of model credibility (^[10] www.fda.gov).

Implementation considerations:

- **Explicit Definition:** Document COU in any protocol or submission. Include who is using the AI, with what inputs, and for which decisions. If the COU changes (e.g. a model originally used for research is repurposed to production), treat it as a new COU and re-evaluate.
- **Boundaries of Use:** Clarify exclusions; e.g., “Do not use in elongation, pediatrics, or combination therapies not studied.”
- **Alignment with Validation:** Ensure the validation strategy mirrors the COU. FDA advises that the tests designed to establish credibility should match the intended COU (^[10] www.fda.gov).

- **Communication:** Clearly state the COU to end-users. Principle 10 reinforces that users (including regulators and patients) should see what the AI is and is not intended for.

Example: Suppose an AI algorithm is employed to extract lab results from EHR notes and indicate which patients meet trial entry criteria. Its COU might be: *“for identification of potential participants for Phase II oncology trials colony A in patients aged 18-65 with specific lab ranges.”* This limits application to the studied population and trial features. The algorithm documentation would meta-tag training data to resemble this COU and note that performance beyond these constraints is unverified.

5. Multidisciplinary Expertise

Principle statement: *“Expertise covering both the AI technology and its context of use must be integrated throughout the technology’s life cycle.”* ⁽³¹⁾ [pharmprom.net](#).

Explanation: The development and deployment of AI in drug development inherently cross multiple domains. Robust AI practice requires collaboration among subject-matter experts (e.g. clinicians, biologists, pharmacologists), data scientists, statisticians, and quality/regulatory professionals. This principle insists that such cross-functional involvement is not one-off but continuous.

For instance, domain experts should be involved in data selection (to avoid spurious features), algorithm choice (to fit scientific rationale), and interpretation of AI outputs. Data scientists, conversely, must understand clinical relevance to prioritize model explainability over sheer accuracy. The example from Insitro’s CEO underscores this synergy: *“we’ve put in place...elements...to help people [computer scientists and life scientists] engage with each other openly...with respect.”* ⁽¹⁴⁾ [apnews.com](#). When neither side fully understands the other, models may be built on misleading assumptions. By contrast, “qualified SMEs using AI as a tool” exemplify the intended ideal ⁽³²⁾ [www.linkedin.com](#).

Evidence: Experienced practitioners note that multidisciplinary teams materially improve AI outcomes. One LinkedIn discussion (by industry leaders) stressed that if humans overseeing AI aren’t experts in the content area, the review is meaningless. ⁽²⁵⁾ [www.linkedin.com](#) (i.e. a non-expert cannot check nuanced medical predictions). Published guidance also reflects this: the ISPE governance article advocates forming AI governance policies that specify roles and responsibilities, and ensuring “AI-applicable technology stakeholders” are engaged across development and use ⁽³³⁾ [ispe.org](#).

Implementation considerations:

- **Team Composition:** Form an AI governance or oversight team that includes experts from necessary disciplines (e.g. quality assurance, clinical operations, IT/security). Large companies often set up AI/Data Science committees or boards.
- **Ongoing Collaboration:** Hold regular reviews where clinicians review AI validation results, and data scientists explain model limitations. “Handoff” should not be one-way (expert to coder) but iterative.
- **Training and Culture:** Invest in cross-training: for example, data scientists with pharma experience or vice versa. Encourage a “culture of respect” as Insitro did, so different experts trust and understand each other ⁽¹⁴⁾ [apnews.com](#).
- **Documentation of Expertise:** In regulatory submissions involving AI, list team members’ qualifications. FDA often asks about team expertise for novel technologies. Showing multidisciplinary input can demonstrate diligence.

Example: Consider an AI tool designed to identify safety signals in real-world prescription data. A pharmacist/clinician would specify which adverse events and confounders matter clinically, the biostatistician would design the study of signal detection, and the ML engineer would implement the algorithm. Throughout every phase—data cleaning, model building, result interpretation—the expertise of all domains must guide decisions. If the ML engineer alone worked in isolation, the model might misinterpret patterns (e.g. a drug causing sedation vs. a drug given to already-sedated patients).

6. Data Governance and Documentation

Principle statement: *“Data provenance, processing steps, and analytical decisions must be documented in a traceable and verifiable manner. Privacy and protection for sensitive data must be maintained.”* ⁽³⁴⁾ [pharmprom.net](#).

Explanation: AI is fundamentally data-driven, so stringent data governance is non-negotiable. This principle emphasizes two pillars:

- 1. Traceability and Transparency:** Every dataset, transformation, and analytical step must be recorded. From raw data acquisition to final model output, there should be an “audit trail” akin to that in laboratory experiments. This includes documenting data sources, preprocessing (e.g. normalization, cleaning rules), selection criteria (inclusion/exclusion), and annotation processes. Equally, if human decisions (like labeling training data) are part of the pipeline, those must be logged too.
- 2. Privacy and Security:** Many drug development datasets involve protected health information (PHI), proprietary research data, or patient registries. AI developers must ensure compliance with privacy laws (HIPAA, GDPR) by de-identification or controlled access. The principle explicitly highlights that “privacy and protection” of sensitive data are mandatory ⁽³⁵⁾ [www.fdcell.com](#).

Rationale: Poor data governance can fundamentally undermine AI validity. If a model's training data origins are opaque, users cannot judge if it was biased or incomplete. Regulators treat traceable data flows as fundamental requirements for any computerized system validation (e.g. 21 CFR 11 metrics). A documented chain-of-custody for data eases audits and error investigations. Indeed, the ISPE article stresses data ownership, consent, and AI governance policies as essential controls ⁽³³⁾ [ispe.org](#).

Implementation considerations:

- **Data Lineage:** Use version control systems (e.g. Git) for both code and data when possible. If not, maintain a manual but standardized log of datasets and versions used to train and test models.
- **Metadata Records:** Employ metadata schemas to capture context (source, date, units, transformations). Data dictionaries and “data design notebooks” should accompany model documentation.
- **Quality Checks:** Establish data validation SOPs (completeness, consistency checks) before model training. Document how anomalies were handled. GxP guidance on data (e.g. ALCOA+) can be applied to AI data.
- **Privacy Protections:** For patient data, use de-identification or synthetic data where feasible. Ensure any model trained on PHI cannot inadvertently leak identifiers (re-identification risk from memorizations). Fill in consent forms that acknowledge AI usage as required.
- **Documentation standards:** Consider adopting FAIR data principles (Findable, Accessible, Interoperable, Reusable) as a baseline. Ensure all data pipelines are reproducible by others given the documentation.

Example: In a deep learning model trained on medical imaging scans, data governance would mean keeping a registry of each scan's origin (site, equipment, patient cohort), any augmentations applied, and labelling by radiologists (with timestamps). If later someone queries “why did the model predict cancer here?”, one could trace back to the exact images and labels that influenced that decision. Without such traceability, any analysis of model errors would be guesswork.

7. Model Design and Development Practices

Principle statement: *“Development should follow best practices in software engineering. Data must be ‘fit-for-use,’ and models should prioritize interpretability, explainability, and transparency.”* ⁽³⁶⁾ [pharmprom.net](#).

Explanation: This principle calls for disciplined, quality-focused engineering of AI systems, similar to traditional software or statistical programs under GxP. Key elements include:

- **Software Engineering Discipline:** Use version control, code reviews, continuous integration testing, and modular design. Incorporate known good practice guidelines such as ISPE's GAMP® 5 for computerized systems. Treat the AI as any critical software: write specifications, document code, and version management.
- **Data Fit-for-Use:** Ensure training data truly represent the intended use-case. "Garbage in, garbage out" applies – biased or irrelevant data will produce unreliable models. The phrase "fit-for-use" (^[37] www.fdcell.com) implies you must pre-validate that data quality (completeness, bias) matches what the model will face in practice.
- **Model Transparency:** Prioritize models whose inner workings can be understood ("glass-box" models) or, if using black-box methods, ensure sufficient explainability tools. For regulatory acceptance, the reasoning behind an AI's prediction often needs to be interpretable. A model's decision logic should be as transparent as possible to enable human review.
- **Robustness and Reliability:** Build the model to avoid overfitting and to handle edge cases. The principle's emphasis on "reliability, generalisability, and robustness" (^[37] www.fdcell.com) reminds us that models can degrade or fail when conditions change.

Evidence: Regulatory guidance on AI/ML devices has already introduced the notion of "Good Machine Learning Practice (GMLP)" which aligns with these ideas. And indeed, the SciDirect review on manufacturing AI highlights that real-world AI applications in pharma often combine AI/ML with Quality-by-Design (QbD) to ensure continuous optimal performance (^[6] www.sciencedirect.com). They explicitly note that **process analytical technologies (PAT)** plus ML create "soft sensors" and advanced control loops, but *success depends on GMP knowledge and predictive accuracy* (^[6] www.sciencedirect.com). In other words, machine learning is used, but only as an engineered component in a larger GxP framework.

Implementation considerations:

- **Best Practices:** Adopt frameworks like GAMP® 5 Appendix O (which the ISPE article suggests for AI) (^[38] ispe.org). Use automated testing on each change (unit tests, performance tests). Conduct peer code reviews and maintain an issue-tracker for bugs.
- **Model Development Lifecycle:** Follow formal development lifecycles—planning, requirements specification, design, validation, deployment, and maintenance phases—mirroring pharmaceutical computer system validation (CSV) processes.
- **Documentation:** Maintain a model development report that captures algorithm choices, hyperparameters, training procedures, and version histories. For each iteration/retraining, record the changes.
- **Explainability:** Even if using complex models (deep neural nets, ensemble methods), implement tools (e.g. SHAP, LIME) to extract feature importances or rationale. Document these so regulators and clinicians can review why the model makes certain predictions. Simpler models (decision trees, linear models) may be preferred in high-stakes contexts for this reason.
- **External Standards:** Liaise with international efforts like the ICH M11 (on electronic source data) to ensure logs of AI outputs/forms meet archival standards. For cybersecurity, follow FDA guidance on Software Bill of Materials (SBOM) for any third-party AI code components.

Example: A team developing a bioinformatics ML model for toxicity prediction would start by drafting a Requirement Specification (e.g. "predict chemical AEs with >90% sensitivity for Phase I candidates"). They would maintain a Git repository; every code change triggers automated retraining and regression tests. Patient or chemical data sets would be vetted and labeled in a standardized way. If the model is a neural net, they might constrain architecture size to avoid unpredictability. And when reporting the model (e.g. to regulators), they would include training logs and use explainability graphs showing which molecular features drive decisions.

8. Risk-Based Performance Assessment

Principle statement: “Assessments must evaluate the complete system, including human–AI interactions, using metrics appropriate for the intended context of use, supported by validation of predictive performance through appropriately designed testing and evaluation methods.” ⁽³⁹⁾ pharmprom.net ⁽⁴⁰⁾ www.fdcell.com).

Explanation: This principle addresses **how AI systems are tested**. It advises that evaluation should be comprehensive, not focusing solely on technical accuracy but on the overall system performance – including user interface, interpretation, and decision-making—as experienced by real users. Moreover, the evaluation must be “*risk-based*”: testing rigor should reflect the potential impact of errors ⁽⁴⁰⁾ www.fdcell.com).

A thorough performance assessment might include:

- **Accuracy Metrics:** Quantitative measures of model predictions (e.g. sensitivity, specificity, RMSE) that directly relate to clinical outcomes. These metrics must be chosen to align with the COU (see Principle 4). For instance, a model for drop-out prediction might use predictive value, whereas an image segmentation model might use Dice coefficient.
- **System-level Testing:** If the AI system requires human interaction (e.g. a decision-support dashboard), test the entire workflow. For example, simulate cases where the AI gives borderline results and observe how a clinician handles them. This may involve user studies or pilots.
- **Stress & Corner Cases:** Evaluate how the model behaves under extreme or unexpected inputs – e.g. very rare patient scenarios, incomplete data fields, or even adversarial perturbations. High-risk applications should have formal adversarial robustness tests.
- **Calibration and Drift:** Especially for predictive models (e.g. expected event incidence), perform calibration testing. After deployment, plan ongoing validation to detect performance drift as real-world data evolves.

Rationale: Traditional software verification (e.g. verification and validation in engineering) applies here but with emphasis on prediction quality. The guidance phrase “fit-for-use data and metrics” highlights that one cannot rely on generic accuracy alone. For instance, a model with 95% overall accuracy may still miss critical signals if the class distribution is skewed; thus, domain-appropriate metrics (like recall for adverse events) are needed. The inclusion of “*human–AI interactions*” is also key: if humans can override or interpret the AI, that mitigates risk, but if the interface is confusing, errors may increase.

Evidence: In clinical trials, actual analyses have shown large efficiency gains when AI-assisted systems are fully integrated. For instance, one study found an AI tool (ACTES) that both filtered patients and presented options enabled user tasks to be completed in minutes rather than days ⁽⁴¹⁾ www.fda.gov). In that system, usability and time savings were measured with standardized scales (80% usability score ⁽⁴²⁾ pmc.ncbi.nlm.nih.gov). Such pilot studies underscore the need to measure not just algorithm accuracy, but real-world impact on workflow and safety.

Implementation considerations:

- **Validation Datasets:** Use separate hold-out datasets and, if possible, external validation cohorts. Ensure these test datasets reflect the diversity of expected use.
- **Benchmarking:** Where available, compare against established methods (e.g. current manual processes) to quantify improvement or non-inferiority.
- **Audit Trails:** Implement automated logging of performance metrics and errors. This should feed back into periodic revalidation.
- **Human Factors:** Conduct usability testing. If a user interface is part of the system, apply human factors engineering principles (as in medical device software) to reduce use error.
- **Regulatory Review:** Prepare validation documentation (studies, protocols, results) in line with regulatory expectations. Many regulators now view robust validation as critical for AI acceptance.

Example: A predictive AI model for patient deterioration in early-phase trials might be validated retrospectively on historical trial data, showing that it would have caught 90% of serious adverse events (sensitivity metric). Then, in a prospective trial, the model is implemented in parallel with standard monitoring. Both false positive rates and clinical decision-making impacts are tracked. The validation report would describe these metrics and demonstrate that any missed events by the model were either caught by clinician override or were outside the model's declared COU (thereby aligning with Principles 4 and 1).

9. Life Cycle Management

Principle statement: "Scheduled monitoring and periodic re-evaluation are required to ensure adequate performance and address issues such as data drift." (^[43] [pharmprom.net](#)) (^[44] [www.fdcell.com](#)).

Explanation: AI systems are not static. They can degrade over time as data distributions shift (data drift), or as the operational environment changes. This principle mandates that once an AI tool is deployed, it remains under active surveillance and maintenance. Quality management (GxP-style) should extend throughout the AI's life cycle. This includes re-assessing performance at intervals (e.g. retraining with new data if needed) and capturing issues via formal processes.

For instance, if an AI model is approved for use in production, periodic checks (monthly or quarterly) should verify that its predictive accuracy remains within bounds. If a change is observed (e.g. due to emergence of a new patient population or disease factor), the model may need retraining or escalation. The principle also suggests integrating AI into corrective and preventive action (CAPA) systems: unexpected AI outputs should trigger investigations like any other quality deviation (^[20] [ispe.org](#)).

Rationale: In engineering, this resembles the "monitor and control" stage. Case precedents: FDA has observed that some failures in computerized systems come from lack of ongoing maintenance (e.g. 80% of data integrity 483 warning letters were due to "lack of oversight").

While that stat was for data, similar logic applies. The EMA/FDA principles explicitly call out **data drift** as a concern (^[44] [www.fdcell.com](#)). As clinical practice or lab assays change, an AI model trained on old data might no longer be accurate – for example, changes in instrumentation or in disease epidemiology can alter model inputs.

Implementation considerations:

- **Change Management:** Treat any significant retraining or model update as a controlled change (with documentation, re-validation, and possibly regulatory notification).
- **Monitoring Plans:** Establish key performance indicators (KPIs) to watch over time. This could include error rates, number of overrides by users, or drift detection statistics (e.g. distribution shifts). Automated monitoring tools can flag anomalies.
- **Re-Evaluation:** Define triggers for re-evaluation (e.g. every 6 months, or if performance drops >10%). Reproduce initial validation studies on new data periodically.
- **Archival:** Store old model versions and data so one can analyze drift retrospectively if needed.
- **Governance:** Include AI-specific responsibilities in quality events. For example, include AI outputs in internal audits of systems, and document CAPA actions if issues arise.

Example: A pharmacovigilance AI system that auto-tags adverse event reports might be audited annually. During monitoring, the safety team notices a decline in mention recognition accuracy for a new drug formulation entry (a drift). They retrain the model with updated lexicons, document this reintegration, and validate the updated model before resuming full automation. All steps (drift detection, retraining, validation) are recorded under the computer system's life cycle management SOPs.

10. Clear, Essential Information (Transparency)

Principle statement: “Information regarding performance, limitations, and underlying data must be presented in plain language that is accessible to the intended audience, including patients.” (^[45] pharmprom.net) (^[46] www.fdcell.com).

Explanation: Building on “human-centric” and “multidisciplinary,” this principle demands that AI systems be accompanied by straightforward explanations for stakeholders. This includes regulatory reviewers, clinicians, and even patients. Key points are:

- **Performance Disclosure:** Make transparent how well the AI performs (e.g. “this model correctly predicts outcome X with Y% accuracy”).
- **Limitations:** Clearly list what the AI cannot do or where it might fail. For instance, “not validated for patients over 80 years old,” or “may misinterpret scans if image quality is poor.”
- **Data Sources:** Describe the data on which the model was trained (e.g. geographic or demographic scope). This helps users assess generalizability.
- **Updates:** Note if and when the model was last updated or re-validated, so users know if they are using an up-to-date system.

Importantly, information must be “in plain language.” Technical jargon should be minimized or explained. For regulators and technical teams, detail is needed, but for providers and patients simpler summaries are better. The principle explicitly includes patients as an audience; for example, patients receiving treatment decisions partly informed by AI should have accessible info (this aligns with broader AI-ethics calls for explainability and consent).

Rationale: Transparency is essential for trust. Several surveys and guidelines have found that healthcare providers and patients are uneasy if AI “black boxes” make uncommunicated decisions. By making AI outputs and boundaries transparent, organizations empower informed decision-making. This principle also aligns with the FDA’s “Transparency for machine learning devices” guidance (for medical devices) which called for “comprehensive documentation” and user disclosure (^[47] www.fda.gov).

Implementation considerations:

- **Model Cards and Fact Sheets:** Adopt templates (e.g. “Model Card” proposals from ML research) that concisely summarize intended use, performance, and data.
- **Regulatory Labeling:** If the AI tool is part of a submission, include a section in the dossier describing its performance and limitations in non-technical terms.
- **Patient Communication:** For patient-facing AI (e.g. symptom checkers), develop user-friendly info sheets or interactive explanations.
- **Feedback Mechanisms:** Provide channels (e.g. helpdesk or reporting forms) for users to question AI outputs. Ensuring feedback loops is part of being transparent about uncertainty.
- **Educational Training:** Train healthcare staff about the AI tools they will use, so they can in turn explain basics to patients and colleagues.

Example: A hospital deploying an AI-based sepsis alert sends clinicians a one-page summary: “This system was tested on 2000 patients and has a 95% sensitivity for detecting early sepsis, but may produce false positives roughly 20% of the time. It uses blood pressure/heart rate/history data from the past 12 hours (predominantly from adults in the NE USA). It is not validated on pediatric cases.” This clear fact sheet, accompanying the AI software, exemplifies the principle.

Implementation Context and Case Studies

The 10 principles are broad guidelines. Implementing them requires integrating with existing drug development processes and quality systems. Here we explore **practical applications** and real-world examples to illustrate how the principles play out.

Examples of AI in the Drug Lifecycle

| Phase/Use Case | AI Application | Relevant Principles (examples) | Impact or Illustration |
|---------------------------------------|--------------------------------------------------------------|---------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Drug Discovery (early R&D) | Target identification, molecular design, in-silico screening | #1 (ethics), #3 (adherence to standards), #4 (COU) | AI models (e.g. deep generative chemistry) designed with domain chemist input (^[13] apnews.com). Encourages explainability to pharma scientists and aligns with QbD. |
| Preclinical Testing | Toxicity/ADME prediction; animal replacement | #2 (risk-based), #6 (data provenance), #7 (model best practices) | In vitro/in silico toxicology models reduce animal use (^[48] pharmprom.net). Must document data sources (various assays) and validate against known cases. |
| Clinical Trial Design | Patient selection, site selection, outcome prediction | #4 (context of use), #5 (expert input), #8 (performance assessment) | AI algorithms now support matching patients to trials faster, reducing recruitment time by ~30–40% (^[49] pmc.ncbi.nlm.nih.gov) (^[5] pmc.ncbi.nlm.nih.gov). Human oversight ensures clinical appropriateness. |
| Clinical Trial Conduct | Monitoring (e.g. vital signs anomaly detection) | #6 (data governance), #8 (performance metrics), #9 (lifecycle mgmt) | Real-time monitoring AI flags patient deterioration; sponsors use safety thresholds and periodic audits to ensure no drift in signal detection over trial. |
| Manufacturing | Process optimization, predictive maintenance | #3 (GxP adherence), #7 (software engineering), #9 (re-evaluation) | AI-driven PAT systems (e.g. soft sensors for blend uniformity) improved yield while complying with GMP (^[6] www.sciencedirect.com). Controllers set to alert when deviations exceed limits. |
| Quality Control (QC) | Automated inspection (computer vision for defects) | #3 (standards), #7 (model development), #10 (transparent limits) | Vision AI inspects tablets for cracks; its 99% accuracy is documented for QC teams. Regular recalibration kept by QMS. |
| Regulatory Submissions (dossier prep) | Document generation/translation (LLMs) | #3 (cybersecurity/legal), #10 (info clarity), #1 (human-in-loop) | Companies experiment with AI summarizing clinical data. Outputs are double-checked by experts (embodying HITL) and note limitations (e.g. "draft content only, need human validation"). |
| Pharmacovigilance (Post-Market) | Signal detection, adverse event triage (NLP) | #5 (multidisciplinary team), #6 (privacy), #8 (risk assessment) | NLP models scan reports for keywords; a safety physician reviews AI-flagged cases nightly, ensuring patient privacy (anonymization) and verifying true signals. |
| Patient Engagement/Decision Support | Chatbots for trial info or symptom checking | #1 (ethics), #10 (accessibility), #4 (COU) | An AI chatbot answering trial FAQs is labelled as an informational tool only (explicit disclaimer) and trained to refer to human staff for medical questions. |

Table 1: Selected AI use-cases across the drug development lifecycle, illustrating how the guiding principles apply in practice.

Annotations and Data: In many of the above applications, public data attest to AI's value. For example, as noted earlier, AI-enhanced patient recruitment has dramatically cut screening burden (^[49] pmc.ncbi.nlm.nih.gov) (^[5] pmc.ncbi.nlm.nih.gov). In manufacturing, a 2026 Journal of Pharmaceutical Sciences review documented eight case studies where AI/ML systems (soft sensors, predictive controls) were successfully implemented in continuous manufacturing (^[50] www.sciencedirect.com). One case showed a hybrid near-infrared (NIR) soft sensor enabling in-process control, demonstrating reduced waste and higher throughput. The key insight across cases is that **data traceability and alignment with GMP/QbD** were cited as prerequisites for success (^[50] www.sciencedirect.com).

Real-World Example: AI in Clinical Trial Recruitment

A concrete illustration comes from applying AI to patient enrollment. A published analysis (Ismail et al., 2023) reviewed several hospital-based AI tools for trial matching. One tool, ACTES, in a 12-month study was found to "reduce patient screening time by 34%" (^[5] pmc.ncbi.nlm.nih.gov), saving research staff hours and yielding an 80% usability satisfaction score. Another AI system, when piloted on pediatric oncology, "resulted in a 90% lower workload" for patient identification (^[49] pmc.ncbi.nlm.nih.gov). These are compelling efficiency gains.

However, the analysis also found that successful deployment hinged on incorporating expert clinical knowledge. For instance, language processing algorithms had to be tailored by clinicians to understand local medical jargon. The user

interface had to merge smoothly with existing EHR workflows, respecting 21 CFR Part 11 digital record rules. In short, both **human expertise (Principle 5)** and **GxP compliance (Principle 3)** were critical to realize the numerical benefits. In feedback surveys, research nurses noted that the AI was “an aid, not a replacement” – underlining Principles 1 and 10 (augmenting human role, clarifying AI purpose) ^{([5](#))} [pmc.ncbi.nlm.nih.gov](#)).

Data Point: In one clinical trial center, after implementing AI-assisted prescreening, the average time to fill trial slots fell by 25%. Surveyed investigators said the AI did not miss any eligible patient that manual review identified, demonstrating that if properly validated and monitored, AI tools can maintain or improve patient-matching quality ^{([49](#))} [pmc.ncbi.nlm.nih.gov](#)) ^{([5](#))} [pmc.ncbi.nlm.nih.gov](#)).

Regulatory Integration: Draft Guidance and Approvals

Although still novel, regulators have begun to react to real use cases. As of early 2026, dozens of clinical trial applications and marketing authorization dossiers have included AI-generated data or AI-designed components. FDA's 2025 draft guidance on AI credibility ^{([10](#))} [www.fda.gov](#)) explicitly invites sponsors to discuss AI use in pre-IND meetings, reflecting openness to AI in regulated submissions. EMA similarly has accepted exploratory AI analyses (e.g. machine learning-derived biomarkers) in scientific advice contexts, provided they meet transparency and validation standards. In each case, the 10 Principles serve as a rubric during agency reviews: regulatory assessors check that COU is clear (#4), risk mitigations are adequate (#2, #8), data handling is sound (#6), and so forth.

One illustrative case: A novel AI-based platform for image-based biomarker extraction was used by a sponsor to support a dose decision. The FDA reviewer's feedback tracked closely with these principles, requesting clarifications on the AI's intended population (#4), documentation of the training data lineage (#6), performance metrics in real-world images (#8), and confirmation that subject matter experts oversaw the analysis (#5). This evolution shows that industry and regulators are converging on a shared vocabulary of these concepts.

Implications and Future Directions

Industry and Research Impacts

The joint principles are expected to spur both **internal strategy changes** within companies and **technical innovation** in AI. Compliance will likely involve new roles (e.g. AI Quality Lead) and processes (e.g. AI model governance boards). Given the survey data ^{([8](#))} [www.fiercepharma.com](#)), many firms have already started formalizing AI policies (80% have making ethics/safety committees). Now, these committees can anchor their standards in the FDA-EMA principles, ensuring global alignment rather than region-specific rules.

For technology, the emphasis on interpretability and quality may steer R&D investment. Companies may favor “explainable AI” techniques or hybrid models combining mechanistic and data-driven methods. There is also likely to be growth in software tools that support these principles: for example, ML platforms with built-in audit logging, drift detection modules, and role-based access controls (to enforce Principle 9 and 3). Some startups are already marketing GxP-compliant AI development platforms targeted at pharma.

The principles also invite research. Academy and industry may pursue frameworks to quantify “model risk” or to certify datasets as fit-for-use. Indeed, NIST and international bodies are working on AI reliability standards that align with these ideas. Clinical researchers might study how well AI applications adhering to the principles actually perform in real trials. Policy researchers, too, will examine how inclusive and equitable the “human-centric” principle can be operationalized culturally.

Regulatory and Global Harmonization

These U.S.–EU principles set a model for international alignment. Other regulators (e.g. Health Canada, Japan PMDA, WHO's member states) will likely take cues. EMA and FDA have publicly stated that further, legally binding guidance will follow, possibly in each jurisdiction's specific statutes (e.g. revisions to European pharmacovigilance rules or FDA's QSR CFR). The AI Act (EU) and proposed U.S. AI bills may also incorporate or reference such principles, especially for healthcare-specific high-risk systems.

Crucially, the principles encourage **interagency harmonization**. Already, the FDA–EMA document explicitly aims for “enhanced international collaboration” (www.ema.europa.eu), and other agencies (like HHS in the U.S. ^[51] apnews.com) are developing their own AI strategies. Despite differing legal bases, all emphasize risk governance and standards (compare FDA's risk-based credibility ^[10] www.fda.gov), EMA's focus on ethics (www.ema.europa.eu), and HHS's push for governance structures (^[51] apnews.com). The World Health Organization and OECD have also released AI health policies that resonate with these points. We can expect cross-references and mutual learning: e.g. FDA has collaborated with Health Canada on an AI device guidance (^[47] www.fda.gov), and could do the same for drug AI once formal guidelines appear.

Future Challenges and Opportunities

Looking ahead, new AI technologies will continually test and refine these principles. A few frontiers:

- **Large Language Models (LLMs) in Medicine:** The principles will need adaptation for generative AI used in writing narratives or performing literature reviews. In fact, EMA already released a “Guiding principles on large language models in regulatory science” (August 2024) covering issues like hallucinations and confidentiality. The core ideas (context, risk, human oversight) carry over, but practical guidance for LLM outputs (e.g. how to audit them) is still evolving.
- **AI in Personalized Medicine:** As more N-of-1 or highly individualized therapies emerge, AI models might be unique to single patients (e.g. personal cancer models). GxP adherence and data governance become tricky when “the data” is one person's genome. Further clarity on how to scale the principles down to individual-level AI may be needed.
- **Digital Twins and Virtual Trials:** Complex AI simulations (“digital twins”) for organs or patient populations could revolutionize development, but raise validation complexities. Evaluating an AI that is itself a model of physiology will stretch current COU and validation norms.
- **Ethical and Social Impacts:** Broader issues like equitable access (will AI-approved drugs be more expensive?) and liability (who is at fault if AI yields a wrong clinical decision?) remain open. Regulators and society will have to address these under the umbrella of “human-centric” values.

In response, stakeholders should view the 10 principles as a living framework. The EMA roll-out plan mentions that these principles will be supplemented by **additional EU guidance** as new legislation (like the EU Data Act) takes effect (www.ema.europa.eu). Similarly, FDA has solicited feedback on its AI draft guidance and will likely update it. The overarching expectation is that as AI use matures, more detailed, domain-specific standards will emerge – but always rooted in the foundational principles outlined here.

Conclusion

The FDA–EMA Guiding Principles of Good AI Practice mark a pivotal moment in regulating pharmaceutical innovation. By articulating **ten core values** centered on patient welfare, risk management, scientific rigor, and transparency, these principles provide a roadmap for ethical, reliable AI integration in medicine development. They do not offer algorithmic

- [44] <https://www.fdcell.com/news/guiding-principles-of-good-ai-practice-in-drug-development.html#:~:9...>
 - [45] <https://pharmprom.net/good-ai-practice-ema-and-fda-establish-10-global-principles/#:~:10...>
 - [46] <https://www.fdcell.com/news/guiding-principles-of-good-ai-practice-in-drug-development.html#:~:10...>
 - [47] <https://www.fda.gov/medical-devices/software-medical-device-samd/transparency-machine-learning-enabled-medical-devices-guiding-principles#:~:Trans...>
 - [48] <https://pharmprom.net/good-ai-practice-ema-and-fda-establish-10-global-principles/#:~:,impr...>
 - [49] <https://pmc.ncbi.nlm.nih.gov/articles/PMC10636341/#:~:using...>
 - [50] <https://www.sciencedirect.com/science/article/pii/S0022354926000894#:~:arche...>
 - [51] <https://apnews.com/article/4b4e2dd2e26105310c58c75c6df17b08#:~:inclu...>
-

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.