

OPINION PAPER

Generative AI and Clinical Development: Exploring the Frontier of AI Governance Utilizing Generative AI -Conceptual Framework for AI Governance in Clinical Development

Munenori Takata^{*†}, Daisuke Ichikawa[‡], Masaharu Harada[§], Makoto Hiraidell,
Mari Sugimoto[¶], Yukikazu Hayashi^{†,¶}, Yasuhiro Yoshimaru[‡], Kenji Fujisawa[‡],
Toshiaki Nojima[‡], Satoru Fukimbara^{†,§}, Hideki Suganami^{†,‡} and Takuhiro Yamaguchi^{*†}

Generative AI is increasingly applied to design and managing quality management pyramid documents, supporting traceability and process integrity. As guidance from authorities such as the FDA, EMA, and WHO continues to evolve, organizations face significant challenges in establishing effective governance frameworks amid rapid technological advances. We review existing regional regulations and guidance, identifying core governance elements such as transparency, data privacy, validation, and risk management. We propose a practical methodology for developing governance documentation, emphasizing the use of retrieval-augmented generation (RAG) techniques to safeguard confidentiality and improve accuracy in documentation. Additionally, we highlight the importance of integrating ethical principles, human in the loop (HITL), and lifecycle monitoring to promote responsible and compliant adoption of generative AI in clinical research. Our comprehensive approach aims to accelerate AI deployment while ensuring patient safety and regulatory compliance.

Keywords: generative AI; clinical development; governance-related documentation; large language model; retrieval augmented generation; human in the loop; quality management system; risk-based approach; clinical data management; clinical data science

Introduction

Generative artificial intelligence (AI) is broadly applied in clinical development, including in the formulation of standard operating procedures (SOPs) and data analysis. It has become a critical tool for ensuring traceability and process integrity in many human actions, thereby supporting the overall quality of research. In this context, various regions and countries have issued notices and guidances regarding the use of generative AI, and regulatory authorities such as the WHO, EMA, and FDA are beginning to issue guidances for its application in clinical

development settings.^{1–3} While the use of generative AI offers the potential to strengthen data-driven approaches and accelerate product development, organizations are actively exploring the most effective frameworks for governance. For example, generative AI is instrumental in creating new insights by integrating and analyzing large volumes of medical data within organizations and healthcare facilities.⁴ Additionally, it assists in maintaining and managing clinical development documents, such as research protocols, investigator's brochures, and quality management system (QMS)-related SOPs, by rapidly extracting useful information through natural language processing techniques, thus enabling smooth and consistent processes.

Given the rapid pace of generative AI development and the regulations being issued by various regions and countries,^{5–8} governance frameworks for the use of generative AI within clinical development organizations often lag behind. Ideally, organizations should establish operational procedures before beginning use; however, the speed of technological innovation and the inability to acquire the necessary operational skills in time mean that organizational responses are urgently needed. This

^{*} Clinical Research Data Center, Tohoku University Hospital, Sendai, JP

[†] Division of Biostatistics, Tohoku University Graduate School of Medicine, Sendai, JP

[‡] Kowa Company Ltd, Tokyo, JP

[§] ONO Pharmaceutical Co., Ltd, Osaka, JP

[¶] Department of Clinical Assessment, Tokyo University of Pharmacy and Life Sciences, Tokyo, JP

[¶] A2 Healthcare Corporation, Tokyo, JP

Corresponding author: Munenori Takata, M.D. Ph.D.
(munenori.takata.a3@tohoku.ac.jp)

paper extracts the core elements for governing generative AI from existing regulations and guidance, and proposes a feasible, concrete methodology for creating governance-related documentation.

Understanding Generative AI Utilization Guidance Issued by Various Regions and Countries

As the use of generative AI is promoted, many regions and countries have issued guidance on its utilization.^{1–3, 5–8} The foundation for developing governance documentation lies in correctly understanding these regulations. For example, clinical trials conducted in the United States must comply with FDA regulatory documents.² Recent FDA guidance on AI includes the issuance of “Use of Artificial Intelligence in Drug Development” in 2025.¹ Its structure encompasses: “Scope and Definitions,” “Dataset Bias and Fairness,” “Model Verification and Validation,” “Transparency and Accountability,” and “Documentation of AI Use in Regulatory Submissions.” Similarly, recent EMA guidance³ on AI outlines general considerations such as a “risk-based approach,” “system lifecycle,” “data processing pipelines,” “integrity of data,” “fitness for purpose,” and the “scientific base.”

Although the Pharmaceuticals and Medical Devices Agency (PMDA) in Japan has not issued notifications specific to generative AI in clinical development, the Cabinet Office (a higher authority) has implemented an “AI Strategy Council” as part of the process to establish the Basic AI Law. Meanwhile, in the manufacturing sector, the Ministry of Economy, Trade and Industry has issued “AI Governance in Japan Ver. 1.1.”⁹ This document references a universal framework designed to describe the overall structure of governance in the Society 5.0 era. It organizes AI governance elements debated across different countries and regions into four categories: (1) the ultimately protected, technology-neutral goals; (2) cross-cutting, intermediate rules for achieving these goals—such as cross-regulatory measures, guidelines, and international standards; (3) rules focused on specific fields or sectors; and (4) layers consisting of monitoring and enforcement.⁹

We present a chart showing the status of laws and regulations issued in each region as of February 2026 (**Figure 1**). We focusing on guidance for clinical development of pharmaceuticals and medical devices. Despite regional differences, common framework elements include: “Transparency and Accountability,” “Data Privacy,” “Data Quality and Bias,” “Validation and Verification,” and “Risk Management.”

Pre-defining Stakeholders Utilizing Generative AI

The types of industries involved in clinical development are diverse, including mega-pharma with global bases, contract research organizations (CROs), academic research organizations (AROs), and third parties providing electronic data capture (EDC) and other clinical development infrastructure, and the clinical development operations of each organization vary. Furthermore, the scope of generative AI utilization is diverse.¹⁰ Since the scope of generative AI use is closely linked to the scope of duties of these professions, the content of the service and the scope of generative AI utilization will be tightly coupled.

When considering the definitions of other specialized professionals beyond clinical data management (CDM), such as principal investigator, clinical research coordinator (CRC), clinical research associates, biostatistician, and pharmacovigilance staff, we utilized the following prompt:

You are an expert well-versed in clinical development. #Define the roles of each of the following professions in clinical development in 2–3 sentences, referencing PubMed articles or official documents. #Also, include the source information, such as PubMed ID codes or official URLs.

We used this prompt to generate role definitions with Google Gemini’s cloud large language model (LLM), including PubMed citations.

For example, the tasks of clinical data managers (CDMs) include designing and maintaining the trial database, overseeing the creation of electronic case report forms, and executing data cleaning processes to identify and




Figure 1: Status of Generative AI Guidance Issuance for Clinical Development.

resolve discrepancies before the database is locked for analysis. CDMs ensure the accuracy, completeness, and integrity of the clinical data, which forms the basis for statistical analysis and regulatory submissions. They manage the entire data lifecycle and filing by generating queries and resolving issues to maintain data quality.

Creating Generative AI-Related Governance Documents and Their Consistency with Existing Organizational QMS Documents

By providing regulatory and professional role information to a cloud-based generative AI powered by an LLM, the items to be included in the highest-level policy of the organizational QMS can be defined with a certain degree of accuracy and precision. However, using cloud LLMs poses two major problems. One is hallucination. For instance, Gemini provided the following incorrect information regarding the ICH-E6 (R3), which has reached the consensus stage at Step 4 (Example: “ICH-E6 (R3) recommends the use of a complete black-box model as a mandatory requirement for AI utilization in the data collection process,” which is factually incorrect). The second problem relates to confidentiality in the research industry. Cloud LLMs are expected to be powerful tools for confirming consistency with the organization's documents, uploading internal documents to a cloud LLM and causing information leakage is extremely dangerous for the organization. While large organizations are internally developing and operating LLMs in a local environment, small CROs and AROs often find it difficult to allocate resources for these initiatives. Local LLMs

enable organizations to keep sensitive QMS and project data on-premises or within trusted clouds, reducing the risk of confidential data exposure and facilitating compliance with data residency rules.

We therefore focused on the retrieval-augmented generation (RAG) approach in local LLMs. The features of RAG include referencing external knowledge bases and generating new text while minimizing hallucinations. RAG can be developed using graphical processing units (GPUs) or run with conventional central processing units (CPUs), and the methods are evolving daily, offering options between large-scale and small-scale language models.¹¹ In the retrieval mechanism, text data are transformed into high-dimensional numerical vectors through a process known as “embedding,” which relies on exact word overlaps. Modern RAG systems utilize dense vector embeddings to capture the semantic context of the text. This capability is particularly crucial for clinical development, where the nuanced interpretation of SOPs and regulatory requirements is essential to maintain procedural integrity.

Our concept, as shown in **Figure 2**, is to first define universal conditions such as regulations and professional role definitions using a cloud-based LLM, then load this information along with the organization's proprietary QMS pyramid documents into a local LLM environment to be used as a knowledge base to create new governance documentation for generative AI utilization. Continuing this cycle will refine the governance related to generative AI. We deployed a local, open-source RAG implementation, successfully generating governance documents for the




Figure 2: Using a Local LLM and RAG methods, Generate the organization's R&D-specific Generative AI Governance Documents.

Set company-wide guidelines, organizational research SOPs, and global/regional regulatory guidance as “knowledge domains,” and produce R&D-focused generative AI governance documents through RAG. This activity is sustainable and will be updated as guidance and SOPs are revised.

*MHLW: Ministry of Health Labour and Welfare, Japan.

organization's QMS pyramid. Benefits include: "Rapid creation of governance documents highly aligned with regulatory requirements and existing QMS," "Generation of SOPs tailored to the organization's operational realities," and "Establishment of a low-cost, secure environment for even small organizations to benefit from generative AI."

Integrating generative AI into the QMS framework enables the transition of CDM roles from manual document creation to strategic oversight. For instance, by referencing clinical protocols and study-specific EDC specifications as a RAG knowledge base, it becomes possible to automate the initial drafting of key documents, such as the data management plan (DMP), EDC development plan, and data review procedures. By automating these drafts via RAG, organizations can significantly reduce the lead time for study startup while minimizing human error in document consistency. Furthermore, this framework facilitates the integration of CDM processes into consistent QMS and risk-based approach (RBA) workflows. The establishment of harmonized, real-time compliance metrics—such as the discrepancy rate between AI-generated drafts and human-finalized versions—allows for continuous process optimization within the QMS. As CDM functions are at the forefront of the digitalization of clinical research, this evolution will directly contribute to enhancing the overall quality and efficiency of the clinical development lifecycle.

Practical Approaches and Common Pitfalls in Using RAG

At Kowa Co., Ltd., incorporating company-wide AI usage guidelines as reference documents enabled alignment between common principles (such as data privacy and information security) and specific clinical trial requirements, thus allowing attempts to generate documents for constructing a QMS pyramid. Creating governance documents tailored for clinical trials demonstrated that guidelines alone were insufficient; it was necessary to combine them with various SOPs and regulatory guidance related to clinical trials. Using RAG enabled efficient generation of governance documents that comprehensively covered key elements such as ethical considerations, data privacy, security, AI reliability, and the use of generative AI in clinical trials. However, it was confirmed that the generated governance documents remained at the high-level framework stage and did not necessarily provide specific cautions or operational decision criteria. Some requirements—such as verification methods and evaluation metrics—were practically difficult for users to implement, and there were limitations to applying the generated outputs directly to operations. In the current implementation of QMS documentation at Kowa Co., Ltd., embedding models specialized for medical terminology, such as BioBERT were not employed for the vectorization process. However, utilizing such models specialized for medical terms could potentially expand the application of this framework to more complex document ecosystems. This includes the integration of comprehensive information from investigator's brochures and safety databases, thereby facilitating the

establishment of a more robust and specialized clinical documentation hierarchy.

Documents produced by RAG should therefore be treated as drafts, not final versions, and should be reviewed and corrected manually. It was also suggested that, when viewpoints are missing or expressions are ambiguous, document quality can be improved incrementally by updating prompts and reviewing reference knowledge.

Establishing Ethical Principles and Human in the Loop (HITL)

Building on the practical challenges outlined earlier it becomes evident that ensuring reliability and safety requires more than just technical solutions. These issues highlight the critical need to establish a solid ethical framework and incorporate human judgment into the governance of generative AI in clinical development.

The application of generative AI must therefore not only preserve scientific integrity but also address ethical and legal considerations. We propose applying fundamental principles of medical ethics—such as Autonomy, Non-maleficence, Justice, and Beneficence—to guide AI use. This involves developing SOPs to proactively identify and mitigate potential biases, especially those related to race and socioeconomic status, originating from training data or model outputs.

Moreover, to effectively manage these risks, it is essential to establish HITL as a core element of governance in clinical development.¹² When AI outputs influence critical clinical or regulatory decisions—such as protocol drafting, data interpretation, or updates to QMS documents—the final responsibility must rest with a human decision-maker. This approach mitigates risks associated with hallucinations or unforeseen errors and clarifies accountability. Consequently, governance documentation should specify the level of human intervention required for each specific use case—whether it involves a simple review after automated generation or making important decisions based on AI outputs, (see **Figure 3**).

Continuous Monitoring and Change Control Throughout the AI Model Lifecycle

Generative AI model performance constantly carries the risk of deviating from its initial validation performance (Model Drift) due to the influx of real-world data different from the training data or the manifestation of biases over time. As this performance degradation has a direct impact on trial reliability and patient safety, especially in the context of clinical development, the governance framework must cover the entire AI model lifecycle.

Our governance documentation concept proposes integrating an "AI Utilization Continuous Monitoring SOP" within the QMS pyramid. This SOP will specifically define the following elements:

1. **Establishment of Performance Metrics and Thresholds:** Clearly define the acceptable performance baseline and the deviation threshold for each generative AI use case (e.g., accuracy of protocol drafts, error rate during data integration).




Figure 3: Conceptual Framework for Utilization of Generative AI Governance.

Based on the global and regional regulation for generative AI implements to the each organizational QMS document pyramid and reviewed by human in the loop methods.

2. **Periodic Audits:** Audit the discrepancy rate between AI output and human judgment, as well as bias indicators against specific attributes (e.g., race, gender) at regular intervals (e.g., quarterly).
3. **Change Control:** Define and integrate into the QMS response procedures—including risk assessment, re-validation, model re-tuning, or temporary cessation of use—if a deviation is detected or if the underlying LLM service (e.g., provider version update) is changed.
4. **Decommissioning:** Establish a safe and traceable retirement process if the model's performance degrades to an irrecoverable level or if the purpose of use ceases.¹³

By incorporating this continuous monitoring and changing control mechanism into the QMS, organizations can benefit from the convenience of generative AI while maintaining audit readiness for regulatory authorities. The level of oversight should be risk-based, with lower-risk tools (e.g., document generation) warranting less scrutiny than higher-risk use cases (e.g., data analysis).

Conclusion

Going forward, governance documentation for generative AI is expected to be developed efficiently and reliably. Operations are likely to employ RAG with local LLMs as a standard approach. However, the implications of this technological integration extend beyond operational efficiency; they signify a paradigm shift from traditional CDM to clinical data science (CDS). As the core competencies of data professionals evolve from manual data management procedures to strategic algorithmic

oversight and the generation of advanced data insights, the necessity for a robust governance framework becomes even more critical. The transition to CDS requires not just technical proficiency, but a QMS-driven approach to monitor and validate AI-driven insights.¹⁴ The framework presented in this paper, which integrates RAG, ethical principles, HITL, and lifecycle management, serves as a set of essential “guardrails.” This governance structure ensures that clinical data scientists leverage AI for complex pattern recognition and predictive analytics while maintaining strict regulatory compliance. Ultimately, this approach offers a practical and responsible pathway for clinical organizations to redefine the value of data functions, fostering an environment in clinical development that is agile and driven by data.

Acknowledgements

We would like to thank Takashi Kawaguchi (Tokyo University of Pharmacy and Life Sciences) for his insightful advice on the potential applications of local LLMs in clinical development. We also thank Junya Kimura (Tohoku University Hospital) for guidance on utilizing Dify.

Competing Interests

The authors have no competing interests to declare

References

1. **World Health Organization.** *Ethics and governance of artificial intelligence for health guidance on large multi-modal models.* World Health Organization; 2024. Accessed February 10, 2026. <https://iris.who.int/server/api/core/bitstreams/e9e62c65-6045-481e-bd04-20e206bc5039/content>

2. **US Food and Drug Administration.** *Considerations for the use of artificial intelligence to support regulatory decision-making for drug and biological products.* Draft Guidance. US Department of Health and Human Services; 2025. Accessed February 10, 2026. <https://www.fda.gov/media/184830/download>
3. **European Medicines Agency.** *Reflection paper on the use of artificial intelligence (AI) in the medicinal product lifecycle.* European Medicines Agency; 2024. Accessed February 10, 2026. https://www.ema.europa.eu/en/documents/scientific-guideline/reflection-paper-use-artificial-intelligence-ai-medicinal-product-lifecycle_en.pdf
4. **Gadde SS, Kalli VD.** Artificial intelligence at healthcare industry. *Int J Res Appl Sci Eng Technol.* 2021;9(2):313–318. DOI: <https://doi.org/10.22214/ijraset.2021.32996>
5. **UK Medicines & Healthcare products Regulatory Agency.** *Impact of AI on the regulation of medical products. Implementing the AI white paper principles.* HM Government; 2024. Accessed February 10, 2026. https://assets.publishing.service.gov.uk/media/662fce1e9e82181baa98a988/MHRA_Impact-of-AI-on-the-regulation-of-medical-products.pdf
6. **Canada's Drug Agency.** *Position statement on the use of artificial intelligence in the generation and reporting of evidence.* Canada's Drug Agency; 2025. Accessed February 10, 2026. https://www.cda-amc.ca/sites/default/files/MG%20Methods/Position_Statement_AI_Renumbered.pdf
7. **Infocomm Media Development Authority of Singapore, Aicadium, AI Verify Foundation.** *Model AI governance framework for generative AI fostering a trusted ecosystem.* IMDA; 2024. Accessed February 10, 2026. <https://aiverifyfoundation.sg/wp-content/uploads/2024/05/Model-AI-Governance-Framework-for-Generative-AI-May-2024-1-1.pdf>
8. **Cyberspace Administration of China, Office of the Central Cyberspace Affairs Commission.** *Interim measures for the management of generative artificial intelligence services.* Published July 13, 2023. Accessed February 10, 2026. https://www.cac.gov.cn/2023-07/13/c_1690898327029107.htm
9. **Expert Group on how AI Principles Should be Implemented.** *AI governance in Japan ver. 1.1.* Ministry of Economy, Trade and Industry; 2021. Accessed February 10, 2026. https://www.meti.go.jp/shingikai/mono_info_service/ai_shakai_jisso/pdf/20210709_8.pdf
10. **Takata M, Sugimoto M, Harada M, et al.** Governance for generative AI in clinical development: A cross-sectional survey in Japan. *J Soc Clin Data Manage.* 2026;6(1). DOI: <https://doi.org/10.47912/jscdm.486>
11. **Lewis P, Perez E, Piktus A, et al.** Retrieval-augmented generation for knowledge-intensive NLP tasks. *Adv Neural Inf Process Syst.* 2020;33:9459–9474.
12. **Lemaire W, Nadolny P, Andrus J, Schaffer C, Cameron S.** SCDM Executive Committee perspective: The future of clinical trials with AI and decentralized, hybrid and data-driven approaches. *J Soc Clin Data Manage.* 2025;5(3). DOI: <https://doi.org/10.47912/jscdm.448>
13. **Liao F, Adelaine S, Afshar M, Patterson BW.** Governance of clinical AI applications to facilitate safe and equitable deployment in a large health system: key elements and early successes. *Front Digit Health.* 2022;4:931439. DOI: <https://doi.org/10.3389/fgdth.2022.931439>
14. **Society for Clinical Data Management.** *SCDM industry competency framework.* Society for Clinical Data Management. Accessed February 10, 2026. <https://scdm.org/cdm-competency-framework/>

How to cite this article: Takata M, Ichikawa D, Harada M, Hiraide M, Sugimoto M, Hayashi Y, Yoshimaru Y, Fujisawa K, Nojima T, Fukimbara S, Suganami H, Yamaguchi T. Generative AI and Clinical Development: Exploring the Frontier of AI Governance Utilizing Generative AI -Conceptual Framework for AI Governance in Clinical Development. *Journal of the Society for Clinical Data Management.* 2026; 6(1): 7, pp. 1–6. DOI: <https://doi.org/10.47912/jscdm.500>

Submitted: 25 March 2026

Accepted: 08 August 2026

Published: 16 September 2026

Copyright: © 2026 SCDM publishes JSCDM content in an open access manner under a Attribution-Non-Commercial-ShareAlike (CC BY-NC-SA) license. This license lets others remix, adapt, and build upon the work non-commercially, as long as they credit SCDM and the author and license their new creations under the identical terms. See <https://creativecommons.org/licenses/by-nc-sa/4.0/>.



Journal of the Society for Clinical Data Management is a peer-reviewed open access journal published by Society for Clinical Data Management.

OPEN ACCESS