Paper SA13

Digital Twins for Clinical Submissions: Simulating Regulatory Review Before Submission

Divya Kapoor PhD, Sycamore Informatics, Irvine, CA, United States

ABSTRACT

Regulatory submissions in the life sciences are often delayed by unforeseen reviewer queries, inconsistencies in metadata, or gaps in documentation quality. This paper introduces a “digital twin” framework designed to replicate the end-to-end regulatory review process before submission. Leveraging AI-driven validation, reproducibility checks, and simulated mock inspections, the system enables sponsors to proactively stress-test datasets, outputs, and narratives against regulatory expectations. Unlike conventional quality control methods, the digital twin model provides a dynamic feedback loop that highlights inconsistencies, flags non-compliance risks, and offers targeted remediation strategies before database lock. Early case studies demonstrate a measurable reduction in submission-cycle iterations and query volumes, alongside improved reviewer confidence in data integrity. By embedding a regulatory perspective into the development pipeline, this approach not only de-risks submission packages but also accelerates approval timelines, offering a scalable pathway to more efficient regulatory interactions across therapeutic domains.

Keywords: ADaM, AI, CDISC, CSR, DT, eCTD, EMA, FDA, Gen AI, GxP, HITL, ICH, LLM, ML, MVP, NCT, NLP, PMDA, SAP, SCE, SDTM, SOP, TLF, QC

1. INTRODUCTION

Regulatory submissions undergo a rigorous review process by agencies such as the U.S. Food and Drug Administration (FDA), European Medicines Agency (EMA), and Pharmaceuticals and Medical Devices Agency (PMDA) to ensure a drug’s safety, efficacy, and quality before approval. Despite the significant effort invested in preparing and reviewing these packages, submissions frequently encounter substantial delays or rejections. These setbacks result in increased costs, extended approval cycles, and delayed product launches for pharmaceutical and biotechnology companies. The common pitfalls in these submissions include:

Clinical Data Inconsistencies: One of the leading causes of submission rejections is the presence of incomplete or inconsistent data e.g., inconsistent value across modules, partial/incorrect datasets, unclear handling of missing data, missing reviewer’s guides, incorrect subject numbers across documents, dataset-define.xml mismatches, dataset values not matching with the CSR results etc.

Lower-Quality Documentation: The other common cause of rejections is the poor quality of submission documents. Lack of completeness and clarity on the results often causes regulatory queries or rejections e.g., documents placed in wrong modules/sections, broken or missing hyperlinks, incorrect file naming conventions, appendix referenced but not submitted, mislabeled outputs etc.

Non-compliance with Regulatory Standards: Submission guidelines are regional—the eCTD structure mandatory for FDA submissions, for instance, does not apply to the EMA or PMDA. As such, all submissions must comply with their respective regional and submission-type requirements. Typical compliance gaps involve a lack of required regional regulatory forms, or errors in formatting and labeling.

The entire pharmaceutical landscape is being rapidly revolutionized by advanced digital technologies, particularly artificial intelligence and machine learning (AI/ML). As the industry targets automation and operational efficiency, the concept of Digital Twins (DT) is attracting significant attention for use in clinical domain.

2. DIGITAL TWIN (DT)

Digital twin is a virtual replica of a real-world physical product, process or system. These virtual replicas differ from simulation by being dynamic and adaptive as they are continuously updated based on real-world data and feedback. Currently, DTs are being explored in clinical research e.g. in drug discovery and development, clinical trial design and simulation, predictive modeling of disease and treatment responses, personalized medicine, process development etc. Digital twin technology is currently being used in several clinical trials to prove its value, and a few examples of notable studies are:

Type 2 Diabetes (NCT05181449): An ongoing trial uses DT to enable precision treatment for type 2 diabetes mellitus patients.

Musculoskeletal Conditions (NCT04849923): An active trial employing DT to develop better, optimized treatment strategies for injuries and diseases related to the musculoskeletal system.

Diabetes Insulin & Meal Prediction (NCT04203823): Completed trials have demonstrated the value of DTs in enhancing diabetes management by improving insulin delivery and accurately predicting meal responses.

Digital Twins are being introduced into regulatory contexts but have not yet achieved full acceptance as primary evidence for formal drug submission packages. Current applications are focused on supporting within manufacturing and device contexts for certain regulatory submissions and pilot programs, with regulatory bodies actively exploring their expanded potential.

Regulatory bodies are actively establishing frameworks to govern the use of AI/ML to support decision-making on a drug's safety, effectiveness, or quality. The U.S. FDA released a Draft Guidance in January 2025 focused on a risk-based credibility assessment framework for evaluating the reliability and trustworthiness of AI models in specific 'contexts of use'. The European Medicines Agency (EMA) published a Reflection Paper in October 2024 and reached a significant milestone in March 2025 by issuing its first qualification opinion on an AI methodology, signaling a readiness to accept evidence generated by validated AI tools. The FDA and the EMA have not yet published explicit formal guidance (as of early 2026) specifying the requirements for the use of digital twins in regulatory submissions.

In this paper, we propose using Digital Twin (DT) technology to create a more accurate method for assembling regulatory submission components. By reducing manual reconciliation, DTs allow organizations to proactively identify and resolve any data inconsistencies or document gaps. This strategy not only improves the overall quality of the submission but also drives greater efficiency and shortens timelines for a faster path to approval.

2.1 Regulatory DT Generation

The success of an effective and reliable regulatory digital twin hinges on the availability of massive amounts of detailed, high-quality data. To construct a reliable and effective data-driven DT, the organizational context is paramount, specifically all historical, high-quality datasets from prior clinical trials, disease registries, and real-world evidence. Training the model also necessitates the organization's codified standards & contextual metadata (such as SDTM, ADaM, and therapeutic area-specific rules) and extensive documentation, including SAPs, Protocols, study reports, and regulatory feedback from historical submissions. Organization-specific data embeds institutional knowledge into DT, making it accurate, reliable, precise, and practically useful.

Each organization has unique risk patterns such as data traceability gaps, metadata inconsistencies, and misalignments between narratives, tables, and datasets. Only organization-specific historical data can empower a digital twin to identify these failure modes, reliably predict submission risks, and prioritize the most critical areas, providing insights that generic data simply cannot. Beyond risk management, this data allows the digital twin to anticipate reviewer queries, simulate realistic review behavior, and stress-test

submissions against historical expectations, effectively creating a “mock review” environment that mirrors how regulators are likely to evaluate.

Therefore, creation of DT using organization-specific data transforms it from a theoretical model into a decision-ready, context-aware, and actionable tool, making it far more accurate, dependable, and effective than one built on generic data.

2.2 Regulatory DT Architecture

The regulatory DT framework comprises of four layers – Data Integration & Harmonization layer, Submission Model Layer, Analytics Layer and Continuous Learning Loop.

2.2.1 Data Integration & Harmonization Layer: The regulatory Digital Twin (DT) is fundamentally built upon an extensive inventory of data and document sources, which forms its core foundation. This layer is used to ingress organization’s clinical trial data - EDC, datasets (SDTM, ADaM), statistical outputs (Tables, Listings, Figures/Graphs), narratives (patient level, study-level, integrated clinical narratives etc.) and key documents such as study reports, protocols, SAPs, CSRs, investigator brochures, informed consent forms etc. The data is harmonized using metadata (controlled terminology) and data quality rules to enforce consistency in units, formats, and coding, thereby ensuring its integrity. The final output of this foundational layer is the standardized, high-quality searchable content ready for use by the DT model.

2.2.2 Submission Model Layer: This layer establishes the foundation for regulatory readiness, serving as the backbone for guaranteeing the completeness and compliance of the submission package according to regulatory expectations. It integrates essential guidelines from multiple authorities such as FDA, EMA, PMDA, mandatory submission standards such as eCTD, ICH, CDISC rules, and supporting checklists. By defining these global and country-specific rules, its primary function is to map specific submission content directly to the correct regulatory modules and section. By systematically linking data, outputs, and submission documents to their respective module/section, this layer provides end-to-end traceability, supports regulatory compliance, and ensures the completeness and internal consistency of the submission package.

2.2.3 Analytics Layer: This layer leverages artificial intelligence and machine learning to simulate regulatory review process scenarios. Compliance checks are performed to ensure submission packages are complete and consistent, identifying any potential gaps or risks or discrepancies. It proactively predicts potential reviewer queries, highlights high-risk areas within submission packages example narrative-table-dataset misalignments, flags data and metadata inconsistencies, detects incorrect file placement across modules/sections and suggests targeted remediation actions. It also includes a "what-if" analysis feature for proactive decision-making, which predicts potential downstream gaps or delays. Ultimately, this ensures the submission package's regulatory compliance and overall readiness, allowing organizations to mitigate issues before they arise.

To maximize the value and trustworthiness of the Digital Twin (DT), a Human-in-the-Loop (HITL) capability is indispensable. The HITL professional serves as a vital safeguard, responsible for the thorough review of all model outputs. This human validation step is critical for ensuring the DT's results are interpretable, compliant, and accountable. By subjecting every generated output to expert scrutiny, the organizations can effectively mitigate systemic errors, prevent the dissemination of non-factual results (hallucinations), and proactively address any risks related to regulatory non-compliance, thereby ensuring DT’s operational integrity.

2.2.4 Continuous Learning Loop: The primary purpose of this component is to keep the DT ‘alive’ by continuously feeding real-world data, regulatory updates and expert feedback. This feedback is sourced from multiple critical sources, including post-submission outcomes such as regulatory authority questions and identified gaps, internal inspection findings, final submission outcomes as well as evolving regulatory guidelines and newly generated data. By systematically incorporating this continuous stream of real-world feedback, the system dramatically improves the predictive accuracy and relevance of its compliance and gap analysis.

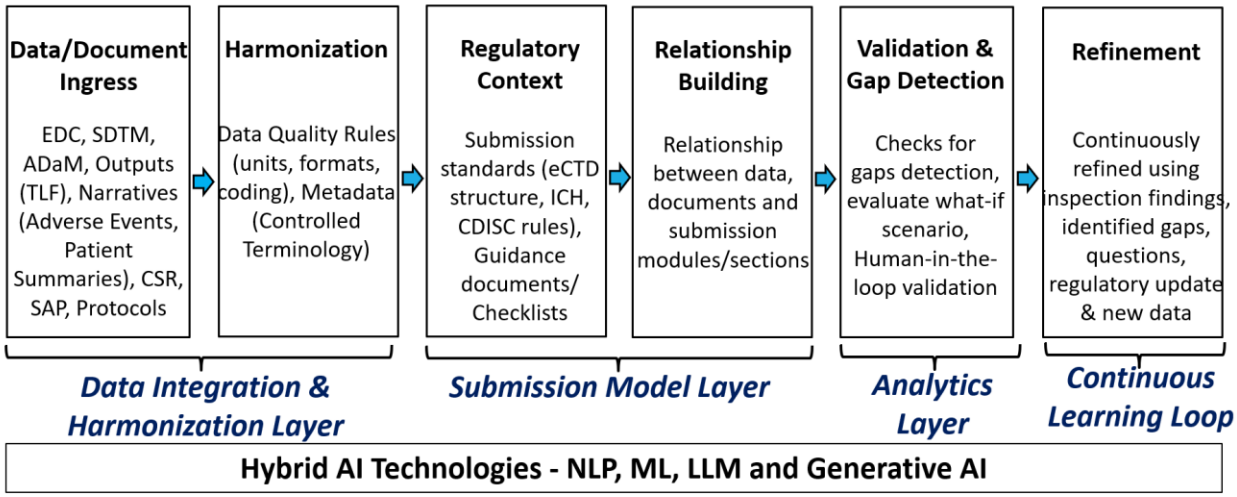

FIG 1: Mapping the Regulatory Digital Twin's process workflow to its architecture layers

The DT framework utilizes a hybrid suite of advanced Artificial Intelligence (AI) technologies, including Large Language Models (LLMs), Generative AI (Gen AI), Natural Language Processing (NLP), and Machine Learning (ML).

Large Language Models (LLMs) are employed to facilitate data harmonization and simulate the regulatory review process. This involves interpreting regulatory language, semantically comparing against data and industry standards to identify potential gaps and inconsistencies across submission modules, and proactively predicting regulatory queries.

Generative AI (Gen AI) is leveraged for sophisticated scenario simulation, enabling the prediction of 'what-if' scenarios. It generates synthetic edge cases to rigorously stress-test submission packages, thereby supporting data-driven decision-making and ensuring comprehensive submission readiness.

Machine Learning (ML) identifies and predicts high-risk areas in submission packages, highlighting inconsistencies or outliers likely to trigger queries.

Natural Language Processing (NLP) autonomously processes and comprehends submission documents, effectively mimicking a reviewer to detect misalignments between textual content and outputs, and identifying missed references.

2.3 Regulatory DT Value Proposition

Integrating a Digital Twin (DT) approach to simulate the regulatory review process is poised to enhance the efficiency of organizational submission processes. The benefits of this approach are:

2.3.1 Enhanced Efficiency and Submission Success

Mitigating Uncertainty and Stress-Testing: The DT framework is designed to replicate the end-to-end regulatory review process *before* the actual submission. This allows sponsors to proactively stress-test datasets, outputs, and narratives against regulatory expectations, which directly mitigates the uncertainty of the submission outcome.

Proactive Risk Identification: It allows for the early identification and mitigation of inconsistencies or risks in submission artifacts. This is achieved through AI-driven validation, reproducibility checks, and simulated mock inspections.

Dynamic Feedback Loop: Unlike traditional QC methods, the DT provides a dynamic feedback loop that highlights inconsistencies, flags non-compliance risks, and offers targeted remediation strategies before database lock.

2.3.2 Accelerated Cycles and Cost Reduction

Reduction in Iterations and Queries: Early case studies of this framework have demonstrated a measurable reduction in submission-cycle iterations and query volumes from reviewers.

Accelerated Approval Timelines: By embedding a regulatory perspective into the development pipeline, the approach not only de-risks submission packages but also accelerates approval timelines. This acceleration directly contributes to reduced trial costs and durations by minimizing delays and manual rework.

2.3.3 Continuous Compliance and Quality Improvement

Continuous Regulatory Compliance: The framework is designed to ensure continuous regulatory compliance by embedding checks and balances throughout the development lifecycle prior to submission.

Improved Quality**:** Quality is improved through the early identification of issues, ensuring greater reviewer confidence in data integrity.

Scalable Pathway: This approach offers a scalable pathway to more efficient regulatory interactions across various therapeutic domains.

These benefits collectively lead to a more efficient and de-risked process, all of which is underpinned by data-driven decisions derived from the DT simulated environment.

2.4 Regulatory DT Challenges

Implementing this initiative presents several key challenges. The most significant impediments, designated as high-impact, stem from Technical and Infrastructural challenges. These primarily revolve around guaranteeing the robust integrity, security, privacy and governance of sensitive operational data, meeting the scalability demands necessary for compute-intensive analytics, and establishing protocols for the reproducibility of generated results. Crucially, this also requires the establishment of comprehensive traceability and audit trails to enable the meticulous monitoring of all updates within the DT environment.

Furthermore, challenges of medium impact are evident in two main areas: Organizational and Ecosystem-Related. On the organizational front, successful deployment necessitates change management and adoption strategies, alongside mitigating the workforce skill deficit in key areas like AI/ML and data science. From an ecosystem perspective, regulatory hurdles include the lack of interoperable data standards and the need for continuous adaptation to evolving regulatory guidelines. Finally, lower impact challenges relate primarily to managing expectations, addressing ethical concerns, and maintaining clear accountability structures for the regulatory DT.

From a technical perspective, the successful implementation of the Digital Twin (DT) framework hinges on robust technology integration. This requires cloud-native, scalable solutions to manage compute-intensive analytics, with a modern, GxP-compliant Statistical Computing Environment (SCE) serving as a key enabler. Since data must be integrated from multiple sources while strictly maintaining integrity, security, governance and privacy through role-based access controls, the DT framework must reside within a fully controlled, regulated GxP environment.

Moreover, unlike static simulations, DTs are dynamic and adaptive, necessitating continuous synchronization and updates based on real-world data or regulatory guidelines. The SCE is crucial for ensuring that all resulting changes are end-to-end traceable, versioned, and documented with comprehensive audit trails.

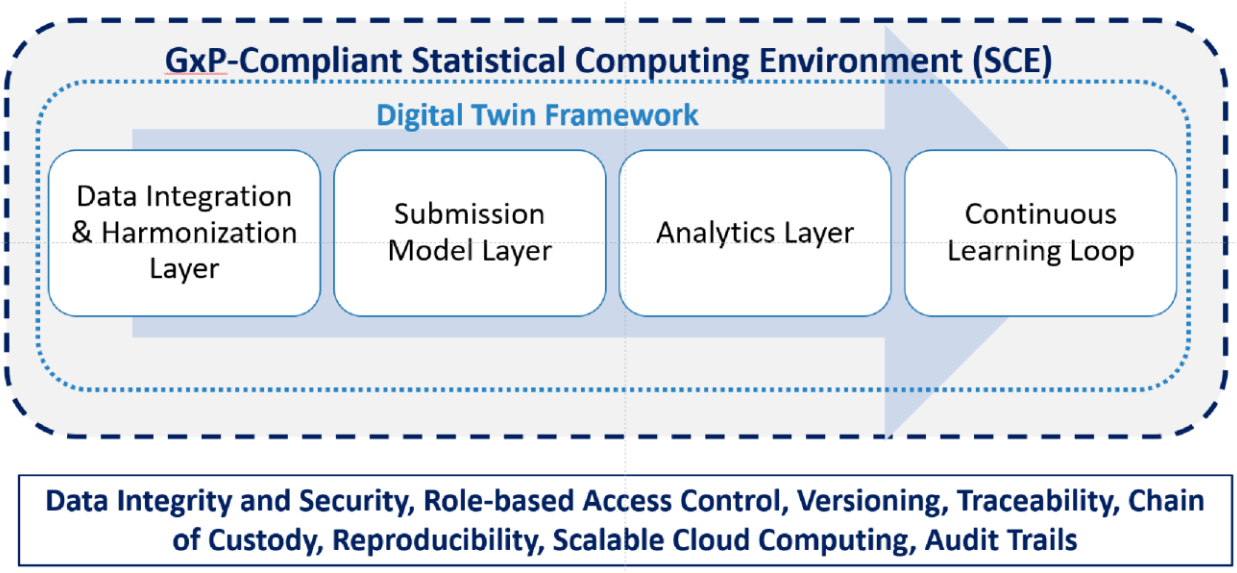

FIG 2: A modern GxP-compliant SCE as a key enabler for Digital Twin Implementation

3. FUTURE DIRECTIONS

DTs hold a promising future in the pharmaceutical industry and the main areas to focus on are:

3.1 Phased Implementation: The strategy for validating and implementing regulatory DT involves a phased approach: initial pilot studies across diverse therapeutic areas need to be conducted to establish the DT's reliability and effectiveness, paving the way for a scaled implementation that progresses from a Minimal Viable Product (MVP) to a full Enterprise-Level DT.

Phase 1: Pilot Studies (Proof-of-Concept): The initial phase focuses on targeted, small-scale pilot studies across multiple therapeutic areas (e.g., oncology, cardiovascular, neurology etc.) to rigorously access the accuracy, statistical reliability and practical impact of the predictions.

Phase 2: Minimum Viable product (MVP): Upon successful validation from the pilots, the projects move to developing a MVP. The objective of this phase is to create a fully functional, end-to-end DT that provides tangible, early value in a live, limited environment. This version is about proving the operational workflow and positive feedback on the MVP's usability and value.

Phase 3: Full Implementation (Enterprise-Level DT): The final phase involves full-scale deployment, transforming the MVP into a comprehensive, Enterprise-level DT solution. The objective of this phase is to create an Enterprise-level DT that is standardized, fully integrated, and adopted across all relevant organizational therapeutic areas.

3.2 Regulatory-industry Partnerships: The core challenge for DT-enabled submissions is lack of interoperable data standards and evolving regulatory guidelines. Without a unified set of guidelines, regulatory bodies (like the FDA, EMA) cannot consistently evaluate submissions, and industry sponsors cannot ensure their DTs are compliant.

Promoting partnerships between regulatory bodies and key industry leaders is a critical next step for the widespread, trusted adoption of Digital Twin (DT)-enabled submissions. These collaborations are essential for moving DT technology from proof-of-concept projects to globally recognized, compliant tools.

3.3 Change Management and Adoption: Achieving the full benefit of a dynamic, transformative DT approach for regulatory submissions requires a dual effort: first, a fundamental cultural and operational shift

from static, traditional Quality Control (QC) to a proactive, predictive compliance strategy; and second, a strategic, systemic investment to close the workforce skill deficit in critical areas like AI, ML, and data science, which is necessary to fully transition the technology from a proof-of-concept to a scalable enterprise solution.

4. CONCLUSION

DTs hold significant promising future in the clinical domain and are increasingly gaining traction across the industry. A key question remains: are DTs intended to replace traditional quality control (QC) methods, or are they better positioned as complementary tools? To answer this, it is important to first examine the limitations of traditional QC.

Traditional QC methods are fundamentally static and reactive. It primarily involves checking submission artifacts *after* they have been produced. This approach is prone to:

Delayed Discovery: Issues like inconsistencies in metadata or gaps in documentation are often discovered late in the development cycle, leading to costly and time-consuming rework before database lock.

Fragmented, point-in time checks: Reviews are performed in silos and lack the holistic assessment.

Unforeseen Queries: Submissions are frequently delayed by unforeseen reviewer queries that arise during the regulatory review process, indicating a lack of alignment with regulator expectations.

Slower approval cycles: Inconsistencies or gaps in data/document lead to extended review cycles.

Manual Reconciliation: Prone to error due to cross-referencing submission artifacts across modules/sections.

Inconsistent Results: Quality of package largely depends on individual reviewer expertise and interpretation of guidelines, leading to variability and inconsistency in the package.

Effort-intensive: QC quality varies by reviewer and hard to scale across therapeutic areas and submission types.

Static: Traditional QC depends on predefined SOPs, checklists and regulatory guidance that are not designed to evolve at the pace of the regulatory updates.

Traditional quality control (QC) approaches are fundamentally reactive in nature, operating through retrospective verification activities such as manual document reviews, checklist-driven assessments, and adherence to predefined Standard Operating Procedures (SOPs) after the submission artifacts have already been produced. While these methods are effective at identifying surface-level errors and ensuring baseline compliance, they are limited in their ability to detect systemic issues, contextual inconsistencies, or emerging regulatory risks early in the submission lifecycle. As a result, deficiencies often surface late, when remediation is costly and timelines are already constrained.

In contrast, the Digital Twin (DT) framework represents a dynamic and predictive paradigm that virtually replicates the end-to-end regulatory review process prior to submission. By leveraging AI-driven validation, automated reproducibility checks, and simulated mock inspections, the DT continuously evaluates data, metadata, analytical outputs, and clinical narratives in an integrated manner. This enables the creation of a closed-loop feedback system that proactively surfaces non-compliance risks, internal inconsistencies, and high-risk areas likely to trigger regulatory scrutiny.

Through iterative simulation and scenario testing, DTs allow sponsors to stress-test submission components against evolving regulatory expectations, reviewer behaviors, and inspection standards. This capability supports earlier risk mitigation, improves submission robustness, and enhances first-cycle review outcomes. While traditional QC can identify errors at a transactional level, its late-stage and static

application often limit its effectiveness in preventing significant submission delays or extensive reviewer inquiries. DTs, therefore, do not merely identify defects but actively support submission readiness by enabling foresight, adaptability, and continuous learning across the regulatory lifecycle.

Consequently, the DT approach yields a measurable reduction in submission-cycle iterations and query volumes, directly translating to an acceleration of regulatory timelines, a superior outcome when compared to the efficiency limitations imposed by static QC methods.

Rather than replacing traditional QC, Digital Twins are best positioned as a complementary capability that augments existing quality frameworks. While QC remains essential for ensuring procedural consistency and baseline compliance, DTs extend these controls by introducing predictive intelligence, system-level validation, and continuous learning. Together, they enable a shift from late-stage error detection to early risk anticipation and prevention.

By embedding DTs upstream in the submission lifecycle and retaining QC as a downstream safeguard, sponsors can achieve greater submission robustness, reduce regulatory friction, and improve first-cycle success. This shift fundamentally changes quality assurance from a final check to a continuous, predictive process.

Appendix: Abbreviations with full form

Abbreviation	Full Form
ADaM	Analysis Data Model
AI	Artificial Intelligence
CDISC	Clinical Data Interchange Standards Consortium
CSR	Clinical Study Report
DT	Digital Twin
eCTD	Electronic Common Technical Document
EMA	European Medicines Agency
FDA	Food and Drug Administration
Gen AI	Generative Artificial Intelligence
GxP	Good Practice
HITL	Human-in-the-Loop
LLM	Large Language Models
ICH	International Council for Harmonisation of Technical Requirements for Pharmaceuticals for Human Use
ML	Machine Learning
MVP	Minimum Viable Product

NCT	National Clinical Trial
NLP	Natural Language Processing
PMDA	Pharmaceuticals and Medical Devices Agency
SAP	Statistical Analysis Plan
SCE	Statistical Computing Environment
SDTM	Study Data Tabulation Model
SOP	Standard Operating Procedure
TLF	Tables, Listings and Figures
QC	Quality Control

REFERENCES

1. U.S. Food and Drug Administration. Considerations for the Use of Artificial Intelligence to Support Regulatory Decision-Making for Drug and Biological Products: Draft Guidance for Industry and Other Interested Parties. January 2025. https://www.fda.gov/regulatory-information/search-fda-guidance-documents/considerations-use-artificial-intelligence-support-regulatory-decision-making-drug-and-biological

2. European Medicines Agency. Reflection paper on the use of artificial intelligence (AI) in the medicinal product lifecycle (EMA/CHMP/CVMP/83833/2023). October 2024. https://www.ema.europa.eu/en/documents/scientific-guideline/reflection-paper-use-artificial-intelligence-ai-medicinal-product-lifecycle_en.pdf

3. U.S. Food and Drug Administration. Digital Twin In FDA Digital Health and Artificial Intelligence Glossary. 2025. https://www.fda.gov/science-research/artificial-intelligence-and-medical-products/fda-digital-health-and-artificial-intelligence-glossary-educational-resource

4. The Cleveland Clinic. Randomized controlled trial of digital twin precision treatment: A novel whole body digital twin enabled precision treatment for type 2 diabetes (ClinicalTrials.gov Identifier NCT05181449). February 1,2022. https://clinicaltrials.gov/study/NCT05181449

5. Swiss Federal Institute of Sport Magglingen. Validation of a Digital Twin to Predict the Adaptation of Muscle Strength and Muscle Size Due to Strength Training (ClinicalTrials.gov Identifier NCT04849923). January 1, 2022. https://clinicaltrials.gov/study/NCT04849923

6. Medtronic Diabetes. Feasibility Study With Personalized Closed Loop (PCL) (ClinicalTrials.gov Identifier NCT04203823). July 20, 2020. https://clinicaltrials.gov/study/NCT04203823

CONTACT INFORMATION

Your comments and questions are valued and encouraged. Contact the author at:

Author Name: Divya Kapoor, PhD
Company: Sycamore Informatics
Address: Irvine, CA, United States
Email: divya.kapoor@sycamoreinformatics.com
Website: www.sycamoreinformatics.com