IntuitionLabs

INTUITIONLABS / RESEARCH & PRACTICAL GUIDANCE

AI Clinical Study Report Automation: SDTM to Module 5

> Build practical AI for pharma and biotech with IntuitionLabs. We help life-science teams turn complex information and workflows into useful software, governed knowledge systems and AI tools.

Published by IntuitionLabs · 2026-04-19 · ai clinical study reports · csr automation · sdtm datasets · adam datasets · module 5 submission · clinical data management · regulatory automation

Original article: https://intuitionlabs.ai/articles/ai-clinical-study-report-automation-module-5

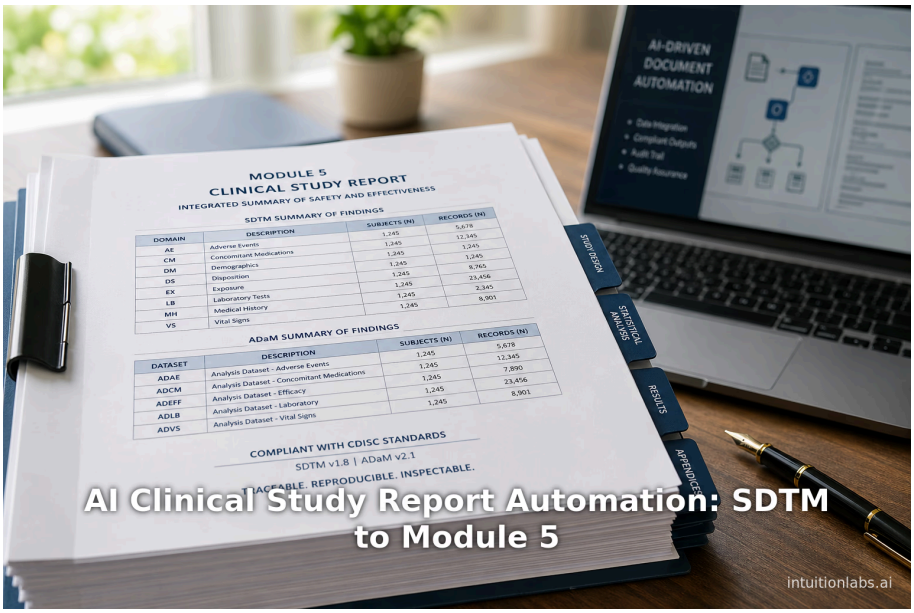

In this article

1. Executive Summary
2. Introduction and Background
3. SDTM and ADaM: The Data Foundation Beneath Every CSR Table
4. ICH E3 and eCTD Module 5: Turning Data Into a Regulatory Document
5. Where Generative AI Fits Today: Drafting, TLF Generation, and Consistency Checking
6. Vendor Landscape, Regulatory Posture, and the Limits of the Current Toolset
7. Implications and Future Directions
8. Conclusion

[Revised July 10, 2026]

Executive Summary

The clinical study report (CSR) is the single largest, most structured piece of writing a drug sponsor produces, and it sits at the exact point where two very different worlds meet: tabular clinical data and regulatory prose. On one side sit the CDISC data standards, chiefly the Study Data Tabulation Model (SDTM) and the Analysis Data Model (ADaM), which discipline how trial data is collected, organized, and analyzed (cdisc.org). On the other side sits ICH E3, the 1996 guideline that still governs the structure and content of a CSR today (fda.gov), and eCTD Module 5, the section of the electronic Common Technical Document where the finished report and its supporting data are filed with regulators (fda.gov).

Generative AI is now being layered onto this pipeline in a handful of concrete ways: drafting narrative text from ADaM-derived tables, generating and quality-checking tables, listings, and figures (TLFs), cross-checking consistency between the numbers in a CSR and the underlying datasets, and triaging regulatory intelligence so writers do not have to manually track every agency requirement. Vendors including Certara, IQVIA, ArisGlobal, Parexel, Cytel, and Formedix have all shipped or announced generative AI features aimed at some part of this workflow, with claimed efficiency gains that range from twenty to sixty percent on specific sub-tasks such as medical writing time or site selection (hitconsultant.net).

None of this changes the underlying regulatory requirement: a CSR must still conform to ICH E3, still be traceable back to SDTM and ADaM, and still survive a human medical writer's and statistician's sign-off. The FDA's January 2025 AI guidance is a nonbinding draft, not for implementation (fda.gov). Its proposed risk-based credibility framework concerns AI producing information or data to support regulatory decisions about drug safety, effectiveness, or quality. It excludes operational drafting or writing of a regulatory submission when that use does not affect patient safety, drug quality, or the reliability of nonclinical or clinical study results (fda.gov), and EMA's parallel reflection paper insists on a human-centric approach to AI across the medicine lifecycle (ema.europa.eu). The practical state of the art in mid-2026 is closer to AI-assisted first drafts and automated consistency checking than to any kind of autonomous CSR generation, and the published literature on large language models in clinical documentation continues to flag hallucination and validation gaps as the binding constraint (pmc.ncbi.nlm.nih.gov).

Introduction and Background

A clinical study report is the document a sponsor files to demonstrate that a specific trial was conducted, analyzed, and reported honestly enough that a regulator can rely on it to support (or fail to support) an approval decision. ICH describes the CSR as the single most important document a sponsor submits to support a new drug or biologic application (gmp-compliance.org). A pivotal Phase 3 CSR for a modestly complex trial commonly runs several hundred pages of narrative plus thousands of pages of appendices, tables, and listings, all of which must trace back to a locked, validated clinical database.

That traceability requirement is why CDISC exists. The Clinical Data Interchange Standards Consortium is a global not-for-profit that develops universal standards for data collected during clinical research (allucent.com), and its foundational standards, principally SDTM and ADaM, are what let a CSR's tables and narrative be regenerated, audited, and cross-checked years after database lock. FDA's initial standardized-study-data timetable applies to studies starting on or after December 18, 2016 for NDA, ANDA, and certain BLA submissions, using the applicable FDA-supported standards specified in its Data Standards Catalog (fda.gov). PMDA's Q&A specifies submission formats by study and analysis type, with exceptions for specified older studies and some pharmacokinetic analyses (pmda.go.jp). ADaM defines dataset and metadata standards that support generating, replicating, and reviewing clinical trial statistical analyses, with traceability among analysis results, analysis data, and SDTM data (cdisc.org). Submission expectations depend on the applicable agency and analysis: PMDA, for example, does not call for ADaM datasets for every analysis described in the statistical analysis plan (pmda.go.jp).

The document that turns those datasets into a narrative regulators can read is structured by ICH E3, a guideline first adopted by ICH in 1995 and by FDA on July 17, 1996, that is still in force (fda.gov). The guideline is explicitly a guideline and not a rigid template: it aims to allow the compilation of a single core clinical study report acceptable to all regulatory authorities of the ICH regions, with a core report and a set of appendices, some of which are regional and submitted only on request (ema.europa.eu). That single-core-report principle is what makes worldwide simultaneous filing practical, and it is also what makes the CSR such an attractive automation target: a standardized structure is exactly the kind of input generative AI systems handle best.

Where the CSR physically lands is eCTD Module 5, the clinical section of the electronic Common Technical Document. Module 5 contains the clinical study reports and related clinical data from the development program, organized per ICH M4E(R2) into a tabular listing of all studies (5.2), the individual clinical study reports themselves categorized by type (5.3), and literature references (5.4) (assyro.com). Within 5.3, section 5.3.5, covering the efficacy and safety studies, is generally considered the single most important section in Module 5, since it contains the full clinical study reports for the pivotal, adequate and well-controlled trials that carry the approval decision (assyro.com). The FDA's M4E(R2) guidance for the CTD's efficacy module explains how the Clinical Overview and the Clinical Summary reference back to those individual ICH E3 study reports (fda.gov), and current practice, rather than filing a CSR as one monolithic document, recommends dividing it into synopsis, report body, and individual appendices so that reviewers can navigate it and sponsors can update pieces without resubmitting the whole thing (allucent.com).

SDTM and ADaM: The Data Foundation Beneath Every CSR Table

Statistical analysis results generated from ADaM should be traceable through the analysis data to their SDTM inputs (cdisc.org). Other CSR content comes from controlled study documents: ICH E3, for example, calls for the planned sample size and its basis (fda.gov). Protocol and statistical-analysis-plan information should remain traceable to those source documents. SDTM organizes raw collected data (from an EDC system, a lab vendor, an ECG core lab, and so on) into standardized domains: demographics, adverse events, concomitant

medications, laboratory results, and dozens more, each following the CDISC Study Data Tabulation Model Implementation Guide (SDTMIG), currently at version 3.4 (cdisc.org). ADaM datasets are then created from SDTM data, frequently drawing on multiple SDTM domains at once, and are named using a consistent ADxxxx convention where xxxx typically echoes the source SDTM domain (quanticate.com). The point of this second layer is that ADaM datasets are meant to be analysis-ready: they drive the statistical programming process that creates tables, figures, and listings more efficiently and with improved traceability (quanticate.com), and the CDISC ADaM Implementation Guide (ADaMIG, currently version 1.3) formalizes traceability as a foundational principle: any analysis result must be explainable by tracing back through the ADaM dataset to its SDTM source (cdisc.org).

CDISC's Define-XML standard, and specifically its Analysis Results Metadata (ARM) extension, is the piece that closes the loop between a specific CSR table or figure and the exact ADaM variables and analysis method that produced it. ARM establishes a standard format identifying key analyses, the analysis methods used, and the ADaM datasets and variables involved, and it explicitly supports regulatory review by linking results, documentation, and data together so a reviewer can verify how and why a given result was generated (cdisc.org). ICON's own technical writing on ARM makes the same point: it provides traceability for selected analysis results back to the data used as input, and when embedded in Define-XML it gives end-to-end traceability from data collection through to the final analytical conclusions reviewers see in the CSR (iconplc.com).

This is why "AI clinical study report automation" is really a data-lineage problem before it is a natural-language-generation problem. A model that drafts a safety narrative sentence like "12 (4.2%) of subjects in the treatment arm experienced a treatment-emergent adverse event of Grade 3 or higher" is only trustworthy if that number can be regenerated deterministically from the ADaM ADAE dataset, and automation should preserve traceability among reported analysis results, analysis datasets, tabulation datasets, and source data (fda.gov). ARM can document links between selected results and their analysis inputs. PMDA says ARM should preferably be included for specified key analyses and permits a separate PDF when inclusion in the dataset definition document is difficult (pmda.go.jp). Metadata repository vendors such as Formedix build their automation precisely on this principle, letting sponsors define SDTM and ADaM mappings, generate datasets, and validate studies against CDISC standards without manual coding (pharmaphorum.com), and Formedix has continued expanding specifically to meet growing demand for this kind of clinical trial automation (clinicalservicesjournal.com).

ICH E3 and eCTD Module 5: Turning Data Into a Regulatory Document

ICH E3 lays out a report structure spanning roughly sixteen major numbered sections: title page and synopsis, table of contents and abbreviations, ethics, investigators and administrative structure, an introduction and study objectives, the investigational plan, the study population, efficacy evaluation, safety evaluation, a discussion and overall conclusions section, and appendices (gmp-compliance.org). ICH's own E3 Questions

and Answers document exists specifically because sponsors kept interpreting the guideline's flexible language inconsistently, and it is periodically updated to clarify implementation questions raised over nearly three decades of use (database.ich.org).

Once a CSR is finished, it does not simply get attached as a PDF. eCTD Module 5 requires every study referenced anywhere in the submission to appear on a single comprehensive tabular listing in section 5.2, covering the study number, title, type, design, patient population, subject counts, dosing regimen, duration, and status, cross-referenced to its eCTD location (assyro.com). Section 5.3 then files the actual CSRs by category, from bioavailability and bioequivalence studies (5.3.1), through pharmacokinetic (5.3.3) and pharmacodynamic (5.3.4) reports, to the efficacy and safety studies in 5.3.5 that carry the approval argument, post-marketing experience reports (5.3.6), and case report forms or individual patient listings (5.3.7) (assyro.com). The FDA's eCTD Technical Conformance Guide sets out the file-format, granularity, and hyperlinking rules that make this navigable at agency scale (fda.gov), and the eCTD v4.0 comprehensive table of contents heading hierarchy is the living reference for exactly where each document type belongs (fda.gov).

Practically, this granularity requirement is a major source of submission defects. Reviewers repeatedly flag hyperlinks that point at a folder instead of an individual document, since hyperlinks can only be made to individual documents, not to section folders (allucent.com), and consolidating sections that should stay separate (for example, combining drug substance sections) means that any future update forces resubmission of the entire consolidated module rather than just the changed piece (allucent.com). FDA-focused regulatory operations guidance is correspondingly specific about hyperlinking hygiene: anchor text should describe what it links to rather than reading "click here," links must be verified functional before submission, and every cross-reference to a document in another module should itself be hyperlinked (pharmaregulatory.in). This is precisely the kind of rules-based, high-volume, low-creativity checking task where deterministic automation and narrow AI classifiers already do reliable work, well ahead of where generative narrative drafting has gotten to.

Where Generative AI Fits Today: Drafting, TLF Generation, and Consistency Checking

The most mature commercial application is AI-assisted first-draft generation of standardized document sections. Certara's CoAuthor product, built around what the company describes as a purpose-built biomedical GPT layered into Microsoft Word with eCTD templates, is marketed explicitly for drafting patient narratives, clinical study reports, protocols, synopses, and toxicology reports, with the company citing greater than ninety percent accuracy for automated analysis and summarization of tables, listings, and figures and efficiency gains of up to forty percent in drafting time (certara.com). Certara separately markets a Pinnacle 21 automation suite aimed at removing manual case report form creation, annotation, data entry, and validation from the clinical data management side of the pipeline (certara.com).

A second cluster of applications sits in regulatory intelligence and consistency checking rather than drafting per se. ArisGlobal's LifeSphere platform layers generative AI, agentic workflows, and a cognitive computing engine (branded NavaX) across the regulatory information management lifecycle, with a chatbot that lets regulatory teams ask what a health authority has recently required and get answers with links back to the specific source document and page for verification, rather than an unverifiable summary (arisglobal.com). In a separate

health-authority-correspondence metadata-extraction use case, ArisGlobal reports extracting 8–16 fields across a sample of 23 languages, with ninety percent accuracy and fifty percent faster processing (arisglobal.com). These vendor-reported figures concern metadata extraction, rather than chatbot-answer or drafting accuracy.

IQVIA and NVIDIA have taken the furthest step toward agentic, multi-step automation, announcing a set of orchestrator AI agents that NVIDIA's chief executive has described as the largest automation platform for clinical trials to date (clinicalresearchnewsonline.com). One of those agents, focused on clinical data review, is claimed to cut review duration from about seven weeks to as little as two through automated checks handled by specialized sub-agents doing speech-to-text transcription, clinical coding, structured data extraction, and summarization, with human experts retaining oversight of the orchestration (clinicalresearchnewsonline.com). IQVIA's broader public framing of AI in clinical operations emphasizes connecting disparate data streams, labs, claims, and EMR data, into insights while explicitly keeping humans in the loop for decision-making (iqvia.com).

Parexel has gone furthest in publishing task-level efficiency numbers specifically tied to CSR production. Its ParexelAI suite, launched under what the company calls a human-in-the-lead model, claims a thirty percent reduction in medical writing time specifically for clinical study reports, a twenty percent reduction in safety literature screening time, and, through a partner integration, up to a sixty percent reduction in IND and NDA authoring timelines, while explicitly positioning the AI as amplifying rather than replacing the judgment of its clinical and regulatory staff (hitconsultant.net). A contemporaneous industry report covered by Forbes cautions that these tools cannot fix clinical trials without the right people in place to supervise them, underscoring that the vendor efficiency claims assume continued expert review rather than a headcount-free pipeline (forbes.com).

On the biostatistics and data-management side, Cytel frames AI's near-term value in narrower, more mechanical terms: spotting outliers, inconsistencies, and missing or duplicate data during cleaning, flagging fraud patterns such as identical dosing times across every patient at a site, and translating plain-text edit-check specifications into programmed edit checks to shorten database build timelines, with medical writing and statistical programming automation described as an emerging rather than mature capability (cytel.com). This is a useful corrective to the more sweeping vendor marketing elsewhere in the sector: most of what is operational today is pattern detection and first-draft generation bounded by a human review step, not autonomous end-to-end CSR authorship.

Vendor Landscape, Regulatory Posture, and the Limits of the Current Toolset

Regulators have been explicit that they are watching this shift rather than blessing it wholesale. FDA's CDER has logged more than one hundred drug and biologic submissions in a single reported year that included some AI or machine learning component, spanning discovery through postmarket surveillance (fda.gov), and its January 2025 draft guidance, "Considerations for the Use of Artificial Intelligence to Support Regulatory Decision-Making for Drug and Biological Products," lays out a seven-step, risk-based credibility framework: define the question the model addresses, define its specific context of use, assess its risk level, build and

execute a credibility-establishment plan, document results and deviations, and determine adequacy for that specific use (goodwinlaw.com). That framework explicitly covers AI models used in the clinical and post-marketing phases where their output is meant to support a safety or efficacy conclusion, while carving out drug discovery and pure operational-efficiency uses that do not touch patient safety, drug quality, or study reliability (federalregister.gov). FDA defines the context of use as the specific role and scope of an AI model in addressing a question of interest; the draft does not endorse a particular AI approach or technique (fda.gov). The joint FDA–EMA principles likewise call for a well-defined context of use (fda.gov).

EMA's parallel reflection paper on AI across the medicinal product lifecycle, adopted after a 2023 draft consultation and developed jointly by EMA's Big Data Steering Group with its human and veterinary medicines committees, takes essentially the same stance: AI tools that draft, compile, translate, or review data in support of an authorization dossier are explicitly in scope, but a human-centric approach should guide all development and deployment of AI and machine learning, and developers whose models could plausibly move the benefit-risk balance of a medicine are told to seek early regulatory dialogue rather than deploy first and ask later (ema.europa.eu). The reflection paper's own scope statement covers everything from drug discovery and non-clinical development through clinical trials, manufacturing, and post-authorization use, including pharmacovigilance case management, which puts CSR-adjacent drafting activity squarely inside the paper's remit rather than in some ungoverned gray zone (ema.europa.eu).

The academic literature backs up the regulators' caution rather than the more optimistic vendor copy. A 2025 mini-review of broader clinical applications of generative AI found that AI systems translating imaging or structured data into narrative text show a persistent tendency to hallucinate clinical findings, describing factual grounding as a critical bottleneck that has to be solved before generative AI can be safely embedded in routine care, and it explicitly calls for robust validation, interpretability tooling, and governance frameworks as prerequisites rather than nice-to-haves (pmc.ncbi.nlm.nih.gov). A separate applied framework for generative-AI-assisted clinical data processing describes large language models functioning as a layer that helps design and deploy extract-transform-load pipelines for behavioral and clinical research data, a more modest, tool-assisted framing than "AI writes the CSR" (mdpi.com). Trade coverage aimed at medical writers and regulatory affairs professionals converges on the same message: generative AI is best deployed today for first drafts, summarization, and literature screening inside a workflow where a qualified human still owns the final sign-off on anything filed with a health authority (appliedclinicaltrialsonline.com).

Two structural facts keep the human-in-the-loop requirement from being merely cautious box-checking. First, ICH E3 itself is not a fill-in-the-blank template; it is a flexible guideline whose intent is a report that is complete, free from ambiguity, well organized, and easy to review, which means judgment calls about what to include, how to frame a discrepant result, and how to phrase a safety conclusion remain inherently interpretive rather than mechanical (ema.europa.eu). Second, the SDTM-to-ADaM-to-CSR chain is only as trustworthy as its weakest traceability link, and CDISC's own ADaM guidance treats traceability, the ability to explain any analysis result by tracing back through the dataset to its SDTM source, as a foundational, non-negotiable principle rather than a best practice (cdisc.org). An AI system that drafts plausible-sounding narrative text without being wired into that traceability chain is producing something that looks like a CSR passage without being one in the regulatory sense.

Implications and Future Directions

The near-term trajectory looks like incremental automation of well-bounded sub-tasks rather than a single leap to autonomous CSR generation. Certara's TLF accuracy benchmark concerns automated analysis and summarization of tables, listings, and figures; the vendor describes this as assistance with interpreting existing tables such as demographics, safety, and exposure tables (certara.com). Narrative drafting for lower-stakes sections, patient disposition summaries, demographic descriptions, protocol deviation listings, is a reasonable next tier, since these sections are more formulaic and less consequential for the ultimate efficacy or safety conclusion than the discussion and conclusions section that ICH E3 places at the center of the report (gmp-compliance.org).

Consistency checking, verifying that ADaM-derived results quoted in the narrative match their analysis data and that planning details match the protocol and statistical analysis plan, that every hyperlink in the Module 5 filing points to a document rather than a folder, that every study referenced in the body appears on the section 5.2 tabular listing, is arguably a better near-term automation target than drafting at all, because it is a verification task with a ground truth (the locked database and the existing filing structure) rather than a generation task with no single correct answer. The recurring, well-documented eCTD defects around hyperlink granularity and section consolidation are exactly the kind of rules-based error that a narrow, deterministic checker (AI-assisted or otherwise) can catch before submission far more cheaply than a health authority review cycle can (allucent.com).

Regulatory acceptance will likely keep pace with, and possibly lag, the FDA's and EMA's own credibility and risk frameworks rather than industry's product roadmaps. EMA already addresses generative-language-model errors in section 2.3.5 of its reflection paper: for medicinal product information documents, it recommends close human supervision and review to ensure generated text is factually and syntactically correct before submission (ema.europa.eu). That section concerns product information, so sponsors should consider its scope carefully when applying its recommendations to CSR drafting. Sponsors experimenting with generative narrative drafting for filed documents should expect to need bespoke validation packages, and EMA's advice to seek early scientific dialogue for any AI use that could plausibly move a benefit-risk assessment applies with particular force here (ema.europa.eu).

The vendor landscape itself is likely to consolidate around the data layer rather than the writing layer. Metadata-repository and automation platforms that already sit inside the SDTM/ADaM pipeline, Formedix, Certara's Pinnacle 21, and comparable tools, have a structural advantage over pure drafting assistants, because they can ground any generative output in the same validated, traceable dataset that a reviewer will eventually check against (pharmaphorum.com). Agentic, multi-tool orchestration of the kind IQVIA and NVIDIA have announced, where a data-review agent, a coding agent, and a summarization agent hand off to each other under human supervision, looks like the more probable near-term architecture for CSR-adjacent automation than a single end-to-end model asked to produce a finished report (clinicalresearchnewsonline.com).

Conclusion

AI-assisted CSR automation in 2026 is real, commercially available, and already producing measurable time savings on specific, well-bounded sub-tasks: first-draft table and narrative generation, regulatory-intelligence lookup, data-review triage, and consistency checking between datasets and text. It is not, and does not yet claim to be, a replacement for the ICH E3 structure, the SDTM-to-ADaM traceability chain, or the human medical writer, biostatistician, and regulatory reviewer who sign off on a submission. FDA's January 2025 framework offers nonbinding draft recommendations, not a universal Module 5 filing obligation (fda.gov). It concerns AI-generated information or data supporting regulatory decisions about drug safety, effectiveness, or quality, and excludes operational drafting that does not affect patient safety, drug quality, or study-result reliability (fda.gov). For sponsors, the more durable investment is probably in clean, well-governed SDTM and ADaM pipelines with documented traceability among source data, datasets, and reported analysis results (fda.gov).

Source links

Links cited in this article, in order of first appearance. Full addresses are provided for readers using extracted text.

1. https://www.cdisc.org/standards/foundational/sdtm
2. https://www.fda.gov/regulatory-information/search-fda-guidance-documents/e3-structure-and-content-clinical-study-reports
3. https://www.fda.gov/media/71610/download
4. https://hitconsultant.net/2026/05/18/parexel-launches-parexelai-clinical-development-suite
5. https://www.fda.gov/regulatory-information/search-fda-guidance-documents/considerations-use-artificial-intelligence-support-regulatory-decision-making-drug-and-biological
6. https://www.fda.gov/media/184830/download
7. https://www.ema.europa.eu/en/news/reflection-paper-use-artificial-intelligence-lifecycle-medicines
8. https://pmc.ncbi.nlm.nih.gov/articles/PMC12620437
9. https://www.gmp-compliance.org/guidelines/gmp-guideline/ich-e3-structure-and-content-of-clinical-study-reports
10. https://www.allucent.com/resources/blog/what-cdisc-and-what-are-cdisc-data-standards
11. https://www.fda.gov/media/82716/download
12. https://www.pmda.go.jp/files/000267939.pdf
13. https://www.cdisc.org/standards/foundational/adam
14. https://www.pmda.go.jp/files/000267935.pdf
15. https://www.ema.europa.eu/en/ich-e3-structure-content-clinical-study-reports-scientific-guideline
16. https://www.assyro.com/blog/ectd-module-5-clinical-guide
17. https://www.fda.gov/media/93569/download
18. https://www.allucent.com/resources/blog/ectd-regulatory-submissions-hints-and-common-mistakes-avoid
19. https://www.fda.gov/media/71271/download
20. https://www.cdisc.org/standards/foundational/sdtmig/sdtmig-v3-4
21. https://www.quanticate.com/blog/bid/90417/exploring-the-analysis-data-model-adam-datasets

22. https://www.cdisc.org/standards/foundational/adam/adamig-v1-3
23. https://www.cdisc.org/standards/foundational/define-xml/analysis-results-metadata-arm-v1-0-define-xml-v2-0
24. https://www.iconplc.com/insights/blog/2018/09/20/analysis-results-metadata
25. https://www.fda.gov/media/153632/download
26. https://pharmaphorum.com/partner-content/using-sdtm-adam-and-send-for-regulatory-submissions
27. https://www.clinicalservicesjournal.com/story/33974/formedix-announces-expansion-to-meet-clinical-trial-automation-demands
28. https://database.ich.org/sites/default/files/E3_Q&As_R1_Q&As.pdf
29. https://www.fda.gov/media/93818/download
30. https://www.fda.gov/media/179699/download
31. https://www.pharmaregulatory.in/module-5-clinical-study-reports-fda-study-placement-and-hyperlinking-best-practices
32. https://www.certara.com/coauthor
33. https://www.certara.com/blog/automation-in-clinical-trials-why-its-essential-for-success
34. https://www.arisglobal.com/blogs/employing-generative-ai-to-drive-connected-regulatory-intelligence
35. https://www.arisglobal.com/blogs/delivering-tangible-genai-use-cases-in-regulatory-accelerated-health-authority-interactions
36. https://www.clinicalresearchnewsonline.com/news/2025/06/11/iqvia-nvidia-build-ai-agents-for-clinical-research
37. https://www.iqvia.com/blogs/2024/02/transforming-clinical-trial-workflows-with-ai
38. https://www.forbes.com/sites/torconstantino/2025/05/16/ai-cant-fix-clinical-trials-without-the-right-people-new-parexel-report
39. https://cytel.com/perspectives/the-use-of-ai-in-clinical-trial-data-management
40. https://www.fda.gov/about-fda/center-drug-evaluation-and-research-cder/artificial-intelligence-drug-development
41. https://www.goodwinlaw.com/en/insights/publications/2025/01/alerts-lifesciences-aiml-fda-publishes-its-first-draft-guidance
42. https://www.federalregister.gov/documents/2025/01/07/2024-31542/considerations-for-the-use-of-artificial-intelligence-to-support-regulatory-decision-making-for-drug
43. https://www.fda.gov/media/189581/download
44. https://www.ema.europa.eu/en/use-artificial-intelligence-ai-medicinal-product-lifecycle-scientific-guideline
45. https://www.mdpi.com/2227-9709/13/4/48
46. https://www.appliedclinicaltrialsonline.com/view/medical-writing-regulatory-affairs-professionals-embrace-deploy-generative-ai
47. https://www.certara.com/wp-content/uploads/2024/11/2024-11-22-Fact-Sheet-CoAuthor-Interactive.pdf
48. https://www.ema.europa.eu/en/documents/scientific-guideline/reflection-paper-use-artificial-intelligence-ai-medicinal-product-lifecycle_en.pdf

THE PUBLISHER

About IntuitionLabs

IntuitionLabs is an AI consulting, custom software development and data engineering firm serving pharmaceutical, biotechnology, medical-device and other life-science organizations. We work with clinical, regulatory, medical-affairs, commercial, quality and IT teams to connect technology decisions with the work people need to accomplish.

AI consulting and adoption

Our AI enablement services cover readiness assessments, use-case selection, governance and policies, team workshops, adoption measurement and ongoing advisory support. We help organizations structure the information layer behind AI: source material, context, permissions and maintained knowledge that make generated answers useful and reviewable. Private LLM inference and hosted AI options support teams evaluating how to operate AI with appropriate control over their data and infrastructure.

Software, data and life-science workflows

IntuitionLabs develops custom software for pharma and biotech, integrates enterprise systems, and builds data engineering and business intelligence solutions. Areas of focus include AI agents, regulatory research, medical writing, medical affairs, CMC information, competitive intelligence and clinical-document workflows. Our eTMF intelligence work includes cross-system reconciliation and inspection-readiness support.

Enterprise platforms and regulated delivery

We provide Veeva services, application support, managed services, integrations and custom applications, alongside enterprise content work involving platforms such as Egnyte. For regulated workflows, our services include GxP enablement, computer-system validation and software development addressing 21 CFR Part 11 requirements. The applicable controls, validation responsibilities and acceptance criteria are defined for each engagement.

Work with IntuitionLabs

Explore AI enablement, pharma and biotech software development, data engineering and BI, and Veeva services. Contact IntuitionLabs to discuss your workflow, information sources and implementation needs.

IntuitionLabs publishes educational research to help life-science teams make informed technology decisions. Coverage of a product or organization does not imply a client relationship, endorsement or partnership.

Website: https://intuitionlabs.ai

This article is educational. Verify consequential medical, legal, financial, regulatory and technical decisions with appropriate professionals and the cited primary sources. Company services are described separately from the article's research findings.