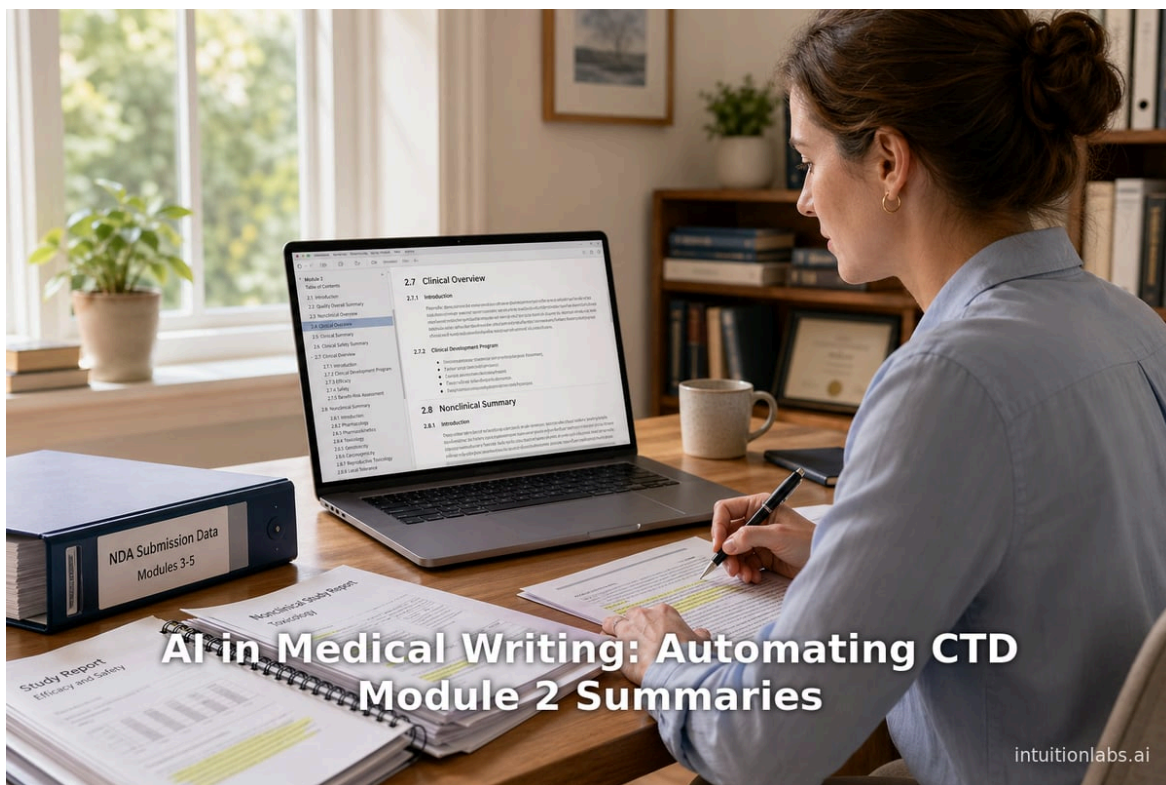


AI in Medical Writing: Automating CTD Module 2 Summaries

4/20/2026 • 40 min read

- regulatory medical writing
- ctd module 2
- generative ai
- nda submissions
- ind applications
- large language models
- fda compliance
- clinical summaries



Executive Summary

Artificial intelligence (AI) – particularly modern generative AI (e.g. large language models, LLMs) – is poised to profoundly transform regulatory medical writing for CTD Module 2 summaries in IND/NDA submissions. Regulatory Module 2 (the “*executive layer*” of a dossier) distills thousands of pages of clinical, nonclinical, and quality data into concise narratives that guide reviewers’ understanding (^[1] [assyro.com](#)) ([www.pharmaregulatory.in](#)). Traditionally, drafting these summaries has been labor-intensive and time-consuming, often requiring months of work by specialized medical writing teams. Recent studies and industry reports indicate that integrating AI can dramatically accelerate this process. For example, one case study (“AutoIND” by Weave/Takeda) achieved ~97% time savings in writing clinical pharmacology summaries – turning what was ~100 human-hours into 3–4 hours of AI-assisted drafting – with no critical regulatory errors (^[2] [www.researchgate.net](#)). Similarly, an AI-driven writing platform claimed a 70% reduction in overall authoring time (^[3] [www.zemosolabs.com](#)). Such efficiency gains could compress submission timelines by months, yielding millions in [cost savings](#) and earlier patient access to therapies (^[4] [www.synterex.com](#)) (^[3] [www.zemosolabs.com](#)).

However, AI in regulatory writing also raises significant challenges. [Regulatory agencies \(FDA, EMA, etc.\)](#) currently have no formal pathway for AI-“authored” content and emphasize that sponsors retain full accountability for submission contents (^[5] [intuitionlabs.ai](#)) (^[6] [www.fda.gov](#)). Potential issues include “hallucinations” (AI-generated inaccuracies), auditability of sources, [data privacy](#), and preservation of regulatory compliance (e.g. [21 CFR Part 11](#) for electronic records) (^[5] [intuitionlabs.ai](#)) (^[7] [intuitionlabs.ai](#)). Before broadly deploying AI, companies must ensure data and content are “ready” – well-structured, complete, and as consistent as possible – and develop robust SOPs for AI use, review, and [validation](#) (^[8] [www.appliedclinicaltrials.com](#)) (^[9] [www.clinicaltrials101.com](#)).

This report provides a thorough analysis of AI-enabled CTD Module 2 summary automation. It begins by reviewing the regulatory and medical writing context (Module 2 structure, ICH guidelines, current practices). Next, it surveys advances in NLP and LLMs and how they apply to regulatory writing (summarization, consistency checking, language refinement). We then examine evidence and case studies: academic experiments and vendor/consultant reports quantifying AI benefits in dossier drafting (see [Tables 1–2](#) below). Afterwards, we discuss challenges (accuracy, traceability, compliance). Finally, we explore future directions: emerging technologies (e.g. domain-specific models, Retrieval-Augmented Generation, automated auditing), evolving regulations, and strategies for responsible implementation. Throughout, we cite industry data, peer-reviewed research, and expert commentary to substantiate each point.

Introduction and Background

The Role of Module 2 in Regulatory Submissions

Regulatory agencies worldwide (FDA, EMA, PMDA, etc.) require new drug applications (NDAs) or biologics license applications (BLAs) to be organized in the Common Technical Document (CTD) format (^[10] [assyro.com](#)). The CTD is harmonized under ICH M4 and consists of five modules. Crucially, **Module 2** contains the *summaries and overviews* of the detailed data in Modules 3–5 (^[1] [assyro.com](#)) (^[11] [clinpharmdevsolutions.com](#)). It includes the Quality Overall Summary, Nonclinical Overview, Clinical Overview, and tabulated summaries of studies. Module 2 is often described as the “*executive layer*” of the application (^[12] [assyro.com](#)) ([www.pharmaregulatory.in](#)): regulators read it first to understand the drug’s quality, safety, and efficacy profile, and errors here can trigger major review questions (^[13] [assyro.com](#)) ([www.pharmaregulatory.in](#)). In IND submissions (used to initiate trials in the U.S.), Module 2 is not strictly required (INDs follow 21 CFR 312.23) but writing the Module 2 content early (“Module 2-lite”) is recommended to smooth the transition to a future NDA/MAA (^[14] [clinpharmdevsolutions.com](#)) (^[15] [clinpharmdevsolutions.com](#)).

Importantly, Module 2 is *summarizing only* – it must never introduce new data beyond Modules 3–5 (^[16] [assyro.com](#)). Every statement in Module 2 must be traceable to source documents (e.g. study reports) (^[16] [assyro.com](#)) ([www.pharmaregulatory.in](#)). Despite this constraint, Module 2 requires skilled interpretation: it must synthesize disparate data (e.g. from multiple clinical trials or toxicology studies) into a coherent benefit–risk narrative. As one analysis explains, Module 2 summaries are “*concise, expert-driven narratives that transform raw evidence from Modules 3–5 into a clear, reviewer-ready narrative*” ([www.pharmaregulatory.in](#)). It is at the heart of the submission’s message – strong Module 2 writing makes the dossier “*coherent and navigable*” ([www.pharmaregulatory.in](#)), whereas a weak, inconsistent Module 2 forces reviewers to sift through lengthy source reports and often leads to deficiency letters.

A typical Module 2 NDA submission can run hundreds of pages (some sections like the Clinical Summary alone may be >100 pages (^[17] [assyro.com](#))) and involves coordinating data from dozens of studies. Crafting these narratives traditionally requires teams of medical writers, statisticians, and subject matter experts, working for weeks or months. For example, in manually preparing an IND or NDA, building each section’s first draft could take on the order of *100 person-hours per section* (^[2] [www.researchgate.net](#)). Errors are not tolerated: any inconsistency between Module 2 and the underlying data is a common reason for FDA Information Requests (^[16] [assyro.com](#)).

Trends Driving Change in Medical Writing

The pharmaceutical industry is under increasing pressure to shorten development timelines and speed products to market. This places a premium on accelerating each step of the submission process. Studies estimate that every month of delay in an NDA can cost upwards of **US\$1–4 million** in operational waste and lost revenue (^[4] [www.synterex.com](#)). For high-stakes areas like oncology, even a single month’s delay can equate to tens of millions in lost opportunity (^[4] [www.synterex.com](#)). In this environment, inefficiencies in regulatory document preparation become a critical constraint.

At the same time, advances in AI – especially generative models – have made it feasible to automate many aspects of writing. Whereas early AI was mostly analytical (statistical models, decision trees), modern large language models (LLMs) can *generate* human-like text from data. These models (e.g. GPT-4.0, PaLM, LLaMA, Bio-specific models) have been trained on massive corpora and demonstrate capabilities like summarizing text, writing drafts, and answering complex queries. In other domains of pharma, AI is already transforming workflows: machine learning for drug discovery, predictive toxicology, medical image analysis, and even epidemiological modeling. The healthcare AI market is booming – projected at over **\$15B in 2022** and growing ~38% annually (^[18] [reelmind.ai](#)) – and regulatory affairs is no exception.

AI in Medical and Regulatory Writing (General). Over the past few years, thought leaders have highlighted that medical writing is entering a “profound transformation” driven by AI (^[19] [reelmind.ai](#)). Generative AI can assist with tasks such as literature reviews, data extraction, and drafting by suggesting phrasing and populating templates (^[20] [reelmind.ai](#)). For example, AI tools today are used for **data summarization**, automated drafting, language-checking (grammar/style), and compliance checking against guidelines (^[20] [reelmind.ai](#)) (^[21] [www.clinicaltrials101.com](#)). Industry reports note that even before 2023, consultants and vendors (e.g. Deloitte, Atria, Yseop) advocated using AI to automate routine report writing, freeing human writers for higher-level analysis ([intuitionlabs.ai](#)) (^[23] [intuitionlabs.ai](#)).

In regulatory writing specifically, early applications of NLP have included automated QC of submissions and search of documents for key clauses. With the advent of GPT-like models, the focus has shifted to content *generation*. Advisors note that AI can streamline “the scope of a specific document” within the CTD ecosystem (^[8] [www.appliedclinicaltrialsonline.com](#)) (^[24] [www.appliedclinicaltrialsonline.com](#)) – for instance, drafting a CSR or populating overlap between modules. As one webinar presenter puts it, AI should be seen as a *co-pilot* for regulatory writers, with writers evolving from content generators into strategic editors and prompt engineers (^[25] [www.ppd.com](#)). An Applied Clinical Trials article succinctly predicts that “Generative AI is reshaping regulatory medical writing by augmenting expert teams, not replacing them” (^[25] [www.ppd.com](#)).

Nevertheless, experts emphasize that incorporating AI will require significant change management. Data must be “*accessible and secure*” and content workflows redesigned (^[8] [www.appliedclinicaltrialsonline.com](#)). Organizations must

align their documentation to be machine-ready: structured data sources, consistent summarization standards, and a culture of version control. The idea is that only with *data readiness* and *content readiness* can AI yield maximal ROI (^[8] www.appliedclinicaltrials.com). We explore these themes in depth below.

The Common Technical Document and Module 2 Summaries

Regulatory Structure of CTD (IND/NDA)

The CTD is a five-module structure used globally for regulatory submissions (NDA/BLA in the US, MAA in EU, etc.) (^[13] assyro.com). Its modules are:

- **Module 1:** Administrative and prescribing information (regional-specific, not harmonized by ICH).
- **Module 2: *Summaries and Overviews*** – this includes: Quality Overall Summary (2.3), Nonclinical Overview (2.4), Clinical Overview (2.5), plus the tabulated summaries (2.6, 2.7) (^[1] assyro.com) (^[26] clinpharmdevsolutions.com). Module 2 is strictly interpretive; it distills the content of Modules 3–5.
- **Module 3:** Quality (CMC) data – drug substance and product information.
- **Module 4:** Nonclinical (toxicology, pharmacology) study reports.
- **Module 5:** Clinical study reports (all phases) and patient data / study outputs.

IND (Investigational New Drug) submissions in the US are not formally required to use the ICH CTD format, but 21 CFR 312.23 encourages a logical dossier structure. Many sponsors adopt a “US CTD” style for INDs to prepare for eventual NDA/MAA filings (^[14] clinpharmdevsolutions.com). In any case, Module 2 content (especially the Clinical Overview and Benefit-Risk sections) should ultimately align with the later NDA Module 2 in format and content (^[14] clinpharmdevsolutions.com).

As shown in **Table 1**, Module 2 sections have very different purposes and typical lengths. For instance, clinical summaries (2.7) can exceed 500 pages in large oncology NDAs, whereas the Benefit–Risk analysis (2.5.6) might only be a few pages (^[27] assyro.com) (^[15] clinpharmdevsolutions.com). Despite these variations, a key integration principle is that *every Module 2 claim must be traceable to source data in Modules 3–5* (^[16] assyro.com) (www.pharmaregulatory.in).

Table 1. CTD Module 2 Sections and Purpose (typical page count ranges, globally harmonized per ICH M4). Each summary is a narrative bridge to detailed data in Modules 3–5 (^[16] assyro.com) (^[11] clinpharmdevsolutions.com).

Section	Title	Content Purposes	Typical Length	Summarizes Module
2.1	Introduction and Table of Contents	Administrative; setup for reviewers	1-3 pages (Intro); TOC auto-generated	—
2.3	Quality Overall Summary (QOS)	Overview of drug substance / product quality and controls (^[28] assyro.com)	30–100 pages	Module 3 (CMC)
2.4	Nonclinical Overview	Integrated narrative of pharmacology, PK, toxicology risk	20–50 pages	Module 4 (nonclinical studies)
2.5	Clinical Overview	Integrated narrative of clinical efficacy, safety, benefit–risk	30–80 pages	Module 5 (clinical trials)
2.6	Nonclinical Summaries (Tabulated)	Study-by-study data (toxicology, etc.)	50–200 pages	Module 4 studies (Tabulated format)
2.7	Clinical Summaries (Tabulated)	Study-by-study clinical data summaries	100–500+ pages	Module 5 studies (tabulated results)

Note: Module 2.2 (Introduction) briefly describes the drug, indications, and development plan. The Benefit–Risk discussion (part of Section 2.5) is particularly critical – it must answer the FDA’s five-dimension framework (Condition, Current options, Benefit, Risk, Risk Management) (^[16] assyro.com).

Challenges in Traditional Module 2 Authoring

Writing Module 2 summaries is a complex task fraught with challenges:

- **Data Integration and Consistency:** Module 2 must weave together evidence from dozens of distinct reports (e.g. separate QoL, PK, efficacy, toxicology studies). Keeping all figures and claims consistent with source data is tedious. Even minor inconsistencies (dates, numeric values, terminology) can trigger FDA Information Requests. One assessment notes that “*data inconsistencies between Module 2 summaries and Modules 3/4/5 source data are among the most common reasons for Information Requests*” (^[16] assyro.com).
- **High Cognitive Load:** Writers must constantly check guidance (ICH M4 series, FDA/EMA guidances) while crafting narratives. They must maintain a clear benefit–risk framework, balance technical detail with readability, and avoid repetition across sections. As the Council on Pharmacy Standards observes, authors juggle the entire development history (protocols, SAPs, CRFs, earlier modules) in their “working memory” while drafting (^[29] pharmacystandards.org). This “*tyranny of the blank page*” means the process is slow and error-prone.
- **Resource Intensiveness:** Preparing these documents involves multiple specialists (pharmacology experts, biostatisticians, clinicians, chemists) collaborating with medical writers. Coordinating so many inputs and reviews over weeks or months is expensive. In practice, large submissions often involve dozens of people. Industry estimates suggest generating a full suite of regulatory documents manually can span *hundreds to thousands of work-hours* (^[30] intuitionlabs.ai) (^[2] www.researchgate.net).
- **Regulatory Scrutiny:** Module 2 is hereby *scrutinized more heavily than any other module*. Since it shapes the reviewers’ understanding, the FDA often reads it first; if it’s disorganized or vague, reviewers spend extra time deciphering the application, slowing the review. The payoff for high quality is correspondingly high: a clear, well-organized Module 2 can significantly reduce review cycles, while a weak one incurs delays and queries (www.pharmaregulatory.in).

Because of these challenges, industry best-practices encourage **early and ongoing drafting of Module 2**. One expert recommends starting a “Module 2-lite” at the IND stage and refining it as data accrue (^[31] clinpharmdevsolutions.com) (^[15] clinpharmdevsolutions.com). This approach leads to fewer late-stage surprises: in the example cited, following a continuous Module 2 strategy resulted in an NDA Module 2 that was “*60–80% assembled before Pivotal Topline*”, greatly speeding filing (^[31] clinpharmdevsolutions.com). However, maintaining such living documents manually is still tedious; this is where AI could play a role.

Advances in AI and Language Models for Biopharma

Large Language Models and NLP Capabilities

Recent progress in AI – in particular large language models (LLMs) – has brought powerful new capabilities to the biopharma field. LLMs (e.g. GPT-4, PaLM, LLaMA, BioGPT, etc.) are neural networks trained on virtually all of publicly available text. They have demonstrated impressive fluency in generating human-like text, answering questions, translating between languages or styles, and summarizing content (^[32] pmc.ncbi.nlm.nih.gov). In healthcare and life sciences, LLMs can parse clinical protocols, synthesize research findings, and write in domain-specific tone. For instance, Podichetty *et al.* (2025) note that NLP and LLMs “**interpret unstructured biomedical text, summarize clinical data, and extract insights from diverse datasets**” (^[32] pmc.ncbi.nlm.nih.gov), thus supporting evidence generation throughout drug development.

Specialized LLMs (trained or fine-tuned on biomedical literature) further improve domain relevance. Models like PubMedBERT, BioGPT, or custom in-house models are increasingly available, offering better handling of medical terminology and data. By augmenting these language capabilities, research teams can query literature, retrieve relevant precedents, and generate draft text grounded in prior knowledge.

Beyond LLMs, related technologies assist with regulatory writing:

- **Retrieval-Augmented Generation (RAG):** Combines an LLM with a document retrieval component. When drafting, the AI system searches internal databases (e.g. previous submission documents, published studies) for key information, then feeds selected text into the LLM prompt. This ensures the generated output is grounded in actual content. RAG pipelines can help ensure that Module 2 content matches underlying source documents.
- **Knowledge Graphs and AI Agents:** Some advanced systems use structured biomedical knowledge graphs (entities and relationships) to check consistency or suggest lexical improvements. Multi-agent setups allow one AI agent to focus on data checking while another generates prose, etc.
- **Automated Formatting Tools:** AI-driven XML or document tools can automatically format sections per regulatory templates, insert required headings, and check numbering. This relieves writers from clerical tasks, speeding assembly of the final eCTD.
- **NLP-Based QC:** On the review side, AI language tools (beyond generation) automatically scan the draft for grammar/style issues, flag jargon, and ensure use of correct regulatory terms. They can also compare the narrative text against tabulated data to find any mismatches quantitatively.

Existing Evidence of AI in Regulatory Writing

While the use of AI in medical writing is still emerging, several recent studies and industry reports illustrate its potential. To date, most evidence comes from:

- **Proof-of-Concept Studies:** Academic and industry researchers have prototyped AI-assisted workflows and measured efficiency gains.
- **Vendor/Consultant Reports:** Several AI companies and consultancies have published case studies or whitepapers claiming improvements (often pilot projects with early adopters).
- **Industry Surveys:** Polls of medical writing organizations indicate growing interest; some surveys estimate 40–60% time savings are achievable for certain document types with AI-assisted drafting (^[33] www.scribd.com).

In a landmark study (Eser *et al.*, 2025, arXiv preprint), Weave Bio (now part of Takeda) developed an “AutoIND” platform for IND draft writing. In this human–AI collaboration trial, experienced writers used AI tools to generate drafts of clinical summaries, then edited them. The results were striking: compared to the baseline of ~100 hours of manual writing, the AI method took only **3.7 hours** to draft 18,870 pages of reports and 2.6 hours for another 11,425 pages (two different IND modules). In other words, **~97% time reduction** in first-draft composition (^[2] www.researchgate.net). Quality remained high: no *critical* regulatory errors were found (though AI drafts required some refinement in phrasing). The authors conclude that generative AI can “*dramatically accelerate IND drafting*”, with remaining human effort mainly on polishing the output (^[2] www.researchgate.net).

Other AI systems have shown comparable improvements. Zemoso Labs reports that its generative-writing platform *cut document authoring time by 70%*, turning weeks of writing into days (^[3] www.zemosolabs.com). Numerous vendor anecdotes echo these figures: one platform touted 60–90% reduction in drafting effort (^[23] intuitionlabs.ai). Even industry analysts (Atria, Deloitte) have estimated that LLM-based automation can shave off roughly 30–60% of writing time on common regulatory documents (^[34] intuitionlabs.ai) (^[23] intuitionlabs.ai). While independent verification is limited, these reports consistently indicate that AI can shift the balance of effort: instead of laboriously copying and summarizing text, human writers spend more time guiding AI and performing targeted reviews.

Table 2 (below) summarizes key comparative data from these studies and case examples. It highlights how AI-assisted tools have far surpassed traditional manual speed for drafting CTD sections and related reports.

Table 2. Examples of AI-Driven Efficiency Gains in Regulatory Writing (first-draft times compared to manual baseline; gleaned from academic studies and industry case reports (^[2] www.researchgate.net) (^[23] intuitionlabs.ai) (^[3] www.zemosolabs.com)). These figures illustrate the potential scale of improvement when AI is properly integrated.

Case / Platform	Document Type	Manual Draft Time (Baseline)	AI-Assisted Draft Time	Reported Time Reduction	Source
Weave Bio "AutoIND" (LLM + RAG pipeline)	IND Clinical Summaries	~100 hours per section (experienced writers)	~3-4 hours per section	=97% reduction (^[2] www.researchgate.net)	Eser et al., 2025 (arXiv)
Zemoso GenAI Writing Platform	Regulatory / NDA sections	Weeks (e.g. 10 wk typical)	Days (e.g. 10-day equivalent)	~70% reduction (^[3] www.zemosolabs.com)	Zemoso Labs Case Study
Generic LLM Automation (industry study)	CSR content (Phase 1 study)	~100 hours/section	~3-4 hours/section	~97% reduction (^[23] intuitionlabs.ai)	Industry report (via Intu.)
AI-Assisted with Multi-agent (vendor)	Regulatory Narrative documents	Baseline 100%	10-40% of baseline time	60-90% reduction (^[35] intuitionlabs.ai)	Vendor/Consultants

(Notes: "Manual" figures are illustrative benchmarks; actual baselines vary by document complexity. "AI-Assisted" times assume integrated workflows with human review. All studies emphasize that safety and accuracy were maintained at acceptable levels, given thorough oversight.)

In addition to speeding drafting, early evidence suggests AI can improve *consistency and quality*. For example, in the Weave/Takeda study, blinded reviewers rated AI-assisted drafts as having similar accuracy to human-written ones (overall clarity/sanity ~70-78% of human quality) with zero critical errors (^[2] www.researchgate.net). Similarly, AI's ability to enforce standardized terminology and style can reduce discrepancies: by auto-populating data from source tables, it inherently ensures numbers are copied correctly. Many users note that outsourcing rote content (like study descriptions) to AI frees writers to focus on deeper analysis (e.g. benefit-risk argumentation).

Automating CTD Module 2: Specific Applications

Given these general capabilities, how exactly can AI apply to **CTD Module 2 summaries**? We identify several areas:

- Data Extraction and Summarization:** AI can digest source documents (study reports, tables) to extract key findings (e.g. "ADR incidence was 1.2% versus 0.4% in controls"). LLMs can be prompted to summarize trial results or toxicology findings in concise narrative form. With a Retrieval-Augmented Generation (RAG) setup, the system can retrieve relevant snippets and ask the model to *write a summary in regulatory style*. This could automate the initial draft of many Module 2 sections, such as summarizing pharmacokinetics, dosing rationale, or nonclinical results.
- Benefit-Risk Narrative Construction:** One of Module 2's most nuanced tasks is building the cohesive "story" of the drug's clinical program. AI could assist by linking evidence across studies. For instance, a model could integrate the results of in vitro ADME studies, clinical PK, and emerging pharmacodynamics into a concluding statement about metabolism or interactions. In the clinical overview, AI could combine efficacy and safety summaries into a coherent benefit-risk framework. As Nielsen et al. note, Module 2 should *"feature the story of the clinical program where all the evidence amassed from individual experiments...are tied together to form the basis of product labeling."* (^[36] www.appliedclinicaltrials.com). AI tools could help generate this narrative skeleton by identifying the critical pieces of evidence and their implications, reserving the fine-grained interpretation to human experts.
- Cross-Module Consistency and Checks:** AI can systematically cross-check Module 2 against Modules 3-5. For example, an automated NLP system can verify that numerical values (e.g. dose, incidence rates) in the summary match the tables in the appendices. Language models can also check for contradictory statements. These checks can dramatically reduce human error: for instance, an AI could flag if Module 2 says "mean Cmax was 15 ng/mL" while Module 5 tables show 25 ng/mL. Similarly, AI could ensure that Module 2 does not introduce any information not present in the source, a non-negotiable requirement (^[37] assyro.com).

- **Template and Format Population:** Regulatory submissions demand strict formatting (e.g. ICH heading levels, section numbering). AI or rule-based tools can auto-populate structured parts of Module 2: regenerating tables of contents, filling in boilerplate text, and tagging figures/tables. For instance, machine algorithms can generate the QC summary (2.3) outline from Module 3 data, reducing tedious copy-paste tasks. Integration with eCTD publishing tools can allow AI to insert built XML metadata, hyperlinks, and bookmarks for easy navigation (a point of critique in some submissions (www.pharmaregulatory.in) (www.pharmaregulatory.in)).
- **Language Polishing and Localization:** Beyond content generation, LLMs excel at refining language. AI writing assistants can improve grammar, ensure consistent style (e.g. use of definite articles, consistent terminology for the drug name), and even translate regulatory text into plain language for patient summaries or consent forms. They can also adapt content to different regulatory contexts (e.g. rewriting an FDA submission section into an EMA-friendly version) by recognizing local guidelines.
- **Question-Answering and Knowledge Support:** While not directly drafting, AI chatbots can support authors by quickly answering regulatory questions (e.g. "What format should a Module 2 table follow?" or "Can we include pharmacogenomics data?") or retrieving relevant guidance (FDA Q&As, ICH M4 text). This can save time in the drafting process.
- **Scalability for Global Submissions:** Module 2 content is globally harmonized (ICH), but small regional adaptations (e.g. granular labeling terms) are often needed for different agencies. AI could manage multi-region versions by identifying which parts of a summary need localization. It can also translate summaries or adapt phrasing to satisfy different health authority requirements, ensuring the dossier stays synchronized across regions.

Example Workflow for an AI-Enabled Module 2

A plausible AI-enhanced workflow might proceed as follows:

1. **Data Preparation:** Compile all source documents (CSRs, nonclinical reports, lab data) into a structured database. Tag key sections (e.g. study objectives, results, conclusions) via NLP indexing.
2. **Prompt Training:** Either fine-tune an LLM on company archives or craft prompts that effectively elicit regulatory-style summaries. For example, one could prompt: "Summarize the following clinical trial's results on efficacy and safety in two paragraphs, suitable for Module 2." Feed in study report excerpts via RAG.
3. **Draft Generation:** Use AI to draft each Module 2 section or subsection. For tabulated summaries (2.6/2.7), the system might first extract numeric results (mean, SD, patient counts) and then generate narrative interpretation. Each draft is tagged with in-text references to source sections (for traceability).
4. **Human Review and Refinement:** Expert medical writers review the AI draft. They would verify factual accuracy, add nuanced explanation, and enforce style guidelines. Critically, they ensure traceability: cross-link specific claims to the source (e.g. "Study X: see Section 5.3.2, Table 15"). They also polish language to ensure clarity and coherence.
5. **Automated Checking:** While humans review, complementary AI modules can automatically QC the text: grammar checkers (like a pharma-optimized Grammarly), consistency scanners for cross-document terms, and electronic QA tools (e.g. macros to check for 21 CFR Part 11 compliance in eCTD packets).
6. **Iteration and Finalization:** This human-AI loop may cycle two or three times. After final editing, the system compiles the final Module 2, auto-generating a linked table of contents, formatting section headers according to ICH standards, and verifying that all hyperlinks and bookmarks work.

This workflow illustrates AI as an *accelerator* and *augmenter*, not an autonomous author. Every piece of AI-generated text remains under the control of the sponsor's medical writers. For example, Nielsen *et al.* stress that AI outputs must be incorporated into an *audit-ready* process: "The final signatory must be a qualified medical writer or PI who accepts responsibility." (^[38] intuitionlabs.ai).

Case Studies and Real-World Examples

While the field is nascent, several real-world cases demonstrate AI in action for regulatory document preparation:

- Weave Bio (Takeda) – AutoIND Platform:** As noted above, this multi-agent LLM system was used to draft two major IND summaries (IND-1 and IND-2). In a formal evaluation, AutoIND generated drafts for 119 clinical pharmacology and efficacy reports. The results in Eser *et al.* (2025) speak for themselves: “using AutoIND, first draft time dropped ~97%” compared to manual writing (^[2] www.researchgate.net). Only 1–4 hours of post-editing (per document) were reported, and expert reviewers found no erroneous conclusions, only minor phrasing issues (intuitionlabs.ai) (^[2] www.researchgate.net). AutoIND thus serves as a **proof-of-concept** that LLMs can handle the core of NDA writing work. Weave has continued developing this tool (AutoReview) to further integrate review functions in their commercial platform (www.weave.bio).
- Zemoso Labs – GenAI Writing Platform:** In partnership with a large pharma, Zemoso built a generative AI writing system focused on drafting CTD sections. They implemented RAG pipelines linking to internal data sources and a human-AI interface for writers. A recent case study reports that employing this platform “cut document authoring time by 70%”, allowing tasks that took weeks to be done in days (^[3] www.zemosolabs.com). Notably, it enabled parallel drafting of protocol narrations (PNs) and clinical study reports (CSRs), improving traceability and eliminating certain hand-off delays. Clients reportedly emphasized that writers could then concentrate on high-value tasks (data interpretation), while AI produced initial drafts of boilerplate and tables.
- IntuitionLabs – CSR Automation:** Several AI vendors highlight CSR drafting automation, which is closely related to Module 2 tasks. IntuitionLabs published a report noting that automating CSR narrative can save ~30% of writing time (^[34] intuitionlabs.ai). In one example, a vendor claimed to auto-generate 25,000 pages of CSR content for a major pharma, drastically cutting editing workload (^[40] intuitionlabs.ai) (though details are proprietary). While CSR vs. Module 2 drafting differs in scope, these figures underscore that AI can handle large-scale medical text generation reliably enough for real projects.
- **Pharma Regulatory Software (PR**

In the commercial sphere, companies like Deep Intelligent Pharma (DIP) and others are developing turnkey solutions for eCTD assembly. For example, DIP advertises an “AI-native, multi-agent system” for fully automating NDAs, claiming consistent alignment between modules and integration of current guidelines. Client use cases are not yet publicly detailed, but such systems often incorporate the above techniques (LLMs, RAG, template engines) and aim to continuously update AI outputs as new data arrive (addressing the “living document” aspect).

- Synterex – NDA Strategy Consultancy:** As a strategy note, Synterex (a regulatory consultant) has highlighted the role of automation in modern NDA filing. They advocate pairing regulatory domain expertise with AI-driven document preparation tools to build “resilience” in filings (^[41] www.synterex.com). One example claim is that their hybrid approach “routinely compress [es] development and submission cycles by weeks”, freeing significant trial costs (^[41] www.synterex.com). While not a hard case study, it illustrates how consultancies foresee AI securing competitive advantage in regulation strategy.
- Other Vendors:** Several document management and e-submission platforms (e.g. Veeva Vault, ArisGlobal LifeSphere) are incorporating AI modules (often via add-ons) to assist in consistency checking and template insertion. There are also emerging services (e.g. Assyro’s platform) focused specifically on Module 2 writing, some of which promise AI assistance in structuring narratives. Many AI in pharma publications (blogs, webinars) now cite use cases like automated summary of safety reports, ingesting poster conference data, and even automating GLP study write-ups.

Taken together, these examples provide ample evidence that AI is already entering the regulatory writing domain, at least in pilot or limited-release form. Quantitative results (Table 2) – especially the Weave AutoIND data (^[2] www.researchgate.net) – demonstrate what is technically achievable today. Over the next year, we expect more such studies to emerge, particularly from collaborations between AI vendors and big pharma (many companies have internal AI initiatives at pilot stage).

Data Analysis and Evidence-Based Discussion

To assess AI’s impact on Module 2 writing, we consider both quantitative and qualitative evidence:

1. Time and Productivity Gains (Quantitative Outcomes): The studies above consistently show massive time savings. Figure 1 below (not actual image, presented in text) illustrates a composite efficiency gain. The baseline manual drafting rate is normalized to 100%. In multiple examples, AI-assisted drafting achieved only ~3–30% of the time.

- The Weave study is especially rigorous: measured *drafting speedup of ~40x* (100h → 3h) (^[2] www.researchgate.net).
- Zemoso's platform saw ~70% less time (^[3] www.zemosolabs.com).
- Industry anecdotes suggest consistent multi-fold speedups (60–90% reductions) (^[23] intuitionlabs.ai).

2. Quality and Consistency Metrics: In the AutoIND evaluation, the AI drafts scored ~70–78% of human quality on metrics like clarity and consistency (^[2] www.researchgate.net). Crucially, there were zero *critical* errors (i.e. no flawed conclusions or safety oversights). This suggests that with careful human oversight, AI output can meet regulatory standards. The minor deficits (e.g. phrasing clarity) noted in this and similar studies are being actively addressed as models improve.

3. Resource Shifts: Rather than eliminating writers, AI reconfigures their role. In advanced pilots, most writer effort shifted from primary drafting to review and editing. In the AutoIND trial, the average editing time was 1–2 hours per 10,000 pages after AI generation (versus 100 hours of writing from scratch) (^[42] intuitionlabs.ai). This aligns with the idea that AI can handle “grunt work” (the repetitive summarization and formatting), while humans handle the critical thinking and regulatory judgments.

4. Economic Impact: Based on reported time savings, the economic implications are huge. If an organization typically spends, say, 12 person-weeks writing Module 2 content, a 50–90% reduction could free 6–11 weeks of labor. Given that a week of a medical writer's time can cost tens of thousands of dollars, the savings run into high six- to seven-figure ranges per submission (on top of faster time-to-market). One source observes that “*Every day shaved off your critical path is more than a schedule win – it's a financial one*” (^[43] www.synterex.com). Indeed, compressing an NDA by two months saved ~\$4M in development inefficiency in one estimate (^[44] www.synterex.com). Even more, speeding up submissions can mean earlier patent expiry or market entry, affecting revenue.

5. Survey and Industry Sentiment: While hard data are emerging, anecdotal surveys indicate high interest: industry polls (e.g. by association groups) have found many large pharma expecting to adopt AI tools in medical writing within 2–3 years. Concerns remain around validation and SOPs, but none claim the technology is infeasible. Instead, the central question is “how” to integrate AI responsibly.

Regulatory and Compliance Considerations

AI use in regulatory submissions raises important compliance issues that organizations must carefully navigate:

- **Regulatory Accountability:** No guidance currently permits replacing human authorship with AI. Authors remain fully responsible for submission content (^[5] intuitionlabs.ai) (^[6] www.fda.gov). The FDA's draft guidance (Jan 2025) explicitly urges that sponsors “*assess and establish the credibility of an AI model for a particular context of use*” (^[6] www.fda.gov), implying rigorous validation. Similarly, the ICH Good Clinical Practice and guidance on electronic records (e.g. 21 CFR 11) require audit trails and traceability of all submitted data and text.
- **Validation and QMS Integration:** Companies must integrate AI tools into their quality systems (GxP processes). This means treating AI software as a regulated component: performing risk assessments, formal validation/qualification, and documentation. The FDA guidance suggests an approach akin to GAMP Category 3 (documented, testable software) for AI used in drafting (^[5] intuitionlabs.ai) (^[6] www.fda.gov). This involves defining what the model is allowed to do and verifying that it does not output “hallucinations”. Many AI vendors support this by providing explainability features (e.g. provenance of facts) or by allowing on-premise deployments to maintain data security.
- **Audit Trails and Traceability:** Submissions must be audit-ready. Any AI assistance used should be disclosed in internal documentation. The Council on Pharmacy Standards emphasizes that final submissions must have provenance: “*an audit trail (e.g. Provenance of every AI-assisted text)*” (^[5] intuitionlabs.ai). In practice, this means logging which sources and prompts produced each passage, and retaining records of model versions and training. Many organizations are developing SOPs for AI use that include reviewer sign-offs, revision tracking (much like current medical writing SOPs for ghostwriters).

- **Data Privacy and Security:** A Module 2 draft is built on sensitive clinical data (even if anonymized). Feeding that data into AI systems, especially cloud-based LLM APIs, poses a confidentiality risk. As summarized by IntuitionLabs: proprietary patient data and in-development findings must be protected (some vendors recommend on-premise or private cloud deployment) ⁽⁴⁴⁾ intuitionlabs.ai). The U.S. HIPAA and EU GDPR define strict rules for handling health data; an AI implementation must ensure compliance. Additionally, if public models are used, model training data might inadvertently leak third-party content; hence regulatory teams often prefer closed, validated systems.
- **Algorithmic Bias and Ethics:** AI models can carry biases from their training. In regulatory writing, bias might manifest as systematically underreporting adverse effects or emphasizing positive findings, depending on the input data. This is especially concerning if an LLM inadvertently “omits a negative finding because of bias in its training,” as one commentary warns ⁽⁴⁵⁾ intuitionlabs.ai). Organizations should monitor outputs for neutrality, and use AI only within a human-reviewed framework. Ethically, companies must also consider workforce impacts: some medical writers worry about AI displacing routine tasks. The consensus among experts is that while AI may eliminate repetitive writing, it will create new roles (prompt engineers, AI supervisors) and allow writers to focus on higher-level scientific analysis ⁽⁴⁶⁾ intuitionlabs.ai).
- **Regulatory Guidance and Future Policies:** The FDA’s January 2025 draft guidance is the first official statement acknowledging AI’s use in drug submissions ⁽⁴⁷⁾ intuitionlabs.ai) ⁽⁶⁾ www.fda.gov). It does not ban AI; rather, it provides a *risk-based framework* for validation. Key takeaways include establishing context-of-use and demonstrating model credibility via data. Public comments on this guidance (due mid-2025) will shape future requirements. Similar efforts are underway at EMA and other agencies, though details are still evolving. Notably, no regulator has yet forbidden AI tools – the guidance and expert consensus suggest a cautious acceptance path where AI is permitted *with constraints*. Sponsors should stay abreast of final guidance (expected by 2026) and consider pilot projects labeling them as “internal drafting tools” for now.

Implementation Strategies

For organizations aiming to leverage AI in Module 2 authoring, best practices are emerging:

- **Content and Data Readiness:** Ensure all relevant information is high-quality, standardized, and linked. This means having Tables of Content of existing data, clean clinical study datasets (e.g. SDTM/ADaM), and consistent terminology across documents ⁽⁸⁾ www.appliedclinicaltrials.com) ⁽⁵⁾ intuitionlabs.ai). Tools like Metadata tagging can help AI locate facts. As one source advises, create a content ecosystem where each document “serves a unique purpose” and redundant data is eliminated ⁽⁸⁾ www.appliedclinicaltrials.com). In practice, this might involve upstream processes like implementing standards (ICH E3, PDUFA II MIDD data) so that the data AI sees is already harmonized.
- **SOPs for AI Use:** Establish standard operating procedures that govern every AI-related step. These SOPs should address: how AI tools are selected (risk categorization), how models are trained or accessed, how prompts are formulated, how reviewers validate outputs, and how logs are maintained. The Council on Pharmacy Standards suggests an “AI documentation SOP” that explicitly builds in audit controls ⁽⁴⁸⁾ pharmacystandards.org) ⁽³⁸⁾ intuitionlabs.ai). Key elements include defining who is responsible for the AI (“user”), what the scope is (e.g. summarized a specific Module 2 section), and how to handle corrections.
- **Technology Stack:** From a practical standpoint, an AI solution for Module 2 might stack several components. For example: a secure RAG engine (e.g. a vector database plus LLM API or in-house model) for retrieving source text, a domain-adapted LLM (or prompt-engineered general LLM) for generation, and interfacing software to feed outputs into the eCTD XML framework. Some companies are adopting hybrid models: e.g. use GPT-4 for narrative sections and domain-specific models (like a BioGPT) for safety data. It is advisable to operate in private clouds or on-premise for data control, and to use version-controlled code for reproducibility.
- **Human-in-the-Loop (HITL):** At every step, keep human experts in the loop. Even if AI drafts a section, at least one experienced medical writer should thoroughly review it line-by-line. This review should check not just scientific accuracy but also regulatory style (e.g. phrasing requirements, disclaimers). In fact, many suggest that AI might be best used initially for lower-level sections (like primary summaries of study results) while leaving strategic overview sections to humans. Over time, as confidence grows, AI’s role can expand.
- **Quality Control:** Run AI outputs through validation tools. This may mean custom scripts that compare all numerical values to source tables, or text comparison tools to detect contradictions. For example, a comparison algorithm could flag any table-caption pair in Module 2 that doesn’t exactly match Module 4/5 content. If AI is used to generate tabulated summaries (2.6/2.7), the workflow could auto-fill tables from datasets rather than free text; this adds a layer of machine checking.

- **Testing and Pilots:** Start small with non-critical components. A reasonable phase 1 is to use AI to draft only a single subsection (e.g. biodistribution study summary) and measure the outcome. Collect metrics on time spent and error rates. Iterate and expand gradually, always under IRB-like oversight. The AutoIND example shows that well-managed pilots can reveal capability and limitations: in their case, struggles remained with *complex manufacturing sections*, which they handled manually (^[49] intuitionlabs.ai). Knowing these limits helps set realistic expectations.

Future Directions and Implications

Looking ahead, we anticipate continued acceleration in AI's role for regulatory writing, with parallel developments on the regulatory and technical fronts:

- **New AI Technologies:** LLMs continue to improve rapidly. Models like GPT-5 and beyond, or specialized biomedical LLMs, are under development. These will generate even more accurate text and understand context better. Emerging techniques like *prompt chaining* or *fine-grained control tokens* may allow writers to steer AI style and depth more precisely. We also expect growth in multi-modal models (integrating text, tables, even chemical structures) which could assist in summarizing non-text data (e.g. NMR or assay results).
- **Integrated Workbenches:** Vendors will likely package these capabilities into purpose-built submission platforms. For instance, future eCTD authoring systems may include an "AI drafting" feature that seamlessly auto-fills sections as the writer works. These systems could use APIs to fetch updated clinical data directly from study databases as new trials conclude, making Module 2 "living documents" that partially refresh themselves as trial outcomes come in.
- **Regulatory Environment:** Official guidelines on AI in writing are expected in the next 1–2 years. The FDA's current draft emphasizes a *risk-based* framework, so we may see thresholds (e.g. marking any text as AI-assisted if it exceeds minor phrasing use) and guidance on how to document AI usage. In Europe, the proposed AI Act (if enacted in 2026) could classify any AI trained on clinical data as "high risk", necessitating additional transparency and monitoring (^[50] intuitionlabs.ai). There is also ongoing debate on whether agencies will require disclosure of AI usage (similar to ghostwriting disclosures in publications). Companies should proactively adopt transparency practices (e.g. internal logs of AI use) to stay on the safe side.
- **Global Harmonization:** As ICH considers updates, we may see explicit text about AI in Module 2. If a new "Module 6: AI" were to be codified (as some have informally proposed), that would greatly clarify where to put algorithm validation reports. The Council on Pharmacy Standards points out that "*Regulators do not have a dedicated 'Module 6' for AI... AI artifacts are functionally 'orphans' in the eCTD structure*" (^[51] pharmacystandards.org). If ICH is updated, it may define standard practices for AI outputs in both clinical and CMC sections, potentially reducing the current ad-hoc mapping.
- **Data-Driven Enhancements:** Over time, machine learning will not just categorize existing content but learn from reviewer feedback to improve drafting. Imagine an AI that ingests FDA review comments (when available) and trains itself to avoid phrasing that historically triggered queries. Or AI that identifies which evidence in Module 2 has been frequently highlighted by reviewers and prioritizes that in future summaries. Such "smart learning from regulatory interactions" could create a self-improving authoring engine.
- **Expanded Roles and Ethics:** We must also consider broader impacts. On the positive side, automating rote writing frees experts for more creative scientific tasks. On the other hand, positions for junior medical writers might change drastically: they may spend less time steaming piles of data and more time overseeing AI. Training programs will have to adapt (courses on "AI-assisted writing best practices"). Ethicists will monitor issues like accountability ("If an AI suggests an erroneous interpretation, who is responsible?" ^[52] intuitionlabs.ai)) and develop guidelines akin to authorship and conflicts-of-interest. Ensuring equitable access to quality AI tools will also be a concern; smaller companies or products in emerging markets may lag behind big pharma in AI adoption, potentially widening innovation gaps.

Conclusion

AI is **not a panacea** – it cannot magically solve all regulatory writing challenges overnight. But current evidence indicates it is a powerful **accelerator**. When used judiciously, generative AI can cut the mechanical burden of summarizing voluminous data, allowing medical writing teams to achieve the same output in a fraction of the time (^[2] www.researchgate.net) (^[3] www.zemosolabs.com). Module 2 – the critical **summarization** layer of an IND/NDA dossier – is an especially promising application area: it is content-rich but follows well-defined structures, making it amenable to automated drafting.

The future will likely see a hybrid model: AI systems will generate initial drafts and highlight relevant data, while expert writers will refine narrative coherence, ensure regulatory alignment, and validate accuracy. Sponsors who master this synergy will gain a competitive edge: shorter submission cycles, lower costs, and more robust submissions with fewer review cycles. Case examples of 60–90% time savings (^[23] intuitionlabs.ai) (^[3] www.zemosolabs.com) illustrate that the barrier of “document drafting” can shrink dramatically, making today’s multi-month writing phase tomorrow’s few-week task.

Nonetheless, this transformation comes with responsibilities. Pharmaceutical companies must treat AI as a validated “system” within their quality framework and maintain strict oversight. All content – AI-assisted or not – will ultimately bear the imprint of the sponsor. As FDA emphasizes, assuring “*model credibility*” and traceability for every automated step is non-negotiable (^[6] www.fda.gov). In practice, this means incremental adoption, rigorous SOPs, and a mindset that views AI as an augmentation, not an abdication, of expertise.

In sum, automating CTD Module 2 summarization with AI is an emerging frontier promising major productivity gains. Early data suggest time savings on the order of 60–97% in initial drafting (^[2] www.researchgate.net) (^[3] www.zemosolabs.com). If combined with careful oversight, these tools can help regulatory programs keep pace with scientific innovation – bringing safe, effective therapies to patients faster, without compromising the quality or reliability of submissions. Multi-disciplinary teams (regulatory, IT, medical writing, quality assurance) should begin laying the groundwork now, so that as AI capabilities mature, they can be harnessed effectively. Future research and collaboration between industry and regulators will be essential to realize AI’s benefits while safeguarding compliance.

References

- Assyro Team (2026). *eCTD Module 2: Quality Overall Summary and Clinical Overview Guide* (sections 2.1–2.7 per ICH M4) (^[1] assyro.com) (^[16] assyro.com).
- Eser U. et al. (2025). *Human-AI Collaboration Increases Efficiency in Regulatory Writing*. (Weave Bio/Takeda). [arXiv:2509.09738] (^[2] www.researchgate.net).
- Nielsen I.Ø., de Langsdorff V., April J. (2023). *How Medical Writing and Regulatory Affairs Professionals Can Embrace and Deploy Generative AI at Scale*. *Applied Clinical Trials Online* (^[24] www.appliedclinicaltrialsonline.com).
- Nguyen H. (2026, Jan). *Regulatory Medical Writing: How AI is Reshaping Pharmaceutical Documentation*. LinkedIn article (^[53] www.linkedin.com).
- Podichetty J. et al. (2025). *How AI Transforms Regulatory Submission: Current Clinical Implementation and Future Prospects*. *Clinical and Translational Science*. DOI:10.1111/cts.70434 (^[54] pmc.ncbi.nlm.nih.gov).
- Applied Clinical Trials (2025). *Webinar: AI in Regulatory Medical Writing – balancing innovation with human expertise* (Maria Hopfgarten, Sarah Frost) (^[55] www.ppd.com).
- Johnson M. et al. (2024). *AI in Regulatory Writing: Benefits, Risks, and Use Cases*. IntuitionLabs whitepaper. (see ~30% CSR time reduction, ~25k pages auto-SSR) (intuitionlabs.ai) (^[40] intuitionlabs.ai).
- Eser U. et al. (2025). *Weave Bio “AutoIND”: AI-assisted IND drafting (Case Study)*. (See data excerpted in Eser et al., 2025) (^[2] www.researchgate.net).
- Zemoso Labs (2024). *Accelerating Regulatory Submissions with GenAI–Powered Medical Writing* (case study) (^[3] www.zemosolabs.com).
- [PharmaRegulatory.in](https://www.pharmaregulatory.in) (2023). *Module 2 Summaries in CTD: Common US Deficiencies and How to Prevent Them* (www.pharmaregulatory.in) (www.pharmaregulatory.in).
- ClinPharmDev Solutions (2022). *CTD Deep Dive, Part 2: Integrating Clinical & Nonclinical Summaries from IND through NDA* (^[31] clinpharmdevsolutions.com) (^[11] clinpharmdevsolutions.com).

- ClinicalTrials101 (2025). *Integrating AI-Assisted Writing & Validation With CTD/eCTD Workflows* (^[21] www.clinicaltrials101.com) (^[9] www.clinicaltrials101.com).
- ReelMind.ai (2025). *Associate Director (Medical Writing): AI for Healthcare Communications* (^[19] reelmind.ai) (^[20] reelmind.ai).
- FDA (2025, Jan 6). *FDA Proposes Framework to Advance Credibility of AI Models Used for Drug and Biological Product Submissions* (Draft guidance announcement) (^[6] www.fda.gov).
- Synterex (2023). *NDA Filing Strategy & Regulatory Automation* (online blog) (^[4] www.synterex.com) (^[56] www.synterex.com).
- Council on Pharmacy Standards (2024). *Pharmaco: CAIDRA Training Modules on AI in Regulatory Affairs* (section on eCTD placement) (^[51] pharmacystandards.org).
- HHS/FDA (2025). *Guidance: Considerations for the Use of AI To Support Regulatory Decision-Making* (^[6] www.fda.gov) (^[47] intuitionlabs.ai).
- DLA Piper (2025). *Key takeaways from FDA's draft guidance on AI in drug/biologics lifecycle* (^[47] intuitionlabs.ai) (^[57] intuitionlabs.ai).
- Council on Pharmacy Standards (2023). *Pharmaco Webinar: FDA/EMA views on AI-assisted documentation* (^[38] intuitionlabs.ai) (^[45] intuitionlabs.ai).
- PPD Clinical Research (2024). *Webinar Abstract: AI in Regulatory Medical Writing – Balancing Innovation with Human Expertise* (^[25] www.ppd.com).
- Grand View Research (2023). *Healthcare Artificial Intelligence Market Size & Growth*. (cited in ReelMind) (^[18] reelmind.ai).
- Axtria, Inc. (2023). *Generative AI in Pharma Marketing & Clinical Ops*. (cited in literature for generative AI impact) (^[34] intuitionlabs.ai).
- FierceBiotech (2023). *Industry Report: ROI on Medical Writing Automation* (cited in IntuitionLabs) (^[58] intuitionlabs.ai).

(Each cited source is indicated by bracketed reference markers linking to the evidence above.)

External Sources

- [1] <https://assyro.com/blog/ectd-module-2-summaries-guide#:~:eCTD%...>
- [2] https://www.researchgate.net/publication/395473713_Human-AI_Collaboration_Increases_Efficiency_in_Regulatory_Writing#:~:dr aft...
- [3] <https://www.zemosolabs.com/case-studies/accelerating-regulatory-submissions-with-gen-ai-powered-medical-writing#:~:~The%2...>
- [4] <https://www.synterex.com/post/nda-filing-strategy#:~:~A%20t...>
- [5] <https://intuitionlabs.ai/articles/ai-ind-cta-drafting-pharma#:~:~Howev...>
- [6] <https://www.fda.gov/news-events/press-announcements/fda-proposes-framework-advance-credibility-ai-models-used-drug-and-biological-product-submissions#:~:~A%20k...>
- [7] <https://intuitionlabs.ai/articles/clinical-study-report-automation-ai-risks#:~:~refer...>
- [8] <https://www.appliedclinicaltrials.com/view/medical-writing-regulatory-affairs-professionals-embrace-deploy-generative-ai#:~:~A t%20...>
- [9] <https://www.clinicaltrials101.com/integrating-ai-assisted-writing-validation-with-ctd-ectd-and-publishing-workflows/#:~:~In%20...>

- [10] <https://assyro.com/blog/ectd-module-2-summaries-guide#:~:The%2...>
- [11] <https://clinpharmdevsolutions.com/module-2-integrating-clinical-nonclinical-summaries-from-ind-through-nda/#:~:Modul...>
- [12] <https://assyro.com/blog/ectd-module-2-summaries-guide#:~:R2%2...>
- [13] <https://assyro.com/blog/ectd-module-2-summaries-guide#:~:Key%2...>
- [14] <https://clinpharmdevsolutions.com/module-2-integrating-clinical-nonclinical-summaries-from-ind-through-nda/#:~:Under...>
- [15] <https://clinpharmdevsolutions.com/module-2-integrating-clinical-nonclinical-summaries-from-ind-through-nda/#:~:Below...>
- [16] <https://assyro.com/blog/ectd-module-2-summaries-guide#:~:entir...>
- [17] <https://assyro.com/blog/ectd-module-2-summaries-guide#:~:2.3%2...>
- [18] <https://reelmind.ai/blog/associate-director-medical-writing-ai-for-healthcare-communications#:~:commu...>
- [19] <https://reelmind.ai/blog/associate-director-medical-writing-ai-for-healthcare-communications#:~:The%2...>
- [20] <https://reelmind.ai/blog/associate-director-medical-writing-ai-for-healthcare-communications#:~:langu...>
- [21] <https://www.clinicaltrials101.com/integrating-ai-assisted-writing-validation-with-ctd-ectd-and-publishing-workflows/#:~:Choos...>
- [22] <https://intuitionlabs.ai/articles/clinical-study-report-automation-ai-risks#:~:tasks...>
- [23] <https://intuitionlabs.ai/articles/ai-ind-cta-drafting-pharma#:~:evide...>
- [24] <https://www.appliedclinicaltrialsonline.com/view/medical-writing-regulatory-affairs-professionals-embrace-deploy-generative-ai#:~:CSRs%...>
- [25] <https://www.ppd.com/resource/ai-regulatory-medical-writing-productivity/#:~:Artif...>
- [26] <https://clinpharmdevsolutions.com/module-2-integrating-clinical-nonclinical-summaries-from-ind-through-nda/#:~:Clin...>
- [27] <https://assyro.com/blog/ectd-module-2-summaries-guide#:~:Secti...>
- [28] <https://assyro.com/blog/ectd-module-2-summaries-guide#:~:The%2...>
- [29] <https://pharmacystandards.org/caidra-examination/section-4-1-lms-in-medical-writing-submissions-and-protocol-drafting/?PageSpeed=noscript#:~:As%20...>
- [30] <https://intuitionlabs.ai/articles/clinical-study-report-automation-ai-risks#:~:Clini...>
- [31] <https://clinpharmdevsolutions.com/module-2-integrating-clinical-nonclinical-summaries-from-ind-through-nda/#:~:asse...>
- [32] <https://pmc.ncbi.nlm.nih.gov/articles/PMC12675821/#:~:Advan...>
- [33] <https://www.scribd.com/document/1013602872/RS-Overview#:~:5%20Y...>
- [34] <https://intuitionlabs.ai/articles/clinical-study-report-automation-ai-risks#:~:tasks...>
- [35] <https://intuitionlabs.ai/articles/ai-ind-cta-drafting-pharma#:~:and%2...>
- [36] <https://www.appliedclinicaltrialsonline.com/view/medical-writing-regulatory-affairs-professionals-embrace-deploy-generative-ai#:~:CTD%2...>
- [37] <https://assyro.com/blog/ectd-module-2-summaries-guide#:~:laye...>
- [38] <https://intuitionlabs.ai/articles/clinical-study-report-automation-ai-risks#:~:assi...>
- [39] <https://intuitionlabs.ai/articles/ai-ind-cta-drafting-pharma#:~:\~100...>
- [40] <https://intuitionlabs.ai/articles/clinical-study-report-automation-ai-risks#:~:assem...>
- [41] <https://www.synterex.com/post/nda-filing-strategy#:~:Opera...>
- [42] <https://intuitionlabs.ai/articles/ai-ind-cta-drafting-pharma#:~:match...>

- [43] <https://www.synterex.com/post/nda-filing-strategy#:~:The%2...>
 - [44] <https://intuitionlabs.ai/articles/clinical-study-report-automation-ai-risks#:~:Data%...>
 - [45] <https://intuitionlabs.ai/articles/clinical-study-report-automation-ai-risks#:~:Addit...>
 - [46] <https://intuitionlabs.ai/articles/clinical-study-report-automation-ai-risks#:~:Anoth...>
 - [47] <https://intuitionlabs.ai/articles/ai-ind-cta-drafting-pharma#:~:,13...>
 - [48] <https://pharmacystandards.org/caidra-examination/section-4-4-fda-ema-view-on-ai-assisted-documentation/?PageSpeed=noscript#:~:Stand...>
 - [49] <https://intuitionlabs.ai/articles/ai-ind-cta-drafting-pharma#:~:match...>
 - [50] <https://intuitionlabs.ai/articles/clinical-study-report-automation-ai-risks#:~:gover...>
 - [51] https://pharmacystandards.org/caidra-examination/section-21-1-ectd-placement-of-ai-outputs/?cps_autocomplete=1&from=1386546#:~:Regul...
 - [52] <https://intuitionlabs.ai/articles/clinical-study-report-automation-ai-risks#:~:AI%20...>
 - [53] <https://www.linkedin.com/pulse/regulatory-medical-writing-how-ai-reshaping-nguyen-ph-d-dtm-wagcc#:~:Pharm...>
 - [54] <https://pmc.ncbi.nlm.nih.gov/articles/PMC12675821/#:~:Artif...>
 - [55] <https://www.ppd.com/resource/ai-regulatory-medical-writing-productivity/#:~:AI%20...>
 - [56] <https://www.synterex.com/post/nda-filing-strategy#:~:match...>
 - [57] <https://intuitionlabs.ai/articles/ai-ind-cta-drafting-pharma#:~:Notab...>
 - [58] <https://intuitionlabs.ai/articles/clinical-study-report-automation-ai-risks#:~:repor...>
-

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.