



INTUITIONLABS / RESEARCH & PRACTICAL GUIDANCE

AI in Regulatory Writing: Benefits, Risks, and Use Cases

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 · 2026-02-14 · regulatory writing · medical writing · generative ai · clinical study reports · compliance automation · ai risk management · large language models · data privacy

Original article: <https://intuitionlabs.ai/articles/ai-regulatory-writing-benefits-risks>



In this article

1. Executive Summary
 2. Introduction and Background
 3. AI Capabilities for Regulatory Writing
 4. Where AI Works: Benefits and Case Studies
 5. Where AI Falls Short: Challenges and Risks
 6. Tools and Technologies
 7. Data Analysis and Evidence
 8. Case Studies
 9. Regulatory, Ethical, and Organizational Implications
 10. Future Directions
 11. Conclusion
-

Executive Summary

AI-assisted regulatory writing is rapidly transforming how organizations prepare compliance documentation in high-stakes industries such as pharmaceuticals and finance. Advanced language models and AI platforms can **dramatically accelerate document drafting**, improve consistency, and reduce routine errors. A disclosed Merck case reports that effort to create a **fully human-reviewed CSR first draft** fell from an average of 180 hours to 80 hours (**Merck disclosure**). This measure includes human review and is distinct from automated draft generation. Observed case outcomes, vendor assessments, and prospective estimates should be evaluated separately. AI also supports **data synthesis tasks**: document workflows can combine source data, draft generation, and expert review.

However, significant limitations remain. Generative models are prone to “hallucinating” plausible-sounding but incorrect statements. In regulated contexts, where **traceability and factual accuracy are paramount**, hallucinations are especially worrying (www.linkedin.com) (www.pharmaceutical-networking.com). No AI system is entirely reliable or self-documenting; regulatory agencies emphasize that *accountability* stays with the sponsor, not the software (www.linkedin.com). Medical writers must still resolve conflicting data and verify results when constructing regulatory narratives (**original Merck/McKinsey case report**). In sum, AI can handle **some aspects** of regulatory writing (first drafts, summaries, consistency checks) quite well, but it **cannot yet replace human expertise** in ensuring scientific rigor, legal compliance, and patient safety. Human oversight, sophisticated data backend architectures, and rigorous validation remain essential.

This report provides an in-depth analysis of the current state of AI-assisted regulatory writing, drawing on case studies, industry data, and expert commentary. We review *where it works* (e.g. drafting, data integration, efficiency gains) and *where it doesn't* (e.g. hallucinations, confidentiality, auditability). We examine tools and platforms, cite specific examples (Merck, Takeda, HSBC, etc.), and discuss the broader regulatory and organizational implications. The report concludes with avenues for future development and recommendations for safely leveraging AI in regulatory documentation.

Introduction and Background

Regulatory writing refers to the creation of documentation required by government and industry regulators to demonstrate compliance with laws and standards. In **pharmaceuticals and medical devices**, this includes lengthy submissions for clinical trials and product approvals (e.g. **INDs**, CTAs, CSRs, protocols, labeling, and safety reports) that can span *thousands of pages* (www.linkedin.com). In **financial services**, it involves drafting policies, regulatory filings, audit reports, and analyses required by agencies like the SEC, FCA, or ECB. These documents demand **exactitude**: adherence to guidelines (FDA, EMA, ICH, IFRS, etc.), precise data reporting, correct terminology, and comprehensive traceability back to source data. Even small errors or omissions can delay approvals, incur fines, or lead to serious consequences for companies and public health. As one expert notes, minor errors in a submission “can trigger significant delays” and **postpone patient access to life-saving treatments** (www.linkedin.com).

Traditionally, regulatory writing is a **labor-intensive, time-consuming** process. Subject-matter experts (scientists, statisticians, medical writers) must sift through vast data sets, interpret complex results, and conform to evolving regulations (www.linkedin.com). For example, preparing a clinical study report (CSR) requires integrating tables, listings, and figures (TLFs) from statistical analysis plans, then writing descriptive narratives for each section (www.axtria.com) (www.axtria.com). This process often involves *multiple rounds of review and revision* among global teams, slowing the path to submission. Even with thorough review, inconsistencies can creep in – a risk unacceptable to regulators.

Recent advances in **generative AI and natural language processing** (NLP) offer opportunities to change this paradigm. Large Language Models (LLMs) such as OpenAI’s GPT series and Google’s Gemini are neural models trained on vast text corpora; they can generate fluent human-like text and perform summarization, translation, and Q&A tasks. In **AI-assisted regulatory writing**, these models (or specialized domain versions) are used to draft documents, summarize scientific content, check consistency, or extract information from raw data. Often they are integrated into platforms that combine LLMs with data warehouses, style guides (e.g. MedDRA dictionaries), and workflow tools to ensure compliance (www.dip-ai.com).

The **context** includes company experimentation and regulatory discussion. Takeda describes developing an in-house GenAI tool ([Takeda](http://www.takeda.com)). The FDA’s January 2025 AI guidance remains a **nonbinding draft, not for implementation**, with recommendations tied to a particular context of use ([FDA draft guidance](http://www.fda.gov)). AI Act high-risk classification depends on the system’s intended use and the conditions in Article 6 and the relevant annexes; working in a regulated industry alone does not establish that classification ([Article 6](http://www.european-council.europa.eu)).

In this environment, case outcomes should be distinguished from prospective productivity estimates. The case studies below separate automated generation from the effort needed for a reviewed draft. The following sections analyze the strengths and weaknesses of AI-assisted regulatory writing, with attention to evidence, human oversight, and source traceability.

AI Capabilities for Regulatory Writing

Modern AI techniques offer several modes of assistance in regulatory writing. Broadly, these include:

- **Generative Drafting:** LLMs can produce initial drafts of documents or sections. By feeding a model with relevant inputs (e.g. a study protocol, data tables, or previous submissions), one can prompt it to generate narrative text. For instance, generative AI has been used to auto-generate Clinical Trial Protocols, Investigator Brochures, or Clinical Study Reports (CSRs) in draft form.
- **Template Filling / Retrieval-Augmented Generation (RAG):** To avoid "hallucinations," many systems use a retrieval approach. Key data and approved language fragments are stored in a structured database ("vault"), and the AI pulls exact figures, study results, and text snippets into predefined templates (www.linkedin.com). This *choreographed* approach treats AI as a sophisticated assembler rather than a freeform improviser (www.linkedin.com). Each statement in the output can be traced back to its source cell or paragraph, satisfying regulators' demand for provenance (www.linkedin.com).
- **Data Summarization:** AI can quickly process voluminous data (e.g. clinical trial tables, adverse event logs) and synthesize them into coherent narratives. An experimental pointer-generator study used TF-IDF "median match" alignment and reported F1 ROUGE-1, ROUGE-2, and ROUGE-L scores of 67.0, 64.0, and 64.0, respectively. These measure text overlap; the authors found that even the best-performing summaries were not consistently comprehensive or informative ([original study](#)). Treat generated medication guidance as experimental and require clinical review. More generally, AI is adept at turning statistical tables or bulleted findings into readable text.
- **Consistency Checking and Quality Control:** Automated tools can cross-check facts and ensure consistency. For example, AI can detect discrepancies between text and datasets, verify that all cross-references are consistent, and flag deviations from style guides. Early pilots have used generative models to **identify data mismatches** and harmonize content across sections (www.linkedin.com). This reduces the mundane burden of manual verification.
- **Language Localization and Style:** AI models can translate or adapt content for multiple markets. For instance, a specialized system may be trained to generate regulatory documents in different languages (English, French, Japanese) while respecting local terminologies and guidelines. NLP can also enforce consistent usage of terms (e.g. MedDRA coding for medical terms) to maintain uniformity across submissions.
- **Regulatory Intelligence and Policy Mapping:** Beyond document text, AI can help track regulatory changes and map obligations. Advanced language models can ingest new regulations or guidance and automatically update internal control documents. LLM-based tools have been shown to **map regulatory text to internal controls** and draft preliminary impact assessments (www.linkedin.com). Similarly, generative AI can create a first-cut policy document aligned with jurisdictional citations and controls (with human review) (www.linkedin.com).

Table 1 summarizes typical tasks and AI roles:

Table 1: AI Use Cases in Regulatory Writing

Task / Capability	AI Role / Functionality	Example / Impact	Sources
Initial Draft Generation	Produces first-draft text from given inputs (protocols, data, etc.)	McKinsey's prospective consultancy analysis describes AI as potentially freeing medical writers to focus on dossier sections requiring more complex clinical interpretation (McKinsey prospective analysis). A safety narrative draft was completed 90% faster with AI (www.linkedin.com).	(www.linkedin.com) (www.mckinsey.com)
Data Synthesis	Summarizes clinical tables or data into prose	Generative models process large datasets and create structured reports. Atria's vendor assessment of potential benefits describes at least a 30% reduction in overall CSR generation time, rather than a measured reduction in initial-draft time (Atria assessment).	(www.axtria.com) (www.mckinsey.com)
Consistency/QC Checking	Cross-checks facts, data consistency, style, references	Early pilots show AI detecting discrepancies between text and underlying data (www.linkedin.com). McKinsey/Merck system reduced CSR errors by ~50% (www.mckinsey.com).	(www.linkedin.com) (www.mckinsey.com)
Formatting / Template Filling	Populates documents based on templates and rules, ensuring compliance	Proprietary platforms (e.g. Deep Intelligent Pharma) claim to automate submission formatting with full audit traceability (www.dip-ai.com). Proprietary "assembly line" AI enforces template rules to avoid hallucinations (www.linkedin.com).	(www.dip-ai.com) (www.linkedin.com)
Regulatory Summaries & Q&A	Drafts responses to agency queries, summaries of regulations	GenAI is used for Health Authority Query (HAQ) responses to reduce manual workload (www.mckinsey.com). LLMs can draft change-impact summaries for updated regulations (e.g. FDA guidances).	(www.mckinsey.com) (www.linkedin.com)
Literature & Guideline Reviews	Extracts and summarizes guidelines and literature for writing	AI chatbots can rapidly retrieve and summarize relevant regulations or prior submission precedent. For example, specialized retrieval can pull policy compliance text for templates (www.linkedin.com).	(www.linkedin.com)
Translation & Localization	Translates or adapts content to other languages/markets	LLM translations (post-edited by experts) speed writing in multi-jurisdiction profiles. Some medical writing teams use AI to draft local-language patient leaflets from English texts.	(Industry practice; no single citation)
Policy Drafting	Generates preliminary policy or control documents	Generative models can create first-cut compliance policies with embedded references; best practice retains human/trace review (www.linkedin.com).	(www.linkedin.com)
Multimodal Integration	Incorporates charts, tables, even images/content via AI support	Agent-based AI systems can ingest figures and replicate visualizations or integrate data charts into word documents. (Emerging capability.)	(Emerging trend; see use of agency AI (www.mckinsey.com))

The above illustrates how AI **augments human workflows**. In pharma companies, these AI capabilities translate to significant efficiency. For example:

- An industry report notes that users of specialized AI platforms created safety narratives in minutes – about **90% faster** than the traditional timetable ([www.linkedin.com](#)).

- McKinsey's prospective consultancy analysis describes potential support for dossier drafting, allowing writers to focus on sections requiring more complex clinical interpretation ([McKinsey prospective analysis](#)). This is a proposed benefit, not a measured study outcome.
- In parallel, a partnership between Merck and McKinsey developed an AI application that **reduced the effort to create a fully human-reviewed CSR first draft from an average of 180 hours to 80 hours**, and concurrently cut drafting errors by ~50% ([www.mckinsey.com](#)).

Financial compliance sees analogous benefits in narrative reports. Banks report that **AI-assisted drafting of Suspicious Activity Reports (SARs)** – with evidence linked into narratives – notably **accelerates case cycles** ([www.linkedin.com](#)). For instance, one source notes LLM tools that “draft SAR narratives with citations to the underlying evidence” allowed investigators to complete reviews more quickly, yielding more consistent case files ([www.linkedin.com](#)). Meanwhile, HSBC's dynamic AI risk-scoring system (built with Google Cloud) reportedly detected 2–4× more financial crime than its legacy approach, illustrating AI's power in *analysis* (though that example is for detection rather than writing) ([www.linkedin.com](#)).

These gains largely arise when AI is applied to well-defined, data-driven writing tasks. Automated systems shine at routine portions of documents (e.g. population PK reports, repeated protocol sections, statistical result summaries) and obvious consistency checks. By offloading template-based drafting and look-up tasks, AI **frees human experts** to focus on high-level content and interpretation ([www.linkedin.com](#)) ([www.linkedin.com](#)). For example, trained writers can spend more effort on scientific messaging and strategic planning, rather than reformatting tables or double-checking every reference ([www.linkedin.com](#)) ([www.mckinsey.com](#)).

Where AI Works: Benefits and Case Studies

AI's strengths in regulatory writing emerge especially in speed, consistency, and handling large data volumes. Key reported benefits include:

- **Time Savings.** A 2025 AutoIND preprint evaluated initial drafting of nonclinical written summaries for two IND examples. It compared directly recorded automated drafting times (3.7 and 2.6 hours) with an approximately 100-hour manual baseline estimated by experienced regulatory writers, reporting an approximately 97% reduction in **initial drafting**, not total submission-ready effort. The authors emphasize that expert writers remain essential to mature the drafts ([original preprint](#)). Estimates and case-specific results should not be presented as a universal 30–60% saving across regulatory-writing workflows.
- **Consistency and Quality.** AI enforces uniformity across documents. By using the same underlying data and style rules, AI-generated sections have **stable terminology and formatting**. For instance, a centralized AI-driven system can apply the same MedDRA coding for adverse events in every periodic safety report. Errors also decline: the Merck/McKinsey CSR project objectively halved data/messaging errors ([www.mckinsey.com](#)). Similarly, AI can adapt learned editorial guidelines (e.g. ICH E3 formatting) automatically, reducing human oversight needed for purely mechanical checks. Case consistency is another benefit: when filing in multiple regions, AI can systematically generate region-specific requirements (e.g. local regulatory forms) from a single source data set, improving global alignment.

- **Streamlined Review and Query Handling.** A major pain point in submissions is handling regulator queries (HAQs – Health Authority Questions). AI can partially automate drafting HAQ responses. McKinsey notes companies are exploring generative AI to “generate HAQ responses” simultaneously for multiple agencies (www.mckinsey.com). Although still in early use, automating these repetitive write-backs can substantially ease CRO/regulatory team workload. In sum, AI helps **focus human experts on strategy** rather than rote letter-writing.
- **Example – Merck Clinical Study Reports.** Merck and McKinsey distinguish rapid automated generation from the effort to create a fully human-reviewed first draft. The detailed case-study section below gives the measures separately ([original case report](#)). Writers review and edit the generated content before it can be relied on.
- **Example – Takeda Prototype.** Takeda describes an in-house prototype for regulatory submissions. Its estimate that an approximately 20-week analysis and tabulation process could be halved is prospective; it should not be presented as an observed reduction in submission-ready authoring time ([Takeda story](#)).
- **Example – Financial Compliance.** HSBC describes AI-assisted financial-crime detection and investigation. These outcomes concern analysis, not evidence of AI-authored SAR narratives. Its original disclosure reports 60% fewer false positive cases ([HSBC disclosure](#)).
- **Improved Uniformity Across Global Submissions.** Maintain approved text and company templates in a controlled library. Check any reused or adapted passages against the current product, region, and source data; assess whether automation reduces rework rather than assuming that an unspecified corporate survey establishes a benefit.
- **Content Enrichment and Storytelling.** McKinsey discusses agentic AI as a **proposed** virtual content challenger alongside a human reviewer to anticipate questions and improve dossier quality. This is a potential use case, not a disclosed Merck deployment ([McKinsey discussion](#)).

In sum, structured drafting, data integration, and formatting are suitable areas to evaluate for AI assistance. Table 1 enumerates potential use cases. Productivity and quality should be measured for the actual workflow, including human review, rather than inferred from unrelated vendor percentages.

Where AI Falls Short: Challenges and Risks

Despite its potential, AI-assisted writing has important limitations in the regulatory context. These weaknesses arise from both technical and regulatory constraints:

- **Hallucinations and Factual Errors.** By design, LLMs predict text based on learned patterns rather than verified facts. A critical drawback is **fabrication of details** – known as “hallucination.” Models can generate convincing but *incorrect* statements or data. In healthcare contexts, hallucinations have led to serious misinformation; one review warns that LLMs may produce false diagnostic or treatment suggestions with high confidence (www.nature.com). In regulatory writing, hallucinated claims (e.g. fake study results, incorrect references) are unacceptable. Regulators emphasize *traceability*: every statement in a submission must be backed by source data (www.linkedin.com). If an AI “hallucinates” a fact that is not directly pulled from the data, it breaks this requirement. An EMA spokesperson warned

that untraceable AI-generated text is fundamentally problematic (www.linkedin.com). Indeed, even minor hallucinations (e.g. citing a non-existent author or misstating a number) could invalidate an entire submission. A 2024 experiment concerned bibliographic references retrieved for systematic reviews, rather than general factual-error rates in scientific prose. It defined a hallucinated reference as one with at least two incorrect fields among title, first author, and publication year. These task-specific findings should not be extrapolated to all regulatory writing ([original study](#)).

- **Lack of Explainability and Audit Trails.** Closely related to hallucination is the **black-box problem**. Pure LLM outputs do not inherently document how an answer was derived. Regulatory agencies demand explanations for claims (e.g. which study or dataset supports each statement). Standard LLM answers lack this: they do not cite sources by default. An LLM would not, for example, say “According to Table 5 of this study...”. This invisibility undermines auditability (www.linkedin.com). Even if a model is trained on internal documents, it typically does not indicate which document fragment generated each phrase. Without heavy post-processing (or RAG-augmented systems), an AI draft is essentially a “blind” assembly and cannot be signed off as verifiable by a regulator. The remedy is often to use retrieval methods that *force* traceability (www.linkedin.com), but those systems are complex. Efforts to “open the black box” (e.g. having the AI footnote its sources) remain immature.
- **Accountability and Liability.** Regulatory submissions are legal documents: errors can have dramatic consequences (pulled approvals, patient harm). Liability ultimately rests with the sponsor (the company submitting the dossier), not the AI vendor. As one regulatory expert put it, regulators like the FDA “will not accept ‘the AI made a mistake’ as an excuse” – the company is fully responsible (www.linkedin.com). This high-stakes accountability means companies must thoroughly validate AI outputs. If an AI system introduces an error that misrepresents trial data, the sponsor (and perhaps individual signatories) are on the hook. Legal frameworks in finance and healthcare reinforce this. For example, banking regulators have made clear that using AI requires robust model validation and governance (www.linkedin.com). In short, AI can assist, but human review is *mandatory*, and final disclaimers or certifications cannot blame the tool.
- **Data Privacy and Security.** Before entering sensitive information, assess the approved deployment, contractual permitted uses, retention settings, and access controls. The governance section below addresses safeguards for HIPAA-regulated data.
- **Intellectual Property Concerns.** Another issue is IP ownership of AI-generated content. Who owns writing produced by an LLM? While less of a writing quality issue, it introduces legal ambiguity. AI training data may itself include copyrighted content; an AI could inadvertently regurgitate a copyrighted passage. Because of opaque data provenance, a generated paragraph might infringe without clear attribution (velvetech.com). Use third-party material only on an appropriate lawful basis, such as permission, a licence, public-domain status, or an applicable statutory exception. In the United States, fair use can permit unlicensed uses in particular circumstances; quotation neither always requires permission nor automatically establishes fair use ([U.S. Copyright Office](https://www.copyright.gov)). Review AI-generated passages for attribution and lawful use.

- **Bias and Fairness.** Large models trained on internet text can inherit biases. In clinical writing, biases might lead to underreporting or overreporting effects in certain patient subgroups if the training data lacked diversity. In financial writing, an AI might mischaracterize entities or events based on skewed media data. While not deeply studied in regulatory writing specifically, generalized concerns are credible: if an AI erroneously infers that a study “failed” because a dataset had selection bias, that’s a factual error. Mitigations include curating training sets and using human review, but no AI is bias-free. Regulatory agencies will scrutinize AI tools for any evidence of systematic error that could disadvantage a protected class or undermine scientific objectivity.
- **Context and Nuance.** Organize source material and review the draft’s clinical interpretation, emphasis, and treatment of uncertainty. Fluent prose should not be treated as proof that a model understands a novel result or the current requirements.
- **Regulatory Acceptance and Ethical Oversight.** The FDA’s January 2025 credibility guidance remains a **nonbinding draft, not for implementation (FDA status)**. It recommends a framework for AI used to produce information or data supporting drug safety, effectiveness, or quality decisions in a defined context of use. It excludes operational drafting or writing that does not affect patient safety, drug quality, or study-result reliability (**scope, section II**). This draft does not establish a blanket obligation to document every writing assistant in a filing. Determine the submission and data-integrity requirements that apply to the actual product and use; treat AI-specific validation and documentation recommendations according to their stated scope.

In summary, AI assistance requires careful review of source accuracy, interpretation, and provenance. Use human review and version control to assess each workflow, and avoid claims that first-pass generation eliminates expert authorship or final sign-off.

Tools and Technologies

The AI landscape for regulatory writing includes a mix of general-purpose LLMs, specialized domain models, and integrated platforms:

- **Large Language Models (General):** General-purpose language models can be evaluated for drafting, summarization, and other text tasks. Assess the specific product’s documented functions and deployment terms rather than assuming a model name establishes suitability for regulated content.
- **Domain-Specific Models:** Domain adaptation can be evaluated against the actual writing task. Test scientific accuracy and source traceability on representative inputs; do not infer regulatory-writing suitability from a biomedical or legal product label alone.
- **Retrieval-Augmented Generation (RAG) Systems:** These systems combine LLMs with a knowledge base. Index prior submissions, internal SOPs, and applicable regulations, then check generated statements against the retrieved evidence. Evaluate factual accuracy and traceability directly rather than relying on untraceable marketing percentages.

- **AI Writing Assistants:** Use primary product disclosures to establish what a platform actually does. Celegence's CAPTIS has a documented Kenvue pilot generating CMC Module 3 documents for FDA and MHRA submissions ([Celegence disclosure](#)). Assess the intended document type, source linking, review workflow, and access controls before selecting a platform.
- **Generic Content-Generation Tools:** Platforms like Jasper, Copy.ai, or Writesonic are not regulatory-specific but provide user-friendly LLM interfaces. They can aid brainstorming and scoping a document. For example, a writer might paste key points and have Jasper generate a draft paragraph, then edit it. These are rule-of-thumb tools: they accelerate authoring but lack built-in compliance checks. Organizations often caution that such tools should only be used to generate **generic text** (e.g. first draft of a background section), never actual clinical data or safety conclusions.
- **Statistical & Data Science Tools:** While not "writing AI", advanced analytics platforms (Python, R) play a supporting role. Machine learning models can preprocess raw trial data, extract key metrics, or prepare TLFs which are then fed into narrative templates. Modern pipeline orchestration (e.g. NextFlow, KNIME) automates these transfers into document generation scripts or into the AI systems. Emerging AI platforms (QuantumBlack/Azure ML) advertise integration of analytical results into narrative text generation.

When selecting tools, organizations must consider **auditability, security, and compliance**. Distinguish a model hosted on premises from a privately networked cloud service. Azure OpenAI is a cloud deployment: its private endpoint controls connectivity to the Azure resource; an on-premises client connects to that cloud resource rather than moving model processing onto the customer's premises ([Microsoft network documentation](#)). Verify processing location, contractual permitted uses, and data-handling settings for the actual deployment.

In practice, a hybrid authoring workflow can keep source documents in a structured repository with version control, retrieve relevant passages, and apply review checkpoints to generated text. Trace each material claim to its source and assess the final draft with qualified reviewers. Evaluate the complete workflow rather than assuming that an enterprise label or model choice establishes regulatory suitability.

Data Analysis and Evidence

Quantitative evidence for AI's impact in regulatory writing is still limited to case studies and consulting reports, but the available data is compelling:

- **Survey and Benchmark Data.** McKinsey's industry benchmarks (2025) show leading pharmaceutical companies have already compressed submission timelines dramatically – some file within 50–65% less time than 2020 standards (www.mckinsey.com). McKinsey attributes part of this acceleration to new technologies like generative AI, alongside process redesign. Their analysis estimates that cutting one month from a \$1B drug's approval can yield ~\$60M NPV in revenue (www.mckinsey.com), highlighting the economic stakes.

- **Time Savings.** Separate reviewed effort from automated generation, and identify the comparator for each result. The Merck case measures a fully human-reviewed CSR first draft separately from generation in minutes ([original case report](#)). Axtria describes its potential efficiency assessment rather than an independent study ([Axtria](#)). The AutoIND preprint concerns initial nonclinical-summary drafting for two IND examples, with an estimated manual baseline and further expert maturation ([original preprint](#)). These different measures do not establish that AI universally halves authoring hours.
- **Error Rates and Quality.** In terms of quality, Merck's experience provides measurable data: their AI-supported drafts had about **50% fewer errors** in categories such as data mismatches, citations, and terminology ([www.mckinsey.com](#)). In finance, AI for transaction monitoring at HSBC reportedly reduced false positives, implying error improvement (though those data are internal). Published studies outside regulatory writing have found LLM-generated scientific text can score lower in certain quality metrics than human authors, indicating room for improvement ([www.pharmaceutical-networking.com](#)). In practice, the positive outcomes reported (fewer red flags per review, higher first-pass acceptance) suggest that AI-checks do improve consistency.
- **Economic Impact.** A common metric cited is the ROI of faster timelines. For example, McKinsey estimates a one-month earlier filing on a \$1B asset yields ~\$60M NPV ([www.mckinsey.com](#)). If AI accelerates multiple projects across a portfolio, the value impact is substantial. Another analysis extrapolates that generative AI in pharma R&D could unlock hundreds of millions in value (by compressing time and reducing cost overruns) ([www.mckinsey.com](#)) ([www.mckinsey.com](#)). Regardless, early adopters uniformly highlight cost reduction (less contractor and OPEX per submission) as a key benefit ([www.axtria.com](#)) ([www.linkedin.com](#)).
- **Productivity Metrics.** Financial-crime detection and investigation metrics should not be relabelled as writing outcomes. Standard Chartered's disclosure describes screening recommendations supported by written explanatory narratives, rather than SAR-authoring performance or a measured tenfold speed gain ([original bank disclosure](#)).

For evaluation, measure total reviewer effort, correction types, source accuracy, and submission readiness alongside generation time. Unidentified unpublished surveys and untraceable vendor percentages cannot establish those outcomes.

Case Studies

Concrete examples of AI-assisted regulatory writing illustrate both its potential and constraints:

- **Merck (Pharmaceutical).** Merck's *LifeSciences.AI* initiative (in partnership with McKinsey's QuantumBlack) is a high-profile case. As reported in mid-2025, Merck built an application to automate CSR drafting for late-phase trials ([www.mckinsey.com](#)). By integrating trial data (tables, listings) and using generative modules, they slashed CSR authoring from 2–3 weeks to **3–4 days** ([www.mckinsey.com](#)). The effort to create a fully human-reviewed CSR first draft fell from an average of 180 to 80 hours ([www.mckinsey.com](#)), while errors were cut by half in key categories. Importantly, this system was *not fully autonomous* – it coupled AI with human experts in iterative review. Merck's

principal scientist noted amazement at producing a full CSR draft in “eight minutes,” a task that previously took months (www.mckinsey.com). This case shows how AI can deliver a **massive speed boost** on a critical path document, while upholding quality with human oversight.

- **Takeda (Pharmaceutical).** Takeda’s 12 December 2024 disclosure describes an in-house GenAI prototype initially tested on the CSR efficacy-analysis section. A writer reported that it could produce readable factual text ([Takeda’s original disclosure](#)). The productivity estimate discussed above is prospective; this prototype account should not be presented as an observed end-to-end authoring result. Generated narratives still need expert review.
- **Celequence/CAPTIS (CMC Module 3).** Celequence’s July 2025 release describes a six-month Kenvue pilot using CAPTIS to generate CMC Module 3 documents for FDA and MHRA submissions. Celequence reports **60% faster draft turnaround** and **80–95% alignment with previously approved submissions** ([Celequence’s primary release](#)). These are vendor-reported results for that CMC pilot; they should not be generalized to medical-device document workflows.
- **Financial Industry.** HSBC’s primary disclosure describes AI-assisted financial-crime detection and investigation, including **60% fewer false positive cases** ([HSBC disclosure](#)). These are detection and investigation outcomes; assess writing performance separately rather than inferring SAR drafting or turnaround from them.
- **Governance and Model Risk.** Though not a “writing case” per se, regulatory pilots such as the UK FCA’s AI sandbox (with Nvidia) illustrate ecosystem support for compliant AI experimentation (www.linkedin.com). In practice, firms have drafted model risk documents (in line with PRA/OCC guidelines) to validate their AI usage. The growing regulatory apparatus around AI also means we should anticipate future case studies of agency-approved AI use: the FDA’s 2025 draft guidance on AI in submissions signals that companies may start including AI methodologies as part of their formal submissions (submitting documentation of model training, validation reports, etc. along with trial data) (www.fda.gov).

Table 2: Case Studies of AI-Assisted Regulatory Writing

Organization / Company	Sector	AI Application	Outcomes / Evidence of Impact	References
Merck & QuantumBlack	Pharmaceuticals	CSR (Clinical Study Report) drafting with LLM and data automation	Effort for a fully human-reviewed CSR first draft reduced from an average of 180h to 80h; 50% error reduction (www.mckinsey.com). CSR authoring time 2–3 weeks → 3–4 days (www.mckinsey.com). High user satisfaction (8-min draft anecdote) (www.mckinsey.com).	(www.mckinsey.com) (www.mckinsey.com) (www.mckinsey.com) (www.mckinsey.com)
Takeda	Pharmaceuticals	In-house GenAI prototype, initially tested on CSR efficacy analysis	12 December 2024 disclosure: a writer reported readable factual text from prototype testing. The prospective analysis-and-tabulation estimate is discussed above; it is not an observed end-to-end authoring result.	(Takeda disclosure)

Organization / Company	Sector	AI Application	Outcomes / Evidence of Impact	References
Celegence (CAPTIS)	Pharmaceutical CMC	Six-month Kenvue pilot generating CMC Module 3 documents for FDA and MHRA submissions	Vendor-reported pilot results: 60% faster draft turnaround and 80–95% alignment with previously approved submissions. These results concern the CMC pilot, not device-document workflows generally.	(Celegence primary release)
HSBC	Finance	AI-assisted financial-crime detection and investigation	HSBC reports 60% fewer false positive cases. This is a detection/investigation outcome, not a measure of SAR-authoring performance.	(HSBC primary disclosure)
Standard Chartered / Others	Finance	KYC and sanction screening with ML/AI	Using ML to pre-check name/sanctions lists, reducing human screening. Silent Eight reports global deployment reducing investigations (third-party claim) (www.linkedin.com). No public metrics, but industry sources cite meaningful time savings.	(www.linkedin.com)
Regulatory Agencies (FDA/EMA)	Regulatory Bodies	AI governance frameworks	FDA 2025 draft guidance acknowledges exploding AI use in submissions; calls for model credibility assessments (www.fda.gov). Joint FDA/EMA principles (2026) emphasize transparency and robust standards (www.linkedin.com). Not an “application” case, but shows official position.	(www.fda.gov) (www.linkedin.com)

These cases underscore two themes:

- 1. Human-in-the-Loop Workflows:** In each example, AI *augmented* rather than replaced human experts. Drafts needed editing, and final sign-off remained with qualified writers. The AI’s role was to handle routine or data-intensive portions, increasing throughput.
- 2. Need for Integration:** Major successes (Merck, CElegence) involved integrating AI into broader systems (data pipelines, company knowledge bases). The “source vault” approach ([www.linkedin.com](#)) was key: raw AI chatbots alone were insufficient. Similarly, secure enterprise deployments were favored (Merck’s solution ran on Merck’s own infrastructure, Celegence’s on secure servers).

By contrast, anecdotal failures (not tabulated above) often stem from isolated use of consumer LLMs without tracked outputs, or attempts to have AI authorship without clear workflow design. Well-documented case studies emphasize that *structured processes and domain data access* are critical enablers.

Regulatory, Ethical, and Organizational Implications

The rise of AI-assisted writing carries significant implications across regulation, ethics, and corporate strategy:

- Regulatory Guidance & Legislation (reviewed 7 October 2026).** The FDA’s January 2025 AI document remains a **nonbinding draft, not for implementation** ([FDA status](#)). Its credibility framework concerns AI-generated information or data used in regulatory decisions about drug safety, effectiveness, or quality, assessed for a defined context of use. It excludes operational drafting or writing that does not

affect patient safety, drug quality, or study-result reliability ([scope](#)). Determine actual submission and data-integrity obligations under the applicable statutes and regulations; the draft does not require every writing workflow to file model documentation.

Under the EU AI Act, assess intended purpose against Article 6's product/safety conditions and Annex III uses. A preparatory or narrow procedural task can fall outside Annex III high-risk classification where the statutory conditions are met; profiling natural persons remains high-risk within Annex III. Pharma or finance use alone does not determine status ([Article 6](#)). The Commission's current implementation page, following the political agreement on the AI Omnibus, gives **2 December 2027** for specified high-risk areas and **2 August 2028** for systems integrated into products. The Commission also distinguishes rules already applicable: prohibitions, definitions, and AI-literacy provisions since 2 February 2025, and governance and general-purpose AI obligations since 2 August 2025 ([Commission implementation FAQ](#)). These are provision-specific dates, not a single deadline for all writing tools ([Commission timetable](#)).

UK banking supervision is separate from EU classification. PRA SS1/23 sets model-risk expectations for regulated UK-incorporated banks, building societies, and PRA-designated investment firms with specified internal-model capital approvals ([PRA scope](#)). It is not a general classification rule for regulatory-writing tools. Assess binding obligations by jurisdiction and function; model cards and early regulator engagement can be governance practices rather than universal deployment prerequisites.

- **Ethical and Quality Considerations.** Ethically, using AI in patient-impacting documents demands STEEP (Scientific, Technical, Ethical, Economical, Patient safety) stewardship. Bias or errors in a submission can harm patient lives. Industry thought leaders stress that AI should prioritize *patient safety and data integrity* over mere speed. For instance, despite efficiency temptations, all AI outputs must be vetted for adherence to the latest clinical guidelines and factual accuracy (www.linkedin.com) (www.nature.com). In practice, companies often establish "AI writing policies" similar to research ethics: requiring peer-review of AI-generated text, checking against source data, and preserving audit trails of all changes. Training analysts to spot likely AI errors is also becoming part of compliance training programs.
- **Workforce Evolution.** The role of regulatory writers is evolving. Many industry commentators (e.g. Hiep Nguyen, Shruti Sharma (www.linkedin.com) (www.linkedin.com)) note that experts can shift focus from drafting to *editorial strategy, data interpretation, and oversight*. Effective AI use requires writers who understand both the domain and the technology. New skills like prompt engineering (crafting effective LLM prompts), AI validation, and data management are emerging as crucial. Several training courses and webinars (e.g. Xtalks, EMWA sessions) now address "AI literacy" for medical writers, covering both how to use tools and how to critically evaluate AI output. The general sentiment is that, while AI will automate many tasks, it will *not replace human experts*. Instead, writers are expected to become *AI co-pilots*: guiding the AI, reviewing its work, and making final judgments (www.linkedin.com) (www.linkedin.com). Evaluate training needs and author responsibilities for the actual workflow.
- **Process and Technology Changes.** On an operational level, companies are rethinking their infrastructure. The Witherell analysis stresses the need for integrated data architecture: if content is not already in a structured repository, implementing AI is an uphill battle (www.linkedin.com). Many

organizations are accelerating investments in Regulatory Information Management Systems (RIMS) and Clinical Data Warehouses precisely to make AI effective (www.mckinsey.com) (www.mckinsey.com). Digitization efforts (like structured authoring, semantic tagging of documents, and APIs linking lab systems) are underway in parallel with AI pilots. The goal is an end-to-end digital submission pathway, where data flows directly from source to final document with minimal manual handoff. Until such integration is in place, many AI experiments remain ad hoc.

- **Security and Data Governance.** Evaluate contractual permitted uses, processing location, access controls, and safeguards before using sensitive data. For HIPAA covered entities and business associates using a cloud provider to process or store ePHI, an appropriate business associate agreement (BAA), risk analysis, and applicable administrative, physical, and technical safeguards are required. The provider's uses and disclosures must comply with the BAA and the Privacy Rule ([HHS cloud guidance](#)). Enterprise branding or a security certification does not replace these arrangements. Review data-retention settings separately from model-training controls; each is a distinct procurement question. Use approved test data and establish review gates before generated content enters a filing.
- **Future AI Regulation and Current Bank Guidance.** Monitor new rules and distinguish proposals from obligations already in force. In the United States, Federal Reserve **SR 26-2**, issued 17 April 2026, supersedes and replaces SR 11-7 and SR 21-8. The letter is expected to be most relevant to Federal Reserve-regulated banking organizations with more than \$30 billion in assets, and emphasizes model-risk management tailored to risk profile and operational size and complexity ([SR 26-2](#)). SR 11-7 is historical guidance; banking supervision should not be extended automatically to every regulatory-writing workflow.

Future Directions

Looking ahead, several trends are likely to shape AI's role in regulatory writing:

- **Improved Models and Data Integration.** We expect more advanced LLMs (larger, better-trained, domain-specific) to emerge. In parallel, RAG will become standard: future systems will likely run on private LLM backends connected to the company's data lakes and document repositories. GenAI tools will add multi-modal inputs (ingesting charts, chemical structures, even audio transcripts of meetings) to incorporate every piece of evidence. Witherell's metaphor of an "assembly line" suggests future platforms will autonomously assemble entire documents from modular pieces — not just produce text ad hoc (www.linkedin.com). This means writers will interact less by writing paragraphs, and more by **tuning the pipeline** (adjusting rules, validating outputs).
- **Regulatory Evolution.** Distinguish current agency activity from forecasts about future disclosure rules. FDA launched Elsa on 2 June 2025 and announced that it was already accelerating clinical-protocol reviews and scientific evaluations and identifying high-priority inspection targets ([FDA announcement](#)). The earlier speculation about an "Inspire" review project is not supported by that disclosure. Future requirements for AI-assisted writing should be evaluated as they are issued, rather than inferred from a single compressed AI Act deadline. The legal overview above identifies the Commission's current provision-specific timetable.

- **Industry Standards and Best Practices.** As the field matures, consortia (e.g. DIA's AI in Pharma initiative) will codify best practices. Standards for validating AI content (test suites, benchmarks for “hallucination” rates, etc.) will emerge. Training certification programs for “AI-savvy medical/regulatory writer” are already appearing. We anticipate checklist-based frameworks to evaluate AI tools (similar to the ASCCC or Purdue rubrics mentioned by vendors (www.dip-ai.com)).
- **Human-AI Collaboration Paradigms.** We may see specialized roles like “Regulatory AI Engineer” or “Document Automation Analyst” join regulatory affairs teams. Job descriptions will incorporate skills in prompt design, prompt reliability testing, and AI auditing. Organizational processes will include mandatory AI review committees for high-risk outputs. The human writer's craft will evolve: emphasis on strategic narrative construction, data interpretation, and reviewing AI suggestions rather than typing every sentence.
- **Ethical and Social Implications.** Beyond workflow, AI in regulatory writing raises broader questions. If approvals can be achieved faster, how will this affect drug safety? Ethical frameworks will need to address issues such as: ensuring AI is only used for legitimate efficiency (not for artificially inflating claims), maintaining public trust in AI-driven science, and guarding against over-reliance on algorithms for critical decisions. Patient advocacy groups may demand transparency on AI's role in drafting materials they see (e.g. patient leaflets).
- **Beyond Drug and Finance.** While currently most focus is on pharma and banking, other sectors could adopt similar approaches. For example, environmental compliance reports, aerospace certification documents, and cybersecurity policy compliance filings are amenable to AI assistance. Lessons from pharma/finance will likely migrate to these fields, with domain-tuned models (e.g. RegTech for environmental law) and analogous auditors.
- **Technological Extensions.** Future AI may integrate with virtual and augmented reality for collaborative authoring, or use voice interfaces to draft on-the-fly. Moreover, AI could proactively *generate insights* or alternatives (e.g. “Given this dataset, what regulatory strategy could minimize risks?”). Agentic AI (AI systems that set subgoals) might come into play to plan entire filing projects automatically, though that lies farther ahead.

In short, AI's role in regulatory writing will deepen but remain **supportive**. By 2030, we anticipate a seamless “smart workspace” where AI tools generate drafts, humans refine them, and automated systems ensure compliance at every step. The future regulatory author will be part engineer, part storyteller – empowered by AI but ultimately the guarantor of accuracy.

Conclusion

AI-assisted regulatory writing is poised to significantly alter how organizations produce compliance documentation. Evidence suggests **notable gains**: first-draft generation times can drop by tens of percent, consistency errors can be reduced, and professionals can redirect effort from typing to analysis. In substance, AI works best on well-defined, data-heavy tasks (figure/table narration, routine sections, initial drafting) and

achieves impressive productivity boosts there (www.linkedin.com) (www.mckinsey.com). Case studies (Merck, Takeda, HSBC, etc.) demonstrate that when properly integrated, AI can compress timelines that would otherwise delay critical approvals or compliance.

However, these benefits require careful review. Check scientific claims, numerical values, and citations against source material; evaluate the intended use and applicable legal requirements for each workflow. Keep human review and clear author responsibilities throughout drafting and final sign-off. Data-handling arrangements require a deployment-specific assessment, rather than assumptions based on a public or enterprise label.

Looking forward, better integration and task-specific evaluation can support AI-assisted writing. The FDA's draft framework assesses credibility for a particular context of use ([FDA draft guidance](#)). Organizations should improve source-data quality and review processes alongside automation, then measure the effort and quality of the reviewed output rather than generation speed alone.

References: Sources are linked inline. Company-reported outcomes, vendor assessments, research findings, and prospective estimates are identified in context.

Source links

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

1. <https://intuitionlabs.ai/articles/regulatory-submission-acceleration>
2. <https://www.merck.com/news/merck-expands-innovative-internal-generative-ai-solutions-helping-to-deliver-medicines-to-patients-faster/>
3. <https://intuitionlabs.ai/articles/llm-clinical-evidence-dossiers>
4. <https://www.linkedin.com/pulse/generative-ai-regulatory-submissions-promise-perils-de-florio-msc-ppuff>
5. <https://www.pharmaceutical-networking.com/a-pharmaconsult-about-ai-hallucinations-in-medical-writing/>
6. <https://www.mckinsey.com/about-us/new-at-mckinsey-blog/with-gen-ai-merck-and-mckinsey-transform-clinical-authoring>
7. <https://intuitionlabs.ai/articles/ectd-submission-ind-nda-guide>
8. <https://www.linkedin.com/pulse/regulatory-medical-writing-how-ai-reshaping-nguyen-ph-d-dtm-wagcc>
9. <https://www.axtria.com/articles/rescuing-clinical-study-reports-30-percent-faster-authoring-with-generative-ai>
10. <https://www.dip-ai.com/use-cases/en/the-best-impact-of-ai-on-regulatory-writing>
11. <https://www.takeda.com/our-impact/our-stories/aiming-to-enhance-regulatory-submissions-with-genai/>
12. <https://www.fda.gov/regulatory-information/search-fda-guidance-documents/considerations-use-artificial-intelligence-support-regulatory-decision-making-drug-and-biological>
13. <https://ai-act-service-desk.ec.europa.eu/en/ai-act/article-6>
14. <https://www.linkedin.com/pulse/part-4-how-choose-your-ai-writing-tool-2026-ingrid-witherell-da7dc>
15. <https://pmc.ncbi.nlm.nih.gov/articles/PMC9950092/>
16. <https://www.linkedin.com/pulse/practical-ai-financial-compliance-what-works-today-wilgenhof-plante-frx4e>
17. <https://www.mckinsey.com/industries/life-sciences/our-insights/generative-ai-in-the-pharmaceutical-industry-moving-from-hype-to-reality>
18. <https://www.mckinsey.com/industries/life-sciences/our-insights/rewiring-pharmas-regulatory-submissions-with-ai-and-zero-based-design>

19. https://www.linkedin.com/posts/alphalife-sciences_aiinlifesciences-regulatorywriting-medwriting-activity-7350797480449495040-RINk
20. <https://arxiv.org/abs/2509.09738>
21. <https://www.hsbc.com/news-and-views/views/hsbc-views/harnessing-the-power-of-ai-to-fight-financial-crime>
22. <https://www.nature.com/articles/s41746-023-00873-0>
23. <https://pmc.ncbi.nlm.nih.gov/articles/PMC11153973/>
24. <https://velvetech.com/blog/generative-ai-risks-regulations/>
25. <https://www.copyright.gov/fair-use/>
26. <https://www.fda.gov/media/184830/download>
27. <https://www.celegence.com/celegence-pioneers-industry-first-automation-of-cmc-module-3-document-creation-with-generative-ai/>
28. <https://learn.microsoft.com/en-us/azure/foundry-classic/openai/how-to/network>
29. <https://www.sc.com/en/press-release/weve-partnered-with-regulatory-technology-firm-silent-eight/>
30. <https://www.fda.gov/news-events/press-announcements/fda-proposes-framework-advance-credibility-ai-models-used-drug-and-biological-product-submissions>
31. <https://digital-strategy.ec.europa.eu/en/faqs/navigating-ai-act>
32. <https://digital-strategy.ec.europa.eu/en/policies/guidelines-ai-high-risk-systems>
33. <https://www.bankofengland.co.uk/prudential-regulation/publication/2023/may/model-risk-management-principles-for-banks-ss>
34. <https://www.hhs.gov/hipaa/for-professionals/special-topics/health-information-technology/cloud-computing/index.html>
35. <https://www.federalreserve.gov/supervisionreg/srletters/SR2602.htm>
36. <https://www.fda.gov/news-events/press-announcements/fda-launches-agency-wide-ai-tool-optimize-performance-american-people>

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.