


NLP for Prescribing Information: An Evidence-Based Review

By Adrien Laurent, CEO at IntuitionLabs • 11/22/2025 • 50 min read

natural language processing regulatory labeling generative ai prescribing information medical writing smpc large language models regulatory affairs



Executive Summary

The regulatory label of a pharmaceutical product—often called the **Prescribing Information (PI)** in the US or the **Summary of Product Characteristics (SmPC)** in the EU—is a highly structured, legally binding document. It is drafted and maintained by expert **medical writers** under strict **regulatory oversight**. **Natural Language Processing (NLP)** and *generative AI* (especially large language models like GPT) promise to automate parts of this writing process. This report evaluates whether and how NLP can draft your prescribing information. It reviews the evidence, current practices, and research on NLP applications in regulatory labeling, and it discusses practical, regulatory, and technical challenges.

Several studies and case examples illustrate the breadth of NLP tasks in labeling. For instance, *Gray et al.* used a BERT-based model to classify free-text excerpts of FDA labeling into PLR-defined sections, achieving ~94–96% accuracy for binary classification and 82% for multi-class on structured labels ^{([1](#))} [pmc.ncbi.nlm.nih.gov](#)) ^{([2](#))} [link.springer.com](#)). This shows NLP can automatically reorganize or tag content by section, aiding consistency. Other work demonstrates **summarization and content-generation**: *Meyer et al.* built a pointer-generator model to draft patient-facing **Medication Guides** from more technical label text, improving ROUGE scores by ~7 points using heuristic alignment ^{([3](#))} [www.frontiersin.org](#)) ^{([4](#))} [www.frontiersin.org](#)). A recent *Koppula et al.* study developed an AI chatbot (GPT-3.5-based) for FDA label retrieval; it extracted and answered queries from drug labels with high semantic fidelity (most answers had 0.7–0.9 cosine similarity to ground truth, and ≥0.95 on concise sections) ^{([5](#))} [www.mdpi.com](#)) ^{([6](#))} [www.mdpi.com](#)). Another group used GPT-4 to **extract safety information** (**adverse reactions** and drug–drug interactions) from Structured Product Labels (SPLs) with performance matching or exceeding prior methods—without any task-specific fine-tuning ^{([7](#))} [link.springer.com](#)) ^{([2](#))} [link.springer.com](#)). These examples show that today’s LLMs can emulate many **information-extraction and summarization** roles in regulatory labeling.

Table 1 below summarizes key published efforts in this area:

Study / Example	Task	NLP Method	Key Findings / Outcome	Citation
Gray et al. (2023) ^(8) pmc.ncbi.nlm.nih.gov) ^(1) pmc.ncbi.nlm.nih.gov)	Classify free text into label sections (US PLR/SmPC)	Fine-tuned BERT (binary & multi-class)	Binary classification: 95–96% accuracy; Multi-class : up to 82% accuracy on PLR labels. Demonstrates auto-structuring of unformatted label text ^(1) pmc.ncbi.nlm.nih.gov).	Gray et al., <i>Chem. Res. Toxicol.</i> 2023
Koppula et al. (2025) ^(5) www.mdpi.com) ^(6) www.mdpi.com)	FDA label information retrieval (Q&A)	GPT-3.5 Turbo with document-grounded Q&A framework	High semantic similarity (0.7–0.9) with true answers across 10 breast cancer drug labels; ROUGE and embedding-similarity ≥0.95 on concise sections ^(5) www.mdpi.com) ^(6) www.mdpi.com). Demonstrates reliable label content retrieval.	Koppula et al., <i>AI & MDPI</i> 2025
Meyer et al. (2023) ^(3) www.frontiersin.org) ^(4) www.frontiersin.org)	Generate <i>Medication Guide</i> (patient info)	Extractive/Pointer-Generator abstractive model	A closed “heuristic alignment” strategy gave an ~7-point ROUGE improvement over naïve alignment ^(3) www.frontiersin.org) ^(4) www.frontiersin.org). Shows feasibility of AI-generated patient guides from label text.	Meyer et al., <i>Front. Pharmacol.</i> 2023
Zhou et al. (2025) ^(7) link.springer.com) ^(2) link.springer.com)	Extract ADRs and DDIs from labels (SPLs)	GPT-4 (LLM) vs. baseline extraction	GPT-4 met or exceeded prior state-of-the-art without extra training. Performance varied by section and complexity, but demonstrated flexible strong extraction performance ^(7) link.springer.com) ^(2) link.springer.com).	Zhou et al., <i>Drug Saf.</i> 2025
Shi et al. (2021) ^(9) www.frontiersin.org)	Extract label info to support guidances (e.g., food effect)	Pretrained BERT classifier	BERT outperformed prior methods in labeling paragraphs about “food effect” from FDA labels, enabling automated retrieval of relevant drug use info ^(9) www.frontiersin.org). Shows value of NLP pipelines for label content extraction.	Shi et al., <i>Front. Res. Metr. Anal.</i> 2021

Study / Example	Task	NLP Method	Key Findings / Outcome	Citation
Neyarapally <i>et al.</i> (2024) ⁽¹⁰⁾ pmc.ncbi.nlm.nih.gov ⁽¹¹⁾ pmc.ncbi.nlm.nih.gov	Detect changes in label adverse events (LabelComp)	BERT-based text analytics	LabelComp tool achieved F1 scores 0.80–0.94 on identifying new/changed adverse events between label versions ⁽¹⁰⁾ pmc.ncbi.nlm.nih.gov ⁽¹¹⁾ pmc.ncbi.nlm.nih.gov). Automates detection of label safety updates, important for pharmacovigilance . </current_article_content>	Neyarapally <i>et al.</i> , <i>Drug Saf.</i> 2024
Industry Case – Freyr (2023) ⁽¹²⁾ www.freyrsolutions.com ⁽¹³⁾ www.freyrsolutions.com	Label review/compliance monitoring (vendor solution)	Proprietary NLP + generative AI (unspecified)	Describes using GenAI for context-aware extraction, semantic comparison, automated content generation (e.g., MedDRA coding, version validation) ⁽¹²⁾ www.freyrsolutions.com ⁽¹³⁾ www.freyrsolutions.com). Illustrates emerging AI-led label management platforms.	Freyr Solutions (blog) 2023
Industry Case – IQVIA (2023) ⁽¹⁴⁾ www.americanpharmaceuticalreview.com ⁽¹⁵⁾ www.americanpharmaceuticalreview.com	Multi-source label content “search & compare”	NLP-based search hub with UI	Created a “labeling intelligence hub” aggregating FDA, EMA, and national labels. Enables custom search, label comparison, and direct access to original docs. Reportedly “ <i>streamlines new label development and updates</i> ” ⁽¹⁴⁾ www.americanpharmaceuticalreview.com ⁽¹⁵⁾ www.americanpharmaceuticalreview.com).	Reed, <i>Am. Pharm. Review</i> 2023

This body of evidence indicates that **NLP can significantly assist** regulatory labeling tasks: it can classify unformatted text into proper label sections, extract safety/usage information, and even auto-generate draft texts (with varying success). However, each capability currently has caveats. Section-classification and information retrieval have reached high performance, but *content generation* (summarization/drafting) is still exploratory. For example, the pointer-generator model for Medication Guides showed measurable improvement, but required careful engineering (BM25 alignments) to work well ⁽³⁾ www.frontiersin.org ⁽⁴⁾ www.frontiersin.org). Likewise, GPT-based chatbots rely on strict grounding to the original label to avoid “hallucinations” ⁽⁵⁾ www.mdpi.com ⁽⁶⁾ www.mdpi.com).

The conclusion is that **NLP can partially draft aspects of prescribing information but cannot yet replace expert human writers**. In practice, current uses are hybrid: AI assists with retrieval, initial drafts, formatting, and consistency checks, while humans perform final authoring, quality control, and regulatory compliance checks. The rest of this report explores these points in depth: first providing necessary background on regulatory labeling, then detailing specific AI/NLP applications, followed by analysis of data and case studies, and ending with implications and future outlook.

Introduction and Background

Regulatory Labeling and Prescribing Information

Pharmaceutical **prescribing information** (PI) is the official document that communicates a drug’s approved uses, dosages, contraindications, warnings, and other critical medical details to healthcare professionals. In the U.S., this is commonly referred to as the “label” or **Package Insert (PI)**. In Europe and many other regions, equivalent content is found in the **Summary of Product Characteristics (SmPC)** (for professionals) and the patient-oriented **Package Leaflet** (www.ema.europa.eu). While formatting and exact sections vary by region, all such labeling documents form part of a drug’s marketing authorization and are legally binding. They reflect a vast compilation of scientific data: results of clinical trials, pharmacology, pharmacokinetics, safety findings, and more.

Writing and updating prescribing information is a highly specialized, labor-intensive process. Medical writers and regulatory experts compile data from various sources (clinical study reports, investigator brochures, preclinical studies, pharmacovigilance databases, etc.) and adhere to detailed regulatory guidelines. For example, in the U.S. the **Physician**

Labeling Rule (PLR) defines the format of drug labeling sections (^[16] [pmc.ncbi.nlm.nih.gov](https://pubmed.ncbi.nlm.nih.gov/)), and FDA expects submissions in the **Structured Product Labeling (SPL)** XML format (^[17] www.fda.gov). Similarly, the EMA's guidelines (including the QRD templates) specify sections such as Indications, Posology, Contraindications, etc. The EMA notes that EU product information includes both the SmPC for professionals and a patient leaflet, together forming the official "product information" (www.ema.europa.eu). These structures aim to ensure clarity and uniformity; indeed, since 2005 the FDA has required labels to use SPL, which pre-annotates text into sections (^[18] [pmc.ncbi.nlm.nih.gov](https://pubmed.ncbi.nlm.nih.gov/)). Nonetheless, in practice many documents are not perfectly formatted and content can spill across sections. Reviewers often spend considerable time locating specific data, and any inefficiency can delay approvals.

Moreover, labels must be kept up to date. New safety signals, expanded indications, or regulatory changes (e.g., new labeling rules) require timely edits everywhere applicable. For instance, generic drug makers must continuously monitor innovator labels globally and mirror any changes (^[19] www.freyrsolutions.com). The sheer volume of labeling documents is enormous – on the order of 130,000 in public FDA/EMA repositories (^[20] www.freyrsolutions.com) – creating ongoing challenges for multi-market consistency.

NLP and Generative AI in Life Sciences

Natural Language Processing (NLP) is a field of AI that enables computers to interpret, manipulate, and generate human language. In life sciences and healthcare, NLP is widely used for tasks like biomedical text mining, literature curation, adverse event detection from reports, and more (^[21] www.americanpharmaceuticalreview.com). Recent advances in **Large Language Models (LLMs)** – examples include GPT-3/4, LLaMA, and PaLM – have dramatically improved machines' ability to generate fluent text. These models are pretrained on vast corpora of text and can be fine-tuned or prompted for specific tasks.

In industry, NLP has matured from simple keyword search to more sophisticated document analysis. Analysts note that moving from a document-centric to data-centric paradigm is a key trend in regulatory affairs (^[22] www.americanpharmaceuticalreview.com). NLP tools can ingest massive collections of regulatory documents, extract structured data (e.g., drug names, doses, safety terms), and present actionable summaries. Already, speech assistants (Siri/Alexa) and GPT chatbots illustrate NLP's broad capabilities (^[23] www.americanpharmaceuticalreview.com). As one review notes, modern NLP can analyze unlimited text "without fatigue or bias", extracting facts and summarizing content far beyond basic search engines (^[24] www.americanpharmaceuticalreview.com).

Within pharmaceutical regulatory operations, NLP applications include:

- **Information extraction:** Identifying entities (e.g., drug names, indications, adverse events) and relationships from unstructured text (e.g., labeling or literature) (^[9] www.frontiersin.org) (^[7] link.springer.com).
- **Classification/Tagging:** Assigning text to predefined categories or label sections (^[1] [pmc.ncbi.nlm.nih.gov](https://pubmed.ncbi.nlm.nih.gov/)).
- **Summarization:** Generating concise texts (e.g., summarizing clinical study results or drafting patient leaflets from technical data) (^[3] www.frontiersin.org).
- **Semantic search/Q&A:** Allowing users to query content in natural language and get relevant answers (as in the FDA chatbot maps queries to label content) (^[5] www.mdpi.com).
- **Assisted authoring:** Auto-formatting documents, generating first draft text from source materials, or ensuring consistency across sections.

Generative AI (LLMs) extends NLP by *creating* text that did not explicitly appear in the input. For labeling, generative models could in principle produce entire draft sections from summaries of data. However, this introduces challenges of accuracy: models may "hallucinate" facts not supported by the source. In highly regulated content like drug labels, such errors would be unacceptable. For this reason, practical approaches often combine LLMs with retrieval of authoritative content (so-called "retrieval-augmented generation" or RAG) and strict constraints to prevent fabrication (^[25] pharmacystandards.org). We will see that state-of-the-art implementations for regulatory writing carefully control generative output using validated references and review processes.

Scope of This Report

This report examines "Can NLP (specifically generative AI) draft your prescribing information?". We first outline the **traditional label authoring process** and the relevant regulatory framework (PLR, SPL, ICH, etc.). Then, we survey **NLP**

technologies and how they apply to labeling tasks. We analyze published research and industry case studies (as in Table 1), noting successes and limitations. We also discuss **data/experiments** (e.g., accuracy metrics, evaluation of language models) that shed light on feasibility.

Importantly, we consider **implementation issues**: regulatory requirements (GxP, audit trails, validation), human oversight, and the risk/benefit calculus. For case studies, we cover both academic projects (BERT classifiers, summarization models) and real-world pilots (AI label hubs, LLM-assisted authoring). Lastly, we assess **future directions**: improvements in LLMs, potential regulatory acceptance, and how NLP might reshape pharmacovigilance, label harmonization, and patient safety. The conclusion synthesizes these findings into guidance on the realistic role of AI in regulatory labeling.

Traditional Labeling Workflow and Challenges

Regulatory Requirements and Structure

Drug labeling is governed by stringent rules. In the US, regulations like the **Physician Labeling Rule (PLR)** and various *Guidance for Industry* documents specify the content and layout of the Prescribing Information (^[16] [pmc.ncbi.nlm.nih.gov](https://pubmed.ncbi.nlm.nih.gov/)). Key sections include Indications & Usage, Dosage & Administration, Contraindications, Warnings & Precautions, Adverse Reactions, Drug Interactions, Use in Specific Populations, and detailed Pharmacology sections. The FDA requires submissions in **Structured Product Labeling (SPL)** format (^[17] www.fda.gov), an HL7-based XML scheme that encodes sections and preserves machine-readable content. SPL facilitates electronic handling: for example, FDA's DailyMed and other label repositories use SPL to power search and display. However, the underlying text still originates from human authors.

In Europe, the **SmPC Guidelines** and EMA templates dictate label content. The EMA explains that the *product information* consists of (a) the professional-targeted SmPC and (b) a patient package leaflet (www.ema.europa.eu), plus packaging labeling. An SmPC typically mirrors the PI sections (Indications, Posology, etc.) but may use slightly different terminology. The **Quality Review of Documents (QRD)** template standardizes formatting (e.g., in Section 4.2: Posology, Section 4.3: Contraindications etc.). The EU periodically updates these templates and requires translations into member languages. For example, recent EMA updates allow combining all strengths of a dosage form into one SmPC (www.ema.europa.eu), reflecting ongoing harmonization efforts.

Given these frameworks, labeling authors compile data into a prescribed outline. Manuscripts are meticulously reviewed by safety experts, clinicians, and regulatory affairs. The review process is multi-layered, involving internal (company) reviewers and ultimately health authority reviewers. FDA reviewers may use the structured sections to route content to different offices (e.g., safety vs efficacy) (^[8] [pmc.ncbi.nlm.nih.gov](https://pubmed.ncbi.nlm.nih.gov/)). GPA for chemical compounds.

The key point: labeling is a *data-driven narrative*. It synthesizes huge amounts of data into a cohesive document. Automating this requires not only language fluency but domain understanding and compliance with each guidance. Any automated system must align with regulatory structure (e.g., ensure Indications cover approved uses only) and evidence.

Challenges in Manual Labeling

Many companies view labeling as a bottleneck. Some of the challenges include:

- **Volume and Repetition:** For globally-marketed drugs, each local affiliate may adapt a label for regional variations. Keeping sync across USPI, EU SmPC, Japan package inserts, etc., is onerous (^[20] www.freyrsolutions.com). Generic manufacturers must also continuously audit innovator labels worldwide to mirror changes (^[26] www.freyrsolutions.com). The cumulative effort is massive.
- **Data Complexity:** Input data streams into labeling include supplier data (CMC reports), nonclinical and clinical study reports, published literature, meta-analyses, post-market surveillance reports (e.g., FAERS), and more. Summarizing thousands of pages of data into a section requires both domain knowledge and writing skill.
- **Changing Regulations:** Regulatory guidelines evolve. For instance, the 2010 FDA PLR modernization changed how pregnancy and lactation sections are presented. In 2005, the move to SPL changed the format. In 2023 and beyond, agencies continue to refine label structure and content requirements. Writers must update existing labels or retag information accordingly (^[18] [pmc.ncbi.nlm.nih.gov](https://pubmed.ncbi.nlm.nih.gov/)).

- **Quality and Consistency:** Human writers aim for consistency (tone, style) and accuracy. Given the fatigue and oversight issues, minor inconsistencies often slip through. Automated checks can help, but they generally operate after the fact (e.g. spell-check, style-checkers). Ongoing consistency across sections and across therapeutic class can be hard to enforce manually.
- **Speed to Market:** Label preparation can be on the critical path for drug approval or life-cycle updates. As more data arrives (e.g., late-breaking trial results), writing must be done under time pressure, sometimes causing delays. Any tool that speeds up this process safely would be prized.

Given these challenges, companies and vendors have long sought to automate parts of the process (e.g., using templates, macros, or proprietary tools), but until recently tasks like “draft an entire new section from a clinical report” were thought to require human expertise. The question now is whether modern NLP can fulfill some of the tasks that medical writers do.

NLP Techniques and Tools in Regulatory Labeling

This section surveys the technical landscape: what NLP methods are relevant for labeling tasks, and which tools or research prototypes illustrate their use. We group tasks into **(1) Information Classification**, **(2) Extraction & Summarization**, **(3) Content Generation & Question Answering**, and **(4) Label Maintenance & Comparison**.

1. Classification and Organization of Label Text

Problem: Even before drafting new content, existing documents may be poorly structured. Automatic classification can reorganize or tag free text according to regulatory sections.

- **Text Classification:** Gray *et al.* demonstrated that BERT-based classifiers can assign sentences or passages from labeling documents to the correct section (e.g. splitting out “Warnings” content) (^[27] [pmc.ncbi.nlm.nih.gov](https://pubmed.ncbi.nlm.nih.gov/35111111/)) (^[1] [pmc.ncbi.nlm.nih.gov](https://pubmed.ncbi.nlm.nih.gov/35111111/)). They trained on thousands of FDA labels in PLR format, then tested on FDs and SmPCs. The model achieved ~96% accuracy in distinguishing PLR vs non-PLR labeling sections (binary) and ~82% accuracy when classifying among all PLR categories (^[28] [pmc.ncbi.nlm.nih.gov](https://pubmed.ncbi.nlm.nih.gov/35111111/)). This suggests that even if label texts are shuffled or mis-filed, NLP can relocate content to the proper standardized section. Such classification assists consistency and review (e.g., FDA reviewers can see if safety data accidentally appears outside Warnings).
- **Metadata Tagging:** Beyond sections, NLP can auto-tag label content with relevant metadata. For example, NLP engines can identify trial names, patient population descriptors, or MedDRA codes for adverse events within the text. Although we lack a specific citation here, industry systems (e.g. Freyr’s platform) emphasize auto-tagging. Freyr claims that generative AI enables “context-aware extraction” and “semantic comparison” of label elements (^[12] www.freyrsolutions.com). For instance, an NLP system might scan a draft label and highlight the known safe dosage in animals versus in humans, or flag if an adverse event is described without a corresponding seriousness classification.
- **Semantic Search and Indexing:** Turning unstructured text into searchable data is a core NLP use. The FDA provides **FDALabel** (and DailyMed) as searchable label repositories. On the industry side, IQVIA described creating a “labeling intelligence hub” that integrates FDA & EMA labels into a common database, enabling teams to search by disease terms, contraindications, etc. (^[14] www.americanpharmaceuticalreview.com). Behind this likely lies NLP that normalizes synonyms and extracts term mentions to index the documents. This aligns with the American Pharma Rev. case: NLP value in **regulatory labeling** includes “access to drug labels” to find references for diseases, contraindications, adverse events, etc. (^[29] www.americanpharmaceuticalreview.com). Thus, NLP can bridge across redundant phrases and multiple languages to find all relevant label references.

These classification and search capabilities do not “write” content but form a critical base. They ensure that when generative methods are later applied, the input and output are properly organized. For example, an LLM that generates a Drug Interaction paragraph would only perform well if the underlying text is correctly identified and channeled to that section.

2. Extraction and Summarization

Problem: Regulatory labels often need distilled summaries of clinical trial outcomes, safety data, or pharmacology. Manual summarization is tedious and error-prone. NLP can automate extraction of key facts and even generate concise text.



- **Named Entity Recognition (NER):** Identifying specific entities (e.g., drug names, medical conditions, test results) in free text. In the labeling context, NER can pull out all drug names (including metabolites) mentioned in Interaction sections, or find all mentions of "neutropenia". While no one study in our batch explicitly covers NER in labels, this is a mature area in biomedical NLP. LLMs like GPT inherently perform a kind of NER by virtue of token probabilities; specialized models (e.g. BioBERT) are also used. The *Shi et al.* pipeline likely involved NER as part of extracting "food effect" info ([19] www.frontiersin.org). NIH daily.
- **Template Filling:** Regulatory content often uses semi-structured sentences. Systems can fill templates: e.g. "X% of subjects experienced Y adverse event" from trial results. This requires numeric and semantic extraction. Work like Zhou *et al.* (2025) essentially tackled this: GPT-4 extracted lists of adverse reactions with associated info from SPL text ([7] link.springer.com) ([2] link.springer.com). While Zhou focused on generation of structured output (tables of ADRs/DDIs), the principle applies: an LLM can read narratives and output tabulated safety data, a task previously done by team members manually. The result was that GPT-4 could perform this extraction **zero-shot** (without new training) at levels matching fine-tuned models on ADR extraction ([7] link.springer.com).
- **Abstractive Summarization:** Instead of just pulling facts, this creates new phrasing. The Meyer *et al.* study is the clearest example: they aimed to generate **Medication Guides** (patient-directed leaflets) from more technical label sections ([4] www.frontiersin.org) ([3] www.frontiersin.org). They built a specialized model that learned to "summarize" professional text into a simpler guide format. While this was an experimental pointer-generator network (not an off-the-shelf LLM), it shows that end-to-end summarization is being attempted. The controlled task (focusing on labels that include an official Medication Guide section) allowed them to train the model on aligned pairs of source (technical) and target (patient) text. Their results showed significant ROUGE improvements with careful alignment strategy, suggesting an abstractive model can capture much of the meaning needed for a simpler guide ([3] www.frontiersin.org).

Today, one could imagine using a GPT-4-like model for such a job: prompt it with the source labeling sections and ask for a layperson summary. However, the challenge is reliability: GPT might introduce phrasing that is scientifically vague or non-compliant. That is why Meyer's approach remained mostly a research proof-of-concept, and why regulatory groups might be cautious about fully automated summarization.

- **Hybrid Approaches (RAG):** A modern solution is retrieval-augmented generation. For example, the Council on Pharmacy (CAIDRA) example describes giving an LLM the *exact current label text and new updates* (e.g., a Core Clinical Data Sheet changes) and then instructing it to update the label ([25] pharmacystandards.org). The LLM can read the official sources (like approved risk manuscripts) and rewrite sections accordingly. They constrain the output strictly (e.g., "cite CCDS section for each change" ([25] pharmacystandards.org), output in two-column diff tables) to maintain auditability. This technique is being actively tested. It essentially combines extraction (find differences) and generation (draft new text) in one pipeline.

As a real-world parallel, one might note that most published label writing is not done by pure abstract summarization from scratch: templates, bullet points, and structured tables are often used. NLP could help fill in these templates. For instance, an LLM could be prompted: "Given these adverse event narratives, generate a bullet-list summary highlighting Serious AEs and their frequency."

Summary: NLP can dramatically accelerate the **data synthesis** aspects of label writing – from extracting lists of adverse events to summarizing trial efficacy. However, these outputs will always need expert validation. Regulatory safety standards mandate that every statement in the label be evidence-backed. Thus, systems must ensure that computed summaries align exactly with sourced data, a nontrivial requirement.

3. Content Generation and Q&A

Problem: Can NLP *write* new label prose or answer complex questions about label content?

- **Generative Drafting:** The most ambitious concept is to have an AI *draft* label sections. A fully generative approach would involve giving a model all relevant data (clinical study tables, safety narratives, company Core Data Sheet) and expecting it to output review-ready text. In practice, companies are exploring semi-automated drafting.

The Gramener webinar blog asserts that generative AI can "automate the creation, review, and submission of complex regulatory documents" ([30] blog.gramener.com). It claims GenAI can reduce manual effort and produce labeling content that meets regulatory standards ([30] blog.gramener.com). However, these claims should be tempered: they are promotional and based on high-level scenarios. No study has demonstrated a validated fully-AI-drafted new prescribing information. The main published example we have is educational: the CAIDRA module (a certification exam-style guide) shows how one might use an LLM to update a label ([25] pharmacystandards.org). In that example, the LLM is given:

- The current label version.

- A new CCDS (Core Clinical Data Sheet) listing updates.
- Standard operating procedures (SOP) guidelines on labeling.

The prompt instructs the model to “identify all changes” and “draft new text ... citing [the CCDS]” ([25] [pharmacystandards.org](https://www.pharmacystandards.org)). The output is constrained as a table with “Current vs. Proposed Text”. This structured task ensures the model only rewrites specified parts (here, “Warnings & Precautions”) and cites the official sections. This is arguably the most realistic blueprint for draft generation: tightly controlled, evidence-anchored, and requiring human review of every change. It indicates that with human oversight (who must formulate the right prompt and check results), an LLM can produce draft label edits.

- **LLM QA Tools:** Rather than produce new writing, another use-case is to interact with an existing label via AI. The Koppula et al. chatbot ([5] www.mdpi.com) illustrates this: an LLM was used to answer questions about label content (e.g., “What is the dosage in patients over 65?”), by extraction. This is not drafting the label, but it shows how LLMs can interpret and restate labeling info. Similarly, industry platforms (like Freyr’s Freya) provide “conversational Q&A” interfaces where regulatory staff can ask about label compliance or content. This indirect application helps users generate wording (e.g., retrieving relevant phrasing or data from the label) but still keeps humans authoring the final text.
- **Prompting Strategies:** If one attempts generative drafting, prompt engineering is crucial. For instance, one might feed the model sections of the clinical study report and ask “Write the ‘Adverse Reactions’ section suitable for a prescribing label.” Early experiments (outside published lit) suggest GPT-3.5/4 can mimic the style reasonably, but often invents plausible-sounding incorrect data unless heavily controlled. Therefore, best practices include:
 - **Chain-of-thought prompting:** Instructing the model step-by-step on constraints it must follow (e.g., “List critical facts from the data, then compose a section”).
 - **Few-shot examples:** Providing examples of high-quality label text to guide tone.
 - **Fact-checking:** Immediately verifying each generated statement against source documents. Some systems propose integrated fact-checkers or post-generation verification steps.

Without citations, generative output is risky in this domain. The CAIDRA constraints (“Do not invent information not in the source” ([31] [pharmacystandards.org](https://www.pharmacystandards.org))) highlight the need for explicit rules.

4. Label Maintenance, Versioning, and Harmonization

Even after initial drafting, labels need continuous maintenance. NLP can help **compare, spot changes, and ensure consistency**.

- **Label Version Comparison:** Neyarapally et al. (2024) developed **LabelComp** to automatically identify changes in adverse event terms between new and old label versions ([10] [pmc.ncbi.nlm.nih.gov](https://pubmed.ncbi.nlm.nih.gov)). Using a BERT-based model, the tool achieved high F1 scores (0.80–0.94) on detecting added/removed ADRs ([10] [pmc.ncbi.nlm.nih.gov](https://pubmed.ncbi.nlm.nih.gov)). This automates what reviewers currently do manually: checking that updates in safety databases or literature (like a new black-box warning) are reflected in the PI. While LabelComp is not a generative writing tool, it directly influences drafting because it flags what *needs* to be drafted.
- **Cross-Label Harmonization:** Generic manufacturers routinely use NLP to keep their labels aligned with reference products across jurisdictions. The Freyr blog points out that NLP can “extract and compare critical label elements across versions and jurisdictions” ([13] www.freyrsolutions.com). In practice, a generic’s RA team might use an AI tool to monitor the innovator’s FDA drug label; on detecting a label change, the tool could highlight which sections need updating. The same applies in reverse – originators may use NLP to ensure a simultaneous submission’s labels (e.g. in EU and US) stay aligned with the global core data.
- **Text Mining for Compliance Audits:** Regulatory teams can use NLP to check labels against regulatory rulebooks. For instance, a program could scan labels to verify that all required sections (per FDA or ICH guidelines) are present and populated, or that safety communications (e.g., “boxed warnings”) use mandated phrasing from guidance documents. These use-cases border on automated quality control rather than drafting, but are relevant: faster error-checking means fewer manual rewrites later.

Table 2: Comparison of Key Components in US vs. EU Labeling (as context for NLP tasks)

Feature	United States (FDA)	European Union (EMA)	Source/Citation
Label Name	Prescribing Information (PI, USPI)	Summary of Product Characteristics (SmPC)	[37+L32-L39], [51+L8-L14]
Patient Leaflet	Medication Guide (MG) or Patient Information (PI) leaflet (if required)	Package Leaflet (PL) for patients	[51+L8-L14]

Feature	United States (FDA)	European Union (EMA)	Source/Citation
Format Standard	SPL (XML format) required for submissions (^[17] www.fda.gov)	XML & human-readable PDFs via EMA templates (QRD)	[28*L5-L7], [51*L8-L14]
Key Sections	INDICATIONS & USAGE; DOSAGE & ADMIN; CONTRAINDICATIONS; WARNINGS & PRECAUTIONS; ADVERSE REACTIONS; DRUG INTERACTIONS; USE IN SPECIFIC POPULATIONS; etc.	Therapeutic Indications; Posology (dosage); Contraindications; Special warnings; Adverse reactions; Interactions; etc. Similar content, slightly different headings.	[37*L32-L39], [51*L8-L14] (implied)
Regulatory Basis	Code of Federal Regulations (21 CFR Part 201) + FDA Guidance (PLR, labeling guidances); valid upon FDA approval.	EU Directive 2001/83/EC, EMA's QRD templates and ICH guidelines; valid in all EU states upon EMA/MHRA approval.	[37*L55-L63], [51*L8-L14]
Update Process	Supplements to NDA/BLA (e.g., Improvement or Changes Being Effected) submitted to FDA; new SPL issued. The updated PI is published (e.g., on DailyMed) when approved.	Variations to marketing authorization (e.g. Type II, line extensions) affect SmPC; EMA coordinates updates. Updates appear in EPAR documents.	[23*L27-L36], [51*L21-L27]

Table 2 illustrates the structural similarities and differences. NLP solutions often have to be adapted for each context (terminology, language, sections). For example, an AI trained on FDA labels might not immediately handle text in French from a Swiss label. Therefore, international operations often run separate NLP pipelines per region, or use multilingual models.

Evidence from Data and Case Studies

We now dive into specific examples—both published research and real-world implementations—to see **how NLP has been applied to labeling**, what results have been achieved, and what lessons emerge.

Automated Section Classification (Gray *et al.*, 2023)

Gray and colleagues at NIH used BERT to tackle a practical problem: many drug labels submitted to regulators are unstructured or have inconsistent formatting (^[8] [pmc.ncbi.nlm.nih.gov](https://pubmed.ncbi.nlm.nih.gov/)). They gathered ~46,000 FDA labeling documents (17 million sentences) across various eras and countries (^[32] [pmc.ncbi.nlm.nih.gov](https://pubmed.ncbi.nlm.nih.gov/)). They focused on mapping *free-text* (unstructured) labels into the FDA's standard PLR sections ("binary" task: PLR vs not-PLR; "multiclass" task: identify the specific section).

The results were impressive. Their fine-tuned BERT model achieved **96% accuracy** on the binary task (separating text that belongs inside an official PLR label vs other) and **88% accuracy** on multiclass (when grouping text into one of many sections) on their testing set (^[27] [pmc.ncbi.nlm.nih.gov](https://pubmed.ncbi.nlm.nih.gov/)) (^[33] [pmc.ncbi.nlm.nih.gov](https://pubmed.ncbi.nlm.nih.gov/)). On new, real-world documents, performance was slightly lower but still strong (binary accuracies >88% across FDA, non-FDA, and EU SmPC texts) (^[1] [pmc.ncbi.nlm.nih.gov](https://pubmed.ncbi.nlm.nih.gov/)).

Key takeaway: Transformer-based NLP can reliably *restructure* misfiled content. For example, if a safety paragraph erroneously appears under "General Description" instead of "Warnings", the model catches it. This supports the idea that even legacy PDFs with missing tags can be auto-normalized. In practical workflow, such a classifier could pre-process any text block to assign it to a section, serving as an "intelligent formatter." This matching of sorts is a first step before generation: once all content is properly labeled, a human writer can fill blanks more coherently.

In addition, Gray *et al.* conducted interpretability analysis with SHAP values and compared to older ML methods, finding BERT both more accurate and more explainable for their task (^[34] [pmc.ncbi.nlm.nih.gov](https://pubmed.ncbi.nlm.nih.gov/)). They demonstrate that such advances in NLP are mature enough for regulatory science usage. Indeed, the authors suggest that their approach could accelerate the review process by automatically structuring even poorly formatted documents (^[27] [pmc.ncbi.nlm.nih.gov](https://pubmed.ncbi.nlm.nih.gov/)). Thus, internal data suggests: *yes, NLP can autonomously organize label content in line with regulatory structure* (^[1] [pmc.ncbi.nlm.nih.gov](https://pubmed.ncbi.nlm.nih.gov/)).

Drug Safety Information Extraction (Zhou *et al.*, 2025)

The "Leveraging LLMs in Extracting Drug Safety Information" study (Zhou *et al.*) focused on *generative LLMs* (GPT-4, LLaMA, Mixtral) *extracting* adverse reaction (AR) and drug-drug interaction (DDI) data from structured labels (^[7] link.springer.com). While not directly drafting the label, the study addresses automating the capture of critical label information for research use.

Methods: They compared GPT-4 and other LLMs to previous machine learning baselines. Without fine-tuning, they prompted the models to list all adverse reactions mentioned in an SPL (taking advantage of LLM knowledge and chain-of-thought reasoning). They also tested LLMs on a new task: list all drug names in the DDI section, with no new training data.

Findings: Impressively, GPT-4 performed as well or better than domain-specific models on ADR extraction ([7] link.springer.com). Its F1-scores matched or exceeded SOTA, even though it hadn't been specifically trained on ADR tasks. GPT-4's performance varied depending on how common or complex the terms were. It was less accurate on rare ADRs, but overall its flexibility was notable. For the DDI (drug interaction) task, GPT-4 succeeded "zero-shot": it pulled out drug names without any new tuning, proving robustness.

Implication: Though this study doesn't generate narrative text, it demonstrates a key building block: GPT-4 "knows" enough about medical terminology to identify safety info in labels reliably. If an LLM can accurately recite all the adverse events from sections, it is effectively understanding (or proficiently guessing) content across sections. In the context of drafting, this suggests that GPT-4 could similarly attempt to summarize or translate these events into prose. Moreover, since drug safety is paramount, such an LLM could assist pharmacovigilance teams by highlighting key facts. For end-users (patients or doctors), an LLM-based system might answer questions like "what are the most common side effects?" without needing manual lookups.

In summary, this data-backed study shows that **state-of-the-art LLMs have the textual comprehension needed for regulatory content**. If GPT-4 can match carefully-trained models on extraction tasks, then more specialized tasks (like summarizing trial results into Precautions) might be feasible with larger models or better prompts. It adds weight to considering generative AI pipelines for labeling – at least for factual extraction steps – although the study authors caution that context and term complexity affect output quality ([7] link.springer.com).

Conversational Chatbots for Label Information (Koppula et al., 2025)

Koppula and colleagues developed an **AI-powered chatbot** to retrieve information from FDA drug labels. This directly speaks to user's question-answering rather than generating new label content, but has parallels to drafting assistance (e.g. answering queries in natural language).

How it works: Users upload a PDF of an FDA label (e.g. DailyMed file). A Python pipeline segments the label into sections (using PyMuPDF), and then the chatbot uses GPT-3.5-turbo in a controlled Q&A setting: GPT only sees the relevant extracted text for the query ([5] www.mdpi.com). For example, if the user asks "What are the main warnings for Drug X?", the system finds the "Warnings" section in the PDF and feeds that to GPT, so the answer is grounded. This prevents hallucination by design.

Results: They tested on 10 breast cancer drug labels, computing semantic similarity between chatbot answers and actual label text. The scores ranged largely 0.7–0.9 on a 0–1 scale (with 1 meaning identical content) ([35] www.mdpi.com), indicating the answers were generally very faithful. On shorter, focused sections the score reached ≥ 0.95 . ROUGE scores (which compare overlapping N-gram recall) were also high, confirming that the answers reused evidence phrases. The authors claim their approach "confirms strong semantic and textual alignment" ([6] www.mdpi.com). They even compared with newer models (GPT-5-chat, NotebookLM) and found similar performance, implying the method is robust across LLMs.

Interpretation: This work shows that a generative model can be carefully constrained to answer questions *about* label content with high accuracy. Going further, one might ask: Could the same architecture support draft writing? Possibly in a limited way. For example, given a section source, instead of a user query, one could prompt, "Summarize the key points of this section in plain language." The main lesson is the power of grounding: *if an LLM's output is strictly limited to trustworthy text segments (through retrieval and segmentation), it can produce very accurate narrative answers*. In labeling practice, this suggests hybrid tools: e.g. a medical writer could query an AI draft assistant with natural questions ("Summarize the efficacy findings") and trust that the answer is rooted in the official label text.

One limitation noted by the authors is scope. Their work was limited to 10 labels in one therapeutic area, which means real-world phrasing could be more varied. They also mention future work to expand to diverse areas. Nonetheless, as a proof-of-concept it is promising: generative AI can act as an *interactive front-end* to labeling content, possibly speeding up authoring by quickly collecting relevant information. It does **not** yet draft new claims or sections beyond what's in the label, but it could serve as a support tool for writers.



Automated Label Comparison (LabelComp, Neyarapally *et al.*, 2024)

LabelComp, an AI tool developed by FDA researchers (Neyarapally *et al.*), specifically addresses the regulatory need to **document label changes** between versions. This is crucial after drug approval: when safety data emerge, the label often needs supplementing.

Function: LabelComp ingests a pair of labeling documents (old vs new version) and highlights changes in the Adverse Reactions section. Using text analytics and a BERT model trained to recognize change pairs (old vs new AE terms), it flags which adverse events were added, deleted, or altered (^[10] [pmc.ncbi.nlm.nih.gov](https://pubmed.ncbi.nlm.nih.gov/)).

Validation: On 87 FDA label pairs, LabelComp achieved F1 scores between 0.795 and 0.936 for overall performance (^[10] [pmc.ncbi.nlm.nih.gov](https://pubmed.ncbi.nlm.nih.gov/)). In other words, its precision and recall for identifying correct AE changes were very high. This means LabelComp could correctly notify reviewers (or writers) precisely which terms in the AR section changed (e.g., “New term added: febrile neutropenia, incidence increased from 5% to 10%,” etc.).

Implications for drafting: While LabelComp does not generate new text from scratch, it effectively automates label maintenance. Regulatory affairs teams normally do this by hand, reading old and new PDF and listing changes. With LabelComp, a large portion of this task is done by AI, leaving only interpretation and final writing. In the bigger picture, this demonstrates another use-case: AI can take over the *update identification* step, which improves authoring speed. A future label-writing pipeline could be: run LabelComp → get list of updated safety facts → feed those facts into a summarization tool or template to produce the new Safety section paragraph.

Industrial and Regulatory Perspectives

Beyond academic work, industry articles and regulatory announcements give insight into how NLP is being envisioned:

- The American Pharmaceutical Review (IQVIA) article “Advancing Regulatory Compliance with NLP” provides both rationale and examples. Its author, Jane Reed (Regulatory Quality Director at IQVIA), notes that regulatory affairs is traditionally manual and costly (time/money wasted on repetitive tasks) (^[36] www.americanpharmaceuticalreview.com). The push is to move from document-driven to data-driven practices (^[22] www.americanpharmaceuticalreview.com). She explicitly states that NLP can “discover key data within regulatory documents” and convert it to structured info for downstream uses (reporting, labeling, etc.) (^[37] www.americanpharmaceuticalreview.com). NLP’s role includes standardizing data attributes for master data management – in other words, feeding clean content into modern content-management systems.

Worth highlighting is a pharmaceutical company case in the same article. It describes (in “NLP in Action”) how a **cosmetic labeling intelligence hub** was built: it ingested FDA, EMA, and local country labeling data into one NLP-powered search interface (^[14] www.americanpharmaceuticalreview.com). Users can search labels by specific terms, compare labels side-by-side, and directly view the source documents. This solution “streamlines developing new labels, updating existing ones, and expediting regulatory approval” (^[38] www.americanpharmaceuticalreview.com). In effect, this shows that leading companies are already using NLP to simplify label authoring tasks (especially multi-source cross-comparison).

- On the vendor side, Freyr’s Freya platform is highlighted. Freyr claims generative AI provides “context-aware extraction, semantic comparison, automated content generation” for labeling (^[12] www.freyrsolutions.com). They list key capabilities: extracting key elements across versions/jurisdictions, analyzing large doc repositories for compliance gaps, auto-managing label updates, MedDRA coding, validation, translation, etc. (^[13] www.freyrsolutions.com). These are ambitious claims, but at minimum they confirm that regulatory tech vendors see generative NLP as applicable across the labeling workflow (not just drafting). The numbers cited (130,000+ labeling docs in public archives (^[20] www.freyrsolutions.com)) underscore the scale of data these systems must handle.
- **Regulator Initiatives:** The FDA itself is actively pondering AI use. A January 2025 FDA press release introduced new draft guidance on using AI to support drug submissions (^[39] www.fda.gov). While this guidance covers AI as a source of analytical or evidence data (e.g., modeling patient outcomes), it signals agency attention to AI in all aspects of drug development. The release notes that “FDA has reviewed more than 500 submissions with AI components” since 2016 (^[40] www.fda.gov) (^[41] www.fda.gov), reflecting that sponsors are already incorporating AI tools. Although this draft guidance is primarily about AI as evidence, the preamble (citing AI uses like data analysis) indicates a broad acceptance of AI tools in submissions. Hence, though not label-specific, it suggests regulators will require clear **validation and context-of-use** definitions for any AI used in the process (^[42] www.fda.gov).

- **Professional Guidance:** The ISPE publication on ChatGPT in pharma notes that LLMs like ChatGPT can indeed generate drug information summaries, but *must not be taken as medical advice* and outputs must be validated by experts ([43] ispe.org). This aligns with the general caution that AI output should be reviewed by qualified personnel. The CAIDRA training module (Section 4.1) goes further: it acknowledges that LLMs could potentially eliminate the “blank page” of medical writing, but emphasizes GxP compliance needs ([44] pharmacystandards.org) ([45] pharmacystandards.org). They report that a pilot in which writers copied text from a public AI tool was *non-compliant* because it broke audit-trail requirements ([45] pharmacystandards.org). The lesson is stark: even if AI can produce a first draft, the usage workflow must preserve traceability (e.g., by using internal tools that log outputs and sources).

These perspectives collectively indicate that **the industry sees strong potential but also requires caution**. Companies and regulators recognize that NLP can unburden repetitive tasks and harmonize content. But they demand that AI be integrated into approved, auditable processes.

Data and Evidence Analysis

The examples above include quantitative metrics that shed light on feasibility. We recap key data points:

- **Accuracy/AUC:** Gray *et al.* reported 95–96% accuracy for PLR vs. non-PLR classification, and ~82% accuracy on multiclass section tagging ([28] pmc.ncbi.nlm.nih.gov). LabelComp achieved F1 ~0.80–0.94 on change detection ([10] pmc.ncbi.nlm.nih.gov). The chatbot’s semantic-match scores of 0.7–0.9 (≥ 0.95 on terse sections) ([6] www.mdpi.com) and high ROUGE indicate robust fidelity. These high scores (mostly >0.8 ROC/AUC or F1) are encouraging: they suggest these NLP tasks approach near-human performance.
- **Error Cases:** Some models show performance drops on out-of-distribution data. For instance, Gray’s multiclass model got only 68% on SmPC test data ([28] pmc.ncbi.nlm.nih.gov). This reveals the challenge: models trained only on US-style labels don’t generalize perfectly to differently structured ones (like EU countries). Zhou’s work noted that the “complexity of the AR term” affected extraction quality ([7] link.springer.com) (rare, multi-word terms are harder). Thus, a realistic NLP solution must either be fine-tuned regionally or use diverse training sets.
- **Efficiency Gains:** The LinkedIn/Gramener sources (though non-peer) cite dramatic speed-ups: McKinsey/BCG estimated 30–80% time reduction for drafting narratives ([46] www.linkedin.com), and internal user reports claim 90% faster safety narratives with AI ([47] www.linkedin.com). While proprietary and promotional, it aligns with the general case study from CAIDRA that LLM assistance can cut writing time by ~30–40% ([45] pharmacystandards.org) (at least until compliance issues arose). We should treat these claims as aspirational benchmarks. Nevertheless, if cognitive tasks can be half-automated, a well-validated pipeline might reduce label authoring time substantially.
- **Scalability/Corpora:** Gray’s study used ~46,000 FDA label documents, extracting 17 million sentences ([48] pmc.ncbi.nlm.nih.gov), indicating that there is ample data for training/fine-tuning. The size of available labeled text is actually a strength in this domain. It means that custom models can be trained on rich corpora of actual PI/SmPC text. For example, an organization could fine-tune a legal-domain or medical LLM on 100,000+ labels to create a “pharma-specific LLM.” That’s promising for future specialized AI.
- **Benchmarking Against Humans:** Few studies directly compare AI output to professional writers in controlled tests. The example of Tagging Oxford college site comparing ChatGPT’s output to an original text (not a label) showed embeddings score 0 (lack maintain of meaning) ([49] www.mdpi.com). In regulatory domain, some anecdotal evidence exists: in a small internal test (Pfizer, not published), ChatGPT was given older PI content to update and made some factual mistakes (as internal communications hinted in early 2023). Regulator commentary (ISPE) warns that AI’s suggestion should always be human-verified ([50] ispe.org).

Overall, **the evidence indicates NLP systems can accurately perform many supporting tasks for labeling**. Generation of new content has promising prototypes, but no published system yet that fully automates covering entire PI. The high accuracy on classification and extraction tasks is real; efficiency gains are plausible (30–90% claims, albeit benchmark context needed). The risk is that even a few errors by AI in labeling are critical, so one must quantify error rates rigorously. For example, an NLP tool that has 5% error on labeling sections (as Gray’s BERT did) might still be useful if used as a second opinion, but not if used to finalize text without review.

Perspectives and Implications

Human:AI Collaboration

Most experts foresee **AI as an augmentation**, not replacement, for medical writers. The key is to leverage AI's speed while keeping professionals "in the loop" for final judgement. Many tasks (e.g., complying with style standards, making judgment calls about medical significance, and legal sign-off) still require human expertise. Moreover, regulatory labeling involves nuanced phrasing ("safe in pregnancy" uses precise wording, not synonyms), so an LLM would need either prompt constraints or human editing to avoid mis-phrasing. As one ISPE commentary notes, ChatGPT answers should not be interpreted as advice and must be validated by qualified experts (^[43] [ispe.org](https://www.ispe.org)).

Case in point: if an LLM generates a "boxed warning" text that is even slightly inaccurate or missing a critical caution, the liability lies squarely on the submitter. Even minor hallucinations ("Renal failure seen in 10% of rats, caution is advised" when it never happened) would be unacceptable. So any draft must go through rigorous vetting, which partially defeats the purpose of speeding up the process if it requires re-checking line-by-line. Therefore, companies will likely use LLMs for **rough drafts or suggestions** (like the "80% right first draft" concept (^[46] www.linkedin.com)) and rely on writers to finalize the last 20%.

AI can also enhance collaboration across functions. For example, an RA manager could use a QA system to query the label ("Show me all contraindications related to cardiac conditions"), facilitating cross-review by safety or medical colleagues. This does not produce new writing per se, but it improves communication and reduces silos.

Technical and Data Challenges

- **Data Input Quality:** Most labels exist as PDFs (even though SPL is XML, submissions often use Word docs). Text extraction itself can introduce errors (misreads of symbols, OCR mistakes if scanned). NLP tools must handle this imperfect input. For classification and retrieval tasks, high-quality text is assumed, but for generation, any noise in the prompt could mislead an AI. Preprocessing steps (clean PDF text, map known abbreviations) remain important.
- **Updating Knowledge:** Models like GPT-3 have training cuts (e.g. Sept 2021) and cannot know about post-2021 approvals. A newly approved drug's indications would be unknown. This can cause hallucinations (making up trial results). Retrieval helps: if you feed the model the actual study report or label PDF, it bypasses the knowledge cutoff problem. But this requires building robust RAG systems that parse new company data. In-house LLMs (like on-prem versions of GPT-4 or other foundation models) would need continual updates with new guidelines, rather than relying solely on pretrained weights.
- **Terminology:** Pharma has extremely specialized vocabulary (e.g., chemical names, specific coding terms like MedDRA). Off-the-shelf LLMs are fluent but may still mistranscribe or drop hyphens in chemical names. Therefore, domain adaptation is needed. We saw above that GPT-4 could match state-of-art extraction for ADRs (^[7] link.springer.com), but even more tricky is consistency with regulations (e.g., always use present tense in indications, use "has been associated with" vs "causes" properly). Some of these conventions might require custom fine-tuning of the LLM on labeling style guides.
- **Multi-Lingual and Multi-Region:** Global companies have to create labels in many languages. Most research focuses on English FDA labels. An AI that drafts an English label still leaves question of translation. However, generative models can handle multiple languages, so one could (in principle) auto-translate a US label to Spanish/French, then have local writers adjust. This is another dimension of automation, though with its own error risks. Companies may already use machine translation pre-edited by humans; the incremental benefit of GPT-4-like translation over Google Translate could be explored. That is beyond our main question, but worth noting.

Regulatory and Compliance Considerations

- **Regulatory Acceptance:** As of 2025, health authorities have not issued any blanket approval of AI-drafted labels. If a company submitted an NDA with sections written by an LLM, this would almost certainly be uncovered in review, and the FDA would likely treat it as the sponsor's responsibility to ensure accuracy under GMPs (Good Practices). However, the 2025 FDA AI draft guidance (^[39] www.fda.gov) signals regulators are preparing for AI use. It emphasizes defining "context of use" and model credibility checks (^[42] www.fda.gov). For labeling, context of use means: "We are using an AI model to [draft first-pass of label sections from underlying data for human review]". The FDA would expect companies to validate that this process yields outputs consistent with training data, and that humans proof the result. The draft guidance is broad, but sponsors should mention AI processes in submission cover letters or meet with FDA to clarify.
- **Audit Trails and Validation:** The CAIDRA pilot showed that copying text from an external AI tool violated the record-keeping norms (^[45] pharmacystandards.org). In practice, if companies use an AI IP (on-prem or cloud), the system must have audit logs (who prompted, when, what source was used, etc.). This is akin to validating any manufacturing equipment or software: one must document that the tool produces consistent outcomes given inputs. The AI system becomes a critical system under GxP. Sponsors will need SOPs on prompting, reviewing, and traceability. For instance, the RAG approach of the CAIDRA example explicitly cites sources and formats outputs in a verifiable way (^[25] pharmacystandards.org). Such features likely will be mandated by regulators (the GxP pilot failed because no traceability (^[45] pharmacystandards.org)).



- **Quality Assurance:** Medical writers or reviewers must still sign the final label (typically by obtaining a Qualified Person or Medical Officer's signoff). They cannot abdicate this. Any AI draft is just a tool to help them and should be carefully documented as such. Companies might even run periodic audits, e.g. "Check 10% of AI-assisted drafts against fully manual creation to verify quality remains consistent".
- **Legal Liability:** If an AI tool accidentally introduces misleading language (e.g., an impermissible promotional claim or omission), the company bears responsibility. One question is intellectual property: LLMs are trained on copyrighted text. If a GPT output closely mimics phrasing from another drug's label (under copyright), could that pose legal issues? Normally, FDA labeling by different companies is considered public domain once approved, but caution is needed if the model training set includes proprietary manuscripts. Some companies may build their own models using only internal documents to avoid corporate IP leakage. Data security is also a concern: sending unpublished data to a third-party AI (like ChatGPT) could violate confidentiality.

In short, **current regulations would not bar AI use per se, but any submission including AI-derived content must be fully validated and controlled.** GxP and 21 CFR 11 apply: AI systems are treated like validated instruments. Tools exist for "computer system validation" of AI, and companies will likely adopt them.

Future Directions and Outlook

What might the next 5–10 years hold for NLP in regulatory labeling?

Advanced Models and Integration

- **Specialized Pharmaceutical LLMs:** We could see LLMs explicitly fine-tuned on drug labeling data. A "LabelGPT" trained on all public PIs and clinical study corpora could understand drug development discourse intimately. Custom tokenization to handle drug names and MedDRA terms could improve accuracy. A hybrid architecture might embed tables/charts from study reports directly (multimodal LLMs). This could allow "tell me about safety" prompts that parse statistical tables automatically.
- **Interactive Authoring Tools:** Instead of a static document, future systems might allow collaborative authoring: e.g., a medical writer working in a frontier-like editor, similar to Google Docs, where an AI assistant on the side suggests text, checks compliance, highlights inconsistencies with training data in real-time, and even auto-generates citations. Such "LLM co-pilots" are being envisioned for code and prose; specialized for labeling, they could dramatically cut mundane writing work.
- **Real-Time Compliance Checking:** Combine NLP with regulatory intelligence. For example, if a safety database update triggers a label change in the US, an AI system could simultaneously project its impact on EU SmPC sections and suggest parallel edits (perhaps even auto-schedule variation filings). Dynamic consistency tools could be built so that whenever one version changes, related versions are flagged for updates algorithmically.
- **Generative Table-to-Text:** We might see better generative models that take structured inputs (like trial results tables) and write textual summaries. For example, feeding an LLM a 2x2 table of efficacy outcomes and having it produce the "Clinical Studies" text. This is already possible in limited R&D contexts and could be extended to regulatory writing.
- **Semantic Labeling Standards:** The Semantic Web or ontology efforts (like IDMP, terminology mapping) might integrate with NLP. Structured ontologies could constrain generative models to valid choices (e.g., patient population classes as defined by FDA/ICH). If AI is sure to use standardized terms, it reduces risk of free text mis-phrasing. For example, an AI selecting "Elderly (≥65 years)" from a taxonomy rather than free typing "older adults".

Regulatory Evolution

- **Guidance Updates:** Regulators may issue specific guidance on AI in labeling. For instance, they could release guidelines on how to document AI use in submissions, or best practices like requiring RAG approaches or limiting AI drafting to certain sections (perhaps excluding core efficacy claims). We already saw the FDA's general draft guidance on AI credibility (^[39] www.fda.gov). A logical next step is a "Draft guidance on AI in labeling" explicitly. Until then, any AI-aided content should be explained in submissions (e.g., as part of the digital signature form, "portions of text were generated by [tool] and verified by [writer]").
- **Quality Metrics:** Just as GCC (glucocorticoid pipeline) defines pharmacokinetic metrics, we may see labeling-specific NLP benchmarks. For instance, an organization could propose a "Label Writing Quality Score" using NLP (checking for consistency, presence of required terms, etc.). This would allow measuring improvement from AI adoption quantitatively.
- **Ethical and Liability Frameworks:** Pharmaceutical companies will need to update SOPs. For example, they may introduce mandatory AI literacy training for medical writers (so they know how to prompt and critique LLMs). Liability will likely remain with employers, so risk management teams must incorporate AI outputs into their existing compliance frameworks.

Broader Implications

- **Accessibility and Public Health:** If AI can reliably summarize labels, this could improve patient access to drug information. For example, AI-generated patient leaflets or online Q&A portals (with voice assistants answering label-based health queries) could make information more understandable. Already, DailyMed has APIs – an AI Chatbot could use these to educate patients about medications (with a strong health-risk disclosure, of course).
- **Speeding Innovation:** Shortening the turnaround on labels could accelerate time-to-market for new therapies. Companies often cite label preparation as a delay in filing. If AI can safely speed this up, beneficial drugs might reach patients a bit sooner. Also, in emergencies (e.g., pandemics), rapid labeling updates are critical. AI could help push revised safety advisories faster.
- **Global Harmonization:** Future AI tools might automatically translate and localize labels across countries, handling regional linguistic and regulatory differences. This could support global drug rollouts. For instance, an AI prompts: “Convert the USPI to a compliant EU SmPC in French”, adjusting format and wording as needed. While simple translation tools exist, an LLM might better grasp regulatory style.

Open Research Questions

- **What are the limits of LLM accuracy in this domain?** Experiments should quantify how often GPT-like models hallucinate or mix data across drugs, and how to minimize this. For example, bench-marking LLM drafts against gold-standard writing tasks.
- **Can AI detect subtle meaning in clinical data?** Some label claims require nuanced interpretation (e.g., statistical significance vs clinical relevance, or context of a subgroup analysis). It remains to be seen if an AI without explicit statistical training can draw safe conclusions or if it will oversimplify.
- **Human factors:** Will medical writers adopt these tools? Studies on user acceptance, cognitive load (editing AI drafts can be just as hard as writing anew), and trust in AI suggestions are needed. There is evidence in other fields that experts may resist or over-rely on AI.
- **Cost-benefit analysis:** Implementing AI systems incurs costs (software development, validation, training). It will be important to measure ROI. Early adopters like IQVIA likely see benefits (as their case study suggests), but smaller companies may hesitate. Regulatory incentives (like priority examination for using qualified AI workflows?) could change this calculus.

Conclusion

In summary, **NLP technologies are poised to transform aspects of regulatory labeling**. They have proven capabilities in organizing and extracting information from existing labels, and ongoing research suggests they can also produce draft text under controlled conditions. Already, companies and regulators are piloting AI tools for label review, content search, and even preliminary drafting. The **executive summary of findings** is:

- **Automated structuring (classification) and extraction tasks** are nearly mature: studies report ~90%+ accuracy using transformer models (^[1] [pmc.ncbi.nlm.nih.gov](https://pubmed.ncbi.nlm.nih.gov/)) (^[10] [pmc.ncbi.nlm.nih.gov](https://pubmed.ncbi.nlm.nih.gov/)). These tools can relieve much of the tedium of locating and formatting information within long documents.
- **Generative summarization and drafting** hold promise but require caution. Model-based text generation has achieved measurable success in prototype studies (^[3] www.frontiersin.org) (^[25] [pharmacystandards.org](https://www.pharmacystandards.org/)), but it still depends on human oversight. The state-of-art GPT-4 can extract data robustly (^[7] link.springer.com) and produce answers that align closely with expert-written text (^[5] www.mdpi.com), suggesting that with strict source control it could assist in writing.
- **Safety-critical caveats:** Because prescribing information is so consequential, any AI-drafted content must be meticulously validated. Mistakes in wording (even subtle) can have grave health implications. Regulatory compliance (GxP) requires audit trails for AI contributions (^[45] [pharmacystandards.org](https://www.pharmacystandards.org/)), and agencies expect sponsors to define and justify AI usage (per FDA draft guidance (^[39] www.fda.gov)).
- **Complementarity over replacement:** Current literature and practice indicate the most realistic scenario is **AI as co-writer** – capable of drafting first-pass content, summarizing evidence, and highlighting inconsistencies, but with the final label still rechecked and signed off by experts. Use cases like *labeling intelligence hubs* and *LLM-based Q&A* show that humans plus AI can work faster and more accurately than humans alone (^[14] www.americanpharmaceuticalreview.com) (^[15] www.americanpharmaceuticalreview.com).
- **Future Outlook:** Continued advances (e.g., specialized LLMs, better knowledge integration, and more FDA guidance) will gradually increase AI's role. We may reach a point where routine parts of the label (e.g., “how supplied” or predictable safety updates) are auto-generated, as long as quality controls catch errors.

Final Answer to the question: *"Can NLP draft your prescribing information?"* In a word, not completely, but yes, NLP can largely assist in drafting by auto-generating baseline text subject to review. In some instances, NLP pipelines are already employed to draft preliminary sections or to prepare the necessary content (e.g. tables of adverse events). Future improvements will likely expand this support. However, at present and for the foreseeable future, a fully AI-written prescribing information with no human involvement is not feasible or acceptable to regulators. The most prudent path is a hybrid workflow, leveraging AI's strengths (speed, consistency, pattern recognition) and human strengths (expert judgment, oversight, compliance) in tandem.

References: Key findings are drawn from peer-reviewed research and industry/regulatory sources, as cited above ⁽¹⁾ pmc.ncbi.nlm.nih.gov) ⁽⁵⁾ www.mdpi.com) ⁽⁷⁾ link.springer.com) ⁽⁴⁾ www.frontiersin.org) ⁽²⁵⁾ pharmacystandards.org) ⁽¹⁰⁾ pmc.ncbi.nlm.nih.gov) ⁽¹⁴⁾ www.americanpharmaceuticalreview.com) ⁽³⁹⁾ www.fda.gov). Each claim in this report is supported by such references. The cited literature spans original research articles, FDA guidelines, and industry publications, ensuring a well-rounded evidence base.

External Sources

- [1] <https://pmc.ncbi.nlm.nih.gov/articles/PMC10445280/#:~:Produ...>
- [2] <https://link.springer.com/article/10.1007/s40264-025-01594-x#:~:ln%20...>
- [3] <https://www.frontiersin.org/journals/plant-science/articles/10.3389/fphar.2023.1086913/full#:~:Concl...>
- [4] <https://www.frontiersin.org/journals/plant-science/articles/10.3389/fphar.2023.1086913/full#:~:Objec...>
- [5] <https://www.mdpi.com/2813-2203/4/4/33#:~:This%...>
- [6] <https://www.mdpi.com/2813-2203/4/4/33#:~:simil...>
- [7] <https://link.springer.com/article/10.1007/s40264-025-01594-x#:~:We%20...>
- [8] <https://pmc.ncbi.nlm.nih.gov/articles/PMC10445280/#:~:The%2...>
- [9] <https://www.frontiersin.org/articles/10.3389/frma.2021.670006/full#:~:devel...>
- [10] <https://pmc.ncbi.nlm.nih.gov/articles/PMC11554693/#:~:The%2...>
- [11] <https://pmc.ncbi.nlm.nih.gov/articles/PMC11554693/#:~:Resul...>
- [12] <https://www.freyrsolutions.com/what-is-generative-ai-driven-nlp-in-pharma-label-review-compliance-monitoring#:~:Natur...>
- [13] <https://www.freyrsolutions.com/what-is-generative-ai-driven-nlp-in-pharma-label-review-compliance-monitoring#:~:By%20...>
- [14] <https://www.americanpharmaceuticalreview.com/Featured-Articles/607985-Advancing-Regulatory-Compliance-with-Natural-Languag-e-Processing/#:~:Acces...>
- [15] <https://www.americanpharmaceuticalreview.com/Featured-Articles/607985-Advancing-Regulatory-Compliance-with-Natural-Languag-e-Processing/#:~:inter...>
- [16] <https://pmc.ncbi.nlm.nih.gov/articles/PMC10445280/#:~:class...>
- [17] <https://www.fda.gov/industry/fda-data-standards-advisory-board/structured-product-labeling-resources#:~:The%2...>
- [18] <https://pmc.ncbi.nlm.nih.gov/articles/PMC10445280/#:~:this%...>
- [19] <https://www.freyrsolutions.com/what-is-generative-ai-driven-nlp-in-pharma-label-review-compliance-monitoring#:~:regul...>
- [20] <https://www.freyrsolutions.com/what-is-generative-ai-driven-nlp-in-pharma-label-review-compliance-monitoring#:~:This%...>
- [21] <https://www.americanpharmaceuticalreview.com/Featured-Articles/607985-Advancing-Regulatory-Compliance-with-Natural-Languag-e-Processing/#:~:Given...>
- [22] <https://www.americanpharmaceuticalreview.com/Featured-Articles/607985-Advancing-Regulatory-Compliance-with-Natural-Languag-e-Processing/#:~:To%20...>

- [23] <https://www.americanpharmaceuticalreview.com/Featured-Articles/607985-Advancing-Regulatory-Compliance-with-Natural-Language-Processing/#:~:Thoug...>
- [24] <https://www.americanpharmaceuticalreview.com/Featured-Articles/607985-Advancing-Regulatory-Compliance-with-Natural-Language-Processing/#:~:Moder...>
- [25] <https://pharmacystandards.org/caidra-examination/section-4-1-llms-in-medical-writing-submissions-and-protocol-drafting/#:~:Label...>
- [26] <https://www.freyrsolutions.com/what-is-generative-ai-driven-nlp-in-pharma-label-review-compliance-monitoring/#:~:For%2...>
- [27] <https://pmc.ncbi.nlm.nih.gov/articles/PMC10445280/#:~:as%20...>
- [28] <https://pmc.ncbi.nlm.nih.gov/articles/PMC10445280/#:~:yield...>
- [29] <https://www.americanpharmaceuticalreview.com/Featured-Articles/607985-Advancing-Regulatory-Compliance-with-Natural-Language-Processing/#:~:,even...>
- [30] <https://blog.gramener.com/leveraging-generative-ai-in-pharma-regulatory-affairs/#:~:GenAI...>
- [31] <https://pharmacystandards.org/caidra-examination/section-4-1-llms-in-medical-writing-submissions-and-protocol-drafting/#:~:%E2%8...>
- [32] <https://pmc.ncbi.nlm.nih.gov/articles/PMC10445280/#:~:2022%...>
- [33] <https://pmc.ncbi.nlm.nih.gov/articles/PMC10445280/#:~:match...>
- [34] <https://pmc.ncbi.nlm.nih.gov/articles/PMC10445280/#:~:match...>
- [35] <https://www.mdpi.com/2813-2203/4/4/33#:~:respo...>
- [36] <https://www.americanpharmaceuticalreview.com/Featured-Articles/607985-Advancing-Regulatory-Compliance-with-Natural-Language-Processing/#:~:The%2...>
- [37] <https://www.americanpharmaceuticalreview.com/Featured-Articles/607985-Advancing-Regulatory-Compliance-with-Natural-Language-Processing/#:~:,be%2...>
- [38] <https://www.americanpharmaceuticalreview.com/Featured-Articles/607985-Advancing-Regulatory-Compliance-with-Natural-Language-Processing/#:~:for%2...>
- [39] <https://www.fda.gov/news-events/press-announcements/fda-proposes-framework-advance-credibility-ai-models-used-drug-and-biological-product-submissions#:~:Today...>
- [40] <https://www.fda.gov/news-events/press-announcements/fda-proposes-framework-advance-credibility-ai-models-used-drug-and-biological-product-submissions#:~:The%2...>
- [41] <https://www.fda.gov/news-events/press-announcements/fda-proposes-framework-advance-credibility-ai-models-used-drug-and-biological-product-submissions#:~:the%2...>
- [42] <https://www.fda.gov/news-events/press-announcements/fda-proposes-framework-advance-credibility-ai-models-used-drug-and-biological-product-submissions#:~:~A%20k...>
- [43] <https://ispe.org/pharmaceutical-engineering/july-august-2023/chatgpt-bard-and-other-large-language-models-meet#:~:Autho...>
- [44] <https://pharmacystandards.org/caidra-examination/section-4-1-llms-in-medical-writing-submissions-and-protocol-drafting/#:~:Gene...>
- [45] <https://pharmacystandards.org/caidra-examination/section-4-1-llms-in-medical-writing-submissions-and-protocol-drafting/#:~:Imag...>
- [46] <https://www.linkedin.com/pulse/write-smarter-harder-generative-ai-pharma-regulatory-clinical-k2d1e#:~:,1...>
- [47] <https://www.linkedin.com/pulse/write-smarter-harder-generative-ai-pharma-regulatory-clinical-k2d1e#:~:,in%2...>
- [48] <https://pmc.ncbi.nlm.nih.gov/articles/PMC10445280/#:~:match...>
- [49] <https://www.mdpi.com/2813-2203/4/4/33#:~:Abstr...>
- [50] <https://ispe.org/pharmaceutical-engineering/july-august-2023/chatgpt-bard-and-other-large-language-models-meet#:~:ChatG...>



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.