Pharmacophore

ISSN-2229-5402

Journal home page: http://www.pharmacophorejournal.com

RETRIEVAL-AUGMENTED AI FOR INTERPRETING ICH QUALITY GUIDELINES AND PRODUCT CONTROL STRATEGIES

Elena Petrova[1]*, Ivan Georgiev[1], Nikolay Stoyanov[2], Petar Kolev[1]

1. *Department of Pharmaceutical AI Systems, Faculty of Medicine, Medical University of Sofia, Sofia, Bulgaria.*
2. *Department of Computational Drug Analytics, Faculty of Engineering, Technical University of Sofia, Sofia, Bulgaria.*

ARTICLE INFO

Received:
19 November 2025
Received in revised form:
03 February 2026
Accepted:
05 February 2026
Available online:
28 February 2026

Keywords: Retrieval-augmented generation, ICH quality guidelines, Pharmaceutical quality, Control strategy, Regulatory intelligence, GMP

ABSTRACT

Pharmaceutical companies face the complex challenge of interpreting and applying ICH quality guidelines while translating general principles into product-specific control strategies, a task made cognitively demanding by variations in modality, dosage form, manufacturing process, and lifecycle stage. Current practices rely heavily on manual document searches, individual expertise, and static templates, which are difficult to scale across diverse product portfolios and evolving regulatory expectations. This article presents a retrieval-augmented AI assistant designed to index ICH quality guidelines, regional regulatory guidance, pharmacopoeial materials, and internal control-strategy documents, enabling natural-language queries about pharmaceutical quality expectations with grounded, citation-backed responses. The system integrates a multi-source ingestion pipeline, a vector database with metadata filtering, a domain-adapted language model, a citation-grounding layer, and a human-review dashboard, emphasizing support for expert interpretation rather than automation of regulatory judgment. By facilitating easier retrieval and comparison of relevant regulatory language, the tool can accelerate the development and review of product-specific control strategies, promote consistency across quality, regulatory affairs, analytical development, and manufacturing teams, and enhance submission readiness while preserving expert oversight, ultimately transforming how pharmaceutical professionals interact with quality guidelines.

To Cite This Article: Petrova E, Georgiev I, Stoyanov N, Kolev P. Retrieval-Augmented AI for Interpreting ICH Quality Guidelines and Product Control Strategies. Pharmacophore. 2026;17(1):72-80. https://doi.org/10.51847/FGAYclSoYD

Introduction

ICH quality guidelines from pharmaceutical development through lifecycle management provide the conceptual backbone for modern quality-by-design practice, including design space, risk-based control, and post-approval change management. Quality-by-design literature has emphasized that development knowledge, process understanding, and control strategy definition must be linked across the product lifecycle [1]. For complex products and evolving manufacturing platforms, the burden is not only knowing the guideline text but applying it consistently to product-specific contexts [2]. Digital quality systems and intelligent information management have therefore become increasingly relevant to how organizations operationalize quality expectations [3].

In many companies, regulatory interpretation still depends on manual review of PDFs, internal expert networks, and inherited templates that may not clearly distinguish general principles from product-specific commitments. Industry 4.0 perspectives on pharmaceutical manufacturing show that smart factories require more connected quality knowledge than static document repositories can provide [4]. Process automation work in continuous and platform manufacturing likewise highlights that control strategies are increasingly embedded in digital workflows rather than isolated narrative documents [5]. When interpretation remains fragmented, teams may face delays, inconsistent justification, and avoidable gaps between development rationale and regulatory submission language.

Retrieval-augmented generation has matured into a practical architecture for grounding language-model answers in retrieved source documents rather than relying only on model memory. Direct applications to drug information, clinical trial protocols, and FDA guidance documents demonstrate how retrieval can support regulatory question answering while retaining traceability to source materials [6, 7]. Pharmaceutical regulatory compliance workflows have also begun to explore RAG-based question-answering systems that combine generative responses with targeted retrieval [8]. Biomedical language models

Corresponding Author: Elena Petrova; Department of Pharmaceutical AI Systems, Faculty of Medicine, Medical University of Sofia, Sofia, Bulgaria. E-mail: elena.petrova@gmail.com.

such as BioBERT, PubMedBERT, and SciBERT further indicate that domain adaptation can improve how technical language is represented for downstream text-mining tasks [9-11].

This article proposes a RAG system designed specifically for pharmaceutical quality and regulatory science, with emphasis on ICH quality guidelines and product-specific control strategies. The system would ingest authoritative regulatory documents and internal quality manuals, retrieve relevant passages in response to natural-language questions, and synthesize answers that remain explicitly grounded in the cited text [12]. It is intended to assist regulatory affairs, quality assurance, pharmaceutical development, and manufacturing science professionals rather than replace accountable human decision-making. Its design is therefore aligned with broader expectations for AI-enabled decision support, including transparency, oversight, and careful evaluation before routine use [13-15].

Background

The ICH Quality Guideline Framework

The ICH quality guideline framework links pharmaceutical development, quality risk management, pharmaceutical quality systems, analytical procedure development, and lifecycle management into a shared model for science- and risk-based control. Quality-by-design approaches describe how formulation, process, analytical, and control decisions should be justified through structured knowledge rather than retrospective testing alone [1]. Control strategy studies show that excipient variability, design-space definition, and predictive modeling must be connected when translating quality principles into practical controls [2]. For emerging manufacturing modes, integrated control strategies also need to accommodate automation, continuous processing, and platform-based production [5].

Challenges in Regulatory Interpretation and Implementation

The difficulty in applying ICH principles lies in translating broad regulatory language into concrete decisions for a specific product, process, analytical method, or post-approval change. Regulatory intelligence studies have shown that organizations seek AI support because interpretation often requires synthesis across guidance, precedent, and internal experience [16]. Reviews of AI in regulatory affairs similarly suggest that regulatory work increasingly involves managing complex information flows rather than simply locating individual documents [17]. In this context, divergent interpretations can arise when teams rely on different templates, informal expert judgment, or incomplete searches of the available guidance.

Retrieval-Augmented Generation for Regulatory Text

RAG is attractive for regulatory text because it separates document retrieval from language generation, allowing an answer to be conditioned on specific source passages. In FDA guidance search, generative AI has been used to make large guidance collections more semantically searchable while still returning evidence from the underlying documents [7]. In pharmaceutical regulatory compliance, QA-RAG approaches propose using retrieval to ground responses to compliance questions rather than allowing a model to answer from latent parameters alone [8]. Clinical and biomedical foundation-model studies further show why grounded retrieval is important in high-stakes domains where fluent language can otherwise mask unsupported inference [18, 19].

Prior NLP and AI Applications in Pharmaceutical Quality

Before RAG became prominent, pharmaceutical and biomedical NLP applications included document classification, semantic similarity, labeling summarization, translation, and domain-specific language representation. Regulatory-document question similarity work has shown that recognizing semantically related questions can support more efficient reuse of regulatory knowledge [20]. Drug-labeling summarization research illustrates the potential of large language models to condense complex regulated text, while also implying the need for careful review of generated summaries [21]. Pharmaceutical regulatory translation and chatbot-based drug-label retrieval further show that natural-language interfaces are becoming practical, although interactive control-strategy interpretation remains less developed [12, 22].

Trust, Auditability, and Human Oversight in Regulatory AI

Trustworthy regulatory AI must be auditable, access-controlled, and embedded in human review workflows because GMP decisions carry accountability beyond the software system itself. Reporting guidance for AI interventions emphasizes transparent description of the AI system, its intended use, and its evaluation context [13]. Early-stage AI decision-support guidance similarly stresses human factors, deployment context, and workflow integration before claims of operational value are made [15]. Regulatory discussions of AI-enabled software also highlight governance, validation, monitoring, and accountability as core expectations for systems used in regulated health environments.

System Architecture Overview

High-Level Design

The proposed system is a secure, on-premise RAG platform that ingests regulatory documents and internal quality manuals, builds a searchable index, and exposes a natural-language interface for regulated users. On-premise or otherwise controlled deployment is important because pharmaceutical quality documents may include confidential development rationales, manufacturing parameters, and product-specific commitments, and AI-enabled pharmaceutical workflows have increasingly

emphasized privacy, governance, and secure implementation [23]. The answer engine would retrieve source passages before generation, following the same source-grounded principle used in regulatory guidance search and drug-label information retrieval [7, 12]. Each generated answer would include mandatory citations to the retrieved documents so that the user can verify the underlying regulatory basis.

Figure 1 presents the proposed retrieval-augmented AI architecture that connects authoritative regulatory sources, internal control-strategy knowledge, citation-grounded answer generation, and human regulatory review.

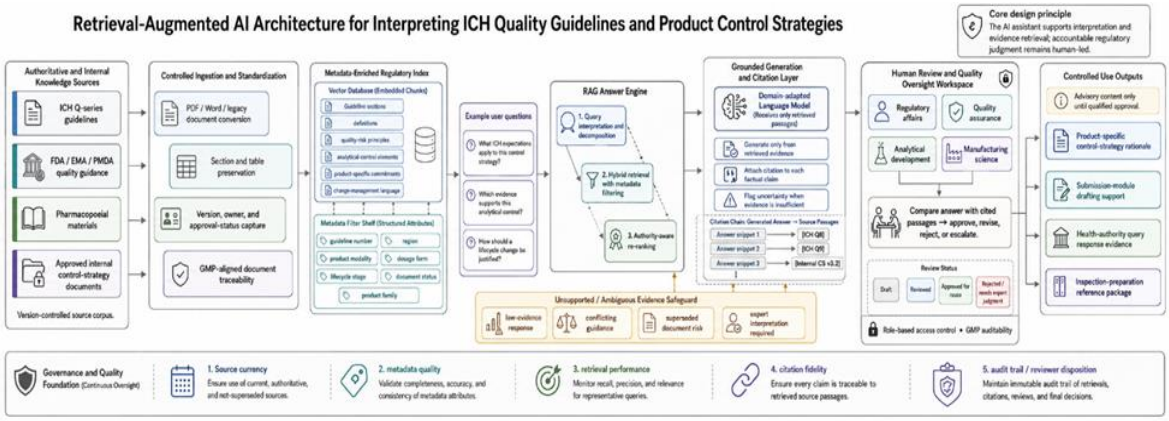

Figure 1. Retrieval-Augmented AI Architecture for Interpreting ICH Quality Guidelines and Product Control Strategies

Core Document Corpus and Data Flow

The core corpus would include ICH Q-series guidelines, regional quality guidance from authorities such as FDA, EMA, and PMDA, relevant pharmacopoeial materials, and internal control-strategy documents. The ingestion workflow would convert source documents into structured, searchable text; attach metadata such as guideline, region, product modality, lifecycle stage, and document owner; and make these units available for retrieval. Intelligent information-management work in pharmaceutical production suggests that quality knowledge becomes more useful when dispersed technical information is organized into reusable digital structures [3]. The output would be a natural-language answer with numbered citations to the relevant document sections, allowing users to move from an answer back to the source evidence.

Design Principles

The system would be designed as an assistive layer rather than an autonomous regulatory decision-maker. It would prioritize source grounding, auditability, privacy preservation, and alignment with GMP data integrity expectations, reflecting broader regulatory concerns around AI transparency and accountability. Human approval would remain necessary before AI-generated language is inserted into a protocol, control-strategy justification, response to a health authority, or submission module. This assistive posture is consistent with clinical and regulatory AI guidance that emphasizes intended use, human oversight, and careful monitoring rather than unrestricted automation [14, 15].

Regulatory Knowledge Base Construction and Indexing

Document Preprocessing and Standardization

Document preprocessing would convert PDFs, Word files, controlled templates, and scanned legacy documents into machine-readable text while preserving section headings, tables, figure captions, and version metadata. Pharmaceutical regulatory translation and labeling summarization studies show that careful preparation of regulated text is essential before language models can reliably operate on it [21, 22]. For quality manuals and control-strategy documents, preprocessing should also preserve document hierarchy so that a retrieved paragraph can be traced to its section, version, owner, and approval status. This standardization would make later retrieval more auditable and reduce the risk that the system answers from superseded or poorly contextualized content.

Chunking and Metadata Enrichment

After preprocessing, the system would split documents into semantically coherent chunks, such as guideline paragraphs, definitions, risk-assessment principles, analytical-control elements, or product-specific commitments. Biomedical language-model work shows that domain terminology benefits from representations trained or adapted to scientific and clinical language, which is relevant when embedding technical regulatory chunks [9, 10, 24, 25]. Each chunk would receive metadata tags for guideline number, topic, dosage form, modality, region, product family, lifecycle phase, and document status. This enriched structure would allow the system to distinguish, for example, a general quality-by-design principle from an approved internal control for a particular biologic, sterile product, or continuous manufacturing process.

Table 1 structures the proposed knowledge base by distinguishing source authority, metadata requirements, retrieval function, and the way each document class supports product-specific control-strategy interpretation.

Table 1. Regulatory knowledge architecture linking source type, metadata logic, retrieval function, and control-strategy use

Knowledge source category	Primary regulatory or quality role	Required metadata for reliable retrieval	Retrieval value for user questions	Control-strategy interpretation contribution	Main risk if poorly indexed	Governance requirement
ICH Q-series guidelines	Establish global quality principles for pharmaceutical development, quality risk management, analytical procedures, lifecycle management, and pharmaceutical quality systems	Guideline number, section heading, topic, version, effective date, quality concept, lifecycle stage	Retrieves authoritative language for broad questions about design space, risk management, analytical control, established conditions, and lifecycle change	Provides the general interpretive foundation for linking development knowledge to quality expectations	Overgeneralized answers that quote principles without explaining product relevance	Maintain current guideline versions and retire superseded text from active retrieval
Regional regulatory guidance	Clarifies jurisdiction-specific expectations from agencies such as FDA, EMA, or PMDA	Region, agency, guidance type, effective status, product scope, modality, submission context	Supports questions where regional expectations may affect submission language, comparability strategy, or control expectations	Helps distinguish global ICH principles from region-specific implementation expectations	Mixing regional expectations without jurisdictional context	Require metadata filtering by region and visible source labels in generated answers
Pharmacopoeial materials	Defines compendial standards, analytical requirements, monograph expectations, and quality test conventions	Pharmacopoeia, chapter or monograph, dosage form, material type, analytical method, revision status	Retrieves standardized test expectations and terminology for analytical-control questions	Supports alignment between product controls and recognized compendial standards	Inappropriate application of general compendial language to product-specific contexts	Link each passage to revision status and applicable material or dosage-form scope
Internal quality manuals	Encodes company-approved quality-system procedures, terminology, and decision pathways	Document owner, approval status, SOP link, department, effective date, review cycle, supersession history	Answers questions about how the organization operationalizes external regulatory principles	Translates regulatory expectations into company-specific procedural expectations	Retrieval of obsolete procedures or uncontrolled templates	Enforce document control, access permissions, and periodic content review
Approved control-strategy documents	Contains product-specific commitments, development rationale, critical quality attributes, process controls, and analytical controls	Product family, modality, dosage form, lifecycle phase, approval status, market, version, document owner	Retrieves product-specific evidence for proposed changes, submissions, inspections, or internal review	Connects general guideline principles to actual control decisions for a product or platform	Leakage of confidential product commitments or use of unapproved draft content	Apply role-based access control and clearly mark draft versus approved source material
Prior regulatory responses and assessment packages	Captures precedent from health-authority questions, company responses, and prior submission rationales	Product, region, submission type, authority interaction type, response date, commitment status, approval outcome	Supports consistency when preparing new responses or interpreting similar authority questions	Provides precedent-based context while preserving the distinction between past response and current decision	Treating historical precedent as universally applicable	Require expert review before reuse and tag commitments by product and jurisdiction
User feedback and reviewer dispositions	Records whether AI-generated answers were accepted, modified, rejected, or escalated	Reviewer role, decision status, comment category, timestamp, answer version, retrieved evidence set	Supports controlled improvement of retrieval, ranking, prompt instructions, and interface guidance	Identifies recurring ambiguity in control-strategy interpretation and evidence gaps	Informal feedback causing uncontrolled model drift or undocumented changes	Govern feedback as quality-system improvement input rather than automatic retraining

Vector Database and Retrieval Strategy

The vector database would store embeddings for each chunk and support hybrid retrieval that combines semantic similarity with keyword matching and metadata filters. Hybrid retrieval is important because regulatory questions often include exact terms such as “design space,” “control strategy,” “established conditions,” or “analytical procedure,” while also requiring conceptual matching to related passages [8]. FDA guidance search and grounded drug-label question answering both illustrate why retrieval must return authoritative source text rather than merely similar-looking language [7, 12]. Metadata filters would

allow users to limit the search to ICH guidance, regional guidance, or approved internal documents depending on the regulatory question and the decision context.

Retrieval-Augmented Answer Engine
Query Interpretation and Decomposition
The answer engine would first interpret the user's natural-language query by identifying the relevant quality concept, product context, regulatory source, and intended use of the answer. A question such as whether an analytical control supports a lifecycle change would be decomposed into concepts related to analytical procedure development, established conditions, control strategy, and change management. Question-entailment methods for regulatory documents suggest that identifying semantic overlap between questions can help route new queries toward prior knowledge and relevant regulatory text [20]. In practice, this layer would convert a broad user question into a retrieval plan while preserving the original phrasing for auditability.

Retrieval and Re-Ranking
The retrieval step would return candidate passages from the indexed corpus, and a re-ranking step would prioritize passages that are authoritative, current, and directly responsive to the user's question. RAG compliance architectures in the pharmaceutical domain show the value of separating retrieval from generation so that the system can select evidence before drafting an answer [6, 8]. The re-ranker would be expected to favor primary ICH text over secondary summaries when the user asks for a guideline interpretation, while internal approved control-strategy documents may be prioritized when the question concerns a specific product. This design would reduce the chance that the model synthesizes a plausible but weakly supported answer from tangentially related material.

LLM-Based Answer Generation with Inline Citations
The language model would synthesize an answer only from the retrieved passages, using instructions that prohibit unsupported extrapolation and require inline citations for each factual or interpretive statement. Domain-adapted biomedical and clinical language models show that specialized pretraining can improve handling of technical terminology, although source grounding remains necessary for regulatory use [9, 10, 18]. In the proposed system, the answer would explain the relevant guideline principle, identify how it applies to the product or control-strategy question, and clearly state any limits in the retrieved evidence. The generated response would remain draft advisory content until reviewed by qualified regulatory or quality personnel, consistent with human-centered AI decision-support principles [15, 19].

Answer Grounding, Citation, and Traceability
Citation Mechanism and Audit Trail
Every generated answer would be stored with the user query, retrieved passages, source-document identifiers, prompt template, model response, reviewer actions, and final disposition. This audit trail is essential because AI-assisted regulatory interpretation must remain inspectable by quality units, regulatory affairs, and potentially auditors, especially when generated text influences GMP documentation or submission rationale. Citation grounding would allow each claim in the response to be traced back to a specific guideline section, internal manual, or approved control-strategy document, similar to how FDA guidance search and grounded drug-label retrieval systems connect answers to source evidence [7, 12]. The system would therefore treat citation fidelity as a design requirement rather than a presentation feature.

Flagging Unsupported or Ambiguous Statements
When retrieved passages do not provide sufficient evidence to answer a question, the system should explicitly state that the corpus does not contain a clear answer. This behavior is important because large language models can produce fluent responses even when the available evidence is incomplete, and biomedical foundation-model work has underscored the need to constrain generated clinical or regulatory language to reliable evidence [18, 19]. In a pharmaceutical quality setting, an unsupported answer about design space, established conditions, or control strategy could create a misleading basis for development or regulatory decision-making. The system would therefore classify low-evidence responses as requiring expert interpretation rather than attempting to resolve ambiguity through model inference alone.

Feedback and Continuous Improvement
The system would include mechanisms for users to mark answers as helpful, incomplete, unsupported, or incorrectly cited, and these signals would be reviewed through a controlled quality process. Human feedback could be used to refine retrieval settings, metadata tagging, re-ranking logic, prompt instructions, and user-interface guidance without allowing uncontrolled model drift in a regulated environment [15]. Reviews of AI-driven pharmaceutical development and manufacturing suggest that continuous improvement must be balanced with governance, validation, and lifecycle oversight when AI tools affect quality decisions [26]. Feedback would therefore be treated as a governed improvement input rather than an informal training signal.

Human-In-The-Loop Review and Quality Oversight
Collaborative Review Dashboard

A collaborative review dashboard would allow regulatory affairs, quality assurance, analytical development, and manufacturing science users to compare each AI-generated answer with the retrieved source passages before approving, modifying, or rejecting it. Early-stage AI decision-support guidance emphasizes that evaluation must consider the interaction between the system, the user, and the workflow rather than the model output alone [15]. In this design, the dashboard would show the answer, cited passages, document metadata, review comments, and approval status in a single controlled workspace. The approved response could then be reused in a regulatory rationale, internal assessment, or health-authority response only after accountable expert review.

Table 2 shows the key elements and functionalities of a collaborative review dashboard for AI-generated responses in regulatory and quality workflows.

Table 2. Key Elements of a Collaborative Review Dashboard

Element	Description / Functionality
AI-Generated Answer	The proposed response produced by the AI system for review.
Retrieved Source Passages	Original document excerpts that support or inform the AI-generated answer.
Document Metadata	Details such as title, authorship, version, and date of source documents.
Review Comments	Feedback provided by regulatory, QA, analytical, or manufacturing science experts.
Approval Status	Indicator showing whether the answer is approved, modified, or rejected.
Workflow Integration	Ensures the answer undergoes accountable expert review before being reused in regulatory work.

Governance and Access Control

Governance would require role-based access control so that confidential product-specific control strategies, manufacturing parameters, and regulatory commitments are visible only to authorized users. AI governance discussions in regulated health environments emphasize transparency, accountability, monitoring, and clear responsibility for outputs that may influence regulated decisions. The system would watermark generated content as draft until a qualified reviewer approves it, preventing unreviewed language from being mistaken for an official regulatory position. Access logs, document-version controls, and review records would support GMP data integrity expectations and make the system suitable for controlled quality workflows.

Integration Into Pharmaceutical Development And Regulatory Workflows

Supporting Control-Strategy Development and Review

During product development, the system could help teams retrieve relevant guideline passages and compare proposed control strategies with established regulatory expectations. Quality-by-design research shows that control strategy development requires connecting formulation variables, process parameters, analytical controls, and risk assessments into a coherent product-specific rationale [1, 2]. Digital manufacturing and process automation studies further suggest that control strategies are becoming increasingly integrated with platform technologies and continuous manufacturing workflows [4, 5]. A RAG assistant would therefore support drafting and review by making the regulatory and internal knowledge base easier to query, while preserving human judgment for final control-strategy decisions.

Regulatory Submission Preparation and Health-Authority Query Response

During submission preparation or health-authority query response, the tool could help users identify precise regulatory language, retrieve internal justifications, and align response drafts with approved product knowledge. Regulatory intelligence studies indicate that AI may support the management of complex regulatory information when its outputs remain transparent and reviewable [16, 17]. Drug-label summarization and regulatory-document question-similarity methods also suggest that NLP can help users navigate large bodies of regulated text, although final interpretation must remain expert-led [20, 21]. In this workflow, the system would function as a traceable evidence-retrieval and drafting assistant rather than an autonomous author of regulatory commitments.

Evaluation Strategy

Answer Accuracy and Relevance

Evaluation should focus on expert review of answer correctness, relevance to the query, completeness of regulatory reasoning, and appropriateness of citations. Reporting guidelines for AI interventions emphasize that evaluation should be explicit about intended use, user population, and the context in which the AI output informs decisions [13, 14]. In this setting, regulatory affairs and quality professionals would compare AI-generated draft answers against manually retrieved guideline evidence and internal documentation. Because the proposed article is conceptual, such evaluation would be described as a necessary future activity rather than reported as completed performance testing.

System Usability and Adoption

Usability evaluation should examine whether the system fits naturally into regulatory research, control-strategy review, submission drafting, and inspection-preparation workflows. Decision-support evaluation guidance highlights the importance

of human factors, workflow effects, user confidence, and safe adoption before broader deployment [15]. Pharmaceutical AI reviews similarly indicate that successful adoption depends not only on model capability but also on integration with quality systems, documentation practices, and organizational governance [23, 26]. The system should therefore be assessed through structured user review, qualitative feedback, and controlled pilot use rather than assumed to be acceptable because it produces fluent answers.

Impact on Regulatory Outcomes

Longer-term evaluation could examine whether the system improves consistency of regulatory references, supports clearer control-strategy rationales, and helps teams prepare more traceable responses to health-authority questions. AI in pharmaceutical regulatory affairs has been discussed as a way to improve information access and regulatory intelligence, but such impact must be assessed in relation to actual regulated workflows and accountable review processes [16, 17]. Quality-by-design and intelligent information-management work suggests that better organization of product and process knowledge can support more coherent pharmaceutical quality documentation [3, 27]. Any claims about regulatory outcome improvement should therefore be framed cautiously and supported only after governed pilot implementation.

Table 3 defines an evaluation and governance framework that separates retrieval performance, citation fidelity, regulatory reasoning quality, workflow usability, and controlled deployment readiness.

Table 3. Evaluation and governance framework for a citation-grounded RAG assistant in pharmaceutical quality workflows

Evaluation domain	Core evaluation question	Suggested assessment method	Evidence standard for acceptable performance	Failure mode to detect	Governance response
Retrieval relevance	Did the system retrieve passages that are directly responsive to the user's regulatory or control-strategy question?	Expert comparison between retrieved passages and manually identified source evidence	Retrieved evidence includes the most relevant authoritative source sections and avoids tangential passages	Semantically similar but regulatory-irrelevant text is retrieved	Refine chunking, metadata tags, hybrid search weights, and re-ranking criteria
Source authority prioritization	Did the system prioritize the correct source hierarchy for the question?	Scenario-based testing across ICH guidance, regional guidance, pharmacopoeial materials, and internal documents	Primary authoritative sources are favored when guideline interpretation is requested; approved internal documents are favored for product-specific questions	Internal template or secondary summary overrides authoritative guidance	Add authority-ranking rules and visible source-type labels
Citation fidelity	Does every factual or interpretive claim trace to the cited source passage?	Claim-level audit of generated answers against retrieved text	Each claim is supported by a specific cited passage, and unsupported claims are absent or flagged	Citation exists but does not support the statement it is attached to	Strengthen citation-grounding prompts and block unsupported answer generation
Regulatory reasoning quality	Does the answer correctly translate broad regulatory language into product-relevant interpretation without overstating certainty?	Blinded review by regulatory affairs, quality assurance, analytical development, and manufacturing science experts	Answer distinguishes general principles, product-specific commitments, regional expectations, and unresolved ambiguity	Fluent but overconfident interpretation beyond available evidence	Require uncertainty labeling, expert escalation, and reviewer disposition tracking
Control-strategy usefulness	Does the response help users connect CQAs, CPPs, analytical controls, risk assessments, and lifecycle decisions?	Use-case testing on control-strategy development, change assessment, and submission-drafting scenarios	Output supports clearer rationale development without becoming an unreviewed regulatory commitment	Answer retrieves guideline language but does not help structure a practical control strategy	Add control-strategy templates, product-context prompts, and reviewer guidance
Handling of insufficient evidence	Does the system recognize when the corpus does not contain enough information to answer?	Low-evidence and conflicting-evidence test cases	System explicitly states uncertainty and recommends expert interpretation when evidence is incomplete	Model fabricates a definitive answer from weak evidence	Implement low-confidence thresholds and mandatory escalation pathways
Usability in regulated workflows	Does the tool fit existing regulatory research, submission preparation, and quality-review processes?	Structured user testing, think-aloud sessions, time-on-task comparison, and qualitative feedback	Users can verify sources, understand answer status, and complete review actions without workflow disruption	Users trust generated answers without checking citations	Improve interface design, reviewer prompts, and draft-status labeling
Access control and confidentiality	Are product-specific documents visible only to authorized users?	Role-based access testing and audit-log review	Users retrieve only documents permitted by	Confidential commitments or manufacturing details	Tighten permissions, document

			role, product, region, and document status	are exposed to unauthorized users	segregation, and access-log monitoring
Auditability and inspection readiness	Can the organization reconstruct how an answer was generated, reviewed, and reused?	End-to-end audit trail inspection	Query, retrieved passages, prompt version, answer version, reviewer action, and final disposition are preserved	Approved text cannot be traced to source evidence or reviewer decision	Enforce immutable audit logs and controlled export procedures
Lifecycle monitoring	Does system performance remain reliable as documents, models, prompts, and regulatory expectations change?	Periodic validation review, drift checks, document-currency audits, and post-deployment issue review	Updated documents are indexed correctly, superseded material is retired, and changes are controlled	Outdated sources remain active or prompt updates change answer behavior unexpectedly	Establish change control, periodic revalidation, and governed improvement cycles

Limitations

Dependency on Document Quality and Currency

The proposed system would be only as reliable as the documents it retrieves, the metadata assigned to them, and the governance process used to keep them current. If internal manuals are outdated, inconsistent, or poorly structured, the system could retrieve technically obsolete passages even if the RAG workflow itself functions as designed [3]. Regulatory AI governance literature indicates that monitoring and lifecycle management are essential when AI tools are used in controlled environments. For this reason, deployment would require disciplined document management, version control, periodic review, and clear procedures for retiring superseded content.

Handling of Emergent and Unstructured Regulatory Questions

The system may struggle with questions that lack explicit guidance, require interpretation of emerging technologies, or depend on business policy and regulatory precedent not present in the indexed corpus. Pharmaceutical innovation reviews show that AI is increasingly relevant across discovery, formulation, manufacturing, quality control, and post-market activities, but these broad opportunities also create interpretation challenges for novel contexts [26]. In such cases, the system should retrieve the closest relevant evidence and identify uncertainty rather than presenting a definitive recommendation. Expert review would remain especially important when the question concerns new manufacturing platforms, novel analytical technologies, or strategic regulatory positioning.

Conclusion

A retrieval-augmented AI system for interpreting ICH quality guidelines and product control strategies would provide pharmaceutical professionals with a structured way to query complex regulatory and internal quality knowledge. By combining document ingestion, metadata-enriched retrieval, controlled generation, and citation grounding, the system could transform scattered regulatory text into a usable decision-support resource.

The principal strength of this approach is that it keeps answers tied to authoritative source documents rather than treating the language model as an independent source of regulatory truth. Its value would be greatest where teams need traceable, reviewable, and product-specific explanations that connect guideline principles to practical control-strategy decisions.

Important challenges remain, including maintaining a current knowledge base, preserving metadata quality, preventing unsupported synthesis, and building trust among users who operate under GMP accountability. The system would also need careful governance to ensure that generated content remains advisory until approved through established quality and regulatory review processes.

Pilot implementations within pharmaceutical quality and regulatory affairs departments would be a practical next step. Such pilots should examine whether the system improves consistency, accelerates regulatory research, and supports clearer preparation of development rationales, submissions, and health-authority responses.

Acknowledgments: None

Conflict of interest: None

Financial support: None

Ethics statement: None

References

1. Yang S, Hu X, Zhu J, Zheng B, Bi W, Wang X, et al. Aspects and implementation of pharmaceutical quality by design from conceptual frameworks to industrial applications. Pharmaceutics. 2025;17(5):623.
2. Kim JY, Choi DH. Control strategy for excipient variability in the quality by design approach using statistical analysis and predictive model: Effect of microcrystalline cellulose variability on design space. Pharmaceutics. 2022;14(11):2416.
3. Zhu Z. Intelligent information management enables quality-by-design in pharmaceutical production. Sci Rep. 2025;15(1):44201.
4. Arden NS, Fisher AC, Tyner K, Yu LX, Lee SL, Kopcha M. Industry 4.0 for pharmaceutical manufacturing: Preparing for the smart factories of the future. Int J Pharm. 2021;602:120554.
5. Schmidt A, Helgers H, Vetter FL, Zobel-Roos S, Hengelbrock A, Strube J. Process automation and control strategy by quality-by-design in total continuous mRNA manufacturing platforms. Processes. 2022;10(9):1783.
6. Waikar S, Bhat AG, Ramanathan M. Retrieval Augmented Generation (RAG) for Evaluating Regulatory Compliance of Drug Information and Clinical Trial Protocols. CPT Pharmacometrics Syst Pharmacol. 2026;15(3):e70201.
7. Proestel S, Jeng LJ, Smith C, Deady M, Amer O, Ahmed M, et al. Semantic search of FDA guidance documents using generative AI. Ther Innov Regul Sci. 2025;59(5):1148-59.
8. Kim J, Hur M, Min M. From rag to qa-rag: Integrating generative AI for pharmaceutical regulatory compliance process. In: Proceedings of the 40th ACM/SIGAPP Symposium on Applied Computing; 2025. p. 1293-5.
9. Lee J, Yoon W, Kim S, Kim D, Kim S, So CH, et al. BioBERT: a pre-trained biomedical language representation model for biomedical text mining. Bioinformatics. 2020;36(4):1234-40.
10. Gu Y, Tinn R, Cheng H, Lucas M, Usuyama N, Liu X, et al. Domain-specific language model pretraining for biomedical natural language processing. ACM Trans Comput Healthc. 2021;3(1):1-23.
11. Beltagy I, Lo K, Cohan A. SciBERT: A pretrained language model for scientific text. In: Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP); 2019. p. 3615-20.
12. Koppula M, Madhulika F, Sreeramoju N, Kolimi P. AI-Powered Chatbot for FDA Drug Labeling Information Retrieval: OpenAI GPT for Grounded Question Answering. Analytics. 2025;4(4):33.
13. Liu X, Rivera SC, Moher D, Calvert MJ, Denniston AK, SPIRIT-AI and CONSORT-AI Working Group. Reporting guidelines for clinical trials evaluating artificial intelligence interventions: the CONSORT-AI extension. Nat Med. 2020;26(9):1364-74.
14. Liu X, Rivera SC, Moher D, Calvert MJ, Denniston AK, Ashrafian H, et al. Reporting guidelines for clinical trial reports for interventions involving artificial intelligence: the CONSORT-AI extension. Lancet Digit Health. 2020;2(10):e537-48.
15. Vasey B, Nagendran M, Campbell B, Clifton DA, Collins GS, Denaxas S, et al. Reporting guideline for the early stage clinical evaluation of decision support systems driven by artificial intelligence: DECIDE-AI. BMJ. 2022;377.
16. Mayer M, Canedo A, Dinh T, Low M, Ortiz A, Garay C. Potential use of artificial intelligence for regulatory intelligence: biopharmaceutical industry's views. Ther Innov Regul Sci. 2019;53(6):759-66.
17. Patil RS, Kulkarni SB, Gaikwad VL. Artificial intelligence in pharmaceutical regulatory affairs. Drug Discov Today. 2023;28(9):103700.
18. Singhal K, Azizi S, Tu T, Mahdavi SS, Wei J, Chung HW, et al. Large language models encode clinical knowledge. Nature. 2023;620(7972):172-80.
19. Moor M, Banerjee O, Abad ZS, Krumholz HM, Leskovec J, Topol EJ, et al. Foundation models for generalist medical artificial intelligence. Nature. 2023;616(7956):259-65.
20. Saraswat N, Li C, Jiang M. Identifying the question similarity of regulatory documents in the pharmaceutical industry by using the recognizing question entailment system: evaluation study. JMIR AI. 2023;2(1):e43483.
21. Ying L, Liu Z, Fang H, Kusko R, Wu L, Harris S, et al. Text summarization with ChatGPT for drug labeling documents. Drug Discov Today. 2024;29(6):104018.
22. Chen T, Mo H, Wang T, Jiang C, Liu Z, Hou F, et al. A lightweight large language model for regulatory affairs translation in pharmaceutical industry. Sci Rep. 2025;15(1):37992.
23. Nene L, Flepisi BT, Brand SJ, Basson C, Balmith M. Evolution of drug development and regulatory affairs: The demonstrated power of artificial intelligence. Clin Ther. 2024;46(8):e6-14.
24. Alsentzer E, Murphy J, Boag W, Weng WH, Jindi D, Naumann T, et al. Publicly available clinical BERT embeddings. In: Proceedings of the 2nd Clinical Natural Language Processing Workshop; 2019. p. 72-8.
25. Huang K, Altosaar J, Ranganath R. Clinicalbert: Modeling clinical notes and predicting hospital readmission. arXiv preprint arXiv:1904.05342. 2019.
26. Huanbutta K, Burapapadh K, Kraisit P, Sriamornsak P, Ganokratanaa T, Suwanpitak K, et al. Artificial intelligence-driven pharmaceutical industry: A paradigm shift in drug discovery, formulation development, manufacturing, quality control, and post-market surveillance. Eur J Pharm Sci. 2024;203:106938.
27. Simões MF, Silva G, Pinto AC, Fonseca M, Silva NE, Pinto RM, et al. Artificial neural networks applied to quality-by-design: From formulation development to clinical outcome. Eur J Pharm Biopharm. 2020;152:282-95.