International Journal of Advanced Artificial Intelligence Research (IJAAIR)
eISSN: 3087-4270
Volume03 Issue11 | Page No: - 11-23

Received: 11 Oct 2026, Revised: 18 Oct 2026, Accepted: 25 Oct 2026, Published: 03 Nov 2026

DOI: https://doi.org/10.55640/ijaair-v03i11-02

Compliance By Architecture: A Human-AI Operating Model For Medical Inquiry Management Conceptual Framework Paper

Bilal Saiyad
Independent Researcher, USA

ABSTRACT: Medical Information functions are exploring artificial intelligence (AI) and advanced analytics to improve the speed, consistency, scalability, and intelligence of Medical Inquiry management. Generative AI can support intake, information extraction, classification, scientific-content retrieval, evidence synthesis, response drafting, documentation, and analytics. These opportunities arise within a controlled environment in which scientific accuracy, response scope, source traceability, privacy, safety reporting, product-quality escalation, documentation, and human accountability are essential. This paper proposes Compliance by Architecture: a practitioner-informed model in which controls are embedded across the complete inquiry pathway rather than concentrated at a final review gate. The framework was developed through a practitioner-informed conceptual synthesis of professional Medical Information guidance, peer-reviewed literature, pharmacovigilance and electronic-record requirements, and responsible-AI principles and frameworks. The model comprises seven interconnected layers: channels and intake; identity, context, and reference data; governed scientific content and knowledge; AI orchestration and retrieval; workflow, routing, and the system of record; analytics and monitoring; and human decision and accountability. It also introduces an automation boundary that distinguishes AI-accelerated activities, human-reviewed activities, and human-owned decisions or escalations. The framework integrates established Medical Information practice, pharmacovigilance expectations, electronic-record controls where applicable, and responsible-AI principles while remaining explicitly conceptual rather than regulator-issued. Its intended contribution is a practical architecture for deploying AI as a governed scientific-workflow capability, not as an autonomous scientific authority. Future empirical work should test the framework across jurisdictions, inquiry types, content models, and human-AI workflows.

KEYWORDS: Medical Affairs; Medical Information; Medical Inquiry; artificial intelligence; generative AI; responsible AI; scientific exchange; compliance; analytics; retrieval-augmented generation; pharmacovigilance

1. Introduction

Medical Information is a scientific interface through which life-sciences organizations respond to unsolicited questions from healthcare professionals and other stakeholders. Published professional guidance characterizes the function as a source of accurate, balanced, timely, and question-specific scientific information and as a point of coordination with safety, quality, regulatory, legal, clinical, and other functions [1-4].

The operating environment is increasingly multichannel. Requests may arrive through contact centers, email, web forms, digital portals, congress interactions, Field Medical personnel, customer relationship management workflows, and other authorized channels [1,5]. At the same time, the scientific evidence needed to answer those requests is distributed across product information, scientific response documents, publications, study data, guidelines, and controlled internal repositories. This combination creates an information-management problem and a decision-governance problem.

Generative AI can assist with transcription, information extraction, classification, retrieval, summarization,

drafting, documentation, and pattern analysis. However, large language models can produce inaccurate, incomplete, or unsupported statements, and users may over-rely on outputs that appear fluent or authoritative [6-8]. The central question is therefore not how much of the inquiry process can be automated, but how automation can be placed inside an architecture that preserves evidence, review, escalation, and accountability.

This paper proposes Compliance by Architecture: a model in which compliance-aware controls operate at every consequential boundary of the inquiry lifecycle. The proposal does not replace jurisdiction-specific legal analysis, pharmacovigilance procedures, content governance, or organizational standard operating procedures. Instead, it offers a vendor-neutral framework for connecting those obligations and practices to AI-enabled workflow design.

The problem addressed is that established Medical Information, pharmacovigilance, electronic-record, privacy, and responsible-AI sources define important but dispersed expectations; they do not provide a single operating architecture for assigning AI and human authority across the end-to-end Medical Inquiry pathway. The objective is to synthesize these expectations into a practical, testable model for governed AI assistance in Medical Inquiry management. The paper makes four specific contributions: (1) a seven-layer architecture that connects intake, data, content, AI, workflow, analytics, and accountability; (2) a task-level automation boundary distinguishing AI-accelerated, human-reviewed, and human-owned activities; (3) provenance and monitoring measures for evaluating both outputs and workflow behavior; and (4) a maturity progression for controlled adoption. Sections 2 and 3 establish the conceptual basis and lifecycle; Sections 4-14 present the architecture, controls, analytics, and design principles; and Sections 15-17 discuss comparison, limitations, future research, and conclusions.

2. Medical Inquiry Management as an Architecture Problem

Medical Inquiry management is often summarized as receive, research, respond, and document. The DRESS scientific process—Define, Research, Evaluate, Synthesize, and Share—provides a more rigorous model for developing Medical Information responses [2]. Operational execution nevertheless requires additional decisions concerning requester context, inquiry classification, safety and product-quality screening, routing, content eligibility, review, fulfillment, recordkeeping, and secondary analytics.

These decisions frequently span separate platforms and teams. An inquiry may be captured correctly but associated with the wrong requester; a scientifically accurate response may rely on superseded content; a potential safety signal may be detected but routed incorrectly; or a completed response may be inadequately documented. The control environment therefore depends on the chain connecting data, content, workflow, AI behavior, and human decisions—not solely on the quality of a final document.

Professional Medical Information guidance supports narrow, request-responsive communications, controlled response documents, documented review and maintenance, accurate records, and processes for detecting and escalating potential adverse events (AEs) and product-quality complaints (PQCs) [1,3-5]. EMA pharmacovigilance guidance similarly emphasizes quality systems, validation, follow-up, duplicate management, controlled terminology, traceability, and clear responsibilities for suspected adverse-reaction information received through medical-inquiry and product-information services [9]. These sources do not prescribe the architecture proposed here; they establish the operational expectations that the architecture is designed to support.

2.1 Conceptual Synthesis and Source Selection

This paper is a practitioner-informed conceptual synthesis rather than a systematic or scoping review. Sources were selected purposively from four categories relevant to the design problem: professional Medical Information guidance and peer-reviewed practice literature; official pharmacovigilance guidance and regulations; electronic-record and privacy requirements where potentially applicable; and recognized responsible-AI and generative-AI risk-management frameworks. The synthesis compared these sources by the operational function they address—scientific response practice, signal interception, records and traceability, AI risk, monitoring, and human oversight—and translated the resulting requirements into design layers and decision boundaries. Professional Medical Information guidance primarily informed the lifecycle stages, governed-content controls, and human review and accountability layer; pharmacovigilance sources informed safety and product-quality

interception and the related workflow and recordkeeping controls; electronic-record and privacy provisions informed access, audit-trail, and system-of-record controls; and responsible-AI frameworks informed the AI orchestration, monitoring, and change-management controls and the automation boundary. No preregistered protocol, exhaustive database search, date-bounded search strategy, formal evidence-grading process, or risk-of-bias assessment was used; the method is therefore intended to support conceptual framework development, not claims of systematic-review completeness.

2.2 Established Foundations and Proposed Integration

Prior work establishes several foundations used in this paper. Medical Information guidance describes question-specific scientific responses, controlled response documents, content maintenance, documentation, service measures, and coordination with safety and quality functions [1-5]. Pharmacovigilance sources establish expectations for receipt, validation, follow-up, duplicate management, recordkeeping, and reporting of suspected adverse reactions [9,10]. Electronic-record and privacy provisions define controls within their respective scopes [13,14]. NIST, WHO, FDA, and EMA sources contribute risk-based governance, context-of-use reasoning, testing, transparency, monitoring, and human-oversight principles [6-8,11,12]. The contribution proposed here is not to replace or restate those sources as a new standard. It is to integrate their dispersed expectations into a Medical Inquiry-specific operating model: the seven-layer architecture, task-level automation-boundary assignments, end-to-end decision lineage, AI-era monitoring measures, and maturity progression are the manuscript's conceptual proposals.

End-to-End Medical Inquiry Lifecycle

The proposed lifecycle extends the DRESS scientific process into a complete operational pathway. It is intended as a vendor-neutral design model rather than a universal regulatory sequence.

Capture and intake: retain the original inquiry, channel, source, date/time, and relevant contextual information.

Identity and context validation: establish requester type, identity or account match, geography, product, language, and applicable privacy context.

Inquiry classification: assign controlled product, topic, request-type, complexity, and priority attributes while preserving the original wording.

Safety, product-quality, privacy, and sensitivity screening: identify candidate signals and route them for qualified assessment.

Routing and triage: assign the accountable queue, function, owner, service expectation, and escalation path.

Evidence and governed-content retrieval: search sources eligible for the requester, jurisdiction, product, topic, and intended use.

Response preparation: organize evidence, prepare a question-specific response, and assemble citations.

Human scientific review and approval: apply the review appropriate to the inquiry, content type, risk, and local procedure.

Delivery and follow-up: communicate through an appropriate channel and capture clarifications or additional information.

System-of-record documentation: retain inquiry, response, sources, decisions, approvals, escalations, changes, and timestamps.

Analytics and learning: monitor operations, quality, scientific-information needs, content performance, and AI behavior using governed data.

Safety screening requires particular precision. Under US postmarketing requirements, applicants must promptly review adverse-experience information obtained from any source and maintain written procedures for surveillance, receipt, evaluation, and reporting [10]. EMA guidance describes minimum information for a valid individual case safety report (ICSR) and expects incomplete cases to be recorded and followed up [9]. An inquiry architecture can support detection and routing, but qualified safety personnel retain assessment, disposition, follow-up, and reporting responsibility.

Compliance by Architecture: The Seven-Layer Model

Compliance by Architecture treats control as a distributed system property. Identity controls operate during matching; content controls operate before retrieval; provenance controls operate during evidence selection and generation; safety controls operate during intake and processing; workflow controls separate draft, review, approval, delivery, and submission; and monitoring controls operate throughout the deployed lifecycle. As shown in Figure 1, these controls span seven connected architecture layers and are governed by an automation boundary that returns consequential decisions to qualified people.

COMPLIANCE BY ARCHITECTURE LAYERS	CONTROL EMPHASIS
1. Channels and intake	AI-accelerated capture; governed provenance
2. Identity, context, and reference data	Suggested matches; human confirmation
3. Governed scientific content and knowledge	Eligibility, version, audience, and use controls
4. AI orchestration and retrieval	Restricted retrieval, citations, abstention, logging
5. Workflow, routing, and system of record	Review states, access, audit trail, escalation
6. Analytics and monitoring	Quality, drift, exceptions, and feedback
7. Human decision and accountability	Human-owned approval, disposition, and submission

Figure 1. Proposed practitioner-informed Compliance by Architecture framework for Medical Inquiry management, presented in tabular form as the seven architecture layers and their primary control emphasis. Row shading indicates the predominant automation-boundary category for each layer: blue denotes AI-accelerated activity; amber denotes human-reviewed activity; red denotes human-owned decisions or escalation. Activity-level boundary assignments are shown in Table 2. The framework is conceptual and not regulator-issued.

Table 1 summarizes the primary purpose of each architecture layer and representative controls embedded at that layer.

Table 1. Seven architecture layers and representative embedded controls

Layer	Primary purpose	Representative embedded controls
Channels and intake	Capture inquiries from authorized channels while preserving origin and context.	Original wording; channel provenance; timestamps; disclosure/consent handling; date of awareness.
Identity, context, and reference data	Establish who is asking, where, about what, and under which use context.	Confirmed versus suggested matches; requester type; geography; product master; controlled vocabularies; data minimization.
Governed scientific content and knowledge	Determine which evidence and response content are eligible for use.	Owner; status; version; effective/expiry dates; audience; market; intended use; localization and supersession lineage.
AI orchestration and	Structure records, recommend	Restricted corpora; retrieval filters; citations;

Layer	Primary purpose	Representative embedded controls
retrieval	actions, retrieve evidence, summarize, and support drafting.	confidence/abstention; prompt/output logging; use-case testing.
Workflow, routing, and system of record	Manage assignment, escalation, review, delivery, follow-up, and final documentation.	Role-based access; mandatory fields; segregation of duties; state transitions; audit trails; attestations where applicable.
Analytics and monitoring	Measure operations, scientific demand, quality, and AI behavior.	Metric lineage; data-quality rules; overrides; unsupported-output review; drift; exceptions; feedback.
Human decision and accountability	Retain scientific, safety, approval, escalation, and submission authority.	Named decision rights; qualified reviewers; escalation authority; final attestation; accountability.

Governed Content, Retrieval-Augmented Generation, and Provenance

Retrieval-augmented generation (RAG) can connect a generative model to selected, updateable knowledge sources. It does not, by itself, guarantee scientific accuracy or appropriate use. Retrieval can return incomplete, outdated, noisy, or contextually inappropriate material, while generation can misstate or overextend the retrieved evidence [7,8].

For Medical Inquiry management, RAG should operate as a controlled evidence pathway. A content-eligibility gate should evaluate source type, approval or authorization status, version, effective and review dates, product, topic, audience, geography, intended use, confidentiality, localization relationship, and supersession status before material becomes available to generation. Professional guidance for scientific and global response documents supports controlled creation, review, localization, maintenance, parent-child linkage, and regular updating of response content [3,4].

Provenance should extend beyond a bibliographic citation. The relevant lineage is: source → version → retrieved passage → generated statement → reviewer modification → approved response → delivered communication → system-of-record entry. This chain is intended to enable reviewers and auditors to determine not only what source was cited, but how the source contributed to the consequential output. NIST's Generative AI Profile recommends source and citation verification, provenance assessment, context-specific testing, monitoring, and documentation of human overrides and changes [7].

The Medical Inquiry Automation Boundary

The automation boundary classifies activities according to the authority that AI may exercise. AI-accelerated activities may operate within validated constraints. Human-reviewed activities permit AI to propose an output but require confirmation before consequential use. Human-owned activities may use AI to organize evidence, but final assessment, approval, escalation, or accountability remains with qualified personnel. This task-level approach is consistent with risk- and context-based AI governance principles from NIST, WHO, FDA, and EMA, while the specific boundary assignments shown in Table 2 are the paper's proposed recommendations [6,7,11,12].

Table 2. Proposed automation-boundary matrix

Lifecycle activity	Appropriate AI contribution	Boundary	Human responsibility
Transcription and structuring	Transcribe and extract explicit fields.	A/R	Verify consequential text and correct errors.
Missing-field detection	Flag absent required information.	A	Determine whether clarification or follow-up is required.
Identity or account matching	Rank candidate matches.	R	Confirm ambiguous identity or account selection.
Duplicate detection	Identify potentially related inquiries or records.	R/H	Confirm linkage, merge, or separate disposition.
Inquiry classification	Suggest topic, product, request type, complexity, and priority.	R	Confirm ambiguous, off-label, complex, or high-risk classification.
Safety/PQC/privacy screening	Detect candidate terms, narratives, or patterns.	A/R	Assess and escalate under qualified safety, quality, or privacy processes.
Routing recommendation	Propose owner, queue, and priority.	A/R	Resolve exceptions and ownership conflicts.
Governed-content retrieval	Search eligible sources and apply metadata filters.	A	Confirm relevance, coverage, and permitted use.
Evidence summarization	Summarize retrieved evidence with source links.	R	Evaluate quality, balance, limitations, and applicability.
Draft-response support	Assemble a question-specific draft.	R	Exercise scientific judgment and approve or edit as required.
Citation assembly	Link statements to retrieved sources.	A/R	Verify that cited evidence supports each consequential claim.
System writeback preparation	Populate a draft record and identify required fields.	R	Confirm data, attestations, and final status or submission.
Final externally delivered response	No autonomous authorization.	H	Approve and deliver under the applicable role and procedure.

Lifecycle activity	Appropriate AI contribution	Boundary	Human responsibility
AE/PQC disposition or regulatory submission	Extract and organize source information only.	H	Qualified personnel make determinations and remain accountable.
Trend analytics and quality monitoring	Cluster topics, sample records, and identify anomalies.	A/R	Interpret significance, investigate exceptions, and authorize corrective action.

Boundary key: A = AI-accelerated (AI may operate within validated constraints); R = human-reviewed (AI may propose an output, but human confirmation is required before consequential use); H = human-owned (AI may organize information, but final assessment, approval, escalation, or accountability remains with qualified personnel). Combined classifications indicate that the boundary depends on the case or output: A/R denotes activities that may be AI-accelerated for routine outputs but require human review where outputs are consequential, ambiguous, or high-risk; R/H denotes activities in which AI output is human-reviewed and the final linkage or disposition decision is human-owned.

Safety and Product-Quality Interception

Medical Information is a potential point of receipt for adverse-event and product-quality information [1,5,9]. The architecture should therefore screen incoming and follow-up content early, preserve source wording, record the date of awareness where applicable, and route candidate signals promptly. AI may improve sensitivity by identifying relevant language in unstructured text, but detection is not determination.

The model should not assign final AI authority over ICSR validity, seriousness, expectedness, causality, product-quality disposition, or regulatory reportability. In the United States, 21 CFR 314.80 establishes postmarketing adverse-drug-experience reporting and recordkeeping requirements for covered applicants and nonapplicants [10]. EMA GVP Module VI provides European guidance on collection, management, validation, follow-up, duplicate detection, data quality, and reporting of suspected adverse reactions [9]. Actual responsibilities and timelines depend on the governing regulation, product, contractual allocation, and company procedure.

A defensible interception pattern is: inquiry → AI-assisted candidate detection → visible workflow flag → qualified human assessment → controlled safety or quality process. Performance evaluation should emphasize missed-signal risk, not only overall classification accuracy, and should include difficult, incomplete, multilingual, and indirect narratives.

Workflow, Electronic Records, and Decision Lineage

The workflow layer converts recommendations into controlled actions. It should make state transitions explicit: draft is not equivalent to reviewed; reviewed is not equivalent to approved; approved is not equivalent to delivered; and delivered is not equivalent to final system submission. This separation is intended to help prevent AI-supported drafting or writeback from crossing an accountability boundary without an authorized decision.

Where an electronic record falls within the scope of FDA predicate rules and 21 CFR Part 11, Part 11 establishes criteria for trustworthy and reliable electronic records and signatures, including validation, authorized access, time-stamped audit trails, operational and authority checks, record protection, training, and controlled documentation [13]. Part 11 should not be described as automatically applicable to every Medical Inquiry record. Applicability requires a record-specific assessment.

Similarly, HIPAA technical safeguards apply to covered entities and business associates handling electronic protected health information within the regulation's scope. They include access control, audit controls, integrity, authentication, and transmission security [14]. A life-sciences organization should determine its data

role and applicable privacy regime rather than assuming that every inquiry workflow is governed by HIPAA.

AI-enabled auditability should preserve consequential decision lineage: the original inquiry; extracted fields; suggested classifications; human changes; sources and versions retrieved; generated content; reviewer modifications; escalations; delivery status; and final writeback. The purpose is not indefinite retention of every technical event, but sufficient reconstruction of decisions under applicable retention and privacy requirements.

Analytics as a Value Layer and a Control Layer

Established Medical Information measures include inquiry volume, service level, turnaround time, topic, response-document use, escalation, quality monitoring, and customer experience [1,4,5]. These measures support resource planning, content maintenance, service improvement, and understanding of external information needs.

AI-enabled inquiry management also requires control analytics. The measures in Table 3 are proposed, not regulator-defined. They should be interpreted together and by context; for example, a high override rate may indicate poor model performance, or it may show that users are appropriately challenging recommendations in complex cases.

Table 3. Proposed AI-era monitoring measures

Monitoring domain	Illustrative measures	Interpretive caution
Retrieval and provenance	Eligible-source coverage; passage relevance; source/version traceability; citation support.	High retrieval volume does not establish scientific relevance or completeness.
Generated output quality	Groundedness; unsupported-claim rate; omission rate; question-scope adherence.	Aggregate accuracy can conceal failures in high-risk inquiry types.
Classification and routing	Precision/recall by class; ambiguous-case escalation; duplicate-detection review.	Class balance and local taxonomies affect performance.
Safety and quality interception	Candidate-signal sensitivity; missed-signal review; escalation timeliness; reconciliation.	AI output does not replace qualified case assessment.
Human-AI interaction	Edit rate; override rate; abstention acceptance; reviewer time; user feedback.	Low override may reflect either strong performance or automation bias.
Operational control	Audit completeness; state-transition exceptions; access violations; unresolved queues.	Targets should reflect risk and workflow context.
Change and drift	Performance after model, prompt, content, taxonomy, language, or workflow changes.	A stable model can still degrade when its environment changes.

Human Decision Rights and Context-of-Use Governance

Meaningful human oversight requires more than an approval button. It should define who reviews, what evidence is visible, what the reviewer may change or override, which cases require escalation, how disagreements are recorded, and who remains accountable. NIST and WHO both emphasize governance, defined tasks, human oversight, monitoring, transparency, and stakeholder participation across the AI lifecycle [6,7].

Risk should be evaluated at the level of the use case and decision consequence. Summarizing an approved response for an internal specialist differs materially from selecting content for an external response, preparing a customized synthesis, or autonomously communicating with a healthcare professional. FDA's January 2025 draft guidance on AI used to support regulatory decision-making for drugs and biological products proposes a risk-based credibility framework tied to a defined context of use [11]. Its formal scope does not make it direct regulation of ordinary Medical Inquiry systems; its context-of-use logic is used here as a design analogy. EMA's adopted reflection paper on AI across the medicinal-product lifecycle likewise provides lifecycle principles, not a Medical Inquiry-specific architecture mandate [12].

Validation, Monitoring, and Change Management

Evaluation should extend beyond model accuracy to the complete sociotechnical workflow. Relevant dimensions include retrieval quality, content eligibility, groundedness, citation correctness, classification and routing performance, candidate safety-signal detection, human factors, workflow integrity, auditability, robustness to incomplete or unusual inquiries, and change impact. NIST recommends testing in conditions representative of deployment, measuring human overrides, verifying sources and citations, monitoring for confabulation and other risks, and reassessing after changes [7].

Generative-AI performance can change when the model, prompt, retrieval logic, content library, taxonomy, product information, language coverage, or workflow changes. Governance should therefore include versioning, change assessment, regression testing, deployment approval, post-release monitoring, incident management, rollback, periodic review, and controlled expansion. Responsible AI is a lifecycle discipline rather than a one-time approval event.

From Inquiry Analytics to Governed Medical Intelligence

Aggregated inquiry patterns can reveal recurring scientific questions, content gaps, customer-experience friction, emerging information needs, and hypotheses about evidence or education gaps [1,4,5]. However, a question is not automatically an insight, and an inquiry trend does not by itself establish causation or justify stakeholder targeting. Secondary use should be governed by defined scientific purpose, appropriate aggregation, privacy, data quality, and human contextual interpretation.

A disciplined learning pathway is: inquiry data → aggregation → pattern detection → contextual review → validated insight → documented action → outcome review. This could create a bridge from inquiry fulfillment to broader Medical Intelligence without converting individual inquiries into automatic promotional or targeting signals.

Maturity Model

The proposed maturity model describes a practical progression rather than an industry benchmark. Organizations may operate at different stages across products, geographies, channels, and use cases. Table 4 summarizes the five-stage progression and the primary achievement associated with each stage.

Table 4. Proposed maturity progression for AI-enabled Medical Inquiry management

Stage	Model	Defining characteristics	Primary achievement
1	Manual and fragmented	Disconnected channels, manual classification and routing, fragmented content, limited analytics, procedure-heavy controls.	Local consistency
2	Digitized workflow	Centralized capture, structured fields, explicit states, controlled repositories, system-of-record discipline.	Traceability
3	Governed automation	Rules-based routing, completeness checks, controlled taxonomies, integrated content, standardized escalation.	Scalable process control
4	Compliance-aware AI assistance	AI-assisted classification, governed retrieval, summarization, drafting, citation support, explicit decision gates, monitoring.	Acceleration without transfer of accountability
5	Monitored human-AI inquiry intelligence	Risk-tiered automation, continuous monitoring, advanced analytics, governed learning, integration with Medical Intelligence.	Adaptive and accountable collaboration

Practical Design Principles

Preserve the original inquiry. AI interpretation should supplement, not replace, the source request.
Separate fact from inference. Distinguish requester-provided data, mastered data, AI suggestions, and human-confirmed decisions.
Govern content before generation. Determine eligibility before information is exposed to retrieval or drafting.
Make provenance reviewable. Link consequential statements to source passages and versions.
Design for abstention. Insufficient evidence or unresolved ambiguity should trigger clarification or escalation.
Intercept regulated signals early. Candidate safety, quality, and privacy signals should enter controlled workflows promptly.
Separate draft from decision. AI-prepared content or records should not silently become approved, delivered, or submitted.
Monitor the model and the workflow. Technical performance, human behavior, exceptions, and operational outcomes all matter.
Define human decision rights. Specify authority, evidence, escalation, and accountability at each consequential boundary.
Learn without overinterpreting. Inquiry patterns require contextual validation before strategic use.

Discussion

The proposed framework reframes AI adoption from a tool-selection exercise into a design of scientific decision rights. Its main contribution is the integration of layers that are often governed separately: data and identity, content, retrieval, workflow, analytics, and human accountability. The automation boundary makes the division of labor explicit, while provenance and control analytics are intended to create the evidence needed to evaluate that division over time.

Compared with established Medical Information frameworks, this model retains their emphasis on accurate, balanced, request-responsive communication, controlled scientific response documents, content maintenance, documentation, customer experience, and coordination with safety and quality functions [1-5]. It

extends those practices by mapping them to an end-to-end technical and operating architecture, including content-eligibility gates, RAG provenance, explicit workflow states, and task-level AI decision boundaries. Compared with recognized AI-risk frameworks, NIST and WHO provide broad lifecycle governance, risk management, transparency, evaluation, monitoring, and human-oversight principles [6-8], while FDA and EMA materials contribute context-of-use and medicinal-product lifecycle perspectives [11,12]. The present framework applies and connects those general principles to the specific work objects, handoffs, escalation points, and accountable decisions of Medical Inquiry management. This comparison is an integration claim, not a claim that the framework is regulator-issued or that the underlying concepts originated here.

The framework does not assume that manual work is inherently safer or that automation is inherently risky. Manual processes can miss signals, use outdated information, or produce inconsistent documentation. Properly designed automation may improve completeness, traceability, consistency, and timeliness. Conversely, automation can scale mistakes and create unwarranted confidence. The appropriate comparison is therefore between complete operating models under defined contexts of use, not between humans and models in isolation.

The likely future is neither fully manual nor fully autonomous. It is a hybrid scientific operating model in which AI performs increasing amounts of information processing while qualified professionals retain responsibility for consequential scientific, safety, quality, privacy, and external-communication decisions. The boundary may evolve as evidence, technology, regulation, and organizational capability mature, but it should evolve deliberately and transparently.

Limitations and Future Research

This paper is conceptual and practitioner-informed. The seven-layer architecture, automation-boundary matrix, monitoring measures, and maturity model have not yet been validated through comparative implementation studies. They should be treated as hypotheses for design and evaluation rather than as established standards.

Requirements vary across jurisdictions, products, requester types, information channels, data categories, and organizational roles. AI regulation and guidance are evolving. Organizations should therefore perform use-case-specific legal, regulatory, pharmacovigilance, quality, privacy, security, and records assessments.

Future research should evaluate: (1) human factors and automation bias in Medical Information review; (2) RAG retrieval, groundedness, and citation-validation methods for scientific response workflows; (3) sensitivity and operational performance of AI-assisted adverse-event and product-quality interception; (4) multilingual and cross-jurisdiction performance; (5) the effect of explicit automation boundaries on quality, timeliness, workload, and trust; and (6) methods for connecting aggregated inquiry intelligence to evidence planning without inappropriate inference or secondary use.

Conclusion

AI has the potential to reduce administrative burden, accelerate retrieval and synthesis, support consistent documentation, and reveal patterns across large volumes of Medical Inquiry data. Its reliability nevertheless depends on the integrity of the complete pathway: the inquiry captured, requester context, eligible evidence, content version, generated synthesis, safety interception, workflow controls, review, and final accountability.

Compliance should therefore not sit at the end of the process. It should operate throughout intake, context, content, AI, workflow, analytics, and human decision-making. Within this architecture, AI retrieves, structures, recommends, summarizes, drafts, and monitors. When uncertainty, scientific judgment, regulated information, external authorization, or consequential decisions arise, authority returns explicitly to qualified people.

The objective is not autonomous Medical Information. It is a traceable, scientifically grounded, continuously monitored human-AI system in which technology is designed to accelerate inquiry management while humans retain judgment, responsibility, and accountability.

Declarations

Funding: No external funding was received for preparation of this conceptual manuscript.

Conflicts of interest: The author declares that this manuscript is an independent practitioner-informed conceptual paper. No company-confidential information, internal metrics, proprietary architecture, named partner workflows, or identifiable operational examples are included.

Ethics approval: Not applicable. No human participants, patient data, or identifiable individual-level data were used.

Data availability: Not applicable. No dataset was generated or analyzed.

Disclaimer: The framework is provided for scholarly and industry discussion. It does not constitute legal, regulatory, pharmacovigilance, privacy, security, quality, or medical advice.

References

1. Medical Affairs Professional Society; Medical Information Leaders in Europe; Pharma Collaboration for Transparent Medical Information. Essential elements of a best-in-class Medical Information organization: standards and guidance. 2021. Available from: https://www.mile-association.org/wp-content/uploads/2023/06/MedInfo-SG-2022-12-1.pdf

2. Albano D, Pragga F, Rai R, Flowers T, Parmar P, Wnorowski S, et al. The Medical Information scientific process: Define, Research, Evaluate, Synthesize, and Share (DRESS). Ther Innov Regul Sci. 2022;56(3):405-414. doi:10.1007/s43441-021-00366-w.

3. Hermes-DeSantis ER, Johnson RM, Redlich A, et al. Proposed best practice guidelines for scientific response documents: a consensus statement from phactMI. Ther Innov Regul Sci. 2020;54(6):1303-1311. doi:10.1007/s43441-020-00151-1.

4. Wnorowski S, Niemoller F, Glaetzer A, Tamburri R, Vine E, Flick F, et al. Best practices for Medical Information global response documents from phactMI and MILE. Pharm Med. 2026;40:87-94. doi:10.1007/s40290-025-00598-3.

5. Guillot P, Shea J, Shoemaker S, Rocco M, DeLuca M, Rai R, Hermes-DeSantis ER. Medical Information delivering improved customer experience: a guide. Pharm Med. 2023;37(2):111-119. doi:10.1007/s40290-023-00469-9.

6. World Health Organization. Ethics and governance of artificial intelligence for health: guidance on large multi-modal models. Geneva: WHO; 2024. Available from: https://www.who.int/publications/i/item/9789240084759

7. National Institute of Standards and Technology. Artificial Intelligence Risk Management Framework: Generative Artificial Intelligence Profile. NIST AI 600-1. Gaithersburg (MD): NIST; 2024. doi:10.6028/NIST.AI.600-1.

8. National Institute of Standards and Technology. Artificial Intelligence Risk Management Framework (AI RMF 1.0). NIST AI 100-1. Gaithersburg (MD): NIST; 2023. doi:10.6028/NIST.AI.100-1.

9. European Medicines Agency. Guideline on good pharmacovigilance practices (GVP): Module VI – Collection, management and submission of reports of suspected adverse reactions to medicinal products. Rev 2. 2017. Available from: https://www.ema.europa.eu/en/documents/regulatory-procedural-guideline/guideline-good-pharmacovigilance-practices-gvp-module-vi-collection-management-submission-reports_en.pdf

10. Electronic Code of Federal Regulations. Title 21, §314.80: Postmarketing reporting of adverse drug experiences. Current electronic edition. Available from: https://www.ecfr.gov/current/title-21/chapter-I/subchapter-D/part-314/subpart-B/section-314.80

11. US Food and Drug Administration. Considerations for the use of artificial intelligence to support regulatory decision-making for drug and biological products. Draft guidance for industry and other interested parties. January 2025. Available from: https://www.fda.gov/media/184830/download

12. European Medicines Agency. Reflection paper on the use of artificial intelligence in the medicinal product lifecycle. EMA/CHMP/CVMP/83833/2023. Adopted September 2024. Available from:

https://www.ema.europa.eu/en/documents/scientific-guideline/reflection-paper-use-artificial-intelligence-ai-medicinal-product-lifecycle_en.pdf

13. Electronic Code of Federal Regulations. Title 21, Part 11: Electronic records; electronic signatures. Current electronic edition. Available from: https://www.ecfr.gov/current/title-21/chapter-I/subchapter-A/part-11

14. Electronic Code of Federal Regulations. Title 45, §164.312: Technical safeguards. Current electronic edition. Available from: https://www.ecfr.gov/current/title-45/subtitle-A/subchapter-C/part-164/subpart-C/section-164.312