American Journal of Healthcare Strategy

DOI 10.61449/ajhcs.2026.8 | E-ISSN 2995-6242

Published in Volume 4, Issue 2, on June 18, 2026

Closing the Insights to Action Gap: How Agentic AI Is Reshaping Medical Affairs and Scientific Strategy in Pharma

Emily Lewis, MS, HIMSS CPDHTS, CCRP

UCB

Disclaimer: The views expressed here are those of the author and do not reflect those of UCB. Any discussion of process is provided for illustrative purposes only.

Abstract

Medical Affairs has spent the better part of two decades trying to reconcile irreconcilable demands: the imperative to move at the speed of science, the obligation to uphold regulatory and compliance standards, and the pressure to demonstrate strategic value beyond commercial support. The arrival of generative artificial intelligence offered partial relief, primarily by accelerating document drafting and literature summarization. The emergence of agentic AI, referring to systems that can plan, execute, and iterate across multi-step workflows with limited human intervention, represents something categorically different. This article examines the strategic implications of agentic AI adoption for Medical Affairs functions within biopharmaceutical organizations. It distinguishes agentic architectures from earlier AI paradigms, anchors the discussion in a five-domain map of where agents are already finding traction across the function (Figure 1), surveys representative use cases within each domain, and identifies the organizational capabilities, governance structures, and risk frameworks that determine whether such deployments succeed or stall. The paper closes with a concrete set of priorities for healthcare leaders tasked with navigating this transition, framed around the

American Journal of Healthcare Strategy

DOI 10.61449/ajhcs.2026.8 | E-ISSN 2995-6242

Published in Volume 4, Issue 2, on June 18, 2026

recognition that agentic systems do not simply augment the workforce but reconstitute how scientific engagement itself is designed, measured, and delivered.

Keywords: agentic AI; medical affairs; pharmaceutical strategy; healthcare leadership; AI governance; scientific exchange; medical science liaisons; evidence synthesis

Introduction: An Inflection Point

Medical Affairs occupies a strange position within the pharmaceutical enterprise. Once regarded as a back-office support function, it has been elevated over the past decade into what industry analysts frequently describe as the third strategic pillar alongside research and commercial.[1] That elevation has not been accompanied by a commensurate investment in how the function operates. Medical information call centers still run on ticketing systems designed in the early 2000s. Medical Science Liaisons (MSLs) log field insights into customer relationship management systems and hope, often in vain, that someone synthesizes those signals into strategic intelligence. Publication plans are assembled in spreadsheets. Congress coverage is compiled by hand.

Against this backdrop, a new technology paradigm has arrived. Agentic AI, broadly understood as artificial intelligence systems capable of autonomous reasoning, tool use, and multi-step task execution, has moved from research curiosity to production reality in non-medical contexts like software engineering and customer service.[2] In Medical Affairs, adoption has been slower and more cautious, for reasons that are both appropriate, given that the stakes of getting it wrong include patient safety, and lamentable, given that risk aversion can itself become a form of inaction.

Healthcare leaders responsible for Medical Affairs at biopharmaceutical companies cannot afford a wait-and-see posture. The functional case for agentic AI adoption is compelling on its own merits. The strategic case becomes urgent when considered against the backdrop of constrained headcount, accelerating therapeutic complexity, and a generation of clinicians who increasingly expect scientific information to be available instantly, accurately, and on their terms.

American Journal of Healthcare Strategy

DOI 10.61449/ajhcs.2026.8 | E-ISSN 2995-6242

Published in Volume 4, Issue 2, on June 18, 2026

The economic context further sharpens the urgency. Industry benchmarks suggest that Medical Affairs budgets have grown only modestly over the past five years while the scope of the function has expanded to encompass real-world evidence, patient engagement, digital health partnerships, and increasingly sophisticated scientific exchange across a widening array of therapeutic modalities.[3] Workforce surveys conducted by major professional societies have consistently shown that Medical Affairs professionals report working well beyond standard hours to keep pace with demand. Agentic AI is among the few interventions that credibly addresses this structural gap without requiring a proportional increase in headcount, which makes its adoption not merely attractive but, for many organizations, financially necessary.

From Generative to Agentic: A Definitional Distinction That Matters

The language around artificial intelligence has grown imprecise. Within Medical Affairs, the terms “generative AI,” “copilot,” “assistant,” and “agent” are often used interchangeably by vendors, executives, and operational teams. The distinction matters more than it might appear.

Generative AI refers to models that produce content, whether text, images, audio, or code, in response to a prompt. The model receives an instruction and returns an output. A user who asks a large language model to summarize a clinical trial publication receives a summary. Whether that summary is accurate, appropriately contextualized, or suitable for a given audience depends entirely on the user’s ability to evaluate it.

Agentic AI introduces three additional capabilities that transform the nature of the interaction.[4] First, agents can plan. Given a goal rather than a single instruction, an agentic system decomposes the goal into sub-tasks, sequences those sub-tasks, and adjusts the plan based on intermediate results. Second, agents can use tools. Where a generative model is confined to the content of its training data plus whatever is included in a prompt, an agent can query databases, search the literature, invoke statistical packages, retrieve clinical trial registries, and access internal systems via application programming interfaces. Third, agents can iterate. An

American Journal of Healthcare Strategy

DOI 10.61449/ajhcs.2026.8 | E-ISSN 2995-6242

Published in Volume 4, Issue 2, on June 18, 2026

agent that produces a draft literature review can critique its own work, identify gaps, retrieve additional sources, and refine the output across multiple cycles before presenting a result to a human reviewer.

These capabilities, individually, are not new. What is new is their integration into systems reliable enough for production use in regulated environments. Frameworks such as LangGraph, CrewAI, and AutoGen, coupled with model-native tool-use capabilities in offerings from Anthropic, OpenAI, and Google, have moved agentic architectures from proofs of concept into software that can be deployed, monitored, and audited.

For Medical Affairs, this distinction changes the strategic question from how to use AI to draft documents faster, to which workflows can be handed off to an autonomous system under appropriate oversight. The answer to the first question is incremental. The answer to the second is architectural.

A Map of the Territory

Before turning to specific use cases, it is useful to orient. Figure 1 offers a map of where AI-assisted methods and autonomous agentic workflows are being applied across the Medical Affairs lifecycle. The figure organizes the function into five interconnected domains that follow the natural flow of medical work: from upstream strategy and planning, through KOL and HCP engagement, into publications and scientific communications, across medical information and content review, and out to real-world evidence and outcomes research. Within each domain, the figure distinguishes between AI-assisted methods that augment a single human task and autonomous agentic workflows that execute multi-step processes end-to-end with human review at the points where judgment matters most.

The recurring patterns summarized at the foot of the figure deserve early attention because they reappear throughout the analysis that follows. Document generation agents tend to follow a draft–self-check–revise–route loop. Insight synthesis agents follow a collect–cluster–theme–route pattern. Continuous monitoring agents follow a

American Journal of Healthcare Strategy

DOI 10.61449/ajhcs.2026.8 | E-ISSN 2995-6242

Published in Volume 4, Issue 2, on June 18, 2026

monitor–detect–investigate–alert pattern. In each case, the agent completes a multi-step task and surfaces output for human review rather than acting unilaterally. This shared architecture is what makes the figure something more than a taxonomy: it is a description of how agentic work actually gets done inside a regulated function.

The five domains in Figure 1 will serve as the structural backbone for the rest of this article. Each subsequent section zooms into one column of the figure, explains why agentic approaches are well suited to that part of the function, and identifies the deployment patterns most likely to deliver durable value. Readers who want to keep their bearings as the discussion proceeds may find it useful to keep the figure visible.

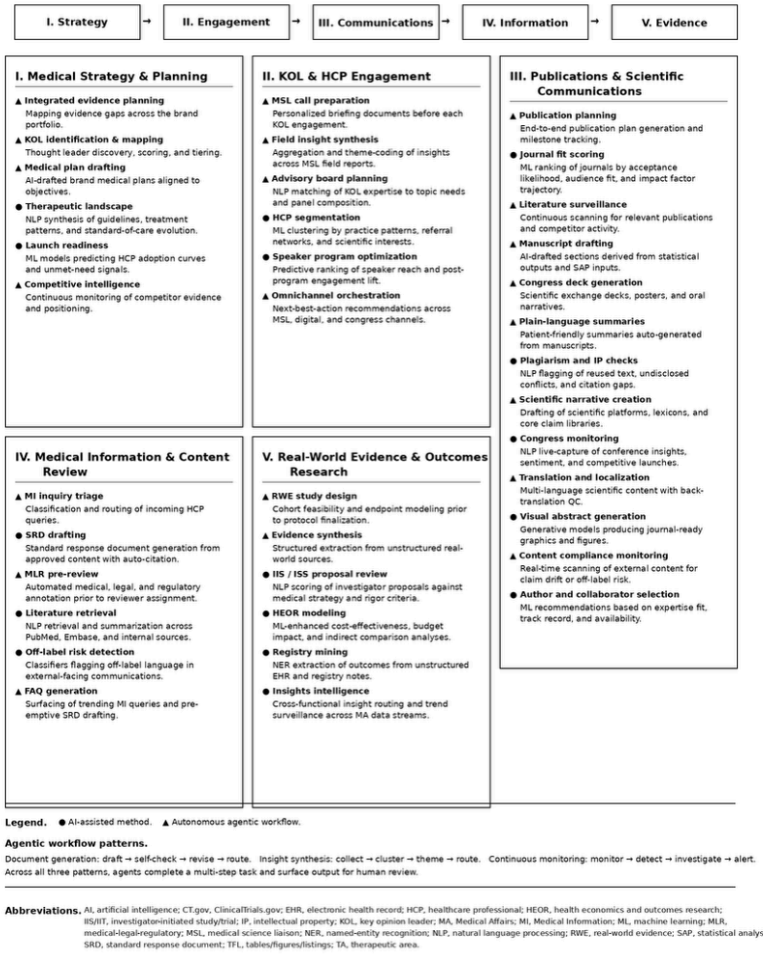

American Journal of Healthcare Strategy

DOI 10.61449/ajhcs.2026.8 | E-ISSN 2995-6242

Published in Volume 4, Issue 2, on June 18, 2026

Figure 1. Applications of artificial intelligence and autonomous agents across the Medical Affairs lifecycle. A schematic of where AI-assisted methods (●) and autonomous agentic workflows (▲) are applied across five functional domains, from medical strategy through real-world evidence generation. Recurring agentic workflow patterns are summarized at the foot of the figure.[5–18]

Why Medical Affairs Is Uniquely Suited for Agentic AI

Several structural features of Medical Affairs work make the function an unusually strong fit for agentic AI deployment.[19]Recognition of these features helps explain why early adopters are already reporting significant productivity and quality gains, and why laggards face a widening gap. These features also explain why every column of Figure 1 has populated faster than skeptics predicted.

The first feature is the density of knowledge work. Medical Affairs is, at its core, a function that produces, curates, and communicates scientific information. The unit of labor is a document, a response, a synthesis, or an insight. These outputs are precisely the kinds of artifacts that modern language models can produce with increasing fidelity.

The second feature is the availability of structured reference material. Unlike many commercial functions, Medical Affairs operates against a corpus of highly structured content: approved product labels, published literature indexed in PubMed, clinical trial registries, congress abstracts, and curated medical information libraries. This structured substrate provides the grounding that agentic systems require to produce outputs traceable to authoritative sources.

The third feature is the strict separation between scientific exchange and promotional activity. This regulatory boundary, often viewed as a constraint, in fact simplifies the design of agentic workflows. The scope of permissible output is well-defined. The audiences are well-characterized. The escalation paths, to medical directors, to pharmacovigilance, to regulatory affairs, are codified in standard operating procedures.

American Journal of Healthcare Strategy

DOI 10.61449/ajhcs.2026.8 | E-ISSN 2995-6242

Published in Volume 4, Issue 2, on June 18, 2026

The fourth feature, and perhaps the most strategically salient, is the chronic mismatch between demand and capacity. Medical information functions routinely field inquiries that require specialist review but cannot, for economic reasons, staff specialists at the volume required. MSL teams are asked to cover therapeutic areas across geographies that no individual could realistically manage. Publications teams operate under perpetual deadline pressure. In each of these cases, an agentic system that can perform a substantial portion of the initial reasoning, content generation, or synthesis, with human review at the points where judgment matters most, addresses a structural problem.

Five Domains of Application

The remainder of this section walks through Figure 1 column by column. Each domain is treated as a distinct territory with its own dominant workflow patterns, its own risk profile, and its own readiness criteria. The intent is not to catalog every possible deployment but to illustrate, with enough specificity to be useful, what high-quality agentic work looks like inside each part of the function.

I. Medical Strategy and Planning

Zooming into the leftmost column of Figure 1, the strategy and planning domain is where agentic AI most directly threatens to upend assumptions about how Medical Affairs work begins. The domain encompasses integrated evidence planning, KOL identification and mapping, medical plan drafting, therapeutic landscape synthesis, launch readiness modeling, and continuous competitive intelligence (Figure 2). Several of these activities have, until recently, depended on consultancy engagements and bespoke analytics projects that took months to deliver and were stale by the time they landed.

American Journal of Healthcare Strategy

DOI 10.61449/ajhcs.2026.8 | E-ISSN 2995-6242

Published in Volume 4, Issue 2, on June 18, 2026

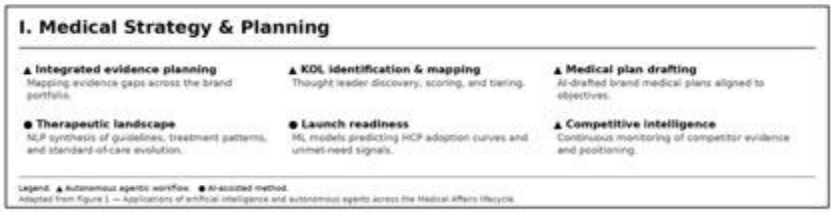

Figure 2. Domain I — Medical Strategy and Planning. Six representative agentic and AI-assisted applications spanning evidence planning, KOL identification, medical plan drafting, therapeutic landscape synthesis, launch readiness modeling, and competitive intelligence. ▲ denotes autonomous agentic workflows; ● denotes AI-assisted methods.

Agentic systems collapse this timeline. A continuous competitive intelligence agent can monitor publications, congress abstracts, regulatory filings, and clinical trial registry updates against a defined therapeutic area, surface signals daily rather than quarterly, and route material findings to the medical strategy team with traceable provenance. A therapeutic landscape agent can synthesize evolving treatment guidelines, real-world treatment patterns, and standard-of-care shifts into living documents that update as new evidence emerges. The shift is not from manual to automated. It is from periodic to continuous.

KOL identification and mapping warrant particular attention within this domain. Traditional KOL analytics relied on structured data from third-party providers and produced ranked lists of clinicians by publication count or clinical trial activity. Agentic systems layer reasoning over this data and answer more nuanced questions: which clinicians have recently changed their treatment patterns in a given geography, which emerging investigators are likely to become influential within two years, or which patient advocacy groups have begun commenting on a specific therapeutic question. The shift is from descriptive analytics to anticipatory analytics.

The ethical dimensions of this use case deserve explicit consideration. Combining publicly available data in novel ways can produce inferences that individual clinicians did not

DOI 10.61449/ajhcs.2026.8 | E-ISSN 2995-6242

Published in Volume 4, Issue 2, on June 18, 2026

anticipate when their data became part of commercial databases. Medical Affairs leaders should ensure that KOL analytics derived from agentic systems are used to inform scientific engagement rather than to profile or manipulate, and that internal governance explicitly prohibits uses inconsistent with the scientific exchange mission. Transparency about the existence and purpose of such analytics, where appropriate in the context of individual engagements, builds rather than erodes trust.

AI-drafted brand medical plans, the third agentic workflow in this column, also deserve commentary. The risk in this application is not that the agent produces low-quality output. Modern agents produce competent first drafts. The risk is that the medical plan, as a strategic artifact, loses its function as a forcing mechanism for cross-functional alignment if the drafting work is offloaded entirely. The most successful deployments observed to date use the agent to produce a comprehensive draft and then deliberately route it through a structured human deliberation process before approval.

II. KOL and HCP Engagement

Moving one column to the right, the engagement domain is where field medical teams generate enormous volumes of qualitative data and where the latency between signal generation and strategic response has historically been measured in months. Typical Medical Affairs organizations systematically analyze less than one percent of field interaction notes.[1] Figure 1 distinguishes within this domain between AI-assisted methods (HCP segmentation, speaker program optimization) and fully agentic workflows (MSL call preparation, field insight synthesis, advisory board planning, omnichannel orchestration); Figure 3 expands the column for closer reading.

American Journal of Healthcare Strategy

DOI 10.61449/ajhcs.2026.8 | E-ISSN 2995-6242

Published in Volume 4, Issue 2, on June 18, 2026

Figure 3. Domain II — KOL and HCP Engagement. Six representative applications spanning MSL call preparation, field insight synthesis, advisory board planning, HCP segmentation, speaker program optimization, and omnichannel orchestration. ▲ denotes autonomous agentic workflows; ● denotes AI-assisted methods.

Multi-agent architectures compress field-insight latency by ingesting interaction notes, classifying them against a predefined taxonomy of scientific themes, cross-referencing the results with external signals from congress coverage and publication activity, and producing trending reports that surface emerging questions, concerns, or misconceptions in the clinician community. In this paradigm, senior medical directors receive weekly synthesis reports where they previously received quarterly ones. The collect–cluster–theme–route pattern called out at the foot of Figure 1 describes precisely this workflow.

Pre-call briefing for MSLs is a more visible application within the same column. The agent retrieves a clinician's publication history, prior interaction notes, institutional affiliations, and relevant competitive intelligence, and produces a briefing document in advance of each scheduled meeting. Well-prepared MSLs gain the bandwidth to build deeper scientific relationships rather than spending the hour before each interaction triaging context.

Human-authored field insights are vulnerable to availability bias, in which recent or emotionally salient interactions receive disproportionate narrative weight. Agentic summarization, properly designed, can weigh insights by frequency, recency, and source diversity, producing a corrected view that strategy teams can trust as a representative reflection of the field rather than a vivid anecdote. Realizing this benefit requires careful design of the

American Journal of Healthcare Strategy

DOI 10.61449/ajhcs.2026.8 | E-ISSN 2995-6242

Published in Volume 4, Issue 2, on June 18, 2026

insight taxonomy and ongoing calibration against human expert judgment, neither of which is trivial.

Omnichannel orchestration, the last entry in this column, is the most ambitious. A next-best-action agent that coordinates engagement across MSL, digital, and congress channels must reconcile inputs from disparate systems and respect the regulatory firewall between scientific and promotional activity at every step. Organizations that attempt this without first putting the lower-risk applications in the same column into production tend to underestimate the architectural lift involved.

III. Publications and Scientific Communications

The third column of Figure 1 covers publications planning, manuscript drafting, congress deck generation, plain-language summaries, congress monitoring, translation and localization, and several adjacent activities (Figure 4). Publication planning has historically been a ceremony of spreadsheets. Agentic systems can bring substantial value here, functioning as operational copilots for publication teams that track the status of each planned output, monitor deadlines, flag authorship conflicts, and draft sections of manuscripts based on clinical study reports under the supervision of named authors.

American Journal of Healthcare Strategy

DOI 10.61449/ajhcs.2026.8 | E-ISSN 2995-6242

Published in Volume 4, Issue 2, on June 18, 2026

Figure 4. Domain III — Publications and Scientific Communications. Thirteen representative applications spanning publication planning, journal fit scoring, literature surveillance, manuscript drafting, congress deck generation, plain-language summaries, plagiarism and IP checks, scientific narrative creation, congress monitoring, translation and localization, visual abstract generation, content compliance monitoring, and author and collaborator selection. ▲ denotes autonomous agentic workflows; ● denotes AI-assisted methods.

These deployments remain controversial, and legitimately so. The International Committee of Medical Journal Editors and several leading journals have issued explicit guidance that AI systems cannot be listed as authors and that human authors remain accountable for every claim in a manuscript.[22] Responsible deployments respect this boundary. The agent drafts; the authors deliberate, revise, and take accountability for what is published under their names.

An emerging and underappreciated use case within this domain is the production of plain language summaries and patient-facing scientific communications. Regulatory requirements in several jurisdictions, including the European Union Clinical Trials Regulation, now mandate lay summaries of trial results.[23] Translating dense clinical content into accessible prose at a sixth-grade reading level is a labor-intensive task that agentic systems perform remarkably well, subject to human review by medical writers and patient advocates. Organizations with active clinical trial portfolios have an opportunity to meet regulatory obligations more efficiently while also serving the broader goal of patient-centric communication.

Congress monitoring deserves brief amplification. The agent ingests presentation abstracts and live session content, captures sentiment and competitive launches in near real time, and produces structured intelligence that reaches strategy teams the same day rather than weeks after the meeting concludes. The monitor–detect–investigate–alert pattern from the foot of Figure 1 governs the architecture here.

American Journal of Healthcare Strategy

DOI 10.61449/ajhcs.2026.8 | E-ISSN 2995-6242

Published in Volume 4, Issue 2, on June 18, 2026

Across this entire column, the dominant pattern is document generation: draft, self-check against compliance and citation rules, revise, route for human approval. The pattern is the same whether the artifact is a manuscript, a poster, a plain-language summary, or a translated regulatory document. What changes is the rubric the agent checks itself against, which is a function of the artifact's audience and risk profile.

IV. Medical Information and Content Review

The fourth column of Figure 1 sits closest to the regulated edge of Medical Affairs. It encompasses medical information inquiry triage, standard response document drafting with auto-citation, medical-legal-regulatory pre-review, literature retrieval, off-label risk detection, and FAQ generation (Figure 5). This column is also where the majority of contemporary agentic deployments are concentrated, in part because the work is high-volume, well-documented, and operates against a corpus of approved content.

Figure 5. Domain IV — Medical Information and Content Review. Six representative applications spanning MI inquiry triage, SRD drafting, MLR pre-review, literature retrieval, off-label risk detection, and FAQ generation. ▲ denotes autonomous agentic workflows; ● denotes AI-assisted methods.

Medical information functions historically operate on a tiered model. Tier-one staff handle routine inquiries using pre-approved standard response documents. Tier-two specialists handle custom inquiries requiring literature synthesis. Tier-three experts handle complex or

American Journal of Healthcare Strategy

DOI 10.61449/ajhcs.2026.8 | E-ISSN 2995-6242

Published in Volume 4, Issue 2, on June 18, 2026

sensitive inquiries. Response time targets typically range from twenty-four hours for standard requests to five business days for custom literature reviews.

Agentic systems reshape this model by collapsing tiers one and two for a meaningful percentage of incoming volume. Modern systems generate draft responses to medical information inquiries, including appropriate literature citations and adherence to organizational response templates.[20] The system retrieves the relevant standard response document if one exists, searches internal and external literature, composes a response, cites sources, checks the output against a compliance rubric, and surfaces the draft for review with traceable provenance for every claim. Once drafted, human medical information specialists review and approve each response before transmission. Time saved is then redirected to tier-three inquiries and proactive content development.

Regulatory inspections of medical information functions examine not only the content of responses but the processes by which those responses were produced. Agentic systems that cannot produce a complete chain of evidence connecting every factual claim in a response to a specific source and a specific reasoning step will fail inspection. Production deployments must therefore treat this chain of evidence as a first-class artifact, persisted alongside the response itself and available for review after transmission. Medical Affairs leaders evaluating vendors would do well to inspect this provenance capability carefully, as the gap between marketing claims and production reality in this area can be wide.

The same column also captures content review activity, a use case sometimes treated as separate but architecturally similar. Medical-legal-regulatory committees operate under significant time pressure. Agents can perform initial compliance screening of materials against regulatory standards, flag potentially problematic claims, and suggest remediations. Human reviewers retain final decision authority, but the volume of material each reviewer can process increases substantially. Off-label risk detection, the natural pair to MLR pre-review, applies the

American Journal of Healthcare Strategy

DOI 10.61449/ajhcs.2026.8 | E-ISSN 2995-6242

Published in Volume 4, Issue 2, on June 18, 2026

same classifier-based approach to outbound communications and external content. Both applications follow the monitor–detect–investigate–alert pattern.

V. Real-World Evidence and Outcomes Research

The rightmost column of Figure 1 covers real-world evidence study design, evidence synthesis, investigator-initiated and -sponsored study proposal review, health economics and outcomes research modeling, registry mining, and cross-functional insights intelligence (Figure 6). This is also the domain where the strategic stakes and the methodological complexity are both highest.

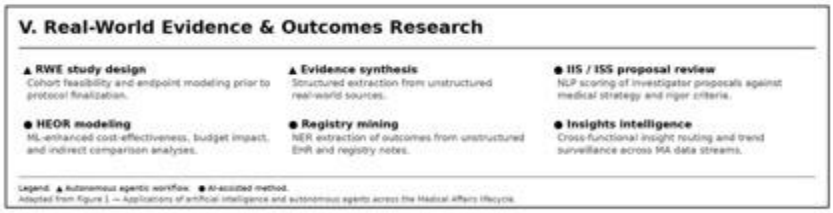

Figure 6. Domain V — Real-World Evidence and Outcomes Research. Six representative applications spanning RWE study design, evidence synthesis, IIS/ISS proposal review, HEOR modeling, registry mining, and cross-functional insights intelligence. ▲ denotes autonomous agentic workflows; ● denotes AI-assisted methods.

Systematic literature reviews and evidence syntheses are labor-intensive exercises that often consume months of time. Agentic systems have demonstrated credible performance on the initial stages of these workflows: search strategy formulation, title and abstract screening, data extraction, and preliminary synthesis.[21] Agentic systems can process in hours what a human team can process in weeks. Importantly, agents like these allow human reviewers to focus on adjudication and critical appraisal, where human judgment remains essential.

Parallel developments are occurring in real-world evidence workflows. Agents that can query structured claims and electronic health record data, generate cohort definitions, execute

DOI 10.61449/ajhcs.2026.8 | E-ISSN 2995-6242

Published in Volume 4, Issue 2, on June 18, 2026

preliminary analyses, and draft results sections of study reports are now possible. Registry mining using named-entity recognition over unstructured EHR notes, called out explicitly in Figure 1, extends the same pattern to data sources that were until recently considered inaccessible at scale. This type of deployment raises important questions about the validation of agent-generated analyses, questions that the regulatory community has not yet fully addressed.

Agentic systems should always require human checkpoints at moments of consequential decision, such as cohort inclusion criteria, handling of missing data, and interpretation of borderline findings. Designing these checkpoint structures is itself a strategic capability that Medical Affairs organizations need to develop, and one in which biostatistics, epidemiology, and data science teams should be deeply involved from the outset rather than consulted after the fact.

The final entry in this column, insights intelligence, deserves a closing note because it sits at the seam between this domain and Domain II. A cross-functional insights agent that surveils all five columns of the figure, identifies trends that span domains, and routes findings to the appropriate medical, commercial, or regulatory leader effectively functions as the connective tissue for the entire function. Organizations that have stood up this layer report that it produces the strategic value that early agentic deployments often promised but rarely delivered: a Medical Affairs function that thinks across its own activity rather than within individual silos.

Organizational Readiness: Governance, Data, and Talent

The gap between organizations that extract value from agentic AI and those that do not is rarely determined by the technology itself. The determining variables are organizational. Three in particular deserve attention.

Governance That Operates at the Speed of Deployment

Traditional governance models in regulated industries are designed for artifacts that change infrequently: a validated system, a locked-down database, a published standard operating procedure. Agentic AI systems do not fit this model. The underlying foundation models are

American Journal of Healthcare Strategy

DOI 10.61449/ajhcs.2026.8 | E-ISSN 2995-6242

Published in Volume 4, Issue 2, on June 18, 2026

updated continuously. The prompts that shape agent behavior can be modified in minutes. The tools that agents can invoke expand over time.

Healthcare leaders who attempt to govern agentic systems using pre-existing computer system validation frameworks will find themselves in one of two failure modes. Either the governance apparatus moves too slowly, and the organization ships nothing, or the governance apparatus is circumvented, and the organization ships without adequate oversight. Neither outcome is acceptable.

The emerging best practice is a tiered governance approach in which the level of oversight scales with the risk profile of the specific agent deployment. The columns of Figure 1 are useful here as a first-cut risk taxonomy. A customer-facing agent in Domain IV producing approved medical information responses for transmission to healthcare professionals warrants the highest level of scrutiny: validated system treatment, formal change control, locked prompts, and documented testing. An internal agent in Domain II producing literature summaries for MSL pre-read warrants lighter governance: documented intended use, sample output review, and periodic audit.

This tiered approach requires a governance body that can move quickly and one that includes medical, regulatory, legal, information technology, and business representation. Organizations that have created dedicated AI governance councils with defined decision rights and service-level agreements for review are moving measurably faster than those relying on ad hoc committees. A practical further refinement is the designation of standing templates for common deployment patterns, so that a new medical information agent does not require the reinvention of a validation approach that has already been approved for a prior, functionally similar system.

Data Infrastructure That Is Actually Usable

American Journal of Healthcare Strategy

DOI 10.61449/ajhcs.2026.8 | E-ISSN 2995-6242

Published in Volume 4, Issue 2, on June 18, 2026

Agentic systems are only as good as the information they can access. Many Medical Affairs organizations may discover that their institutional knowledge is locked in formats that agents cannot effectively reason over: scanned PDFs, legacy document management systems, email archives, individual MSL notebooks.

The work of preparing the institutional knowledge base for agentic access is unglamorous and substantial. It involves document extraction, metadata tagging, ontology alignment, and integration with modern retrieval systems. Organizations that treat this work as a one-time information technology project invariably underestimate it. Organizations that treat it as an ongoing capability, with dedicated product ownership and continuous investment, create a durable advantage that their competitors will find difficult to replicate.

Talent That Can Partner With Agents

The workforce implications of agentic AI in Medical Affairs are significant but not primarily a story of displacement. The story is closer to reconstitution. Tasks that once constituted the core of a junior medical information specialist's job will increasingly be performed by agents. Tasks that emphasize judgment, stakeholder communication, cross-functional problem solving, and oversight of automated systems will become more central to every Medical Affairs role.

The capability gap is real. Few Medical Affairs professionals have been trained to write effective prompts, evaluate agent outputs against quality rubrics, or diagnose the failure modes of autonomous systems. Organizations that invest in structured reskilling programs, including hands-on work with production tools rather than conceptual webinars, will see faster adoption and better outcomes.

Risks and Rational Guardrails

No serious discussion of agentic AI in Medical Affairs can omit the risks. Four deserve particular attention, and each maps cleanly onto specific columns of Figure 1.

DOI 10.61449/ajhcs.2026.8 | E-ISSN 2995-6242

Published in Volume 4, Issue 2, on June 18, 2026

The first is hallucination, the phenomenon in which language models produce fluent but incorrect content.[24] Mitigations include retrieval-augmented generation, grounding every claim in retrieved sources, and automated fact-checking against authoritative databases. None of these mitigations eliminate hallucination entirely. Human review at appropriate points in the workflow remains essential, particularly for any output that will reach an external audience. The risk concentrates in Domains III and IV, where outputs are most likely to be transmitted externally.

The second is bias. Foundation models reflect the distribution of their training data. A model trained predominantly on English-language publications from North American and European institutions may underperform on therapeutic areas or populations underrepresented in that corpus. Medical Affairs organizations with global remits must evaluate agent performance across the populations they serve, not only the populations well represented in the training data. This risk is most acute in Domain V, where evidence syntheses and real-world analyses can perpetuate underrepresentation if not actively interrogated.

The third is privacy. Protected health information and personally identifiable information cannot be processed through systems that do not meet applicable regulatory standards, including the Health Insurance Portability and Accountability Act in the United States and the General Data Protection Regulation in the European Union. This is not a new risk, but agentic systems expand the attack surface by giving models access to more data and more tools. Careful architectural decisions about which data flows through which systems, and under what access controls, are essential. The exposure here is greatest in the registry mining and EHR-adjacent applications in Domain V.

The fourth is over-reliance. An agent that produces consistently plausible-sounding output conditions its human reviewers to approve rather than to scrutinize. This is a well-documented cognitive phenomenon sometimes referred to as automation complacency. Organizations must design their review workflows to preserve human engagement, including through randomized audits, structured critique prompts, and rotation of reviewers. This risk is

American Journal of Healthcare Strategy

DOI 10.61449/ajhcs.2026.8 | E-ISSN 2995-6242

Published in Volume 4, Issue 2, on June 18, 2026

universal across all five domains of Figure 1 and arguably the hardest of the four to engineer around.

An Implementation Framework for Healthcare Leaders

For healthcare leaders considering how to move from interest to impact, the following framework offers a pragmatic sequence.

Begin with a clear articulation of the strategic question the function is trying to answer. Agentic AI is not a strategy. It is a capability that may or may not advance a particular strategy. The right question is not how to use AI. It is what the organization can accomplish with AI that it could not otherwise accomplish, and at what cost.

Prioritize deployments where the business case is strong, the risk is manageable, and the organizational readiness is sufficient. Within Figure 1, Domain IV use cases such as medical information response are common starting points precisely because they meet these criteria. Domain III use cases such as publications planning are another common starting point.

Invest in the infrastructure before investing in the visible applications. A well-curated knowledge base, a governance framework that can make decisions in days rather than months, and a talent base with baseline agentic AI literacy will determine the ceiling of what the organization can accomplish across all five domains.

Measure outcomes, not activity. The relevant measures are those that matter to the function’s strategic mission: speed and quality of medical information response, depth and timeliness of field insight, productivity and quality of evidence synthesis.

Partner aggressively across functional boundaries. The organizations making the fastest progress have established close working relationships between Medical Affairs, information technology, compliance, legal, and data science.

American Journal of Healthcare Strategy

DOI 10.61449/ajhcs.2026.8 | E-ISSN 2995-6242

Published in Volume 4, Issue 2, on June 18, 2026

Start narrow and go deep rather than broad and shallow. A common failure pattern involves launching a dozen simultaneous pilots across disparate workflows, each staffed with minimal resourcing and sponsored by a different executive. The resulting portfolio generates ambiguous signals about what works, consumes disproportionate governance bandwidth, and leaves the organization without a genuinely production-grade deployment it can point to. The inverse pattern, selecting a single workflow within a single column of Figure 1, investing heavily, and achieving genuine operational excellence, tends to produce both measurable business value and the organizational learning required for responsible expansion to additional domains.

Key Takeaways for Healthcare Leaders

The following priorities distill the preceding analysis into a set of near-term actions for healthcare executives responsible for Medical Affairs functions or for their oversight.

Establish a strategic thesis before evaluating vendors. The market for agentic AI applications in Medical Affairs is growing. Evaluating vendors without a clear internal view of the workflows most worth automating, mapped against the five domains in Figure 1, guarantees incoherent purchasing and disappointed users. Leaders should require that every proposed deployment articulate the specific workflow, the expected outcome, the human oversight model, and the measurement plan before procurement conversations begin.

Build governance in proportion to risk. A single governance framework that treats every AI deployment identically will be simultaneously too slow for low-risk applications and too lax for high-risk ones. Tiered governance, supported by a responsive review body with named accountability and informed by the column-specific risk profiles in Figure 1, is a prerequisite for responsible speed.

Invest in structured medical knowledge infrastructure. The durable competitive advantage in agentic AI is not the model. Foundation models are increasingly commoditized. The advantage lies in the organization's ability to bring its own proprietary knowledge, curated

DOI 10.61449/ajhcs.2026.8 | E-ISSN 2995-6242

Published in Volume 4, Issue 2, on June 18, 2026

literature, historical inquiry data, and internal scientific platforms into effective contact with those models. This investment should be resourced as a core capability rather than a side project.

Redesign roles around human-agent collaboration. The Medical Affairs workforce of the future will emphasize oversight, judgment, escalation, and strategic content development rather than routine response composition. Organizations that wait for natural attrition to reshape their workforce will lag those that proactively redesign role descriptions, career paths, and performance expectations.

Measure what matters. Leaders should resist vanity metrics and insist on outcome metrics. The right questions are whether medical information responses are faster and more accurate, whether field insights are reaching decision-makers in time to shape strategy, and whether evidence synthesis is enabling more timely engagement with the external scientific community.

Plan for iteration rather than completion. Agentic AI capabilities are improving at a rate that will render today's architectural decisions obsolete within two to three years. Leaders should budget not only for initial deployment but for continuous refresh, re-evaluation, and, where appropriate, retirement of systems that have been superseded.

Engage with policy and standards development. Regulatory frameworks for AI in life sciences are evolving rapidly, with consequential contributions from the Food and Drug Administration, the European Medicines Agency, and the International Council for Harmonisation.[25] Medical Affairs leaders who engage proactively with these processes, through industry associations, public comment, and direct participation in working groups, will help shape a regulatory environment more favorable to responsible deployment.

Center the patient in every deployment decision. Medical Affairs exists to ensure that patients benefit from the appropriate use of therapies. Every agentic AI deployment should be evaluated not only against metrics of efficiency or scientific quality but against its consequences

American Journal of Healthcare Strategy

DOI 10.61449/ajhcs.2026.8 | E-ISSN 2995-6242

Published in Volume 4, Issue 2, on June 18, 2026

for the patients the function serves. Deployments that accelerate access to accurate scientific information, improve the quality of clinician decision-making, or make trial participation more accessible deserve priority over deployments whose benefits accrue exclusively to internal stakeholders. This orientation is not a matter of sentiment but of strategic coherence: the long-term defensibility of Medical Affairs as a function depends on its credibility as a scientific interlocutor, and that credibility is built by demonstrable service to patient outcomes.

What This Moment Asks of Leaders

Medical Affairs is at an inflection point. The convergence of mature foundation models, practical agentic frameworks, and accumulating operational experience has created conditions under which autonomous AI systems can meaningfully contribute to the function's strategic mission across all five domains mapped in Figure 1.

Realizing the full potential of this moment will require healthcare leaders to move beyond the framing of AI as a productivity tool and toward the recognition that agentic systems are becoming part of the operational fabric of Medical Affairs. That shift requires new governance, new infrastructure, new talent strategies, and, most fundamentally, a new relationship between human expertise and autonomous systems.

The organizations that navigate this shift thoughtfully will find that their Medical Affairs functions become faster, more responsive, and more strategically consequential. The trajectory of agentic AI in Medical Affairs will not be determined solely by technology vendors or regulatory agencies. It will be shaped, in meaningful part, by the choices of the Medical Affairs leaders who commission, govern, and ultimately own these deployments. Those choices will determine whether the function evolves into a more effective scientific engine or whether it is reshaped by forces it neither understands nor directs. The moment calls for intellectual engagement with the underlying technology, pragmatic experimentation within appropriate guardrails, and a willingness to revise long-held assumptions about how Medical Affairs work is organized and

American Journal of Healthcare Strategy

DOI 10.61449/ajhcs.2026.8 | E-ISSN 2995-6242

Published in Volume 4, Issue 2, on June 18, 2026

performed. Leaders who bring curiosity, rigor, and patient-centric values to this work will find that agentic AI does not diminish the strategic importance of their function, but amplifies it.

American Journal of Healthcare Strategy

DOI 10.61449/ajhcs.2026.8 | E-ISSN 2995-6242

Published in Volume 4, Issue 2, on June 18, 2026

References

1. Doshi K, Jouannic E, Vogel D, et al. A vision for medical affairs 2030: Five priorities for patient impact. McKinsey & Company.2022. Available at: https://www.mckinsey.com/industries/life-sciences/our-insights/a-vision-for-medical-affairs-2030-five-priorities-for-patient-impact

2. Gartner. Gartner Predicts 40% of Enterprise Apps Will Feature Task-Specific AI Agents by 2026, Up from Less Than 5% in 2025. Press release. August 26, 2025. Available at: https://www.gartner.com/en/newsroom/press-releases/2025-08-26-gartner-predicts-40-percent-of-enterprise-apps-will-feature-task-specific-ai-agents-by-2026-up-from-less-than-5-percent-in-2025

3. Beals E, Foggitt V, Jain A, et al. Agentic AI advantage for pharma. McKinsey & Company. 2025. Available at: https://www.mckinsey.com/featured-insights/week-in-charts/agentic-ai-advantage-for-pharma

4. Chakraborty C, Bhattacharya M, Pal S, Lee SS. The Agentic Era: Why Biopharma Must Embrace Artificial Intelligence That Acts, Not Just Informs. Curr Res Biotechnol. 2025. Available at: https://pmc.ncbi.nlm.nih.gov/articles/PMC12048886/

5. Fröling E, Rajaeean N, Hinrichsmeyer KS, et al. Artificial Intelligence in Medical Affairs: A New Paradigm with Novel Opportunities. Pharmaceut Med. 2024. Available at: https://doi.org/10.1007/s40290-024-00536-9

6. IQVIA. Harnessing the Power of Artificial Intelligence in Medical Affairs. White paper. 2025. Available at: https://www.iqvia.com/library/white-papers/harnessing-the-power-of-artificial-intelligence-in-medical-affairs

DOI 10.61449/ajhcs.2026.8 | E-ISSN 2995-6242

Published in Volume 4, Issue 2, on June 18, 2026

7. Medical Affairs Professional Society (MAPS). GenAI in Medical Affairs: Use Cases. Available at: https://medicalaffairs.org/gen-ai-use-cases/

8. Ng KKY, Zhang PC. Advancing medical affair capabilities and insight generation through machine learning techniques. J Pharm Policy Pract. 2023. Available at: https://doi.org/10.1186/s40545-023-00670-w

9. Mijatović A, Žuljević MF, Uršić L, Marušić A. Responsible Use of Large Language Models in Manuscript Preparation. Curr Protoc. 2026. Available at: https://doi.org/10.1002/cpz1.70300

10. Tao Y, Li X, Yisha Z, et al. LitAutoScreener: development and validation of an automated literature screening tool in evidence-based medicine driven by large language models. Health Data Sci. 2025. Available at: https://doi.org/10.34133/hds.0322

11. International Society for Medical Publication Professionals (ISMPP). Position statement and call to action on artificial intelligence. Curr Med Res Opin. 2024. Available at: https://doi.org/10.1080/03007995.2023.2273139

12. Singhal S, Hegde B, Karmalkar P, et al. Weakly supervised learning for categorization of medical inquiries for customer service effectiveness. Front Res Metr Anal. 2021. Available at: https://doi.org/10.3389/frma.2021.683400

13. Lau J, Bisht S, Horton R, et al. Creation of scientific response documents for addressing product medical information inquiries: mixed method approach using artificial intelligence. JMIR AI. 2025. Available at: https://doi.org/10.2196/55277

14. Scherbakov D, Hubig N, Jansari V, et al. The emergence of large language models as tools in literature reviews: a large language model-assisted systematic review. J Am Med Inform Assoc. 2025. Available at: https://doi.org/10.1093/jamia/ocaf063

DOI 10.61449/ajhcs.2026.8 | E-ISSN 2995-6242

Published in Volume 4, Issue 2, on June 18, 2026

15. Gartlehner G, Kahwati L, Hilscher R, et al. Data extraction for evidence synthesis using a large language model: a proof-of-concept study. Res Synth Methods. 2024. Available at: https://doi.org/10.1002/jrsm.1710

16. EVERSANA. EVERSANA ORCHESTRATE MLR: an AI-powered solution to transform the medical, legal and regulatory process. 2025. Available at: https://www.eversana.com/2025/04/10/eversana-unveils-eversana-orchestrate-mlr-an-ai-powered-solution-to-transform-the-medical-legal-regulatory-process/

17. IQVIA. Medical Information Teams Embrace Artificial Intelligence. White paper. Available at: https://www.iqvia.com/library/white-papers/medical-information-teams-embrace-artificial-intelligence

18. Warner J, Prada Jardim A, Albera C. Artificial intelligence: applications in pharmacovigilance signal management. Pharmaceut Med. 2025. Available at: https://doi.org/10.1007/s40290-025-00561-2

19. Evers M, Fleming E, Ghatak A, et al. Artificial Intelligence in Medical Affairs: A New Paradigm with Novel Opportunities. Pharmaceut Med. 2024. Available at: https://pmc.ncbi.nlm.nih.gov/articles/PMC11473552/

20. Boston Consulting Group. Agentic AI in Biopharma: Game-changing Efficiency. 2025. Available at: https://www.bcg.com/publications/2025/agentic-ai-in-biopharma-game-changing-efficiency

21. Gilson A, Safranek CW, Huang T, et al. LLM-assisted systematic review of large language models in clinical medicine. Nat Med.2026. Available at: https://www.nature.com/articles/s41591-026-04229-5

DOI 10.61449/ajhcs.2026.8 | E-ISSN 2995-6242

Published in Volume 4, Issue 2, on June 18, 2026

22. International Committee of Medical Journal Editors. Recommendations for the Conduct, Reporting, Editing, and Publication of Scholarly Work in Medical Journals. Updated January 2024. Available at: https://www.icmje.org/recommendations/browse/roles-and-responsibilities/defining-the-role-of-authors-and-contributors.html

23. European Parliament and Council. Regulation (EU) No 536/2014 of 16 April 2014 on clinical trials on medicinal products for human use. Official Journal of the European Union. Available at: https://eur-lex.europa.eu/eli/reg/2014/536/oj/eng

24. Hatem R, Simmons B, Thornton JE. A Call to Address AI "Hallucinations" and How Healthcare Professionals Can Mitigate Their Risks. Cureus. 2023. Available at: https://pmc.ncbi.nlm.nih.gov/articles/PMC10552880/

25. U.S. Food and Drug Administration, Center for Drug Evaluation and Research. Artificial Intelligence for Drug Development. Available at: https://www.fda.gov/about-fda/center-drug-evaluation-and-research-cder/artificial-intelligence-drug-development

Disclosures: The author is an employee of UCB and holds stock in the company.