

White Paper

Harnessing the Power of Artificial Intelligence in Medical Affairs

An exploration of practical applications and responsible adoption

NIAMH MCGUINNESS, Director, Pharma Solutions, IQVIA Applied AI Science



Table of contents

Introduction	1
Navigating Medical Affairs challenges	1
Responsible adoption is essential for effective transformation	2
Defensible AI applications for Medical Affairs	4
Literature review: Unlocking efficiency and insight	5
Revolutionizing workflows with AI agents	5
Consolidating RWD for actionable insights	7
Chart review: Reducing manual burden through AI-assisted information extraction	8
Case study: Streamlining chart review with NLP	8
KOL identification and profiling: Enhancing precision and engagement	9
Medical communications: Reshaping scientific dissemination	10
Conclusion: Cutting complexity and creating value, responsibly	11
References	12
About the authors	13
Acknowledgements	13

Introduction

The rapid proliferation of Artificial Intelligence (AI) carries transformative potential across the pharmaceutical and life sciences sector, and Medical Affairs is no exception.

From Natural Language Processing (NLP) and machine learning to the emergence of Large Language Models (LLMs), generative AI tools and autonomous AI agents, the ability to strategically harness these technologies is quickly becoming a powerful differentiator. Organizations that embrace AI transformation in a responsible and reliable way will be empowered to enhance performance and unlock new opportunities for growth.

This white paper outlines some of the key challenges and imperatives facing professionals in the Medical Affairs space and considers how responsible AI adoption can equip teams to solve for them. We will examine several of the most compelling defensible AI applications that promise Medical Affairs impact, showcasing proven, real-world examples of what is already possible today.

Navigating Medical Affairs challenges

As Medical Affairs has evolved from a support function to become the third strategic pillar between commercial and R&D, professionals in this space find themselves navigating increasingly complex, interconnected challenges. The growing importance of Medical Affairs has brought with it a requirement for greater strategic focus and, as a result, additional expectations.

Consequently, Medical Affairs professionals are now grappling with an explosion in data volume, both structured and unstructured, flowing from disparate sources. To serve the function's expanded strategic role, these data must be translated into medical insight and converted into action points for commercial and R&D functions, ultimately benefiting

the patients we serve. There is also increasing pressure to communicate the medical impact and return on investment (ROI) for these activities across the organization.

The complexity and scale of the challenges require smart solutions that can enhance efficiency and impact. With this in mind, AI is poised to bring tremendous value to Medical Affairs, across the full breadth of the function's initiatives and sub-domains.

By effectively harnessing AI, teams can more easily convert large scale and disparate data into actionable insights, boost operational efficiency through streamlined workflows, and improve decision-making processes.

Responsible adoption is essential for effective transformation

AI transformation is no simple task for any organization and must be underpinned by a range of practical and ethical considerations.

To truly drive value in the healthcare space, it is essential to start with a clear understanding of the risks and how to mitigate them. The key challenges associated with AI use, particularly LLMs, include:



RELIABILITY OF DATA

- Biases in input data can be amplified and lead to biased outputs
- As models are only as reliable as the data on which they are trained, incomplete training data can cause models to “hallucinate” or fill in the blanks with plausible, but ultimately false, outputs
- Learning techniques can result in responses that are unverifiable
- Version control leads to often outdated training inputs



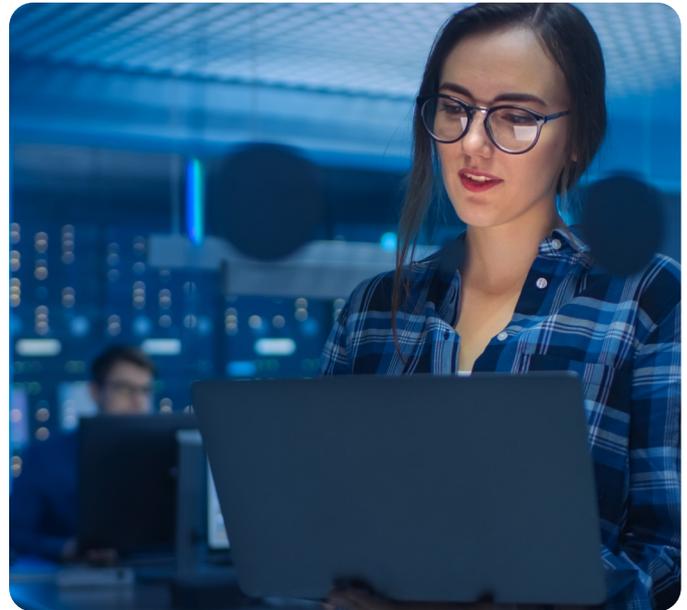
SCALE

- Enormous computational scale leads to both extensive monetary and environmental cost
- In Medical Affairs, this scale is particularly significant due to the complexity, sensitivity, and diversity of biomedical data



REGULATORY CONSIDERATIONS

- The speed of innovation can outpace regulatory frameworks
- Despite global frameworks, there are no harmonized global AI regulations — with conflicting AI laws across over 70 countries

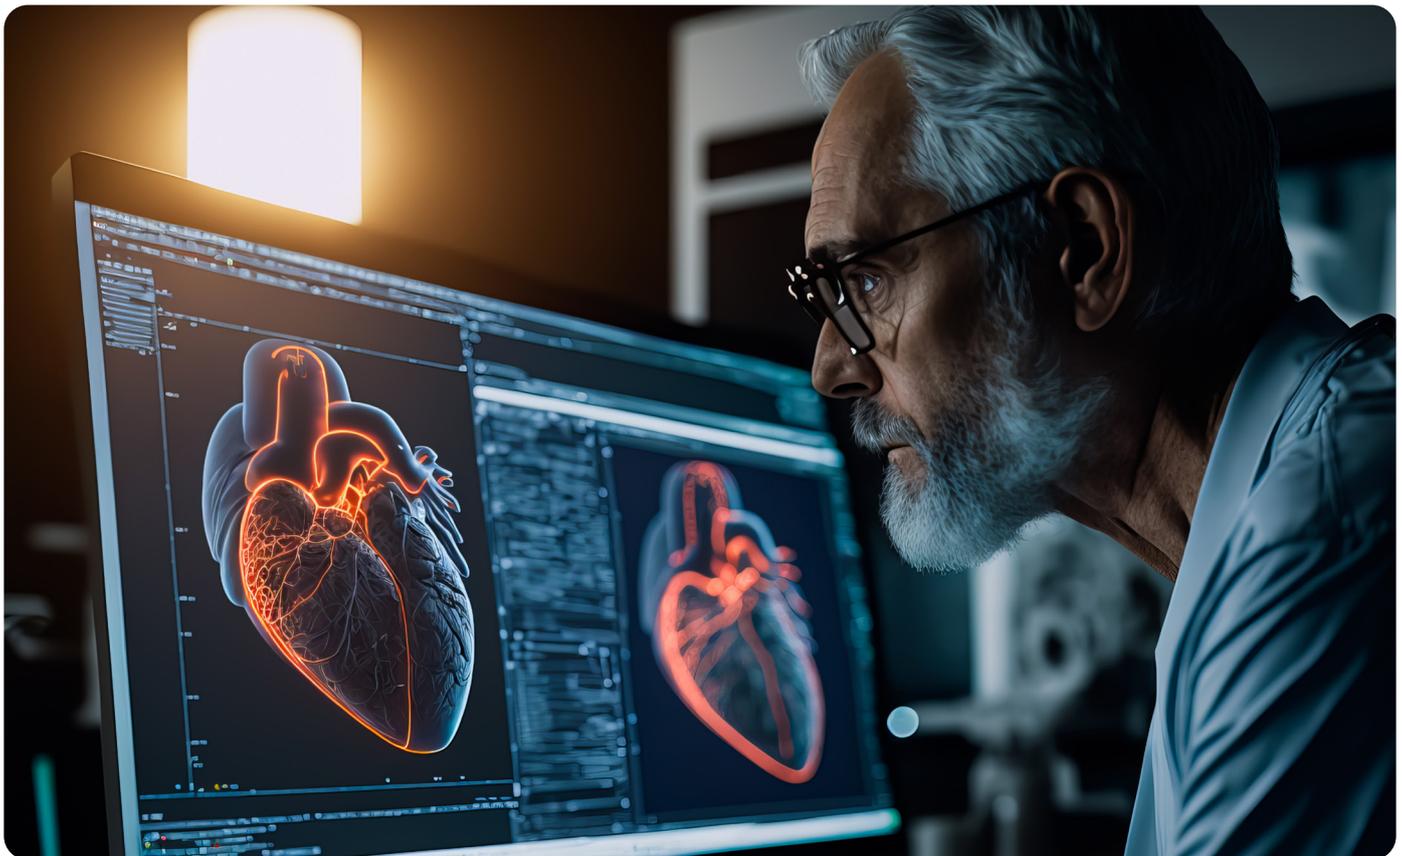
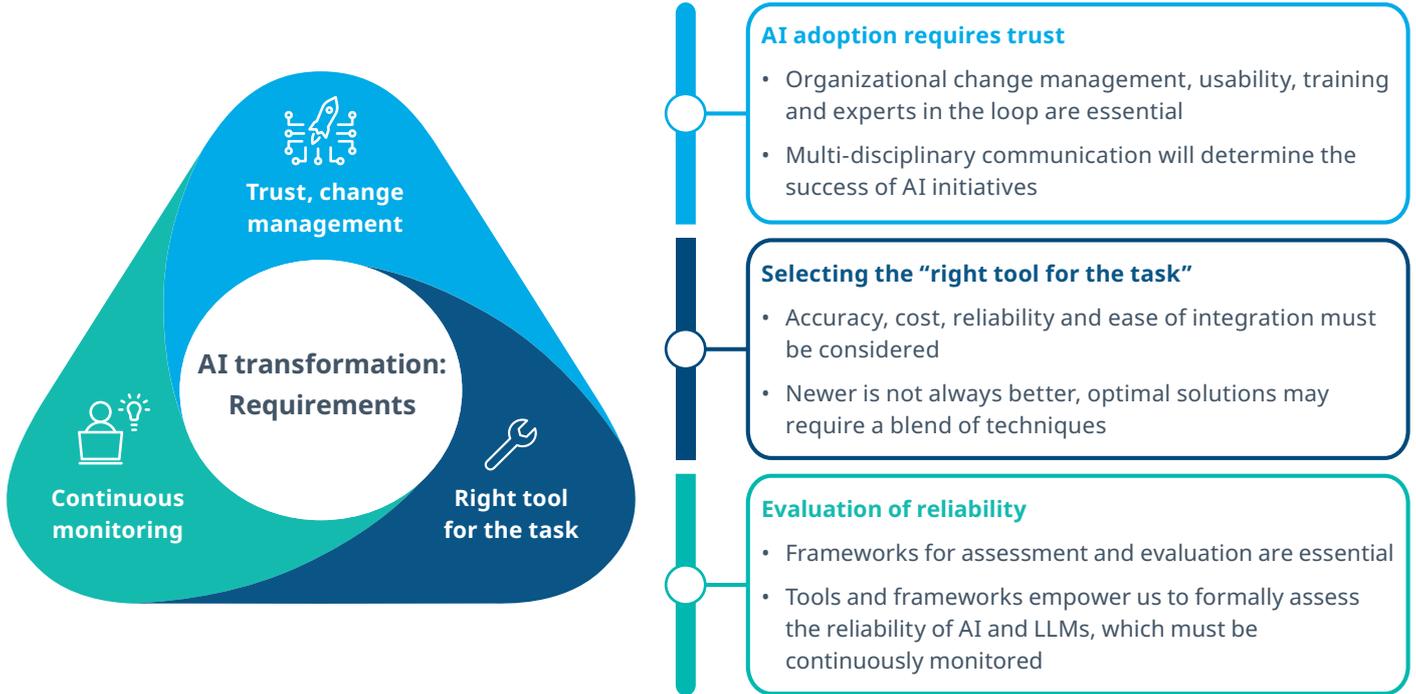


To harness AI safely and reliably in our industry, models must be built and trained on high-quality health data, benchmarked and optimized, and fit into purpose-built applications that are monitored by experts to ensure trustworthiness and regulatory compliance.

Broadly, efforts to secure user trust and understanding, appropriate tool selection, and establishing a framework to evaluate reliability will form the bedrock of effective AI adoption at enterprise level, including the Medical Affairs function (see Figure 1).

If handled responsibly, AI transformation will help to unlock the full potential of Medical Affairs, enabling professionals to deliver enhanced insights and value for the benefit of patients, Healthcare Professionals (HCPs), and the organizations they operate within.

Figure 1: Requirements for AI transformation in the enterprise



Defensible AI applications for Medical Affairs

We have discussed how AI, with appropriate governance and quality assurance, has the potential to bring great value to Medical Affairs. So where can these technologies be best deployed to deliver impact?

There are a large and multifaceted range of defensible applications in this space, converting complexity into

value by supporting streamlined processes, enhanced data analysis and deep insight generation.

While AI holds significant promise across various sub-domains of Medical Affairs, this paper concentrates on use cases in four key areas:



LITERATURE REVIEW

A ubiquitous requirement which underpins decision-making in Medical Affairs, and across the drug development lifecycle more broadly. Literature review is typically the first step in any scientific endeavor, which, especially in the case of Systematic Literature Reviews (SLRs), must be conducted with a high degree of methodological rigor. The insights extracted from literature reviews can also complement evidence generated from other, disparate Real-World Data (RWD) sources, such as information from call center logs, social media, news feeds and Medical Science Liaison (MSL) notes.



CHART REVIEW

Provides rich, patient-level data on treatment patterns, outcomes, and disease progression, making it one of the most powerful sources for Real-World Evidence (RWE) generation in Med Affairs and other domains.



KEY OPINION LEADER (KOL) IDENTIFICATION AND PROFILING

Essential for ensuring strategic engagement with the most relevant experts across therapeutic areas, enabling personalized scientific dialogue and amplifying the company's medical voice throughout the product lifecycle.



MEDICAL COMMUNICATIONS

Critical for translating complex scientific data into clear, credible, and impactful messaging that supports product value and patient access.



The following sections showcase key AI use cases in these four areas, as well as the tangible benefits they offer.

Literature review: Unlocking efficiency and insight

In Medical Affairs, literature reviews serve as a foundational tool for evidence-based decision making, strategic planning, scientific communication, stakeholder engagement, as well as compliance with regulatory expectations.

While literature reviews, especially targeted or systematic reviews, can be tremendously valuable for Medical Affairs teams, such efforts are not without their challenges:



METHODOLOGICAL RIGOR

Selecting the appropriate review type requires a deep understanding of the research question, the available evidence, and must consider the regulatory context. Each type of review has distinct methodological and scope-related requirements. Especially in the case of systematic literature reviews, there is a need for objectivity and robust scientific rigor, adhering to frameworks like the Preferred Reporting Items for Systematic Reviews and Meta-Analyses (PRISMA) and guidance such as the Cochrane Handbook, which is considered the gold standard for conducting SLRs in healthcare.



DATA VOLUMES

The large quantity of published literature can be overwhelming, especially in rapidly evolving therapeutic areas. The research question must be precise and unambiguous, as broader search strategies will generate an unmanageably large corpus. Critical appraisal of available data by experts is essential to identify the most relevant and current studies from a variety of databases and sources.



TIME AND RESOURCE CONSTRAINTS

Literature reviews require meticulous planning and deep domain expertise, as well as significant manual effort from experts. This reliance on highly trained personnel, tasked with conducting repetitive tasks like data extraction, can make reviews difficult to scale. Furthermore, as medical literature evolves so rapidly, reviews can become outdated quickly unless regularly updated.

Revolutionizing workflows with AI agents

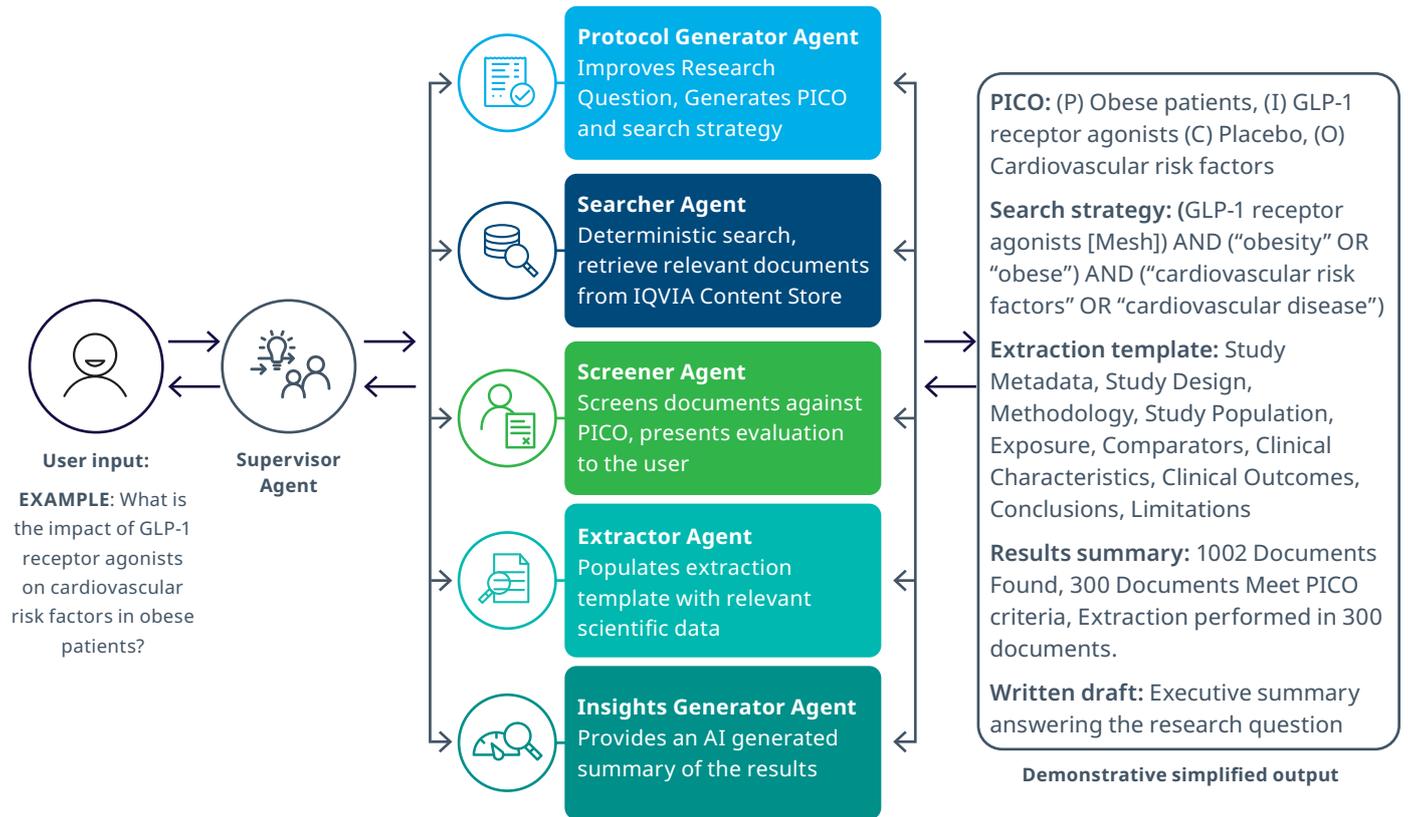
AI holds transformative promise for literature reviews. Recently, the integration of agentic AI into literature review workflows is making it possible to accelerate this essential and necessarily rigorous task, enhancing efficiency and accuracy.

Agentic AI involves the use of autonomous AI agents able to understand and perform complex processes with limited direct human supervision. However, in using this technology, it is critical to maintain the

“expert-in-the-loop” and take care to ensure the validity of the outputs.

With these principles at its core, IQVIA has created an agentic framework for literature review, leveraging a series of AI agents (Figure 2). The approach unburdens the user from executing every task manually, while allowing them to provide input to shape the outcome and check for accuracy.

Figure 2: IQVIA's agentic framework for AI-assisted literature review



Searcher Agent proven to match the manual gold standard of capture and reduce the number of documents requiring manual screening by 87%¹ in an SLR focused on non-small-cell lung cancer.

The process is set into motion by a human prompt; the expert poses one or more research questions in natural language to a "Supervisor Agent", which in turn calls upon a team of discrete agents to guide the workflow to its evidence-generating conclusion. The first is a Protocol Generator Agent, which will seek input to refine the research question(s) before generating PICO criteria (Population, Intervention, Comparators and Outcomes) and creating a documented search strategy.

Once the expert has validated the search strategy, a Searcher Agent will retrieve relevant documents of interest from IQVIA's indexed scientific Content Store, containing continuously updated scientific literature and conference proceedings.

A Screener Agent then qualifies the retrieved documents by comparing them against the PICO criteria, and presents the evaluation to the expert, who may assess whether to accept, reject, or refine the suggested corpus.

Next comes data extraction, which is typically a manually intensive task requiring experts to exhaustively search through documents to extract and normalize the relevant variables. Using the agentic AI framework, an Extractor Agent clarifies with the expert which variables are required and performs the extraction autonomously, populating an extraction template.

Extractor Agent proven to achieve the same degree of recall as manual extraction, with a 2.5-3x reduction in manual effort² in a highly complex living review.

Once verified, these extractions may be used to create a summary of the results, which the user can instruct an Insights Generator Agent to shape into the required format.

In this way, collaboration between AI agents and the ever-present expert-in-the-loop ensures trustworthy implementation of AI to solve complex, multi-step problems like scientific literature review.

Consolidating RWD for actionable insights

The insights generated from AI-assisted literature reviews can also be used to complement other RWD sources, enabling the creation of a living “data lake” for organizations looking to get a 360° view of what the world is saying about the medications they have taken to market.

In one case from 2019³, even before the more global adoption of generative AI, a combination of AI techniques was used to collate evidence from literature with insights extracted from social media, call center logs, news feeds and MSL notes to generate a centralized repository.

This AI-enabled knowledge base was migrated to a cloud-based global big data and analytics platform with visualization and summarization of collated RWD insights through interactive dashboards (see Figure 3).

This provided business users and product teams across the globe with enhanced access to RWD for real-time derivation of actionable insights. Additionally, the use of an AI-enabled approach unburdened global experts from manual and repetitive tasks, allowing them to focus their efforts on other critical, value-adding activities.

Figure 3: AI-driven insight generation from disparate RWD sources⁴

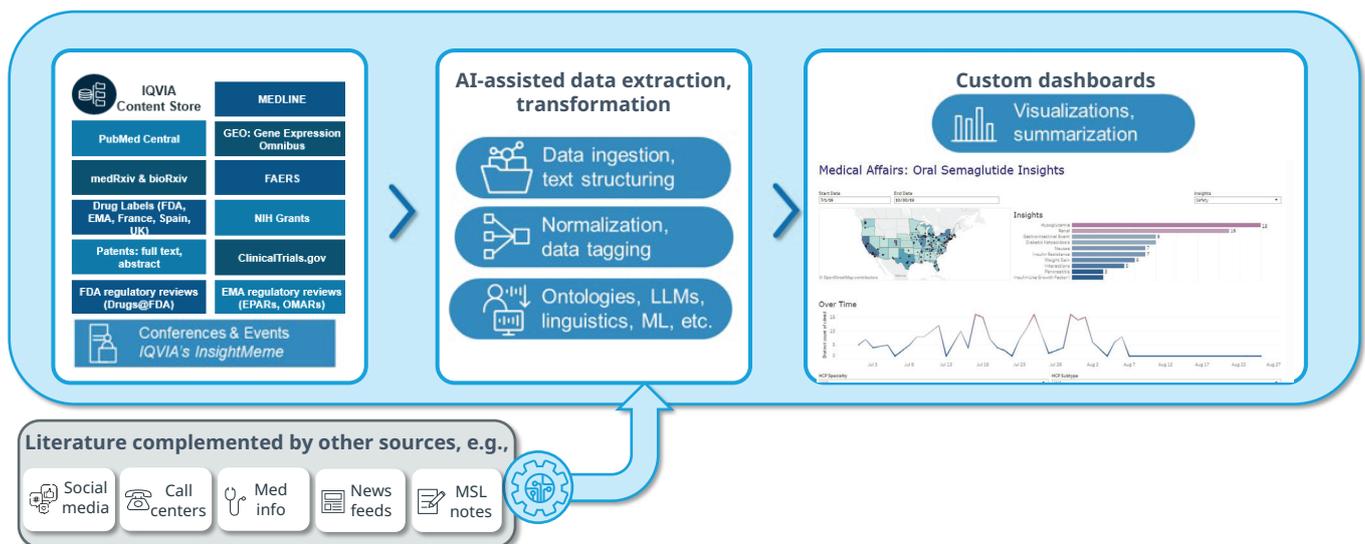


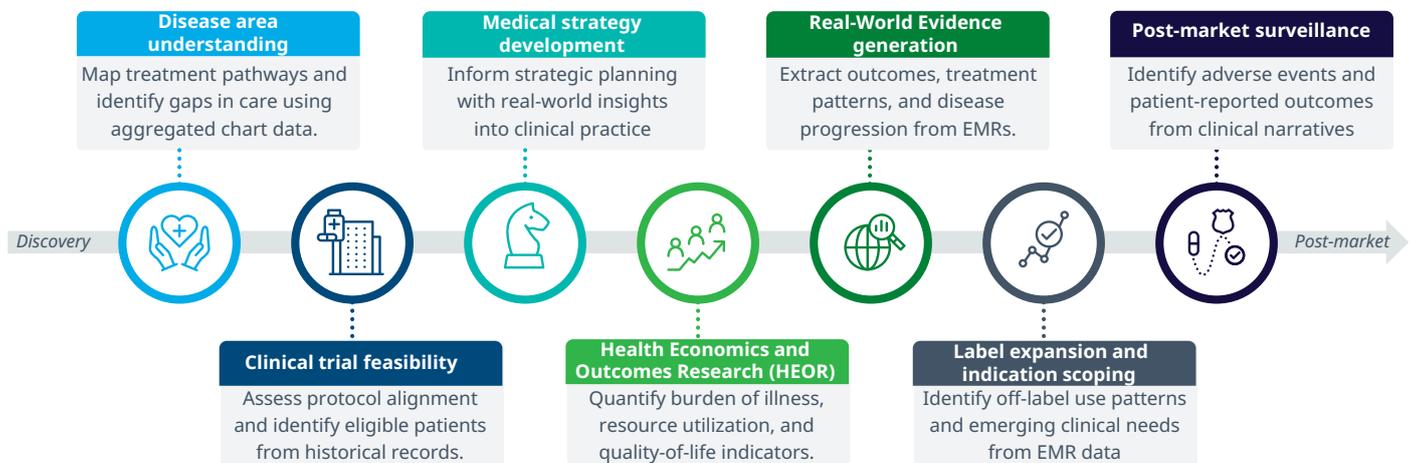
Chart review: Reducing manual burden through AI-assisted information extraction

It is estimated that approximately 80% of Electronic Medical Record (EMR) data exists in unstructured formats such as clinical notes, scanned documents, and free-text entries.⁵ These sources often contain rich, nuanced information about treatments, outcomes, and disease progression, details that are rarely captured in structured fields.

As such, chart review has broad applicability in Medical Affairs (Figure 4) as it enables deep insights into real-world patient data, supporting RWE generation, identifying unmet needs, and informing strategic decision-making.

Increasingly, chart reviews are being augmented with AI to expedite insight generation from these rich data sources while reducing the chart review burden on sites.

Figure 4: How AI chart review can support Medical Affairs — selected examples from discovery to post-market



Case study: Streamlining chart review with NLP

A health system in the midwestern United States partnered with a life sciences company to provide insights on real world outcomes in over 100,000 patients with congestive heart failure. Key data points included ejection fraction (a measure of cardiac function), New York Heart Association (NYHA) classification (a four-tier severity scale), and symptoms such as dyspnea, fatigue, and dizziness: all embedded in free-text clinical notes.

Manually reviewing this volume of data and capturing the pertinent information in a structured format would have been impractical and prohibitively time-consuming.

To streamline the task, IQVIA's AI platform was installed within the organization firewall and domain-specific NLP was used to identify, extract, normalize and collate critical unstructured information from over 34 million clinical notes, with review and validation by the expert-in-the-loop.

This resulted in a 100x reduction in manual effort for chart review by site staff, who could instead focus their time on verifying and validating AI derived insights — which were proven to achieve accuracy of 95-99%, including critical endpoints.

AI efficiently and accurately unlocked insights from 34 million clinical notes from over 100,000 patients⁶

KOL identification and profiling: Enhancing precision and engagement

Identifying and understanding Key Opinion Leaders is a cornerstone of Medical Affairs strategy. These individuals influence scientific dialogue, guide treatment adoption, and shape the broader understanding of disease areas across the clinical community.

Traditionally, KOL identification and profiling have relied on manual, time-intensive reviews of publication lists, clinical trial registries, and congress participation records. While effective to a degree, these approaches often provide only a snapshot in time, can miss emerging voices, and require significant human effort to interpret and update.

AI is redefining this process. By automating the ingestion and analysis of large, diverse data sources, AI helps Medical Affairs teams to identify experts more comprehensively and with greater precision. Leveraging NLP, entity recognition, and network analytics, AI systems can analyze millions of records across structured and unstructured datasets, including scientific publications, conference abstracts, trial registries, and social media.

This provides a multidimensional view of KOL influence that extends beyond publication counts to capture context, collaboration, reach, and emerging thought leadership. AI-powered KOL mapping reveals how experts connect, what topics they shape, and where influence is emerging. Advanced network analytics allow teams to visualize professional relationships and uncover previously hidden connections between investigators, institutions, co-authors, and clinicians.

This level of insight enables Medical Affairs professionals to distinguish between global scientific leaders, regional and local peer leaders, clinical treatment leaders, and Digital Opinion Leaders (DOLs). It also makes it possible to identify how targeted engagement within these networks can accelerate the spread of new scientific insights into clinical practice.

From an operational perspective, AI dramatically reduces the manual burden of sifting through dense



scientific and clinical data. Automated entity extraction, topic modeling, and sentiment analysis streamline the profiling process by dynamically organizing experts according to therapeutic focus, publication velocity, network importance, and engagement relevance. This shift allows Medical Affairs engagement teams to devote more time to strategic planning and personalized outreach rather than data collation.

IQVIA's AI-enabled KOL identification and profiling solution⁷ exemplifies this next generation of Medical Affairs capability. It combines vast data assets, including publication and congress activity, RWE, clinical trial participation, and digital presence metrics, with advanced network analytics that illuminate the scientific and clinical communities shaping disease understanding and influencing care. This integrated dataset enables Medical Affairs teams to prioritize engagements that best align with scientific strategy, product lifecycle stage, and geographic focus.

Ultimately, the adoption of AI for KOL identification and profiling enhances not only efficiency but also the strategic precision of Medical Affairs operations. By integrating vast, heterogeneous data into actionable insights, AI empowers professionals to better understand the evolving landscape of scientific, clinical, and digital influence.

The result is more targeted and meaningful engagement, fostering stronger, more authentic scientific exchange aligned with organizational goals and ultimately contributing to improved patient outcomes.

Medical communications: Reshaping scientific dissemination

AI is also revolutionizing how scientific communication is developed, personalized, and disseminated.

Generative AI in particular promises benefits through the acceleration of content creation, driving enhanced engagement strategies with healthcare professionals and patients alike.

Similarly to previous examples, AI enables Medical Affairs teams to access and assess large volumes of structured and unstructured data to inform Integrated Medical Communications Plans (IMCP), from clinical trial results and RWE to market insights and literature reviews. By combining human expertise with AI-driven speed, teams can produce plain language summaries, manuscripts, posters, and publication extenders tailored to diverse audiences and geographies.

Neural machine translation engines can also facilitate multilingual scientific dissemination, with validation by medical translation experts, before they are provided to journals and local affiliates. This helps to ensure accessibility across global markets.

AI is also reshaping how Medical Affairs engages HCPs. By curating insights from behavioral data, AI can assist in the development of personalized communication strategies that reflect individual preferences and clinical interests. This digital engagement approach supports omnichannel delivery, whether through peer-

reviewed articles, virtual symposia, or AI-powered chatbots, ensuring that scientific information reaches HCPs in the most effective format and context.

Beyond HCPs, AI is being used to deepen patient engagement. By mapping the patient journey and identifying 'moments of meaning', generative AI helps create evidence-based personas and personalized educational materials. These outputs are designed to resonate with patients' lived experiences, improving understanding and adherence while supporting better health outcomes.

Although much is promised by AI in this space, its application in medical communications must be governed by responsible design. All AI-generated content should pass through expert-in-the-loop checkpoints to ensure scientific validity, tone, and compliance with regulatory standards. This safeguards against hallucinations, bias, and inaccuracies, preserving the integrity of medical information.

AI is redefining scientific communication in Medical Affairs, enabling faster, more personalized, and globally accessible content. As adoption grows, especially among HCPs, success will depend on combining AI's speed with human oversight to ensure accuracy, relevance, and trust.



Conclusion: Cutting complexity and creating value, responsibly

The use cases presented in this paper — while only the tip of the iceberg — provide compelling evidence for the transformative potential of artificial intelligence in Medical Affairs. From strategic decision making to evidence generation and dissemination, this versatile technology is already making high-impact contributions across the function and its sub-domains.

By cutting through the overwhelming volumes of data facing Medical Affairs professionals, AI is bringing enhanced efficiency through automation, but not only that — it is enabling professionals to work smarter, communicate more effectively, and deliver improved value. Even greater impact can be anticipated for the future as the technology continues to advance.

However, enthusiasm to embrace AI transformation and grasp the benefits must come alongside appropriate checks and balances.

It is crucial to remember that not all AI solutions are created equal. Selecting the right tool for the task is

essential; not every problem requires an LLM or AI agent, and sometimes a hybrid approach may deliver the best results, combining traditional methods with AI to ensure methodological rigor and relevance.

Ensuring reliability is also paramount. For AI to deliver value, its outputs must be dependable. Continuous monitoring, keeping human experts in the loop, and applying tools and frameworks to formally assess the performance and trustworthiness of AI solutions will be essential components of any successful adoption. It is not just about what AI can do, it is about how responsibly we use it.

As we harness the power of AI in Medical Affairs, we must remain vigilant and committed to maintaining the highest standards of reliability and ethics. By doing so, professionals will be empowered to fully realize the potential of AI to revolutionize Medical Affairs and drive vital advancements in medical science and patient care.



References

1. Rath, N. et al (2023). MSR112 Assessing the Feasibility of Applying Natural Language Processing for Systematic Literature Reviews: A Case Study in Non-Small-Cell Lung Cancer. Value in Health, 26(12) S414. <https://doi.org/10.1016/j.jval.2023.09.2171>
2. Raghuram A, Ssentongo A, Pereira LM, Kuang Y, Uyei J, Park P (2023). 2378. Use of Natural Language Processing to Extract Published Real World Data on a COVID Vaccine and Antiviral Treatment. Open Forum Infect Dis. 27 (10) (Suppl 2) <https://doi.org/10.1093/ofid/ofad500.1999>
3. Borfitz, D (2019, July 29). Novo Nordisk: Natural Language Processing On A Steady Diet Of Real-World Data. Bio IT World. <https://www.bio-itworld.com/news/2019/07/29/novo-nordisk-natural-language-processing-on-a-steady-diet-of-real-world-data>
4. Reed, J. and Breyette, T (2019). Using Natural Language Processing to transform real world data. <https://www.lexjansen.com/phuse/2019/rw/RW05.pdf>
5. Au Yeung, J., Shek, A., Searle, T. et al (2024). Natural language processing data services for healthcare providers. BMC Med Informatics and Decision Making, 24(356). <https://doi.org/10.1186/s12911-024-02713-x>
6. Miliard, M (2018, June 8). How Mercy is using NLP with its Epic EHR to improve analytics for cardiac care. Healthcare IT News. <https://www.healthcareitnews.com/news/how-mercy-using-nlp-its-epic-ehr-improve-analytics-cardiac-care>
7. IQVIA (2025). Expert Ecosystem. <https://www.iqvia.com/solutions/real-world-evidence/health-data-apps-and-ai/expert-ecosystem>

About the author



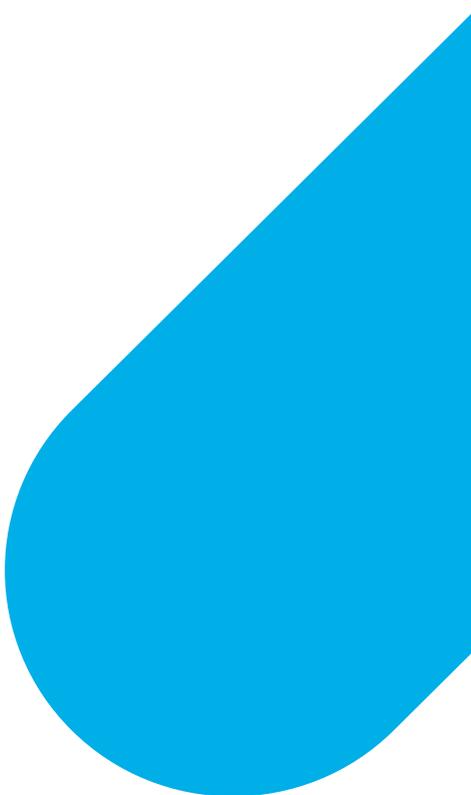
NIAMH MCGUINNESS

Director, Pharma Solutions, IQVIA Applied AI Science

Niamh leads a diverse group of experts who build AI solutions and utilize a suite of award-winning technology to help pharma customers protect participant privacy and surface critical insights from unstructured data and literature, for a vast array use cases from bench to bedside, including R&D, drug safety and medical affairs. Niamh has been with IQVIA for 9 years. She received her PhD in neuroimmunology from Trinity College, Dublin.

Acknowledgements

The author would like to thank Barbara Arone, Steve Eichert, Ben Gallarda and Calum Yacoubian for their support and valuable contributions to the development of this white paper.



CONTACT US
iqvia.com