White Paper

Promising New Developments in Artificial Intelligence:

Innovating with Generative and agentic AI to drive efficiency and quality in clinical trials

RAJNEESH PATIL, VP, Digital Innovation and Strategy, IQVIA

Part 1 of a 4-part series highlighting key innovations in clinical trials

Table of contents

Overview

As implied by the article title — *Promising New Developments in Artificial Intelligence* — these are extraordinary times for the clinical research arena. Each year, with ever-increasing anticipation, thought leaders and industry professionals speculate and debate where and how to best unlock AI's potential to help to make clinical trials more efficient and inclusive — and ultimately move healthcare forward by improving site and patient experiences.

Within the vast and complex clinical trial ecosystem, there has been no dearth of opportunities to harness these remarkable technologies — whether in design and planning, patient engagement, trial delivery, clinical monitoring, regulatory submissions, or well beyond into the commercial drug development space.

Recently, however, the landscape has transformed, thanks in no small part to the actual, successful *application* of GenAI tools into clinical trials — most of which were considered to be within the realm of 'blue sky ideation' not long ago. And while these advancements are proving to enhance both efficiency and quality in actual use cases, the FDA and other regulatory agencies are beginning to provide guidance on proper guardrails, context and validation requirements alongside other key considerations to incorporate Responsible-AI into clinical studies.

"Recently, the landscape has transformed, thanks in no small part to the actual, successful application of GenAI tools into clinical trials — most of which were considered to be within the realm of 'blue sky ideation' not long ago."

In this paper, we will review several use cases that leverage IQVIA Healthcare-grade AI® engineered to meet the level of precision, speed and trust required in life sciences and healthcare, provide an overview of the technological and operational drivers of their success, and specifically address how each of these new AI methodologies substantially enhanced efficiency and quality.

How the bar was raised: ensuring safe and efficient use

Before delving into the more technical aspects of these AI use cases, it is important to understand the distinction between GenAI and agentic AI, which are both subsets of artificial intelligence.

Fundamentally, GenAI is designed to generate responses, content, reviews, and outputs, including reports or visuals, based on data it has processed during model training. It tends to be transactional in nature, such as providing a generic response to a specific question, which can then be improved by the science of prompt engineering. Agentic AI is built to independently handle complex, multi-step problems and operations, such as controlling systems (think of the self-driving car). Generative AI turns vast amounts of data into actionable knowledge, and AI agents can build on this capability to execute specific process steps and actions — allowing for the development of innovative solutions that can impact quality, compliance, speed and productivity.

A holistic, philosophical approach was taken when developing these particular AI frameworks — which, in some ways, bears resemblance to parenting. Like children who enter the world capable of untold possibilities, unlocking the potential of AI requires proper training, a strong ethical framework, model integrity and transparency, as well as safeguards to protect these tools from going down sub-optimal paths. A critical aspect of proper grooming is human oversight — with several data science approaches that focus on iterative training and augmented learning for GenAI models (more on this later).

But the first order of business in these successful applications was ensuring that they received proper contextualization and training.

Selectivity of the training data

The accuracy and reliability of artificial intelligence hinges on the integrity, relevance, and — of critical importance — the *selectivity* of the input data used in its training. Open-source consumer Large Language Models (LLMs), such as *ChatGPT*, *DeepSeek* and *Grok*, cast an extraordinarily wide net — typically scouring vast datasets containing billions of words of human-generated text.

While the capacity to coalesce and process such an enormous breadth of data is technologically remarkable, and useful for many purposes, casting too wide a net in the scientific realm risks generating poor-quality or hallucinatory responses.

For scientific inquiries, in fact, it is not uncommon for certain open source GenAI agents to hallucinate — providing false citations, making presumptuous conclusions, and even offering counterproductive medical recommendations. Without robust modifications, these open source LLMs are unfit for scientific or medical applications.

A range of vital safeguards are therefore required to ensure there are limitations, vetting, and proper verification of training data to ensure GenAI agents can be used safely and efficiently in clinical trials.

Taking a multi-pronged approach to safeguarding efficiency and quality

It is important to consider 5 critical categories of safeguards when developing and applying Generative and agentic AI solutions in clinical studies — irrespective of whether they take the form of a chatbot, summary generator, or automated workflow assistant (see *figure 1*). These principles help ensure efficiency, quality and maximize safety. These 5 categories of optimization are listed as follows:

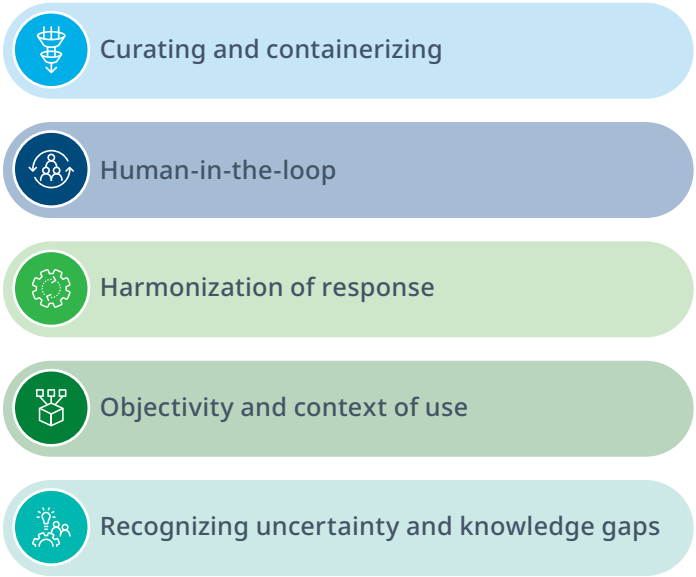

"Like children who enter the world capable of untold possibilities, unlocking the potential of AI requires proper training, a strong ethical framework, model integrity and transparency, and safeguards to protect these tools from going down sub-optimal paths."

Figure 1. Key categories of safeguards that optimize efficiency and quality of GenAI and agentic AI in clinical trials.

1. Curating and 'containerizing' data

Alexander Pope's long-held adage that "a little bit of knowledge is a dangerous thing" characterizes how erroneous conclusions (and overconfidence) can result from having insufficient knowledge. While true for humans, putting prudent limits on knowledge — accomplished by *curating* — is critical in ensuring the quality and efficiency of GenAI.

In other words, to generate reliable, clinical operations-grade responses to queries, training data must be hermetically sealed from tangential or speculative information, which risks ambiguous responses and/or opining.

Therefore, any GenAI/agentic AI solution used in clinical trials must be deployed within the pre-determined boundaries of the data ecosystem, operational process and human-in-the loop touchpoint.

2. Integrating "human-in-the-loop"

Since its inception, artificial intelligence has catalyzed an understandable debate about the future of human involvement, raising concerns that workers could be displaced on a global scale. At IQVIA, we steadfastly maintain — and have consistently observed in successfully developing novel applications — that a healthy balance between human engagement and AI is indispensable in optimizing quality and preempting hallucinatory responses.

Again, the classic example is the self-driving car. Despite its autonomous decision-making capability through agentic AI, safety is optimized for these vehicles because a human is seated at the controls to override the automated decisions when necessary. This "*human-in-the-loop*" component — paired with agentic AI — helps ensure safe and ethical responses are generated when encountering a broad range of scenarios.

The same holds true, of course, in the clinical trial arena. Human-in-the-loop can optimize efficiency and quality by being integrated into complex, multi-step processes — such as end-to-end data flow and monitoring of several streams of data-generated (via) disparate data collection systems utilized on a clinical trial.

Agentic AI can be used autonomously for certain functions within a given workflow, particularly administrative or automated tasks (such as document gathering and compliance review), while human-in-the-loop can be reserved for steps requiring greater sensitivity — such as critical data reviews and decision-making. This integrated approach can significantly streamline and accelerate processes, while providing greater assurances of quality.

3. Harmonization of response

Another critical success factor is the ability to generate consistent, quality responses independent of how a question is asked. A query can be expressed in any number of ways — depending on the word choices, idioms and figures of speech chosen by the user. Different personalities, backgrounds and cultures can also impact the way a question gets phrased.

To meet the impeccable standards required by medical science, natural language processing must be highly sensitized to reliably detect context, unearth the core question at hand, and provide consistent, accurate responses each and every time. While a simple concept, producing globally-harmonized responses requires highly sophisticated expertise, technologies and in-depth experience in best practices.

"To maximize efficiency and quality, agentic AI can be used autonomously for certain functions within a given workflow, particularly administrative or automated tasks, while human-in-the-loop can be reserved for steps requiring greater sensitivity."

4. Objectivity and context of use

While objectivity is critical in any medical or scientific engagement, it requires even greater rigor in the application of GenAI. Consider, for example, diagnostic imaging, where different specialists may look at the same scan, yet analyze and interpret findings very differently. While this is not "bias" in the traditional sense, the reality is that healthcare professionals — like all of us — have unique backgrounds, skills and experiences that may influence conclusions.

GenAI can be used to objectify the diagnosis of medical images and accurately assess the severity of Serious Adverse Events (SAEs) based on standardized grading and well-established guidelines.

5. Recognizing uncertainty and knowledge gaps

An underlying cause of incorrect or hallucinatory responses from many GenAI models is the lack of a built-in mechanism to recognize uncertainty. Like human beings, artificial intelligence benefits from self-awareness and recognizing its own limitations. The novel GenAI agents being applied in medical science must therefore be rigorously trained to refrain from partial answers or guesswork.

Now that some of the important principles that optimize efficiency and quality for AI in clinical research have been outlined, let's delve into some of the AI agents currently being applied within this ecosystem.

"Like human beings, AI benefits from self-awareness and recognizing its own limitations. The novel GenAI agents being applied in medical science are therefore rigorously trained to refrain from partial answers or guesswork."

Putting principles into practice: use cases of successful utilization

Now that we've provided an overview of the core, multi-pronged approach to leveraging innovative AI methodologies to drive efficiency and quality, let's review some of the exciting and successful models where IQVIA leveraged these principles (see *figure 2*).

Figure 2. Novel AI tools currently being applied in medical research.

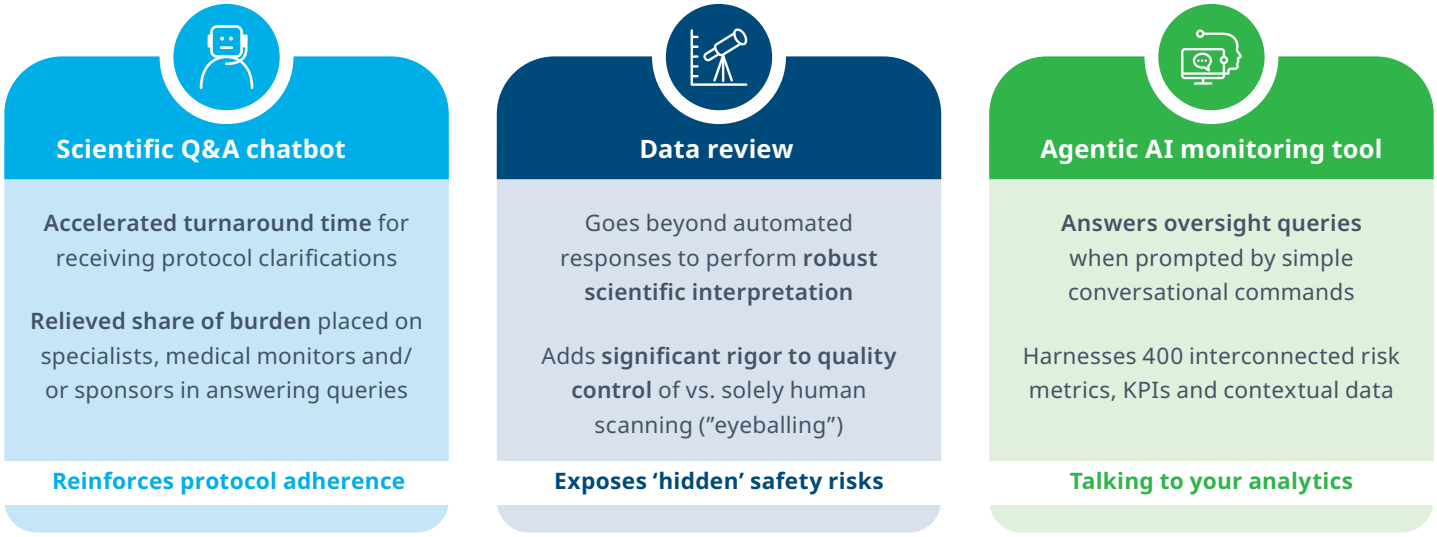

Successful utilization of a scientific chatbot in a Phase III trial

A key advancement in leveraging GenAI was the implementation of a conversational Q&A chatbot that effectively answered protocol-related queries. Utilized by site staff and monitors in a large Phase III study in Q4, 2024, the chatbot quickly and accurately addressed a variety of scientific, protocol-specific questions — ranging from IP administration, screening procedures for enrollment, schedule of events for patients, compliance-driven documentation of safety and adverse events and other critical controls impacting quality of compliance of processes at clinical trial sites.

By providing this diverse spectrum of rapid, informative responses, the chatbot accelerated the turnaround time for receiving protocol clarifications while relieving a share of the burden placed on specialists, medical monitors and/or sponsors in answering queries.

The keys to success of this chatbot aligned with the principles set forth in the previous section:

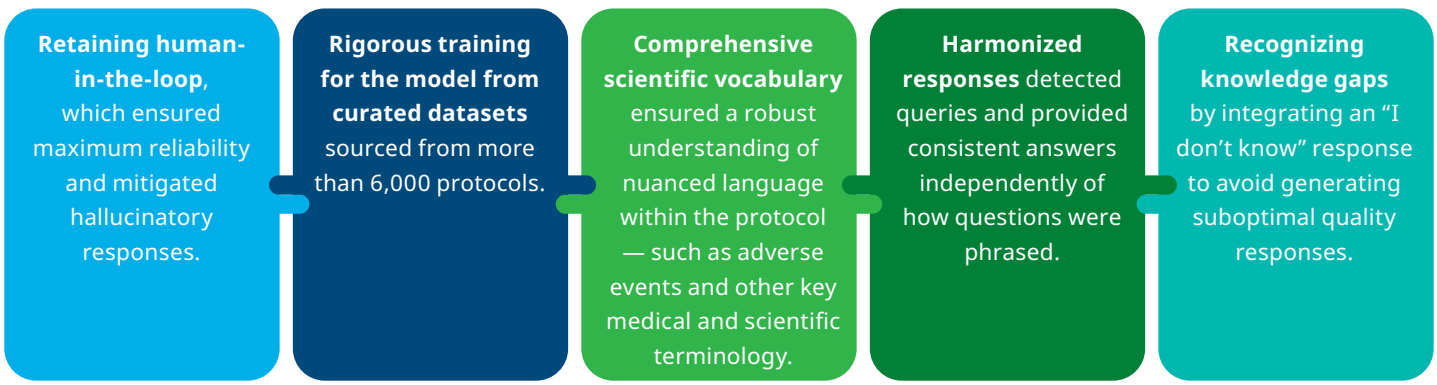

As one would expect, user experience was a key factor in the successful adoption of this chatbot, which was characterized by users as being intuitive, simple to understand and reliable. By making information easily accessible, this tool was proven to enhance real-time decision-making — helping to ensure clinical processes remained fully compliant with protocol guidelines.

Because of its success, the core framework of this chatbot (the LLM foundational model) is now being scaled to support a wide range of data-based inquiries — including IP management and processes, a 'virtual-assistant' for investigational brochures and beyond.

"User experience was a key factor in the successful adoption of this chatbot, which was characterized by users as being intuitive, simple to understand, and reliable."

Data review: empowering central monitors and CRAs

Another cutting-edge GenAI breakthrough was the capability to enhance data reporting by performing critical reviews across a range of safety risks. This application detects and analyzes a wide range of signals, such as patient eligibility issues, medications and administration risks and adverse events — and flagging inconsistencies and errors (such as incorrectly graded AEs). It then promptly generates a well-organized summary that can expose "hidden" risks before they can occur.

What makes this data review tool an important milestone is its ability to go beyond automated responses to perform legitimate scientific analysis and interpretation. As a result, it adds significant rigor to the quality control of data reporting versus solely human scanning (or "eyeballing") to identify data patterns and signals.

As alluded to earlier, different medical experts — such as radiologists — may review the identical chart but draw different conclusions or prognoses. Another significant feature of this technology is its ability to objectively analyze data visualizations.

A pertinent example is the grading of SAE severity — a critical issue in therapeutic areas like oncology, where patients are often in advanced stages of illness and may be more likely to have SAEs occur. Having been rigorously trained on curated data sourced from the Common Terminology Criteria for Adverse Events (CTCAE), this data review tool can "see" the laboratory data and accurately grade SAE severity based on these well-established guidelines.

Of course, human-in-the-loop remains a key safeguard, as this tool enhances — not replaces — human engagement and analysis. That being said, this capability increases the rigor of the analyses and helps mitigate redundant reviews from board-certified doctors, medical reviewers or centralized monitors.

"What makes the data review tool an important milestone is its ability to go beyond automated responses to perform legitimate scientific analysis and interpretation."

Talking to your analytics: an agentic AI platform for monitoring support

One of the more remarkable GenAI applications emerging in clinical trials today is the agentic AI platform monitoring tool, which answers highly specific oversight queries when prompted by simple conversational commands.

With access to more than 400 interconnected risk metrics, KPIs and contextual data, this application "answers" rapidly by generating granular, harmonized and highly accurate reports.

Key risk signals, predictive analytics and a range of insights critical to safety are presented simply, intuitively and without human error or bias — eliminating the need to retrieve or analyze multiple reports.

The game-changing process simplification and user experience has been enhanced by a careful calibration of its prompts to align with a clinical oversight function — enabling the technology to mimic a human-like interaction. For example, by tracking dialogue history and using it to inform each new answer, it provides a "step-by-step" feature that enables users to ask a series of questions to generate custom reports or data analysis.

Looking ahead: supporting a shared vision to improve patient lives

In summary, the principles and use cases presented in this article represent significant advances in harnessing GenAI/ agentic AI to drive efficiencies and quality in clinical trials.

But we are still just scratching the surface.

At IQVIA, we are continuing to pioneer agentic AI and other advanced technologies that empower sponsors to innovate with confidence and make better informed decisions.

In keeping with these objectives, a strategic collaboration between IQVIA and NVIDIA was recently announced to advance Healthcare-grade AI® through agentic automation of complex and time-consuming processes across the therapeutic life cycle. Combining IQVIA's vast data, analytics and Connected Intelligence™ expertise with NVIDIA's AI Foundry service, the collaboration aims to automate complex workflows, enhance precision, and ensure trust and compliance via AI solutions.

So stay tuned, as we will have much more to report on how artificial intelligence evolves to improve efficiency, innovation and patient outcomes in healthcare and life sciences.

About the author

RAJNEESH PATIL
Vice President, Digital Strategy & Innovation, IQVIA

Rajneesh Patil is Vice President, Clinical Operations where he leads Digital Strategy & Innovation with a team focused on transformation of operating models, risk-based monitoring and AIML, and GenAI implementation in clinical development. Rajneesh has 16 US patents in clinical technology and analytics. He is a University Gold Medalist in Dental Surgery, studied Public Health in Australia and certified in Digital Strategy from Harvard Business School, USA.

CONTACT US

iqvia.com/rds

© 2025. All rights reserved. IQVIA® is a registered trademark of IQVIA Inc. in the United States, the European Union, and various other countries. 03.2025.RDS. BCS2025-0799-03MAR