Medable

Whitepaper

Is data readiness slowing down AI in clinical trials?

How agentic AI enables immediate impact

Table of contents

Executive summary

According to seven different market research organizations, the “Agentic AI in clinical trials market” is expected to grow at a compound annual growth rate (CAGR) of anywhere between 12.5% to 43%. While the estimates obviously vary depending on who you ask, one thing remains crystal clear, agentic AI in clinical trials isn’t slowing down.

As agentic AI capabilities continue to come to market, pharma, biotech, and CRO leaders keep circling back to the same question.

Is our organization’s data ready? However, that’s the wrong question.

Clinical data doesn’t need to be perfectly clean, standardized, or consolidated before AI agents can start delivering value. Organizations are already ready with their existing datasets and SMEs. Instead, they should ask where and how exactly to start.

1. The data readiness paradox

Clinical development generates large volumes of data. On average, a trial contains a myriad of different systems ranging from electronic data capture systems, clinical trial management systems, electronic trial master files, to safety systems, laboratory platforms, interactive response technologies, patient-facing applications etc. This fragmentation has made data readiness a persistent challenge for clinical technology programs. It can also create the impression that organizations must first solve their entire clinical data architecture before deploying AI.

Data readiness is not a binary state. An organization does not become universally “AI ready” at a specific point in time. Instead, readiness exists relative to a defined task.

Consider two potential applications. An AI agent designed to identify eligible patients for a clinical trial may require highly accurate clinical and eligibility information because incorrect inputs could contribute to patient misclassification.

In contrast, an agent designed to support clinical data review is expected to encounter incomplete, inconsistent, or unresolved data. In that case, imperfect data is not necessarily an impediment. Identifying those imperfections may be part of the agent’s purpose.

The implication is important. Organizations should evaluate data readiness at the level of the use case rather than waiting for every data source across the enterprise to reach an arbitrary standard of completeness.

---

2. Data curation may matter more than perfect data

A distinction should also be made between data cleanliness and data curation.

An AI agent requires access to information that is relevant to its assigned task. Providing more data does not necessarily improve performance. In some circumstances, excessive or poorly differentiated information may make it more difficult for an agent to identify the appropriate evidence.

Take for instance, audit trail review, which serves as a useful example.

Clinical data may show the current state of a record. Audit trail data may show how that record changed over time, including whether information was removed, restored, or modified. If an agent is asked to identify deleted adverse events, providing only the current clinical dataset may produce a materially different result from providing the appropriate audit trail.

The problem is not necessarily that either dataset is "dirty." The problem is whether the agent has been given the correct dataset for the question being asked. This suggests a more functional definition of data readiness. Data should be curated so that an agent can reliably access the information required to perform its defined task without being confused by irrelevant sources.

For clinical organizations, this is often far more achievable than attempting to create a universally pristine data environment.

---

3. Accessible data is not necessarily usable data

Technical accessibility represents only one dimension of readiness. An agent may be able to retrieve a data point without understanding its clinical or operational significance. In other words, usable data requires context.

For example, an agent performing clinical data review might identify an elevated number of open queries at a particular site. The number of queries is accessible information. Determining what that number means may require additional knowledge. The site may have recently implemented a protocol amendment. It may be experiencing an operational issue. Enrollment may have increased rapidly. The age and severity of the queries may also matter.

Human experts routinely combine these contextual signals when interpreting clinical information. Agentic systems must increasingly be capable of doing the same.

This is why the evolution from data access to contextual understanding is fundamental. Data becomes more useful when it is standardized, related to relevant metadata, and connected with the operational circumstances in which it was generated.

The objective is not simply to give an agent more information. The objective is to give it the appropriate information and enough context to interpret that information correctly.

4. A three-layer architecture for agentic clinical workflows

A scalable agentic architecture can be conceptualized across three interconnected layers, infrastructure, data, and orchestration.

The infrastructure layer

The infrastructure layer contains the underlying clinical and operational information. This may include EDC data, protocol deviation data, coding information, audit trails, laboratory data, third-party vendor data, and other sources generated during trial execution.

Historically, organizations have often connected directly to these systems through point-to-point integrations or APIs. This approach can support individual applications, but complexity increases as the number of systems, data sources, and AI applications expands.

The data layer

The data layer provides structure and context between source systems and downstream applications.

Within this layer, information can be standardized, curated, categorized, and made available in consistent formats. Clinical data can be distinguished from audit data, metadata, coding data, and protocol deviation data. Terminology can also be harmonized across systems.

This is particularly important when different platforms use different terms for equivalent concepts. One system may use “subject ID.” Another may use “subject number.” Another may use “participant ID.” A human user can often recognize these as related concepts immediately. An agent requires a mechanism for establishing the same semantic relationship.

A shared ontology can provide this common language.

The orchestration layer

The orchestration layer sits above the data foundation and determines how information is used.

This layer may contain analytics applications, visualization tools, workflow systems, AI models, and autonomous or semi-autonomous agents. It can also govern how agents interact with underlying systems and when human approval is required.

The three layers are interdependent.

An organization may be able to build an isolated agent without a mature data layer. Scaling dozens of agents across clinical workflows, however, becomes substantially more difficult without a consistent architecture for accessing, interpreting, and acting on clinical information.

5. Why interoperability becomes a strategic requirement

Interoperability is one of the most important prerequisites for scalable agentic AI.
Clinical systems have historically been designed around specific functional requirements. Data management, monitoring, safety, statistics, regulatory operations, and trial management frequently operate through different platforms and data models.

An AI agent intended to work across these boundaries requires a reliable mechanism for retrieving and interpreting information. APIs provide part of that mechanism and semantic models provide another.

Emerging approaches such as Model Context Protocol, or MCP, can further extend this architecture by giving agents standardized mechanisms for interacting with external tools and systems. In practical terms, these connections can function as the operational interface through which an agent retrieves information or initiates actions.

However, the ability to take action must be separated from the authority to act autonomously. An agent may technically be capable of modifying information in an EDC or another source system. That does not mean unrestricted autonomous modification is appropriate.

For regulated clinical workflows, organizations may require specific actions to be reviewed and approved by qualified personnel before they are committed to a source system. Other low-risk activities may support greater automation.

Agentic architecture therefore requires both technical interoperability and configurable governance.

6. Digital protocols illustrate the opportunity

Protocol digitization demonstrates how standardization can enable downstream agentic applications.

A traditional clinical protocol contains the information required to define a study, but much of that information historically exists in a document-oriented format. Digital protocol standards such as the Unified Study Definitions Model, or USDM, create a structured representation of protocol information. Once protocol information is standardized and machine-readable, numerous downstream workflows become possible.

Structured protocol information could support database design, specification development, statistical programming workflows, schedule-of-events interpretation, and other operational activities. Agents may ultimately facilitate transformations between the protocol and downstream systems.

Digitizing the protocol is the first key step.
The value of structured information depends on interoperability. A highly structured digital protocol has limited operational impact if downstream systems cannot consume, interpret, or act on its information.

This illustrates the broader principle for clinical AI. Standardization creates a foundation. APIs create connectivity. Semantic context creates understanding. Orchestration turns that foundation into operational action.

7. Trust is a technical and organizational requirement

For agentic AI to move into regulated clinical workflows, technical performance alone is insufficient. Users must trust the system. A useful framework for establishing this trust consists of three characteristics, transparency, predictability, and accountability.

Transparency

Users should be able to understand the information an agent used and, where appropriate, the reasoning or evidence supporting its output. This is particularly important when AI replaces or augments processes that were previously deterministic and readily inspectable. Traditional statistical programming, for example, allows reviewers to inspect code and understand how an output was generated.

AI systems can appear less transparent. Agentic platforms therefore require mechanisms that expose source information, configuration, activity history, and other evidence necessary for review.

Predicatbility

AI systems are probabilistic. The same input can sometimes produce variations in output.In regulated workflows, organizations need to understand whether these variations remain within acceptable performance boundaries.

Evaluation frameworks can help establish those boundaries. Human-generated or human-approved evaluation sets can define expected behaviors and allow organizations to continuously assess whether an agent continues to perform appropriately.

Predictability requires consistent performance against defined scientific and operational expectations.

Accountability

Responsibility cannot disappear when AI is introduced. Organizations should define who owns the final decision associated with an agent-assisted workflow. Audit trails should allow relevant actions and outputs to be traced. Human review should be incorporated where risk warrants it.

The fundamental principle is straightforward: AI can support a decision, but accountability for regulated clinical activities must remain clearly defined.

8. Uncertainty should be visible

One of the most consequential considerations for trust is how an agent communicates uncertainty.Clinical researchers already operate in disciplines where uncertainty is explicitly quantified. Statistical analysis does not simply provide an estimate. It also characterizes uncertainty around that estimate, allowing decision-makers to interpret the strength of the evidence.

Agentic AI benefits from a similar philosophy. A system that provides an incorrect answer with apparent certainty may damage trust more substantially than a system that identifies ambiguity and requests review.

Where technically and scientifically appropriate, agentic workflows should therefore provide signals that help users understand confidence, ambiguity, missing information, or conditions requiring escalation.

The objective should be to make limitations visible enough that qualified humans can make informed decisions.

9. The human in the loop must be a subject matter expert

Human oversight is frequently discussed as a safeguard for clinical AI. However, the presence of a human reviewer alone does not guarantee meaningful oversight. The reviewer must possess sufficient domain expertise to identify when an agent has made an error.

This creates an emerging workforce challenge.

Historically, clinical professionals developed expertise by performing many of the activities that AI may increasingly automate. A statistical programmer learned by writing code. A data manager learned by reviewing clinical data. A monitor developed judgment through direct experience evaluating sites and study information.

If entry-level professionals increasingly receive AI-generated outputs rather than producing those outputs themselves, organizations must determine how future experts will develop the knowledge required to supervise AI effectively.

This question has implications for training, competency models, career development, and organizational design. The most effective human-agent model is likely to pair automation with genuine subject matter expertise rather than treat human review as a procedural checkbox.

---

10. A practical framework for selecting initial use cases

Organizations beginning their agentic AI journey should resist the temptation to automate everything simultaneously. A focused use case provides a more controlled environment for evaluating value, trust, governance, and organizational adoption. The initial use case should satisfy several criteria.

First, the task should be clearly defined.

Agents perform best when their objectives and boundaries are explicit. Broad agents attempting to perform many unrelated activities create greater complexity and may produce less reliable results.

Second, the required data should be identifiable.

Organizations should know which information the agent requires, where that information resides, and whether it provides sufficient context for the task.

Third, human accountability should be defined.

The organization should determine who reviews outputs, who can approve actions, and who remains responsible for the resulting decision.

Fourth, performance should be measurable.

Evaluation criteria should be established before deployment. Measures might include time saved, percentage of outputs requiring correction, sensitivity for identifying relevant issues, reduction in manual effort, or other task-specific indicators.

Fifth, the use case should have a path to expansion.

A narrowly scoped project can still contribute to a larger strategic architecture.

For example, an agent supporting one component of statistical programming could eventually become part of a broader ecosystem supporting specifications, code generation, metadata, annotated case report forms, analysis datasets, and reporting workflows.

This allows organizations to start small without thinking small.

11. The three P framework for evaluating AI value

One useful way to evaluate the potential value of clinical AI is through three dimensions: predictability, personalization, and productivity.

Predictability

AI may improve the ability to anticipate future clinical or operational events. Potential applications include identifying emerging safety signals, forecasting enrollment performance, predicting site-level risks, and anticipating data quality issues.

Personalization

AI may support increasingly individualized approaches to clinical development. This includes applications related to digital twins, treatment selection, patient-specific risk assessment, and other methods that use multidimensional data to improve understanding at the individual participant level.

Productivity

AI agents can reduce the manual burden associated with repetitive and information-intensive clinical activities. Potential applications include document authoring, clinical data review, monitoring preparation, statistical programming, and other workflows where professionals spend substantial time collecting, synthesizing, or transforming information. Organizations do not need to pursue all three dimensions simultaneously. They should identify where AI can contribute most directly to their scientific and operational objectives.

12. What leading organizations can do in the next 90 days

The most significant risk may not be imperfect data. It may be prolonged inaction. Organizations can begin with three practical actions.

Identify an AI champion

Smaller organizations may benefit from establishing an individual or cross-functional group responsible for evaluating AI opportunities, coordinating strategy, and translating emerging capabilities into relevant clinical use cases. This role should connect technology with clinical, operational, quality, and regulatory expertise.

Modernize workflows and standard operating procedures

Many existing clinical processes were designed before agentic AI became viable. Organizations should evaluate where current SOPs assume that activities are performed entirely by humans and determine how AI-assisted workflows can be incorporated while preserving appropriate control, validation, and accountability. Technology cannot be successfully integrated into a process that does not recognize how the technology will be used.

Implement one controlled use case

Organizations should select one meaningful application and deploy it within a defined scope. The first implementation does not need to represent the final state. Its purpose is also to generate organizational learning. Teams can observe how users interact with agents, identify governance requirements, refine training, understand validation expectations, quantify performance, and determine where architecture needs to evolve.

A successful initial deployment can create a value and trust flywheel

As users see measurable benefits, organizational confidence can increase. Increased confidence can support additional use cases. Each implementation can strengthen the underlying data, integration, governance, and evaluation infrastructure required for subsequent deployments.

13. From data readiness to organizational readiness

The clinical research industry is unlikely to reach a moment when every dataset is clean, every system is standardized, and every workflow is fully integrated. Waiting for that moment is therefore not a viable AI strategy.

Organizations instead need to determine what level of readiness is appropriate for each application. Some agents will require highly controlled inputs. Others can derive value specifically because the underlying data contains inconsistencies requiring review.

Across both scenarios, the same foundational requirements remain important: appropriate data curation, contextual understanding, interoperability, semantic standardization, governance, evaluation, and qualified human accountability.

This shifts the conversation from "Is the data ready?" to a more useful question. "Is the organization ready to define a controlled use case, give an agent the right context, measure its performance, and learn from its deployment?"

For clinical research organizations, that distinction may determine how quickly agentic AI moves from an emerging technology to a reliable component of clinical development.

The organizations most likely to benefit will not necessarily be those with perfect data. They will be those that identify the right problems, establish the right controls, develop the right expertise, and begin learning while the technology and the operating model continue to evolve.

To learn more about Medable and our Agentic AI capabilities, visit medable.com