

Beyond the Algorithm: Why Trust is the Cornerstone of Agentic AI in Pharma



Agentic AI—autonomous systems capable of planning, acting and adapting—are rapidly moving from experimental pilots to core roles in pharmaceutical research, development and patient engagement. As these systems become more influential, the question of trust is no longer academic; it is central to their adoption and impact.

AI is reshaping how Pharma works. Case studies show that agentic AI can deliver measurable gains: a 30% increase in productivity, reduced error rates and faster drug discovery cycles. However, these benefits are only realized when trust is established through transparency, oversight and ongoing evaluation (DigitalDefynd, 2025).

Building trust in agentic AI begins with explainability and transparency. As highlighted by an expert in the field, Nate Nichols, healthcare professionals need agentic AI systems to justify their recommendations with clear reasoning and references to medical evidence. Doctors

want to see the logic, the sources and the process behind an AI's suggestion, much as they would expect from a human colleague. For patients, however, the need is different: they seek intuitive, accessible explanations that make sense to them personally. This dual requirement means agentic AI must be able to communicate effectively with both audiences, adapting its explanations to the context and the recipient's needs (ARF Pharma Council, 2025).

There are challenges in fulfilling these needs. Large language models, the backbone of many agentic AI systems, sometimes “bluff” their reasoning—constructing plausible explanations after the fact that don't reflect the actual decision process. Ongoing research is focused on ensuring that AI systems can provide genuine, auditable accounts of their actions. Regulatory frameworks are emerging to support this, enabling adverse outcomes to be traced and responsibility assigned when mistakes occur (ARF Pharma Council, 2025; Clozel, 2025).

A promising model for safe adoption is to treat agentic AI as an “intern”—capable of making a first pass at a problem, but always subject to review by a qualified healthcare provider. Before any recommendation reaches a patient, a human must validate it. This approach reassures patients that a human has reviewed the advice and gives providers confidence that the agent is not acting beyond its competence. This model is already reflected in real-world deployments, such as Atropos Health’s Evidence Agent, which generates real-world evidence at the point of care but always within a framework of human oversight and auditable decision-making (Landi, 2025).

Comfort with agentic AI varies by generation and context. Younger patients, accustomed to interacting with AI in other domains, may be more willing to share information and accept advice from AI agents. However, when it comes to high-stakes decisions—such as medication changes or diagnoses—most patients and providers still seek human validation. Industry surveys echo this: while a majority of pharma executives see agentic AI as a top priority, only about half of end-users trust AI to consistently deliver correct answers, and even fewer trust it with critical decisions like drug pipeline management (Clozel, 2025).

For pharma, trust in agentic AI is built on several pillars. Data integrity and security are foundational. Agentic AI’s value is clearest in data-heavy tasks: harmonizing, cleaning and analyzing multimodal datasets. Trust hinges on robust data governance, secure access and clear audit trails. Companies must invest in unified, orchestrated data platforms to ensure agents have the right information and that every action is traceable (AWS, 2025).

Regulatory compliance is equally critical. Pharma is among the most regulated industries, and agentic AI systems must be designed with compliance in mind. Every recommendation must be defensible, and adverse outcomes must be traceable to their source. This requires collaboration between AI developers, pharma companies and regulators to establish clear standards for transparency, auditability and human oversight (Coulter, 2025).

Experts recommend starting with low-risk, well-defined tasks—such as automating routine data analysis or patient

Sources:

ARF Pharma Council (Host). (2025, September 22). Agentic AI—Part 1: What is it and how could it change health care? (No. 26) [Audio podcast episode]. In *Future Health*. [ARF Pharma Council](#).

AWS. (2025, October 24). *Three critical data shifts your company needs to thrive in the agentic AI era*. [Business Insider](#).

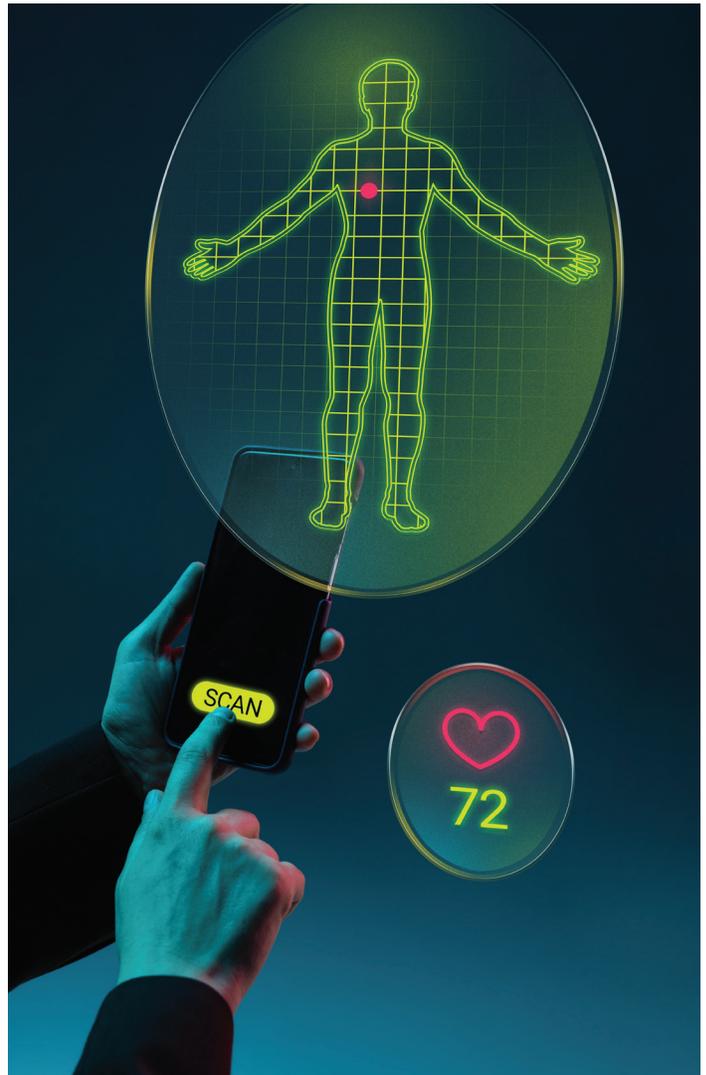
Balvair, A. (2025, July 11). *Agentic AI revolutionizes pharma efficiency and trust*. [HealthPoint](#).

Clozel, T. (2025, October 15). *Can pharma trust the new wave of agentic AI?* [Fast Company](#).

Coulter, B. (2025, July 3). *From hype to impact: How pharma can use agentic AI to drive efficiency and trust*. [MedCityNews](#).

DigitalDefynd. (2025). *Agentic AI in the pharma industry [5 case studies]*. [DigitalDefynd](#).

Landi, H. (2025, October 28). *Atropos Health launches agentic AI for clinical evidence at Stanford Health Care*. [Fierce Healthcare](#).



scheduling—before expanding agentic AI’s scope. Strong governance structures, including clear escalation paths and human-in-the-loop review, are essential for building trust and ensuring safe adoption (Balvair, 2025).

Ultimately, trust in agentic AI is not a given—it must be earned through explainability, transparency and robust human oversight. For pharma, the path forward involves strategic alignment, thoughtful implementation and a commitment to measurable outcomes. As agentic AI matures, those who invest in building trust will be best positioned to realize its transformative potential.