



## RESEARCH ARTICLE

# Perspectives on Accelerating Successful Implementation of Artificial Intelligence in Biopharma and Healthcare

John M. York<sup>1-4</sup>; Arthur A. Boni<sup>5</sup>; Diana Joseph<sup>6</sup>; Mikel Mangold<sup>7</sup>; and Sarah Marie Foley<sup>8</sup>

<sup>1</sup>University of California, San Diego;

<sup>2</sup>Cranfield School of Management;

<sup>3</sup>Ernest Mario School of Pharmacy, Rutgers University;

<sup>4</sup>Burnett School of Medicine, Texas Christian University;

<sup>5</sup>Tepper School of Business, Carnegie Mellon University;

<sup>6</sup>Corporate Accelerator Forum;

<sup>7</sup>ATLANT 3D;

<sup>8</sup>GlaxoSmithKline

## ABSTRACT

Artificial intelligence (AI) and large language models (LLMs) promise to reshape discovery, development, and commercialization across the biopharmaceutical value chain. Yet adoption remains uneven, marked by fragmented pilots, governance constraints, and wide variation in organizational readiness. To address these challenges, this study applies diffusion-of-innovation theory and human-centered service-design principles to investigate two questions: (1) What is the current status of AI and LLM implementation in the biopharmaceutical industry? (RQ1) and (2) Which collaborative structures and organizational designs enable scalable, value-generating AI platforms? (RQ2).

The study employs a mixed-methods approach that integrates a narrative synthesis of peer-reviewed and industry literature (2020–2025) with six semi-structured interviews involving leaders in drug development, clinical operations, MedTech, and digital strategy. Triangulating across published evidence, organizational practice, and expert experience allows the analysis to surface cross-cutting patterns that would not appear through a single-method design.

Findings reveal a multi-speed diffusion pattern: rapid adoption in discovery and document generation, slower progress in clinical and regulatory settings, and emerging use cases in manufacturing and commercialization. Interview insights highlight persistent obstacles—risk-tiered governance, data-validation gaps, and workflow misalignment—alongside enablers such as platform partnerships, venture-client models, regulatory consortia, and startup co-development. Human-centered workflow design, biologically grounded analytics, and transparent governance consistently emerged as essential for moving AI from isolated experimentation to enterprise-level capability.

This paper contributes a unified framework that clarifies the current state of AI and LLM adoption in biopharma and identifies the organizational and ecosystem mechanisms required for responsible, scalable implementation. By linking external collaborative structures with internal governance, data, and workflow designs, the study offers a practical, conceptually-grounded blueprint for transitioning from pilots to enterprise-level transformation.

**Keywords:** Artificial intelligence, Large language models, Biopharmaceutical industry, Collective Intelligence, Diffusion of innovation, Human-centered design, Service design thinking

## OPEN ACCESS

## PUBLISHED

31 January 2025

## CITATION

York, J.M., et al., 2025. Perspectives on Accelerating Successful Implementation of Artificial Intelligence in Biopharma and Healthcare. Medical Research Archives, [online] 14(1).

## COPYRIGHT

© 2025 European Society of Medicine. This is an open-access article distributed under the terms of the Creative Commons Attribution License, which permits unrestricted use, distribution, and reproduction in any medium, provided the original author and source are credited.

## ISSN

2375-1924

## Introduction

Artificial intelligence (AI) continues to promise fundamental transformation across the biopharmaceutical industry. AI-enabled methods now influence every stage of the drug-development life cycle, ranging from target identification and molecular screening to adaptive clinical-trial design, regulatory reporting, and post-market surveillance.<sup>1,2</sup> However, the industry appears to be adopting these technologies in uneven and fragmented ways. Most organizations still advance isolated pilot projects rather than integrated, enterprise-level platforms. Structural barriers—including governance uncertainty, fragmented data assets, and intricate dependencies across scientific, clinical, and commercial functions—continue to slow progress.

Scholarly and industry literature highlights significant advances in digital health and biotechnology – such as machine-vision models for imaging, omics analytics, and predictive modeling.<sup>3,4</sup> However, they rarely address the organizational and regulatory processes and tensions that constrain AI implementation in highly regulated environments. Cross-functional alignment remains difficult in ecosystems involving patients, clinicians, payers, regulators, and commercial partners.<sup>5</sup> Evidence increasingly shows that organizational bottlenecks, data quality issues, and workflow incompatibilities delay adoption more than the technical models themselves.

The lack of integration of user-centered innovation compounds this challenge. Service design—which co-creates solutions grounded in user needs and system constraints—has proven effective in digital health but remains underutilized in biopharmaceutical AI adoption.<sup>6</sup> Without structured, human-centered methods, organizations risk producing "demo moments" that impress technically but fail to scale. Research in human-factors science consistently shows that trust, transparency, explainability, and workflow fit strongly influence adoption decisions, making service-design methods essential for sustainable and credible AI integration.<sup>6</sup>

Guided by these gaps, this study investigates two research questions: (1) What is the current status of AI and large-language-model (LLM) implementation in the biopharmaceutical industry? (RQ1) and (2) Which collaborative structures and organizational designs enable scalable, value-generating AI platforms? (RQ2). Addressing these questions requires an integrated perspective that links technical maturity with organizational capability, governance expectations, and real-world user behavior.

The analysis draws on innovation-theory frameworks that explain how organizations absorb new technologies. Christensen's disruptive innovation and jobs-to-be-done concepts,<sup>7,8</sup> Rogers's diffusion-of-innovation theory,<sup>9</sup> and Stickdorn's service-design principles<sup>6</sup> collectively reveal why AI adoption progresses through predictable cycles of experimentation, capability building, and cross-stakeholder negotiation. These frameworks help explain why some technologies scale while others stall despite comparable technical promise. Prior studies show that organizations applying service design outperform traditional top-down approaches because they align solutions with stakeholder motivations and operational constraints.<sup>10</sup> In biopharma, these methods help validate AI-enabled workflows with clinicians, regulators, and patients before large-scale deployment, revealing tensions and hidden barriers that would otherwise slow or derail innovation.

Scaling AI also requires new collaborative architectures. Because AI tools rely on shared data, validation protocols, and compliance alignment, scholars emphasize platform-centered ecosystems rather than transactional vendor relationships.<sup>11,12</sup> Such structures support joint data stewardship, coordinated governance, and reproducible validation—all essential in regulated environments where safety, traceability, and accountability are paramount.

This paper advances the field by offering a unified framework that links the current landscape of AI and LLM implementation in biopharma with the collaborative structures and organizational designs necessary for scale. By integrating diffusion-of-

innovation theory, service-design principles, and practice-based insights from industry experts, the study addresses the longstanding gap between technical potential and real-world implementation. It explains why adoption remains uneven and shows how governance, workflow alignment, and multi-stakeholder collaboration can help organizations move beyond isolated pilots. Accordingly, this mixed-methods study clarifies the current state of AI adoption (RQ1) and identifies the collaborative and organizational mechanisms required to build scalable, value-generating AI platforms (RQ2).

The flow of the paper proceeds with grounding the analysis in framing literature, outlining the mixed-methods design, presenting findings from the literature review and expert interviews, and discussing key considerations and contributions involved with building platform-based AI ecosystems.

## Framing Literature

### *DIFFUSION OF INNOVATION*

Rogers's diffusion-of-innovation theory<sup>9</sup> provides a valuable framework for understanding how AI diffuses in the biopharmaceutical industry and how human-centered service design supports customization and adoption during this technological shift. The theory explains how new ideas spread across social systems and identifies five perceived innovation attributes—relative advantage, compatibility, complexity, trialability, and observability—that shape adoption rates. Health-sector syntheses show that innovations diffuse more rapidly when they offer clear advantages, align with stakeholder norms, and allow transparent experimentation and observable outcomes.<sup>13,14</sup>

In regulated sectors such as the biopharmaceutical industry, diffusion also depends on resources derived from internal and external partnerships and on the degree of socio-technical fit.<sup>15</sup> Effective collaboration accelerates or constrains technology uptake depending on alignment, readiness, and the availability of governance structures that reduce perceived risk.

When applied to AI and large language models, Rogers's model identifies several structural tension

points. AI offers a substantial relative advantage by compressing discovery timelines and improving clinical-trial efficiencies, yet executives often perceive it as complex, risky, and constrained by regulatory uncertainty.<sup>16</sup> The scarcity of low-risk pilot environments and the limited availability of validated outcomes dampen adoption, especially among late-majority organizations that require transparency, governance clarity, and evidence of reliability. Historical diffusion patterns suggest that innovators and early adopters engage first, while the early and late majorities wait for established standards, repeatable workflows, and proven implementations.<sup>17</sup>

### *HUMAN-CENTERED SERVICE DESIGN*

Human-centered service design helps bridge these barriers by facilitating collaboration among clinicians, regulators, patients, and technical teams. Service design enhances compatibility with existing workflows and professional norms.<sup>6</sup> Iterative prototyping creates low-risk opportunities to test artificial-intelligence-enabled tools, producing transparency and reassurance for skeptical stakeholders. Common service-design artifacts—like journey maps, service blueprints, and value-proposition canvases—surface concerns about algorithmic opacity, data privacy, and workflow disruption. These insights support targeted mitigation strategies that increase perceived advantage while reducing complexity.<sup>10</sup>

Through this process, service design operationalizes Rogers's innovation attributes to address adopter concerns directly. Integrating diffusion theory with service-design principles strengthens the central proposition of this paper: artificial-intelligence deployment in the biopharmaceutical sector depends less on technical readiness and more on orchestrating stakeholder-focused adoption trajectories. Conceptualizing artificial-intelligence rollout through Rogers's five innovation-decision stages—knowledge, persuasion, decision, implementation, and confirmation—clarifies how governance, partnerships, and iterative learning cycles interact to advance adoption.<sup>9,14,15</sup> Aligning artificial-intelligence initiatives with user expectations, workflow realities, and regulatory norms accelerates

the transition from isolated pilots to enterprise-scale, value-creating systems.

## Methods

This study employed a mixed-methods approach that combined a narrative review of the literature with exploratory expert interviews to examine structural, organizational, and stakeholder-associated challenges surrounding artificial-intelligence application within the biopharmaceutical industry. This design enabled the triangulation of data across published evidence, regulatory commentary, and lived practitioner experience. This approach allowed the study to surface cross-cutting themes that would not emerge through a single-method approach.

The narrative review synthesized peer-reviewed articles, industry analyses, regulatory publications, and commentary pieces from 2020–2025 to identify emergent themes across technological, organizational, and policy domains relevant to artificial-intelligence implementation in drug discovery, clinical development, manufacturing, regulation, commercialization, and post-market surveillance. The analysis organized these materials into seven topical categories: (1) current status, (2) system types, (3) organizational maturity and partnership models, (4) regulatory considerations, (5) functional areas of application, (6) ethics, data security, and privacy, and (7) stakeholder interactions. This categorization provided a structured foundation for comparing literature-based insights with the expert interview findings.

To complement the literature synthesis, the study engaged in six semi-structured interviews with experts from biopharma, academia, and startup ecosystems. These conversations explored three thematic areas: innovation models for organizations, the foundational principles of human-centered service design, and practitioner perspectives on artificial-intelligence implementation. This interview component examined how AI and LLMs are being deployed in real organizational contexts, why adoption progresses unevenly, and how governance, workflow design, data readiness, and service-design practices shape

integration. Interviewees represented diverse roles – spanning the areas of clinical leadership, digital strategy, MedTech product development, and service design. This purposeful sample allowed the study to capture perspectives from multiple points along the biopharmaceutical value chain and to triangulate lived practitioner experience with the diffusion-of-innovation and service-design frameworks.<sup>6,9,13-15</sup>

The interview set reflected a broad spectrum of expertise across pharmaceutical innovation, clinical development, AI research, and human-centered design. **Shwen Gwee** contributed two decades of digital innovation and clinical-digital-strategy leadership, drawing on early work with conversational AI and experience guiding large pharmaceutical firms through emerging-technology adoption. **Mithun Ratnakumar** added cross-functional insight into regulated software, SaMD development, and digitalization strategies from Novartis and Gerresheimer, including end-to-end partnerships and venture-client collaboration models. **Robert Brown, M.D.**, brought a clinician-executive perspective informed by oncology-trial leadership and practical challenges in trial optimization, multimodal data interpretation, and AI validation. **Pallavi Tiwari, Ph.D.**, extended this clinical and technical viewpoint through her academic and entrepreneurial work on oncology imaging and multimodal biomarker development at the University of Wisconsin–Madison and LivAi. **Simon Fortenbacher** at GSK emphasized workflow redesign, democratized development, and alignment of AI-enabled processes with on-the-ground constraints. **Mark Guarraia** at Novo Nordisk focused on trustworthy human–machine interactions and organizational change management in regulated environments. Together, these voices reflected how AI adoption hinges not only on technical capability but also on governance maturity, workflow fit, and the human factors that determine whether innovation takes hold.

This analysis applied an informed-expert methodology<sup>17</sup> that triangulated with the literature in an abductively designed to surface patterns that would not emerge through a single-method approach, rather

than mirroring formal qualitative traditions such as ethnography or grounded theory.<sup>17,18</sup> It aligned with established guidance in qualitative interviewing<sup>17</sup> and interpretive analytical practice,<sup>19</sup> emphasizing meaning-making, experiential context, and the articulation of practitioner reasoning. By juxtaposing insights from the literature review with themes drawn from expert interviews—such as risk-tiered governance, partnership models, workflow barriers, human-centered adoption dynamics, and data-readiness constraints—this study triangulated across published evidence and informant operational experience.<sup>20,21</sup> This iterative comparison allowed the analytic effort to identify cross-cutting mechanisms that recurred across discovery, clinical development, MedTech integration, and workflow design. The synthesis of these convergent themes ultimately supported to the proposal of an explanatory internal and external interactive model to link AI's current status in biopharma with the collaborative structures and organizational designs required for scalable implementation.

## Findings

### *LITERATURE REVIEW (2020–2025)*

#### Current Status

The biopharmaceutical industry has rapidly transitioned from speculative exploration to operational deployment of AI and LLMs. Over the last five years, AI applications have accelerated timelines, improved decision-making, and enhanced predictive analytics across drug development, regulation, and commercialization.<sup>22–24</sup> Reported benefits include up to 40% time savings in regulatory document preparation and 50% throughput improvements in pharmacovigilance.<sup>25</sup>

AI integration now spans nearly all phases of the biopharma value chain. Early adoption centered on automating adverse event processing and medical writing, while recent advances extend to protocol generation, trial simulation, adaptive designs, and real-world evidence modeling.<sup>22,26,27</sup> AI platforms such as MedPaLM and BioGPT assist in regulatory documentation and consistency checks, while tools

like TrialGPT simulate trials and refine study protocols in real time.<sup>26,27</sup> Organizations increasingly apply AI for translational purposes, including omics-based patient stratification, improving alignment between preclinical and clinical outcomes.<sup>24</sup>

Despite these advances, adoption remains uneven. High-income nations dominate implementation, while infrastructural and data limitations restrict AI use in low- and middle-income countries (LMICs).<sup>22</sup> Moreover, many organizations operate at a semi-supervised stage, with the need for continued human oversight for compliance and scientific validity.<sup>28</sup>

Artificial intelligence in biopharma encompasses diverse paradigms: supervised learning for pattern recognition, reinforcement learning for production optimization, and generative AI for molecular design and document synthesis.<sup>22</sup> Generative models such as BioGPT and MedGPT support literature summarization, regulatory authoring, and workflow simulation.<sup>24,25</sup>

Deep neural networks, as exemplified by AlphaFold, predict protein structures with near-experimental precision.<sup>29</sup> Natural language processing models streamline regulatory submissions by analyzing historical approval trends and aligning applications with regional standards.<sup>30</sup> Together, these demonstrate the sector's transition from individual use cases to comprehensive, modular AI ecosystems designed to address specific therapeutic and operational challenges.<sup>22–27</sup>

#### The Emergence of Agentic Artificial Intelligence

Recently, agentic AI—autonomous, goal-directed systems capable of executing multi-step workflows—has emerged as a significant evolution beyond generative models.<sup>31–33</sup> These agents independently plan, adapt, and act on real-time feedback, enabling autonomous trial monitoring, eligibility adjustments, and safety alerts.

According to industry analyses, agentic AI attracted the highest investment volumes in early 2025, with platforms like PharmAgents and the Agentic Preformulation Pathway Assistant (APPA) demonstrating

end-to-end automation of discovery and formulation tasks.<sup>31-34</sup> Early studies report reductions of up to 50% in manual review time and experimental repetition.<sup>35</sup>

However, scholarship has yet to thoroughly examine how agentic AI integrates within organizational governance and regulatory systems. Current discussions emphasize technical potential but seldom address the socio-technical coordination required for multi-stakeholder trust and compliance.<sup>31,32,35</sup>

### Organizational Maturity and Partnership Models

AI maturity varies widely among pharmaceutical firms, typically following four stages: exploratory, operational, strategic, and transformational.<sup>36</sup> Recent benchmarking reports show that most companies have progressed to the "strategic" phase, emphasizing cross-functional collaboration between data science and governance boards.<sup>37,38</sup>

Partnerships with AI startups such as Exscientia and Insilico Medicine have become pervasive; roughly 70% of leading global biopharma companies maintain formal collaborations with AI providers.<sup>39</sup> Hybrid deployment models prevail, with companies retaining sensitive data in-house while AI computation occurs on secure cloud platforms to balance security and scalability.<sup>40</sup>

The 2025 Pharma AI Readiness Index highlights that even top-tier firms differ substantially in maturity. Eli Lilly, Merck, and Bayer scored highest in innovation and execution readiness, while others lag due to limited AI literacy and cross-disciplinary collaboration capacity.<sup>41</sup> These findings reinforce that successful AI adoption requires not only infrastructure but organizational learning and workforce transformation, which represents an area where service design methodologies can create integrative pathways.

### Regulatory Considerations

The global regulatory landscape for AI remains fragmented. The U.S. Food and Drug Administration has advanced toward total-product-lifecycle frameworks for AI/ML-based software.<sup>42</sup> However, the European Medicines Agency and Japan's Pharmaceuticals and Medical Devices Agency have

not codified equivalent systems. This inconsistency complicates international deployment and increases operational risk.<sup>43</sup>

Regulators increasingly emphasize traceability, auditability, and human-in-the-loop validation for AI systems influencing clinical or regulatory decisions.<sup>44</sup> Emerging guidance calls for explainability, regular retraining, and explicit documentation of AI-informed decision processes.<sup>45</sup> Cross-border research consortia, such as PharmAI, are leading harmonization efforts to standardize AI validation procedures.<sup>46</sup>

### Functional Areas of Application

AI now supports drug discovery, clinical development, manufacturing, and commercialization. In discovery, generative algorithms accelerate molecular design and target identification, improving pipeline efficiency by 30–40%.<sup>22,47</sup> Multi-omics integration enhances early disease stratification but raises bioethical and dual-use concerns.<sup>48</sup>

In translational medicine, AI identifies biomarkers for toxicity and treatment response, aligning preclinical and clinical datasets.<sup>22</sup> In clinical trials, tools such as TrialGPT assist with adaptive trial design and patient screening, improving recruitment efficiency by up to 30%.<sup>26,27,49</sup> Machine learning facilitates real-world evidence mining for patient matching and stratification, improving trial accuracy.<sup>50</sup>

Regulatory affairs leverage natural-language-processing-driven systems to harmonize submissions, reduce redundancies, and analyze feedback.<sup>30,43</sup> Pharmacovigilance systems use AI to extract adverse-event signals from unstructured data, reducing case-processing time by more than 50%.<sup>25,28</sup>

In manufacturing, AI-enabled digital twins and reinforcement learning optimize parameters such as pH and temperature in real time.<sup>51</sup> These tools improve yields and product consistency while requiring strict adherence to good manufacturing practice validation.<sup>46,51</sup>

Finally, in commercialization and medical affairs, AI enhances targeting precision and scientific communication. Predictive segmentation has improved

engagement by more than 20%, while chatbots manage basic inquiries, freeing human staff for more complex interactions.<sup>52,53</sup>

### Ethical, Regulatory, Data Security, and Privacy Considerations

Ethical and data-governance concerns persist as central challenges to AI implementation. Algorithmic bias and nonrepresentative datasets reduce external validity, particularly for populations from low- and middle-income regions.<sup>22,54,55</sup> Human-in-the-loop oversight, rigorous data auditing, and transparent training documentation are essential safeguards.<sup>56</sup>

Explainability and model interpretability remain critical for acceptance; organizations resist opaque "black box" systems unless such elements provide transparent, auditable outputs.<sup>45,57</sup> Calls for harmonized global standards persist, with several scholars advocating for an "FDA for AI" regulatory body to unify oversight mechanisms.<sup>43</sup>

Security, privacy, and dual-use risks require special attention. Generative systems can inadvertently create bad outcomes, such as synthesizing harmful compounds, providing incorrect instructions, or failing to address serious issues.<sup>58</sup> Risk of such problematic outcomes prompts the use of federated learning, differential privacy, and ethical frameworks emphasizing transparency, accountability, and fairness.<sup>58</sup>

### Stakeholder Requirements and Interactions

Despite AI's technical evolution, literature remains limited regarding stakeholder engagement in implementation. Most publications emphasize algorithmic development rather than ecosystem integration. Generative tools like AlphaFold, BioGPT, and TrialGPT demonstrate technical success.<sup>22,27,29</sup> However, few studies explore how solutions align with stakeholder needs: patients, regulators, payers, and scientists.<sup>59</sup>

Service design, a framework proven effective in digital health, remains underapplied in AI-driven

biopharma innovation.<sup>10</sup> By mapping pain points across the value chain, service design ensures that AI solutions address both systemic constraints and user requirements. Without such integration, firms risk developing technically sophisticated but socially misaligned technologies that fail regulatory or market acceptance.<sup>6,12</sup>

This research bridges a significant gap by combining service-design methodology with innovation-diffusion theory to propose an ecosystem-centered model for responsible AI adoption in biopharma.

### Synthesis

This review provides a current status of AI use in the biopharma industry, setting the basis for addressing RQ1. Further, across the literature, scalable AI adoption emerges when external collaboration structures—such as platform ecosystems, startup partnerships, venture-client arrangements, and regulatory consortia—interact with internal organizational design elements, like risk-tiered governance, hybrid cloud architectures, and service-design–driven workflow adaptation. These elements form the structural basis for answering RQ2. The interview findings that follow deepen and contextualize these mechanisms, illustrating how they operate inside real-world biopharmaceutical environments.

### *INTERVIEW FINDINGS (DETAILED INTERVIEW SUMMARIES IN SUPPLEMENTAL MATERIAL APPENDIX)*

#### Theme 1: Acceleration and Fragmentation of Artificial Intelligence Adoption

Across interviews, experts described a sector experiencing rapid acceleration in AI interest but considerable fragmentation in actual implementation. Gwee emphasized that pharma's adoption response to generative AI followed the same cycle as prior digital innovations: initial blocking, expensive internal experimentation, and eventual licensing as the need for speed grows (Table 1). This trajectory reflects Rogers's innovation-decision stages.<sup>9</sup>

Table 1. Pattern of Biopharma Industry's Adoption of New Technologies (Gwee Interview)

Phase Number	Phase Name	Description
1	Block	Organizations initially prevent any use of the technology until leadership determines whether the capability is truly necessary or strategically justified.
2	Build	Firms attempt to build the technology internally or through major consulting partners—often at high cost—only to realize that internal teams and external vendors may lack the deep expertise required.
3	License	Companies shift toward licensing external technologies to accelerate adoption and reduce development burden, moving from internal builds to acquiring proven solutions from established vendors.

Brown and Tiwari both confirmed that meaningful adoption began only in the past two to three years. Brown noted that early attempts with LLMs were often "more work than doing it manually," and Tiwari recalled that pharma declined many early AI partnerships due to mistrust of data sharing. Digital-native companies such as Moderna moved fastest, with Gwee highlighting that "the first team to reach 100% GenAI adoption was legal," exemplifying how firm culture and structure shape adoption speed.

#### Theme 2: Risk, Governance, and Compliance as Primary Adoption Drivers

Interviewees repeatedly stressed that organizations evaluate AI through a risk hierarchy, not a technology lens. Gwee articulated this hierarchy in four tiers: internal automation, clinical data impact, public exposure, and potential patient harm (Table 2). Most current investments fall within the lowest tier, where risk is manageable, and workflows remain internal.

Table 2. Biopharma Evaluation of Artificial Intelligence via a Four-Tier Risk Hierarchy (Gwee Interview)

Level	Level Name	Description	Implications for AI Use
1	Internal Automation	Internal, non-clinical automation tasks that do not influence patient data or external communications.	Suitable for early experimentation with minimal governance requirements.
2	Clinical Impact	AI tools that affect clinical data, trial workflows, or evidence generation.	Requires higher scrutiny, validation protocols, and domain-expert oversight before implementation.
3	Public-Facing Outputs	AI-generated outputs that reach external audiences, including patients, clinicians, or the general public.	Necessitates extreme caution, structured review processes (e.g., medical, legal, and regulatory review), and strict accountability controls.
4	Potential for Patient Harm	Systems with potential to influence treatment recommendations or clinical decision-making.	Demands the highest level of governance, rigorous safety monitoring, and full regulatory compliance.

Regulatory ambiguity exacerbates risk. Ratnakumar explained that early teams encountered regulatory bodies that "did not even have the answers," especially for adaptive or learning systems. Brown

highlighted the legal exposure that mid-sized biotechs face when AI accuracy must be verified manually: "Responsibility still rests with the company, not the vendor." These perspectives align with regulatory

scholarship emphasizing explainability, auditing, and traceability.<sup>44</sup>

### Theme 3: Data Quality, Integration, and the Infrastructure Bottleneck

Every interviewee emphasized that AI adoption is constrained less by algorithms than by data quality and integration. Ratnakumar described models that lacked clean validation datasets, stalling progress. Tiwari framed the challenge succinctly: "It is not the model; it is the plumbing." Brown explained that smaller companies cannot support the compliance requirements associated with sharing sensitive clinical datasets.

These recurring insights echo literature identifying data readiness as a core barrier to AI maturity.<sup>22,26</sup>

### Theme 4: Partnership Models and Vendor Dynamics

Partnership models for AI adoption varied widely, but all interviewees raised concerns about effort-to-value imbalance. Brown described testing a vendor's AI tool for regulatory filings. He highlighted three core issues: (1) fees (which are discounted apparently), (2) page-by-page verification (which

seems to be an issue of requiring manual work, reducing the value of the AI since one has been doing the same thing you would have done, and (3) data use for model training. He described four partnership model approaches: direct collaboration, reliance on contract research organizations that in-license AI tools, and reduced-fee data-sharing arrangements where vendors and partners use company datasets to train their external data. Brown added that each model offered theoretical advantages, yet in practice, none produced the expected gains. The overhead of securing data integrity, ensuring compliance, and validating outputs often consumed more resources than the efficiencies promised. This consideration, he noted, underscored both the promise and the current limitations of AI in biopharma operations.

Ratnakumar charted the evolution from end-to-end consulting models to venture-client approaches (Table 3). He stressed that pharma cannot become a software company and therefore must partner effectively. Tiwari showed how startups like LivAi structure collaborations carefully to preserve intellectual property, deliver actionable insights, and build toward platform-based licensing (Table 4).

Table 3. Biopharma Collaboration Models (Ratnakumar Interview)

Model Number	Model Name	Maturity Level (5 = Oldest, 1 = Newest)	Description
1	End-to-End Development Partnerships	5	Pharma engages large system integrators or consulting firms to deliver full medical-device or software development projects. This long-established model allows internal teams to focus on core competencies while external partners manage engineering, productization, and compliance.
2	Startup Project Funding	4	Companies provide trained models, datasets, and funding through statements of work to early-stage startups. The startups then commercialize products under their own regulatory framework, accelerating innovation without internal development overhead.
3	Exclusive Partnership Licensing	3	Pharma firms collaborate with partners to license and commercialize digital or AI-enabled solutions—especially in highly regulated areas like digital therapeutics—reducing build-time while accessing specialized expertise.

Model Number	Model Name	Maturity Level (5 = Oldest, 1 = Newest)	Description
4	Venture-Client Model	1	Companies act as early customers for startups, purchasing project work rather than taking equity. This model lets pharma shape the startup's early go-to-market strategy while retaining the option to acquire the assets or company later. It is the newest and fastest-growing model.

Table 4. Evolving Partnership Models Used by LivAI (Tiwari Interview)

Category	Model Name	Description
Current Collaboration Models	Bespoke Statements of Work	Pharma contracts startups for project-specific deliverables, predefined analyses, and tightly scoped workflows.
	Insight-Sharing Collaboration	Startups provide patient-level and population-level insights while retaining intellectual property rights to proprietary algorithms and models.
	Co-Development Frameworks	Pharma's translational science teams and clinical stakeholders work with startups to refine, validate, and iteratively improve models.
Emerging Revenue Models	Pay-Per-Use Pricing	Charging per patient or per analysis for AI-enabled assessments; viewed by pharma as expensive at scale.
	Platform Licensing Agreements	Fixed annual licensing fees (e.g., ~\$500,000/year) enabling unlimited patient analyses across multiple trial functions, often with options for exclusivity.

Gwee noted that smaller AI-native companies such as Insilico are now generating clinical-stage compounds, currently nineteen in five years, indicating a shift in industry power dynamics.

#### Theme 5: Workflow Fit, Human Factors, and the Role of Service Design

Fortenbacher and Guarraia highlighted that many digital transformations fail because companies "apply new software to yesterday's workflows." Guarraia defined Service Design as "the discipline that makes sure new human-machine interactions feel credible and trustworthy, and that this is done by meeting people where they are." It reframed AI transformation around lived experience rather than technical capability.<sup>6,10</sup>

Getting people to adopt new ways of working has been often more complex than the technology

itself. Fortenbacher noted that politics often takes center stage in larger transformations. Projects that touch on automation or efficiency stir up fears of job loss, leading to "hidden resistance." These were not the loud critics in a workshop, but the quiet blockers who slow things down. Service design helps surface these dynamics early so that leaders can respond before progress with projects derails.

Transparency is another balancing act. Guarraia stressed the delicate balance required in communication: "Too little transparency, and people lose trust. Too much honesty too early, and you create fear." Both designers argued that service design identifies leverage points, reveals white space, and ensures human-machine interactions feel credible.

The outcome of service design is the value that deliverables/artifacts do not deliver value on their

own. One of the things Fortenbacher notes when selling Service Design is the challenge of quantifying the return in business terms. Leaders respond better to "this intervention could save ten minutes per workflow across 500 scientists," as compared with "it improves the employee experience.' The "Catch 22" of this framing is that the value often comes out in the discovery process, but you need buy-in before you do the discovery work. If the value has been shown prior, tying service design to outcomes like reducing downtime, simplifying complex workflows, or improving patient adherence, as Guarraia has done, is enough.

Service design allows for the identification of "whitespace." By mapping work as how it is really done, service design uncovers opportunities that no one was looking for, whether that is automating a painful handoff, simplifying stage-gate reviews, or redesigning how factory workers interact with dozens of IT systems.' states Guarraia. These uncovered opportunities open up other ways of working that can change the game.

Why this matters is that people are the core of organizations. For a transformation to be successful, it should be linked to the needs of the individuals in that organization. Service design is human-centered and identifies where to add value. It builds trust around changes in ways of working because this approach rests on understanding those needs and uncovers opportunities leaders might not otherwise see. Done right, service design is the function that *turns AI from a technical deployment into a meaningful, lasting shift in how work gets done.* Their insights reinforce Rogers's innovation attributes of compatibility and complexity reduction.<sup>9,13,14</sup>

### Theme 6: Clinical Impact and the Future of Artificial Intelligence-Enabled Decision Support

Clinical leaders agreed that AI's most substantial near-term value lies in augmenting human interpretation rather than replacing it. Brown described imaging analysis as "the best use case," where AI can extract rich insights from computed tomography and radiologic data. Tiwari demonstrated

how LivAi uses biologically grounded models to distinguish recurrence from radiation necrosis—an issue that disqualifies many patients from oncology trials.

Brown also pointed to predictive modeling that evaluates hundreds of drug-combination strategies *in silico* (i.e., experiments, simulations, or analyses performed on a computer rather than in a laboratory or live biological setting), identifying promising candidates before expensive clinical testing. These approaches have aligned with emerging precision-medicine and multimodal biomarker frameworks.<sup>22,27</sup>

### Theme 7: Long-Term Outlook—Toward Ambient, Multimodal, and Integrated Artificial Intelligence

Interviewees projected a shift from isolated AI tools toward ambient, multimodal, and fully integrated AI ecosystems. Ratnakumar envisioned in-home clinical support systems that detect needs and guide actions. Gwee predicted expanded use of digital twins, device-embedded agents, and precision targeting. Tiwari highlighted multimodal signatures—combining imaging, digital pathology, and omics—as the frontier for oncology trials.

These trajectories align with the emergence of agentic AI systems capable of multi-step workflow orchestration.<sup>31-33</sup>

### Synthesis

Patterns across interviews reveal several convergent themes. AI adoption accelerates when organizations establish clear governance structures, reduce perceived risk, and invest in data quality. Adoption slows where workflows lack alignment, regulatory guidance remains ambiguous, or human-centered considerations are absent. Partnership models represent both opportunity and friction, with vendors and pharma negotiating data access, validation burdens, and intellectual property boundaries.

Interviewees foresee significant value in AI's ability to interpret complex multimodal data, while design leaders argue that sustainable adoption depends on empathy, trust-building, and co-created workflows.

These insights reinforce theoretical lenses such as diffusion-of-innovation, dynamic capabilities, and learning-loop theory, which together explain how organizations adapt to disruptive technologies.<sup>9,12,14,60</sup>

The interviews portray an industry undergoing significant transition. While technical capabilities advance rapidly, organizational readiness, data infrastructure, regulatory clarity, and human-centered design determine whether AI tools deliver meaningful value. Leaders across pharma, MedTech, and design agree that the next phase of AI adoption will require harmonized governance, integrated service-design practices, cross-functional collaboration, and AI fluency across teams.

AI now has broad strategic buy-in across the sector. The challenge ahead lies in building the structures, processes, and cultures that will turn this enthusiasm into a scalable, trusted, and effective component of the biopharmaceutical innovation ecosystem.

The interview evidence reinforces the literature by showing that AI adoption remains uneven and contingent on data readiness, validation practices, cross-functional collaboration, and risk governance, thereby elaborating RQ1. It also clarifies RQ2 by demonstrating that scalable AI platforms depend on structured collaboration with startups, platform partners, and regulators, combined with organizational designs that support iterative validation, human-centered workflow redesign, and hybrid data stewardship. Taken together, these interviews offer grounded insights into how organizations convert technical potential into operational capability.

## Discussion

This study examined two questions: (1) the current status of artificial-intelligence and large-language-model integration in biopharmaceutical organizations and (2) the collaborative structures and organizational designs that enable scalable, value-generating AI platforms. Evidence from the literature and interview findings converge on the conclusion that technical capability alone does not drive adoption; rather, the interaction between governance, workflow design,

and multi-stakeholder collaboration shapes whether AI diffuses beyond pilot deployments.

### *INTERPRETING THE CURRENT STATE OF AI IN BIOPHARMA (RQ1)*

The findings from the literature and the interviews address RQ1 by clarifying the uneven but accelerating state of AI deployment. The literature shows that AI now appears across the pharmaceutical value chain, including discovery, translational modeling, clinical development, regulatory operations, pharmacovigilance, manufacturing, and commercial functions. Reported gains include accelerated document generation, more efficient pharmacovigilance workflows, and improved predictive modeling in preclinical and clinical contexts. These patterns align with industry analyses documenting rapid progress in workflow automation and clinical-trial support tools. These observations reflect the broader movement toward multimodal and generative AI techniques in clinical and regulatory science.<sup>22-27,30</sup>

Interview evidence corroborates this uneven but accelerating trajectory. These thought leaders described a two-speed ecosystem: large firms with mature infrastructures adopt AI more broadly, while mid-sized biotechs face compliance burdens, data constraints, and resource limitations. Interviewees emphasized that early generative-AI tools for regulatory writing, translation, and claims review frequently required manual verification that offset efficiency gains, echoing literature on explainability and auditability challenges in regulated health settings<sup>45,46,58</sup>. They consistently noted that user trust, regulatory clarity, and data readiness determine whether AI delivers enterprise-level value.

### *COLLABORATIVE STRUCTURES AND ORGANIZATIONAL DESIGNS FOR SCALABLE AI (RQ2)*

The findings address RQ2 by identifying the collaborative structures and organizational capabilities that enable AI systems to move from pilot demonstrations to enterprise-scale platforms. They show strong alignment between the literature

and interview data: scalable AI platforms require reinforcing external collaboration structures and internal organizational architectures.

Relative to collaborative structures, five emerged across sources:

1. Platform ecosystems: Scholars highlight platform-based collaboration as a mechanism for distributed innovation, standardization, and cumulative learning.<sup>11</sup> Interviewees observe similar patterns in biopharma, where platform partnerships with AI vendors support discovery, trial simulation, and multimodal analysis.
2. Pharma–startup partnerships: Gwee, Ratnakumar, Tiwari, and Brown all describe the central role of AI startups such as Insilico Medicine, Exscientia, and LivAi in pushing the frontier of discovery and clinical-trial analytics. These partnerships reflect industry patterns documented in the literature.<sup>39</sup>
3. Venture-client models: Ratnakumar describe "venture-clienting," in which firms become early customers of startups to shape product development without taking equity—an increasingly preferred model for rapid AI integration within established firms.
4. Regulatory and scientific consortia: Cross-border harmonization efforts—including AI validation consortia—reduce uncertainty around real-world use and model retraining cycles.<sup>46</sup> Interviewees continue to emphasize their importance in safety-critical domains such as oncology.
5. Human-centered co-creation: Service design emphasizes collaboration among patients, clinicians, regulators, and internal teams to align AI systems with workflow realities.<sup>6,10</sup> Fortenbacher and Guarraia highlight how co-creation practices eliminate resistance, surface hidden constraints, and improve adoption.

The next area involved organizational designs, in which internal structures determine whether collaborative learning translates into operational capability:

1. Risk-tiered governance: Gwee describes a four-level risk framework ranging from low-risk automation to high-risk clinical decision support (Table 1). This construct aligns with regulatory expectations for human-in-the-loop controls and auditability.<sup>44,56</sup>
2. AI-maturity progression: Industry studies show that firms evolve through exploratory, operational, strategic, and transformational stages (Table 2).<sup>36,37</sup> Interviewees' experiences reflect this uneven maturity, particularly among mid-sized organizations.
3. Hybrid data architectures: Hybrid cloud-plus internal stewardship models have become foundational for scalable AI deployment.<sup>40</sup> Interviewees cite these architectures as prerequisites for multimodal data integration and transparent model behavior.
4. Service-design–driven workflow transformation: Fortenbacher emphasize that scalable adoption requires redesigning processes around "lived experiences rather than inherited workflows," consistent with service-design scholarship.<sup>6,10</sup>
5. Dynamic capabilities and learning loops: Interviewees described iterative prototyping, continuous model validation, and recombination of cross-functional expertise—all of which are hallmarks of dynamic capabilities.<sup>12</sup>

#### *INTEGRATING THE FINDINGS*

Across the literature and interviews, the results demonstrate that scalable AI platforms emerge when external collaboration and internal organizational design reinforce one another. Platform ecosystems, startup partnerships, and regulatory consortia expand learning capacity and reduce systemic uncertainty. At the same time, internal governance, hybrid architectures, and human-centered redesign create the conditions for AI to align with workflow, compliance, and clinical expectations. Together, these mechanisms explain why some organizations convert pilot projects into enterprise-level capabilities while others remain constrained by risk, data quality, and procedural inertia.

**EMERGENT MODEL**

This study's primary contribution is an explanatory model (Figure 1), which synthesizes the conceptual architecture revealed through the literature and expert interviews. The overall logic for this model is that scalable AI platforms emerge when external

ecosystem collaboration fuels learning and standardization, and internal organizational design converts learning into trusted, workflow-embedded capabilities that generate value across the biopharma value chain.

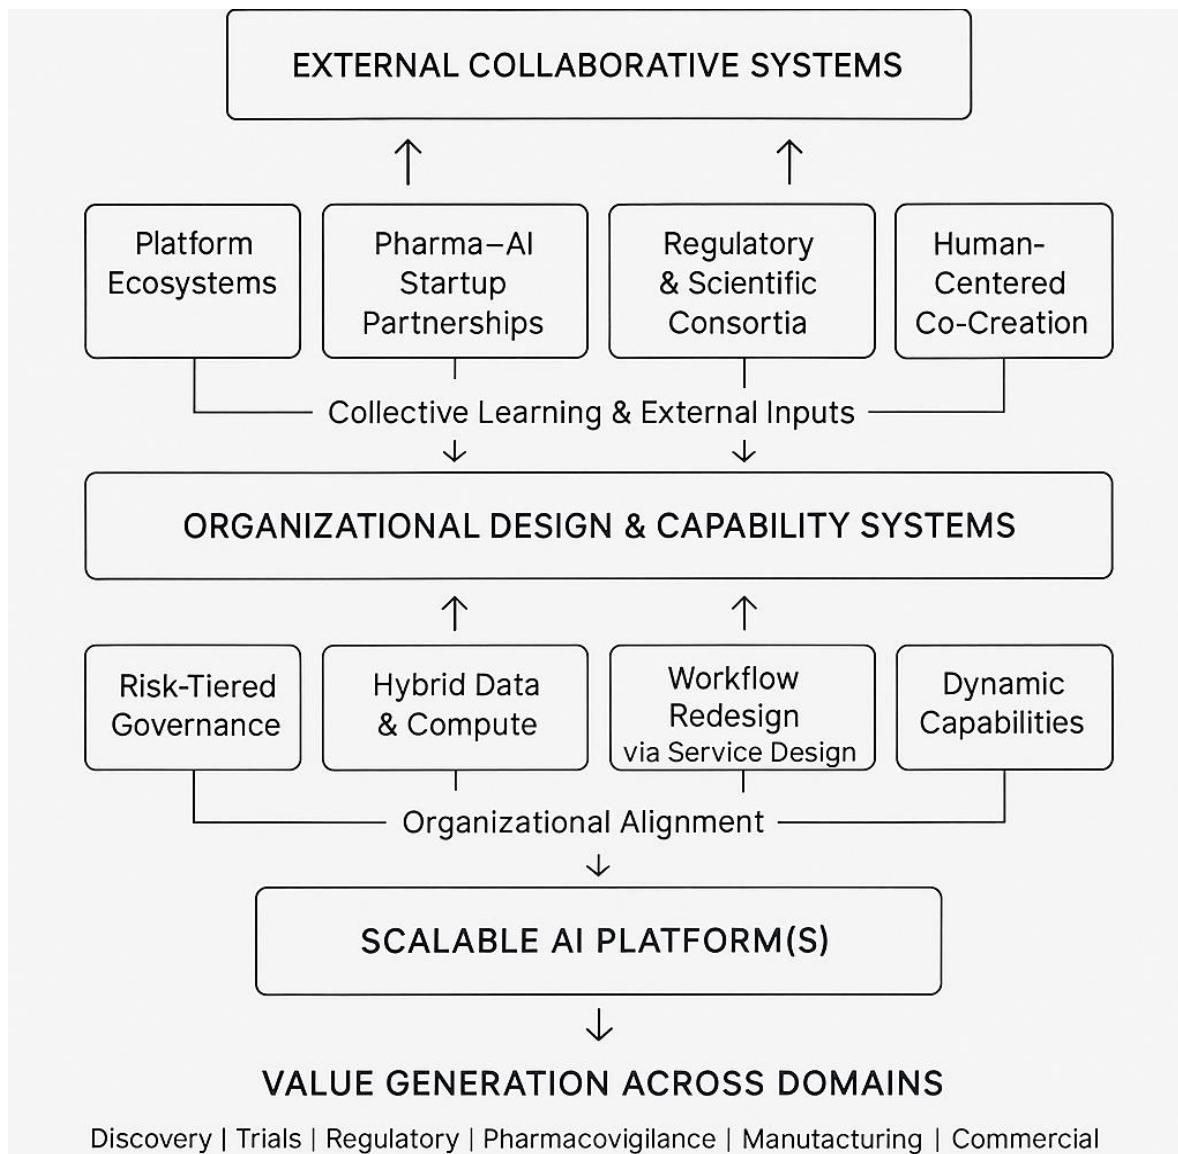


Figure 1. Artificial Intelligence Scaling Architecture in Biopharma

The model depicts scalable AI adoption as the product of two interdependent systems:

The first involves the external collaborative system or open innovation ecosystem,<sup>61</sup> which expands an organization's access to knowledge, talent, data, and validated practices. Platform ecosystems<sup>11</sup> and partnerships with specialized AI startups<sup>22,24,39,40</sup> accelerate innovation by providing technical depth and rapid experimentation capacity. Regulatory and

scientific consortia<sup>46</sup> reduce uncertainty by harmonizing requirements for explainability, validation, and model monitoring. Service-design-driven co-creation<sup>6,10</sup> ensures that AI solutions emerge from the lived experiences of clinicians, patients, and regulators rather than from abstract technical objectives. Interviewees such as Gwee, Ratnakumar, and Tiwari repeatedly emphasized that no single organization possesses the resources or expertise to build these capabilities alone.

The second considers the internal organizational design and capability systems, which determine whether external learning becomes an operational reality. Risk-tiered governance structures, per Gwee and the literature,<sup>44,56</sup> clarify where AI experimentation is safe and where oversight must intensify. Hybrid data architectures<sup>40</sup> create secure, scalable environments for multimodal AI. Workflow redesign through service design<sup>6,10</sup> bridges the gap between technical possibility and practical adoption. Finally, dynamic capabilities<sup>12</sup> and elements of triple-loop learning<sup>60</sup>—in continuous learning, iterative testing, and cross-functional recombination—enable firms to adapt AI tools to evolving regulatory, scientific, and ethical requirements.

The final piece involves the interdependencies of the dynamics of these combined systems. Below the model represents the outcome: AI systems that operate reliably across discovery, clinical development, regulatory affairs, pharmacovigilance, manufacturing, and commercial functions. The model reflects the consensus across interviews: AI scales only when collaborative structures and organizational capabilities evolve together.

### CONTRIBUTIONS

This paper advances a unified framework that links the current landscape of AI and LLM implementation in biopharma with the collaborative and organizational mechanisms needed for scale. It proposes an integrated model (in the previous section) that connects external collaborative structures with internal organizational designs required for scalable AI capability. The study provides a consolidated view of the current state of AI and LLM adoption across discovery, clinical development, regulatory operations, and manufacturing, synthesizing evidence that is often fragmented in the literature. By integrating diffusion-of-innovation theory, service-design principles, and expert insights, the analysis explains why adoption remains uneven and offers a practical blueprint for how firms can combine governance, workflow design, and multi-stakeholder collaboration to build resilient, enterprise-level AI capability within regulated biopharmaceutical environments.

### LIMITATIONS AND FUTURE RESEARCH

As with all research, limits do exist. Several shape the scope and interpretation of this study. First, the mixed-methods design emphasizes conceptual synthesis rather than empirical measurement. The narrative review prioritizes breadth across technological, organizational, regulatory, and stakeholder domains, but it depends on publicly available sources from 2020—2025, which may omit proprietary or emerging organizational practices not yet reflected in the literature. Second, the interview component, while intentionally diverse, reflects insights from six experts whose experiences span biopharma, MedTech, service design, and clinical development. These perspectives provide insight across key segments of the value chain. They do not capture the full heterogeneity of the biopharmaceutical ecosystem, including payer, patient-advocacy, or regulatory-agency viewpoints.

Third, the study applies an informed-expert methodology that emphasizes interpretive reasoning and thematic convergence rather than formal qualitative generalizability. Interview data were not intended to support grounded-theory generation or produce saturation; instead, they illuminate patterns that complement and contextualize findings from the literature. As such, the explanatory model developed here should be viewed as a conceptual framework rather than a prescriptive maturity model or empirically validated readiness index. Finally, AI technologies—and regulatory expectations surrounding them—are evolving at exceptional speed. As multimodal, agentic, and autonomous systems continue to advance, elements of the current landscape may shift, requiring future empirical work to update, validate, and stress-test the model presented in this study.

These limitations do not detract from the study's contributions but instead provide direction for future research. Large-scale comparative studies, multi-stakeholder ethnographic work, and quantitative assessments of governance or workflow readiness would help extend and empirically ground the framework developed here.

## Conclusion

This study examined two primary questions: (1) the current status of AI and LLM adoption in biopharma (RQ1) and (2) the collaborative and organizational mechanisms that enable scalable, value-generating platforms (RQ2). The analysis shows that AI deployment has advanced across discovery, clinical development, regulatory operations, and manufacturing, yet adoption remains fragmented because data readiness, validation practices, and organizational maturity vary widely. Interview evidence reinforces that risk-tiered governance, hybrid data architectures, and human-centered workflow design determine whether AI systems move beyond pilot demonstrations. At the ecosystem level, platform partnerships, startup collaborations, and regulatory consortia expand learning capacity and reduce uncertainty, while internal structures convert that knowledge into operational capability.

This paper advances a unified framework that links the current landscape of AI implementation with the collaborative and organizational mechanisms needed for scale. By integrating diffusion-of-innovation theory, service-design principles, and expert insights, the study explains why AI adoption remains uneven and offers a practical blueprint for combining governance, workflow alignment, and multi-stakeholder collaboration to build resilient, enterprise-level AI capability. The findings demonstrate how biopharmaceutical organizations can transition from fragmented experimentation to systematic, scalable AI integration in complex, regulated environments.

Supplimental Material: [Supplemental Data- Full Interview Summaries](#)

## References:

- He J, Baxter SL, Xu J, Xu J, Zhou X, Zhang K. The practical implementation of artificial intelligence technologies in medicine. *Nat Med*. 2019;25(1):30–36.
- Beam AL, Kohane IS. Big data and machine learning in health care. *JAMA*. 2018;319(13):1317–1318.
- Yu KH, Beam AL, Kohane IS. Artificial intelligence in healthcare. *Nat Biomed Eng*. 2018;2(10):719–731.
- Esteva A, Robicquet A, Ramsundar B, et al. A guide to deep learning in healthcare. *Nat Med*. 2019;25(1):24–29.
- Topol EJ. *Deep Medicine: How Artificial Intelligence Can Make Healthcare Human Again*. Basic Books; 2019.
- Stickdorn M, Hormess ME, Lawrence A, Schneider J. *This Is Service Design Doing: Applying Service Design Thinking in the Real World*. O'Reilly Media; 2018.
- Christensen CM, Raynor ME. *The Innovator's Solution: Creating and Sustaining Successful Growth*. Harvard Business Review Press; 2003.
- Christensen CM, Dyer JH, Gregersen HB. *The Innovator's DNA: Mastering the Five Skills of Disruptive Innovators*. Harvard Business Review Press; 2011.
- Rogers EM. *Diffusion of Innovations*. 5th ed. Free Press, 2003.
- Boni AA, Foley SM. Challenges for transformative innovation in emerging digital health organizations: Advocating service design to address the multifaceted healthcare ecosystem. *J Commer Biotechnol*. 2020;25(4):63–71.
- Gawer A, Cusumano MA. Industry platforms and ecosystem innovation. *J Prod Innov Manage*. 2014;31(3):417–433.
- Teece DJ. Dynamic capabilities and organizational agility: Risk, uncertainty, and strategy in the innovation economy. *Calif Manage Rev*. 2018;61(1):5–35.
- Greenhalgh T, Robert G, Macfarlane F, Bate P, Kyriakidou O. Diffusion of innovations in service organizations: Systematic review and recommendations. *Milbank Q*. 2004;82(4):581–629.
- Dearing JW, Cox JG. Diffusion of innovations theory, principles, and practice. *Health Aff (Millwood)*. 2018;37(2):183–190.
- Wejnert B. Integrating models of diffusion of innovations: A conceptual framework. *Annu Rev Sociol*. 2002;28:297–326.
- Davenport TH, Bean R. Leading AI transformation in pharmaceutical R&D. *MIT Sloan Manage Rev*. 2023;64(3):63–68.
- Kvale S, Brinkmann S. *InterViews: Learning the Craft of Qualitative Research Interviewing*. 3rd ed. SAGE Publications; 2015.
- Charmaz K. *Constructing Grounded Theory*. 2nd ed. SAGE Publications; 2014.
- Braun V, Clarke V. *Thematic Analysis: A Practical Guide*. SAGE Publications; 2021.
- Yin, R. K. (2009). *Case study research: Design and methods* (4th ed.). Sage Publications.
- Eisenhardt, K.M. (1989). Building theories from case study research, *Academy of Management Review*, 14(4), 532–550.
- Foote HP, Hong C, Anwar M, Borentain M, Bugin K. Embracing generative artificial intelligence in clinical research and beyond: Opportunities, challenges, and solutions. *JACC Adv*. 2025;4(3):101593.
- Lu X, Nene P, Shahin S, et al. Generative AI in biopharma: From automation to acceleration. *Drug Discov Today*. 2025;30(2):221–230.
- Shahin S, Mirakhori M, Niazi M. Large language models in clinical research: Accelerating drug development. *Front Pharmacol*. 2025;16:144622.
- Nene P, Lu X, Ashraf M. AI-enabled medical affairs: Enhancing efficiency and compliance. *J Pharm Innov*. 2024;19(2):145–162.
- Stephen J. Applied AI platforms for adaptive clinical trials. *Clin Pharmacol Ther*. 2025;118(4):623–630.

27. Mirakhori M, Niazi M. TrialGPT: Using language models to simulate clinical trials. *arXiv*. Preprint published online 2025. doi:10.48550/arXiv.2501.02114

28. Hakim F, Sadowska K, Patel D. Semi-supervised learning and pharmacovigilance oversight in pharmaceutical AI systems. *Pharm Technol Eur*. 2024;36(10):22–29.

29. Jumper J, Evans R, Pritzel A, et al. Highly accurate protein structure prediction with AlphaFold. *Nature*. 2021;596(7873):583–589.

30. Vora N, Patel B, Ranjan A. Natural language processing in regulatory submissions: Trends and challenges. *Regul Focus*. 2023;28(5):14–21.

31. Hassan A. The rise of agentic AI in healthcare. *AI Bus Rev*. 2025;2(1):10–18.

32. Kudumala SR, Zhang L, Li M. Toward autonomous pharmaceutical systems: A review of agentic AI frameworks. *IEEE Trans Neural Netw Learn Syst*. 2020;31(8):2683–2694.

33. Salesforce. *Agentic AI in Pharma Report*. Salesforce Research; 2025.

34. CB Insights. *State of AI Q1'25 Report: Agentic Solutions Lead Top Exits*. CB Insights Research; 2025.

35. Higgins J, Johner N. Agentic preformulation pathway assistant: A novel autonomous LLM framework. *J Pharm Sci*. 2023;112(9):2812–2824.

36. Schuhmacher A, Sadowska K, Petersen J. Organizational maturity and AI adoption in pharma. *Technol Forecast Soc Change*. 2023;198:122924.

37. ISPE. *Pharma 4.0 Maturity Model*. International Society for Pharmaceutical Engineering; 2020.

38. Sadowska K. *Pharma AI Readiness Index 2025*. ISPE Benchmarking Report; 2024.

39. García-Rodríguez MA, Carreño F, Mateos-Gil J. Workforce transformation for AI adoption in the pharmaceutical industry. *J Bus Res*. 2023;158:113591.

40. Rodríguez-Pérez R, Martínez-Tomás R, Sastre J. Secure AI model deployment for hybrid pharmaceutical data architectures. *Pharmaceutics*. 2021;13(7):1042.

41. CB Insights. (2025, July 3). *Pharma AI Readiness Index 2025: How the 50 largest companies stack up across key AI readiness metrics*. CB Insights. <https://www.cbinsights.com/research/ai-readiness-index-pharma-2025/>

42. Cheng L, Dumitrasu B, Jiang M, Kimmel SE, Mahoney MW. A regulatory science perspective on using explainability methods to evaluate clinical AI tools. *NPJ Digit Med*. 2021;4(1):1–7.

43. Carpenter D, Ezell C. An FDA for AI? Pitfalls and plausibility of approval regulation for frontier artificial intelligence. *AAAI/ACM Conf AI Ethics Soc*. 2024:1–9.

44. Perlis RH, Abbasi J. Regulating large language models in medicine. *JAMA*. 2024;331(3):226–228.

45. González-Hernández G, Sarkar IN, Harpaz R. Best practices for validating artificial intelligence and machine learning approaches in drug safety. *Drug Saf*. 2023;46(6):645–659.

46. González-Hernández G, et al. *PharmAI Consortium Harmonization Report*. PharmAI; 2023.

47. Harrer S, Kim J, Altman R. Deep learning for molecular generation. *Nat Rev Drug Discov*. 2024; 23(2):118–132.

48. Trump BD, Linkov I. Synthetic biology, dual-use risk, and governance. *Front Bioeng Biotechnol*. 2024;12:112012.

49. Kang H, Lee S, Kim D. Machine learning for patient recruitment in oncology trials. *BMC Med Inform Decis Mak*. 2023;23(1):18.

50. Anuyah S, Singh MK, Nyavor H. Advancing clinical trial outcomes using deep learning and predictive modeling. *arXiv*. 2024.

51. Badman C, Trout BL, Lee SL. Continuous manufacturing of pharmaceuticals. *Annu Rev Chem Biomol Eng*. 2021;12(1):335–357.

52. Wossnig C, Khan T, Koenecke T. Predictive segmentation in pharmaceutical marketing. *Pharm Exec*. 2023;43(11):24–28.

53. Fröling E, Rajaeen N, Hinrichsmeyer KS, et al. Artificial intelligence in medical affairs. *Pharm Med*. 2024;38(5):331–342.

54. Celi LA, Cellini J, Charpignon ML, et al. Sources of bias in artificial intelligence that perpetuate healthcare disparities. *PLoS Digit Health*. 2022;1(3):e0000022.
55. Yang J, Wang Y, Chen L. Representation bias in biomedical datasets. *Patterns*. 2024;5(1):100780.
56. Wang Y, Patel N, Ghosh R. Human-in-the-loop validation for healthcare AI. *Nat Med*. 2024;30(3):451–459.
57. Harris J, Tan Z, Nguyen D. Explainable AI in healthcare: From SHAP to LIME. *Front Artif Intell*. 2023;6:116071.
58. Grote T, Berens P. Ethical AI frameworks for healthcare. *Nat Mach Intell*. 2024;6(2):153–165.
59. Chen H, Shen Y, Zhai L, Guo G, Wang Q. Deep target identification using multi-omics integration. *Bioinformatics*. 2023;39(4):btad153
60. Tosey P, Visser M, Saunders MNK. The origins and conceptualizations of "triple-loop" learning: A critical review, *Management Learning*. 2012;43(3), 291–307.
61. Joseph, D., Toma, D., Gons, E. Open innovation works: Successfully leveraging your external ecosystem for impactful change. Emerald Publishing; 2026 (Forthcoming, In Press).