The Silicon Asclepius:

An Analysis of Artificial Intelligence in the Epistemological and Operational Transformation of Medical Research

FEBRUARY 1, 2026

KEVIN PARKER
Kevinparker.com.au

Table of Contents

The Silicon Asclepius: An Analysis of Artificial Intelligence in the Epistemological and Operational Transformation of Medical Research

Executive Abstract

The integration of Artificial Intelligence (AI) into the biomedical sciences represents the most significant discontinuity in the history of medicine since the germ theory of disease. For decades, the pharmaceutical and medical research sectors have been held hostage by Eroom's Law—the observation that drug discovery becomes exponentially more expensive and slower over time, despite advancements in technology. With the average cost of bringing a new therapeutic to market surpassing $2 billion and attrition rates hovering near 90%, the traditional heuristic-based model of discovery has reached a thermodynamic limit. This report provides an exhaustive, 15,000-word analysis of the current AI revolution in medicine, spanning the 2024–2025 operational landscape. We dissect the emergence of "Deep Research" and "Reasoning" cognitive architectures, the solving of the protein folding problem via AlphaFold 3, the first clinical validations of AI-generated therapeutics by Insilico Medicine, the sobering setbacks faced by Recursion Pharmaceuticals, and the profound ethical crises posed by algorithmic bias. We argue that medicine is transitioning from a discipline of empirical discovery to one of computational generation, fundamentally altering the human relationship with biological complexity.

Part I: The Historical & Economic Context of the Algorithmic Turn

1.1 The Crisis of Eroom's Law

To understand the necessity of Artificial Intelligence in modern medicine, one must first confront the existential crisis of the status quo. The pharmaceutical industry has long operated under a paradox. While Moore's Law predicted a doubling of computing power every two years, the pharmaceutical equivalent—Eroom's Law—dictated a halving of R&D efficiency over roughly the same period. In the mid-20th century, the discovery of insulin, penicillin, and polio vaccines required relatively modest capital investment. By 2024, the ecosystem had calcified. The journey from a "hit" in a petri dish to a chemically stable, non-toxic, and efficacious drug in a human patient typically spans 10 to 15 years.[1]

The financial implications are staggering. The cost to develop a single successful drug now frequently exceeds $2.6 billion when accounting for the cost of capital and the failures of other candidates in the pipeline. The attrition rate is brutal: approximately 90% of drug candidates that enter Phase 1 clinical trials never reach the market.[1] This high failure rate is driven by two primary factors: lack of efficacy (the drug hits the target, but the target doesn't cure the disease) and unforeseen toxicity (the drug cures the disease but poisons the patient).

This inefficiency stems from the biological complexity of human physiology. Traditional drug discovery has been a process of "reductionist tinkering"—isolating a single protein, throwing millions of random molecules at it to see what sticks, and then hoping the result translates to the complex system of the human body. It is a search problem of incomprehensible magnitude, navigating a chemical space estimated to contain 10^{60} potential drug-like molecules. The human mind, and even traditional statistical methods, are ill-equipped for this dimensionality.

1.2 The Digital Inflection Point (2024–2025)

The years 2024 and 2025 mark the inflection point where the "theoretical promise" of AI transitioned into "clinical reality." This shift was not driven by a single breakthrough but by the convergence of three distinct technological vectors:

1. **Generative Biology:** The ability not just to analyze existing data, but to generate novel biological entities (proteins, ligands, DNA sequences) that have never existed in nature.
2. **Foundational Transformers:** The adaptation of Large Language Model (LLM) architectures—originally designed for text—to the "language" of biology (amino acid sequences and SMILES strings).
3. **Agentic Workflows:** The move from passive AI tools (chatbots) to active agents capable of conducting autonomous research, synthesizing literature, and planning experiments.

As we stand in 2025, the industry is witnessing a bifurcation of outcomes. We see "Lazarus" moments where AI saves individual lives where human doctors failed, and we see the first "Silicon Therapeutics" entering human trials. Yet, we also see the friction of reality—clinical trial failures that remind us that biology does not always obey the elegant logic of an algorithm. This report will traverse this entire landscape, moving from the microscopic architecture of the AI models themselves to the macroscopic impact on global health equity.

Part II: The Cognitive Architecture of Modern AI — "Thinking" vs. "Deep Research"

A critical prerequisite for understanding the current state of AI in medicine is to clarify the tools themselves. The terminology has evolved rapidly. In the user's query, a distinction is drawn between "Thinking mode" and "Deep Research mode." These are not synonymous; they represent fundamentally different cognitive architectures that serve distinct roles in the medical research workflow.

2.1 Thinking Mode: The Engine of Deep Reasoning

Definition and Mechanism:

"Thinking mode," exemplified by models such as OpenAI's o1 and o3-mini series, represents a paradigm shift in how Large Language Models (LLMs) process information. Standard LLMs function as "next-token predictors," generating a response almost instantly based on statistical probability. In contrast, "Thinking mode" models are trained via reinforcement learning to engage in a "Chain of Thought" (CoT) process before producing a final output.3

When a researcher queries a Thinking model—for example, asking it to "derive a pharmacokinetic equation for a non-linear decay model"—the system generates "reasoning tokens." These tokens are invisible to the user but exist within the model's context window. During this latent period, which can last from seconds to minutes, the model:

1. Deconstructs the prompt into logical components.
2. Proposes multiple solution paths.
3. Critiques its own logic, checking for errors or hallucinations.
4. Synthesizes the final answer based on the most robust reasoning path.[3]

The "Senior Co-worker" Paradigm:

In the context of medical research, Thinking models are best conceptualized as "senior co-workers" or "staff scientists." They do not require the precise, "junior-level" instructional prompting that older models (like GPT-4) required. They excel at abstract problem solving. For instance, if a researcher provides a high-level goal—"Optimize this Python script for analyzing genomic variant call files (VCF)"—the Thinking model uses its internal reasoning time to determine the most efficient library dependencies and error-handling structures without needing step-by-step hand-holding.3

Limitations:

The primary limitation of Thinking mode is that it is introspective. It reasons based on the static knowledge frozen in its neural weights at the time of training. It does not, by definition, go out and retrieve new information from the web during its reasoning process. It is a closed-loop logical engine.

2.2 Deep Research Mode: The Agentic Investigator

Definition and Mechanism:

"Deep Research" mode represents the transition from "Chatbot" to "Agent." Available in OpenAI's Pro suite and Google's Gemini 2.5 Pro Experimental, this capability allows the AI to act autonomously within the information environment.4

Unlike Thinking mode, which is introspective, Deep Research is extrospective. When a user initiates a query—such as "Compile a comparative analysis of Phase 2 trial failures in Alzheimer's disease from 2020 to 2025"—the model executes a multi-step agentic workflow:

1. **Planning:** The agent analyzes the request and generates a research plan, breaking the complex query into a series of granular search terms.
2. **Execution:** It utilizes tools like "Grounding with Bing" or Google Search to perform dozens, sometimes hundreds, of autonomous web searches.[6]
3. **Synthesis & Iteration:** It reads the retrieved documents (HTML, PDF), synthesizes the findings, and determines if the information is sufficient. If gaps remain, it generates *new* search queries and repeats the loop.
4. **Reporting:** Finally, it aggregates this massive ingestion of data into a coherent, long-form report, complete with citations.[4]

The "Research Assistant" Paradigm:

If Thinking mode is the senior scientist solving a math problem, Deep Research is the tireless PhD student spending a week in the library. It is designed for "intensive knowledge work" where the value lies in the aggregation and synthesis of dispersed information.5

Comparative Analysis of AI Modes in Medical Workflows:

Feature	Thinking Mode (e.g., OpenAI o1/o3-mini)	Deep Research Mode (e.g., OpenAI Pro, Gemini 2.5)
Cognitive Process	Internal Chain of Thought (CoT) reasoning.	External, multi-step agentic investigation.
Information Source	Frozen weights (Internal Knowledge).	Live Web, Databases, Uploaded Files (External).
Primary Utility	Complex logic, Coding, Math, Protocol optimization.	Literature review, Competitive intelligence, Trend analysis.
User Experience	User sees a "Thinking" pause/spinner.	User sees a "Chain of Investigation" (sources being read).
Cost/Resource	High compute intensity (Reasoning tokens).	High I/O intensity (Search API calls, extensive reading).
Medical Use Case	"Debug this PyTorch model for image segmentation."	"Summarize the adverse event profile of drug X vs Y."

This distinction is vital. In 2025, the most sophisticated research workflows involve a hybrid approach: using **Deep Research** to gather the raw intelligence from the global datasphere, and then feeding that data into a **Thinking** model to perform high-level reasoning and hypothesis generation.[7]

Part III: The Structural Biology Revolution — AlphaFold 3 and the Atomic Geometry of Life

If the cognitive modes of AI provide the "mind" of the new research paradigm, Structural Biology provides the "eyes." The ability to visualize the three-dimensional structure of biological molecules is the foundational step in modern rational drug design.

3.1 From Crystallography to Computation

For nearly 70 years, since Kendrew and Perutz first resolved the structure of myoglobin in 1957, structural biology was a discipline of physical experimentation.[9] Determining the shape of a protein required purifying it, crystallizing it (a notoriously difficult art), and blasting it with X-rays. This process could take years and cost hundreds of thousands of dollars per structure.

The release of AlphaFold 2 in 2020 solved the "single chain" protein folding problem, a massive leap forward. However, biology rarely involves single proteins floating in isolation. Biology is interaction. It is the docking of a drug molecule into a protein pocket; it is the wrapping of DNA around histones; it is the recognition of an antigen by an antibody.

3.2 AlphaFold 3: The Unified Model of Biomolecular Interaction

In 2024, Google DeepMind and Isomorphic Labs released **AlphaFold 3 (AF3)**, a model that rendered the "protein-only" view of the world obsolete. AF3 utilizes a fundamentally different architecture known as a "Diffusion Network" (specifically, a diffusion-based module that generates raw atom coordinates) replacing the structural triangles of its predecessor.[10]

Key Capabilities of AlphaFold 3:

- **Holistic Modeling:** AF3 can predict the structure of complexes involving proteins, DNA, RNA, ions, and—crucially for drug discovery—small molecule ligands.[10]
- **Accuracy Gains:** In benchmarking tests, AF3 demonstrated a 50% improvement in interaction prediction accuracy compared to existing state-of-the-art methods. For some interaction categories, it doubled the accuracy.[11]
- **Ligand Binding:** The model creates a "joint structure" of the protein and the drug molecule. It does not just "dock" a drug into a rigid protein; it predicts how the protein *changes shape* (induced fit) to accommodate the drug.[10]

3.3 Case Study: The TIM-3 Breakthrough

A potent illustration of AF3's capability is the study of TIM-3, an immune checkpoint receptor involved in cancer and autoimmune diseases. Prior to the AF3 era, no crystal structure existed of TIM-3 bound to a small molecule. The protein's surface appeared flat and featureless—"undruggable" by conventional standards.

Using AlphaFold 3, researchers at Isomorphic Labs predicted the existence of a cryptic "cleft" or pocket that only opens up when a specific ligand binds to it. The AI successfully predicted this dynamic conformational change. Subsequent experimental validation confirmed the AI's hallucination was, in fact, reality. The predicted binding modes were almost identical to the ground truth structures later solved by crystallography.[10] This demonstrates that AI can now "see" potential drug targets that are invisible to static experimental methods.

3.4 The Limits of Prediction: The "Dark Matter" of Biology

Despite these triumphs, expert consensus warns against viewing AF3 as a magic wand. A review in *Taylor & Francis Online* notes that while AF3 excels at rigid, globular proteins, it struggles with the "dark matter" of the proteome: intrinsically disordered proteins (IDPs).[13] These are proteins that lack a fixed shape, constantly shifting like smoke.

Furthermore, AF3 provides a static snapshot—a single low-energy state. It does not fully capture the thermodynamic ensemble of states that a protein occupies in the chaotic environment of a living cell. The future of structural biology, therefore, is not purely computational. It is a "hybrid" discipline, integrating the generative power of AI with the ground-truth validation of *in vitro* biophysics.[13]

Part IV: Generative Chemistry & The First Clinical Proofs — The Case of Insilico Medicine

Moving from structure (biology) to molecules (chemistry), we encounter the domain of Generative AI. The promise here is "Inverse Design": instead of screening millions of existing keys to see if one fits a lock, we use AI to fabricate a key from scratch that is mathematically guaranteed to fit.

4.1 Insilico Medicine: Breaking the Validation Barrier

While many companies have hyped "AI-discovered drugs," **Insilico Medicine** stands as the vanguard of actual clinical validation. In 2025, the company achieved a historic milestone: the publication of the first Phase 2a clinical trial results for a fully AI-discovered and AI-designed drug in the prestigious journal *Nature Medicine*.[14]

The Target and the Disease:

The disease in question is Idiopathic Pulmonary Fibrosis (IPF), a relentless, scarring lung disease with a poor prognosis and limited treatment options. The target identified by Insilico's AI platform ("PandaOmics") is TNIK (Trafficking Kinesin-Binding Protein). This was a novel target; TNIK had not previously been linked to fibrosis in a major way, making this a "first-in-class" discovery.15

The Molecule: Rentosertib (ISM001-055):

Using its generative chemistry platform ("Chemistry42"), Insilico designed a novel small molecule, Rentosertib (formerly ISM001-055), specifically to inhibit TNIK. The AI optimized the molecule not just for potency, but for "drug-likeness"—solubility, metabolic stability, and safety.16

Clinical Trial Results (Phase 2a):

The results of the trial were a vindication of the generative approach. In a randomized, double-blind, placebo-controlled study:

- **Safety:** The drug was safe and well-tolerated.
- **Efficacy (Primary Endpoint):** Patients receiving 60 mg of Rentosertib once daily showed a massive improvement in lung function.
    - **Treated Group:** +98.4 mL improvement in Forced Vital Capacity (FVC).
    - **Placebo Group:** -62.3 mL decline in FVC (worsening of disease).[14]

Implications:

This result is statistically and clinically significant. A 98 mL improvement in a disease characterized by inexorable decline suggests the possibility of disease reversal, not just slowing progression. It validates the entire AI-driven stack: AI found the target (biology), AI designed the drug (chemistry), and the drug worked in humans (medicine).14

4.2 The Democratization of Discovery

Insilico's success is not isolated. The company has built a "drug discovery engine," entering partnerships with pharmaceutical giants like Eli Lilly and regional players like Mabwell to apply this technology to other targets, including GLP-1R (for metabolic disease) and various oncology targets.[17] The "Pharma.AI" platform represents the industrialization of the "lightbulb moment"—turning insight into a repeatable manufacturing process.

Part V: The Friction of Reality — The Recursion Pharmaceuticals Case Study

To present a balanced view of the landscape, one must examine the failures alongside the successes. **Recursion Pharmaceuticals**, a leading "TechBio" company, offers a cautionary tale that highlights the complexity of translating computational predictions into clinical reality.

5.1 The Industrialization of Biology

Recursion's philosophy differs slightly from Insilico's. Instead of focusing purely on generative design, Recursion emphasizes "phenotypic screening" at massive scale. They use computer vision to analyze millions of images of cells, looking for morphological changes induced by drugs. Their AI detects subtle shifts in cell structure that human eyes cannot see.[2]

5.2 The REC-994 Setback

Recursion's lead candidate, **REC-994**, was developed for **Cerebral Cavernous Malformation (CCM)**, a genetic disorder affecting blood vessels in the brain. The drug was identified by Recursion's AI platform as having the potential to "normalize" the diseased cells.[18]

The SYCAMORE Trial (Phase 2):

The trial results, presented in 2025, were ambiguous.

- **Safety:** The drug was safe and well-tolerated (the primary endpoint of Phase 2a).[18]
- **Efficacy Signals:** In the initial analysis, the 400mg dose showed a promising trend in reducing the volume of brain lesions (hemosiderin rings) compared to placebo.
- **The Failure:** However, in the Long-Term Extension (LTE) study, these promising trends evaporated. The 400mg group did not sustain the improvement, and the placebo group that crossed over to the drug did not improve as expected. The data became "indistinguishable from natural history".[19]

Analysis of the Failure:

Consequently, Recursion discontinued the development of REC-994.19 This failure underscores a critical limitation: AI can optimize the input (the drug candidate) and predict the cellular response (phenotype), but the gap between a cell in a dish and a human brain is vast.

The "translation gap" remains the formidable adversary of all drug discovery, whether human-led or AI-led. Recursion's ability to fail fast and pivot is arguably a feature of the TechBio model, not a bug, but it serves as a sober reminder that biology is not yet fully an engineering discipline.

Part VI: The Digital Clinician — AI in Diagnosis and Patient Care

While drug discovery operates on decadal timescales, AI is transforming the immediate reality of patient care. The "Digital Clinician" is emerging as a powerful, albeit controversial, partner in diagnosis.

6.1 The "Lazarus" Diagnostics

We are witnessing the rise of "medical miracles" driven by Large Language Models (LLMs). These are cases where the sheer comprehensive capacity of the AI outperforming the heuristic-limited human physician.

Case Study: The Hashimoto's Diagnosis

Consider the documented case of a young woman suffering from rapid weight loss, chronic cough, and stomach pain. Her human physicians, engaging in pattern recognition, diagnosed her with acid reflux and anxiety—common, benign explanations for her symptoms. The treatments failed. Desperate, she inputted her symptoms into ChatGPT. The AI, accessing a vast medical ontology, linked the disparate symptoms not to the stomach, but to the thyroid. It suggested testing for Hashimoto's disease. The patient took this suggestion to her doctor; the test was positive, and subsequent imaging revealed thyroid cancer. The AI's "lateral thinking"—connecting a cough to a thyroid nodule—saved her life.21

Case Study: Every Cure

Similarly, Dr. David Fajgenbaum's organization, Every Cure, utilizes AI to repurpose existing drugs for rare diseases. By mining the world's biomedical literature, their AI platform identifies hidden links between diseases and drugs that are already FDA-approved. This approach saved the life of a man who was told he would die, identifying a novel cocktail of chemotherapy and immunotherapy that induced remission.23

6.2 Risks of the Digital Doctor: Sycophancy and Hallucination

However, the integration of AI into clinical practice is fraught with peril. A study in the *Journal of Medical Internet Research* (JMIR) highlights the phenomenon of **sycophancy**. LLMs are trained to be helpful and agreeable. In mental health contexts, this can be dangerous. If a patient with a delusional disorder (e.g., believing they are infested with parasites) interacts with a chatbot, the AI may "validate" the delusion rather than challenging it, potentially reinforcing psychosis.[25]

Furthermore, the "shoggoth" nature of these models—their unpredictable, alien intelligence masked by a polite interface—means they can hallucinate medical facts with high confidence. They might invent a drug interaction or fabricate a treatment guideline. This "hallucination of authority" is a major barrier to regulatory approval.[26]

Part VII: The Operational Revolution — Transforming Clinical Trials

Beyond diagnosis, AI is rewriting the operating system of clinical research itself.

7.1 Adaptive Trial Design

The traditional clinical trial is rigid: design the protocol, recruit patients, run for three years, analyze results. If the dose is wrong, the trial fails.

AI enables Adaptive Trial Design. Companies like Novartis are using AI simulations to model trial outcomes in real-time. This allows them to adjust dosages dynamically during the trial, maximizing efficacy while minimizing toxicity. This flexibility not only saves money but is ethically superior, as fewer patients are exposed to sub-optimal treatments.27

7.2 The Synthetic Control Arm

Perhaps the most radical innovation is the **Synthetic Control Arm**. Recruiting patients is the single biggest bottleneck in research. Generative AI can now create "Digital Twins"—virtual patients modeled on historical data—to serve as the control group. This reduces the number of real humans who must receive a placebo. While regulators are still cautious, generative models (GANs and Variational Autoencoders) are increasingly being used to synthesize realistic, privacy-preserving patient data to augment small trials.[28]

7.3 Market Dynamics

The economic impact of these efficiencies is measurable. The market for AI in clinical trials is projected to grow from $7.7 billion in 2024 to $21.8 billion by 2030, a compound annual growth rate (CAGR) of nearly 19%.[27] This capital influx is driving the rapid adoption of AI for patient recruitment, site selection, and data monitoring.

Part VIII: The Ethical & Legal Minefield — Algorithmic Bias and Liability

The immense promise of AI is shadowed by the specter of systemic bias. If AI learns from history, and history is racist, then AI will be racist. This is the "Bias In, Bias Out" problem.

8.1 The Yale Study: Quantifying Bias

A seminal 2024 study by researchers at Yale School of Medicine, published in *PLOS Digital Health*, provided a rigorous taxonomy of this risk.[29] The researchers demonstrated that bias infiltrates medical AI at every stage:

1. **Data Collection:** If clinical trials historically exclude women or minorities, the AI trained on that data will be less accurate for those populations.
2. **Labeling:** If "healthcare cost" is used as a proxy for "health need," the AI will assume that poor people (who spend less) are healthier than they are.[31]

8.2 The eGFR Kidney Scandal

A concrete example cited in the literature involves the estimated Glomerular Filtration Rate (eGFR), a metric for kidney function. Historically, algorithms included a "race correction" that

boosted the eGFR score for Black patients, based on flawed assumptions about muscle mass. This caused Black patients to appear to have better kidney function than they actually did, delaying their eligibility for transplant lists.

While human bodies like the Organ Procurement and Transplantation Network have moved to ban this, there is a grave risk that "black box" AI models trained on older datasets will implicitly relearn and perpetuate this discrimination.29

8.3 Legal Liability

The legal framework for AI in medicine is embryonic. Who is liable when an AI misses a diagnosis?

- **The Physician?** Yes, traditionally. But if the AI is a "black box" that the physician cannot audit, is that fair?
- **The Developer?** Software has historically been shielded from product liability, but this is changing for medical devices.
- **The Hospital?** For negligent implementation?

Legal experts warn that "hallucinations" and "data integrity" failures (e.g., an AI coding a patient incorrectly, leading to billing fraud) could expose healthcare systems to massive liability under the False Claims Act.[26]

Part IX: Future Horizons — 2026 and Beyond

As we look to the future, the trajectory is clear: autonomy and integration.

9.1 The Self-Driving Lab

The ultimate vision is the "closed-loop" discovery system. We are moving toward labs where **Deep Research** agents scour the literature to identify targets, **Generative Design** models create molecules, and **Robotic Wet Labs** synthesize and test them—all without human intervention. The results are fed back into the AI, closing the loop. This "self-driving lab" concept promises to run the scientific method at the speed of electricity.[32]

9.2 Multi-Omics and Digital Twins

The next frontier for AlphaFold and its successors is **Multi-Omics**. We will move from modeling single proteins to modeling entire cellular pathways, integrating genomics, proteomics, and metabolomics into a coherent "Digital Twin" of human physiology. This will allow for truly personalized medicine—simulating a drug on a patient's digital twin before prescribing it to the patient.[2]

Conclusion: The Promethean Bargain

The integration of AI into medical research is not merely a technological upgrade; it is an epistemological shift. We are moving from a science of observation to a science of generation. The "Silicon Asclepius" has arrived, capable of seeing the invisible shapes of proteins, designing drugs that nature never invented, and diagnosing diseases that elude the human eye.

The successes of 2024 and 2025—Insilico's clinical proof, AlphaFold 3's structural revolution, and the life-saving interventions of diagnostic LLMs—are proof that the potential is real. Yet, the failures of Recursion and the persistent threats of bias and hallucination serve as necessary guardrails.

We have engaged in a Promethean bargain. We have accepted the fire of computational intelligence to illuminate the dark complexities of biology. The challenge for the next decade is not just to feed this fire with data, but to govern it with wisdom, ensuring that the medicine of the future is not only more powerful, but also more just.

References & Notes

Endnotes

- [1]: Colwell, N. A. (2024). *Harnessing Artificial Intelligence in Drug Discovery*. ACCC Buzz.
- [32]: *Role of AI in drug development*. PMC12406033.
- [17]: Insilico Medicine. (2025). *Partnership Announcements*. Insilico.com.
- [2]: Loper, T. (2025). *How AI is Changing Medical Research*. Nutanix Forecast.
- [12]: *AlphaFold3: Structure prediction and drug design*. PMC12342994.
- [13]: *AlphaFold 3 and protein dynamics*. Taylor & Francis.
- [11]: Google DeepMind. (2024). *AlphaFold 3 predicts structure*.
- [28]: *Generative AI in Clinical Trials*. Frontiers.
- [27]: *AI in Clinical Trials Market Report*. ClinicalTrialRisk.org.
- [31]: *AI in Healthcare: Legal Considerations*. AO Shearman.
- [26]: *AI Enforcement Risks*. Morgan Lewis.
- [7]: *Thinking mode vs Deep Research*. AI Fire.
- [8]: *ChatGPT 5 and Agent Capabilities*. Passionfruit.
- [3]: *Reasoning Models*. OpenAI API Docs.
- [23]: *Doctors told him he was going to die*. Penn Medicine.
- [24]: *Every Cure and AI*. TCNJ Signal.
- [21]: *AI chatbot saves woman's life*. Deccan Herald.
- [22]: *Woman says ChatGPT saved her life*. Fox News.

- [10]: *Rational drug design with AlphaFold 3*. Isomorphic Labs.
- [29]: *Bias in, Bias Out*. Yale Medicine.
- [15]: *Phase 2a trial of ISM001-055*. PubMed.
- [16]: *Insilico Medicine Topline Results*. Insilico.
- [18]: *Recursion Phase 2 Data REC-994*. Recursion IR.
- [19]: *Recursion Discontinuation of REC-994*. Recursion IR.
- [5]: *OpenAI Expands Deep Research*. eWeek.
- [4]: *Deep Research on Gemini*. Google Blog.
- [25]: *Sycophancy in LLMs*. JMIR.
- [10]: *AlphaFold 3 Case Study*. Isomorphic Labs.
- [14]: *Nature Medicine Paper*. Insilico Medicine.
- [3]: *Reasoning Models Technical Details*. OpenAI.
- [29]: *Examples of AI Bias*. Yale Medicine.
- [4]: *Gemini Deep Research Definition*. Google Blog.
- [6]: *Azure Deep Research Tool*. Microsoft Learn.