AI vs Cancer

How AI Can, and Can't, Cure Cancer

By **Emilia Javorsky MD, MPH**

March 2026

Contents

AI vs Cancer

How AI Can, and Can't, Cure Cancer

Every year cancer kills over 600,000 Americans: our friends, colleagues, and family members. It can strike suddenly and fatally, even in those who have done everything right for their health. We know woefully little about how to prevent it. We have few effective weapons once it spreads. The statistics are worsening, as cancer is increasingly killing young people. We must move beyond hollow promises, light on specifics, which simply promise a cure for cancer. Fundamentally, a promise without a plan is a lie. We need a plan as urgent and unrelenting as the disease itself: a plan with the scale, coordination and resolve to end it.

Major tech companies are racing to create artificial superintelligence (ASI). Current artificial intelligence already operates beyond human capabilities in specific domains with well-defined boundaries, such as in chess or image classification. By contrast, ASI would be AI that substantially exceeds human cognitive capability across the vast majority of domains .

Tech companies have pitched ASI as the answer, making cancer cures their flagship promise. **The pitch is seductive: summoning superintelligent AI genies to grant unlimited wishes like unimaginable economic growth, breakthrough treatments for devastating diseases, and reversing climate change.** "Think of the children!" they implore, pulling heartstrings by invoking illness, suffering, and hope. Curing cancer is a big promise, and one that is universally considered one of the most noble and good things we can fund. But rarely is the substance of their promises examined.

The logic appears airtight: If we can create systems of superior intelligence to humans across all domains, surely they will solve what has eluded our brightest scientists for decades. In 2024, Anthropic CEO Dario Amodei suggested that "AI-enabled biology and medicine will allow us to compress the progress that human biologists would have achieved over the next 50-100 years into 5-10 years," calling this the "compressed 21st century." The potential for infinite benefit, we're told, justifies near-infinite risks and infinite investment. Any regulatory or resource constraint that might slow this race becomes unconscionable when cancer cures hang in the balance.

This approach makes a dangerous assumption: that insufficient intelligence is the primary barrier to new cancer therapies. Somehow raw computational power alone is the thing that can overcome the complex landscape of data gaps, biological complexity, regulatory constraints, and misaligned incentives that have caused billions in previous healthcare investments to fail.

In fact, exponential growth in biomedical knowledge is already here. The doubling rate of medical knowledge was 50 years in the 1950s and by some estimates was down to every 73 days by 2020. Yet this intelligence explosion has not significantly moved the needle on cancer mortality or greatly increased annual new drug approvals. **An abundance of knowledge and an oversupply of brilliant scientists have not moved the needle on more cures.** The ASI narrative shapes capital allocation, policy priorities, and public expectations. The National Cancer Institute's 2025 budget, which funds most fundamental cancer research in the United States, was $7.2bn, a mere 1.3% of the $540bn projected total spend by private markets to build out ASI in 2026. The opportunity cost for medical progress is significant. The share of biotech funding is at a 20 year low, while unprecedented sums of VC dollars are directed to ASI development - the markets at least believe the hype.

The obsession with the pursuit of multibillion dollar superintelligence obscures what AI can already do, nearly for free. Today's AI capabilities already deliver real medical value, not through Big Tech's pursuit of god-like machines, but through targeted solutions for specific problems. Even Google DeepMind's AlphaFold succeeded by focusing on one well-defined challenge: protein folding. Across pharma companies, biotech startups, and academic labs, AI is already cracking concrete bottlenecks in cancer treatment: drug discovery, toxicity prediction, and clinical trial efficiency. This is where investment belongs.

But developing practical AI tools to solve problems and remove friction is fundamentally different from chasing ASI genies that assume raw computing power will magically cure all ills. Every complex problem demands we ask: is this bottlenecked by insufficient intelligence? If not, what's actually blocking progress? If so, can current AI solve it? For cancer, today's AI already addresses real intelligence and efficiency gaps in drug development. **The deeper question remains: is intelligence truly the fundamental barrier to curing cancer or are we misdiagnosing the problem entirely?**

Who Am I and Why Did I Write This?

I'm Emilia Javorsky, MD, MPH and I am the Director of the Futures Program at the Future of Life Institute. Throughout my career I've had the opportunity to work "bench-to-bedside" from basic science, co-founding startups, conducting clinical trials, medicine, regulatory compliance, to public health. In 2017 I became motivated to work on how to ensure AI advances human progress and to ensure that we're not taking on risk without benefits. At this point many in the biosecurity community, an area that overlapped with my work in public health, were thinking about the potential risks emerging from artificial intelligence.

Since the launch of consumer-facing LLMs and multi-billion dollar AI fundraising rounds, these two previously disconnected worlds—medicine and AI—have collided with the "promise" of Artificial Superintelligence. This is the magic genie which many companies suggest will grant us cancer cures. The corporate promise is unsurprising, but what is strange is that the claim has gone largely unexamined. I have seen how a new therapy is developed, and intelligence, super or otherwise, was definitely not the bottleneck. I believe we need to break down the promise and examine both the real and exciting ways current AI tools can advance medicine and also flag what the holdups actually are to medical progress.

Ultimately, my choice to write this was not my professional background but an examination of my own intense, visceral response to the cancer-curing promise. I lost my father to esophageal cancer shortly after starting medical school in 2011. He went to the doctor for a pesky cough he'd developed over the past year. There were no other symptoms. When the usual suspects were eliminated, he had an endoscopy. He went in with a cough and came out with a terminal diagnosis. I scoured my medical literature, and still seared in my memory is the table of five-year survival statistics that showed "20%" staring back at me.

We were lucky to live in a hub of leading academic medical centers, visiting each one to understand and compare what treatment options were available. Medicine did not have any great answers: a savage surgery that rendered you unable to eat, toxic chemotherapy to carpet bomb your body, and radiation therapy to shrink the unresectable tumor. Even the combination of all three was unlikely to even get him into that fortunate 20%. My father ultimately chose quality of life with the time he had left and to have radiation therapy. He lived a wonderful two years, feeling tired, but overall quite well. One day in November he started slurring his speech.

We went to the ER at a major academic medical center. The doctors said the cancer had metastasized to his brain, but they had state-of-the-art proton therapy and he'd be home by Christmas. He passed away in the hospital a week later.

I share this story because it is almost everyone's story. Many of us have loved someone who has died of cancer. We know survivors whose lives have been forever changed by the disease, and how they often live with the perpetual fear that it could return. We all have felt the promise and betrayal of hope. Receiving a diagnosis, assuming medicine must have good tools to help, only to discover that for many cancers, there is little that can be done.

One of the most intense memories I have around my father's passing was the, ultimately false, promise that I'd have one more Christmas with him.

Flippant promises by tech CEOs that their technology will cure cancer must not go unexamined. At best, the promise reflects a bias towards over-optimism and naivety about the state of medicine. At worst it's leveraging our collective human hope as a tool to raise funds and to shield their technology from criticism. Market reality dictates that given the unprecedented dollars flowing into ASI, money is flowing away from biotech and medical innovation, which now struggle to raise funds for promising, breakthrough ideas.

The unprecedented mega-rounds being raised by AI companies raise a societal conversation: if we are spending all of this capital to cure cancer, is superintelligence the right bet? In writing this I revisited the survival rates for esophageal cancer 14 years later, the number staring back now is 21.9%, largely unchanged. Yes, let's make curing cancer an urgent national priority, but let's couple it with an honest analysis of why progress is stalled, what are the most promising solutions, how AI can really help us and how to best allocate finite funds and start saving lives.

In this essay, I leverage my background in both medicine and AI to examine the battle of AI vs Cancer and understand who will win.

AI Excels At Problems That Are Rules-Based or Data-Rich, Biology Has Neither

AI excels in domains with complete rules, objective win conditions, instant feedback, and no physical constraints. Medicine has none of these. It offers incomplete information, stochastic outcomes, physical constraints, and delayed consequences. Win conditions are subjective and can take years. AI can help, but history shows us that intelligence alone only suffices in domains specifically structured to reward it.

> *In the human body, rapid cellular changes are more likely to be cancer than progress.*

Silicon Valley's mental model for progress is fundamentally shaped by Moore's Law. This creates an entire culture that expects, plans and budgets for exponential improvement. When tech leaders look at medicine, they instinctively assume similar dynamics should apply. But biological systems are not semiconductors. Evolution optimized tradeoffs in human biology over billions of years for robustness and redundancy. Safety mechanisms that prevent harmful mutations can also slow beneficial changes. Unlike software where updates change systems instantly, biological interventions operate on systems that evolved to resist rapid change. In the human body, rapid cellular changes are more likely to be cancer than progress.

Further, biology imposes fundamental limits on the compressibility of time. You cannot speed up a pregnancy with more engineers, and you cannot compress clinical research beyond the rate at which disease progresses in human bodies. Human biology is bounded by the timescales of cells, organisms, and populations.

Even as computational power and medical knowledge increased exponentially, life expectancy gains have been linear at best, and FDA drug approvals have remained flat for decades. AI companies hang the promise of medical progress on ASI amounting to a "country of geniuses in a data center," yet there is already an oversupply of human genius scientists.

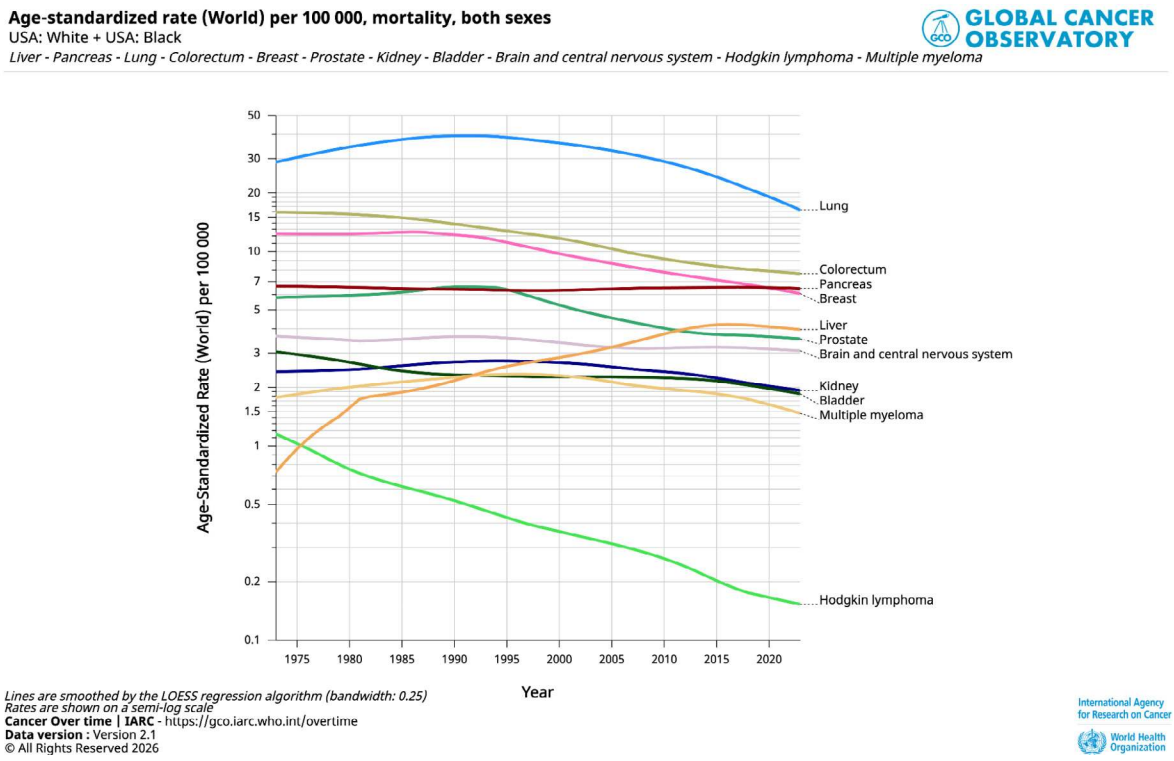

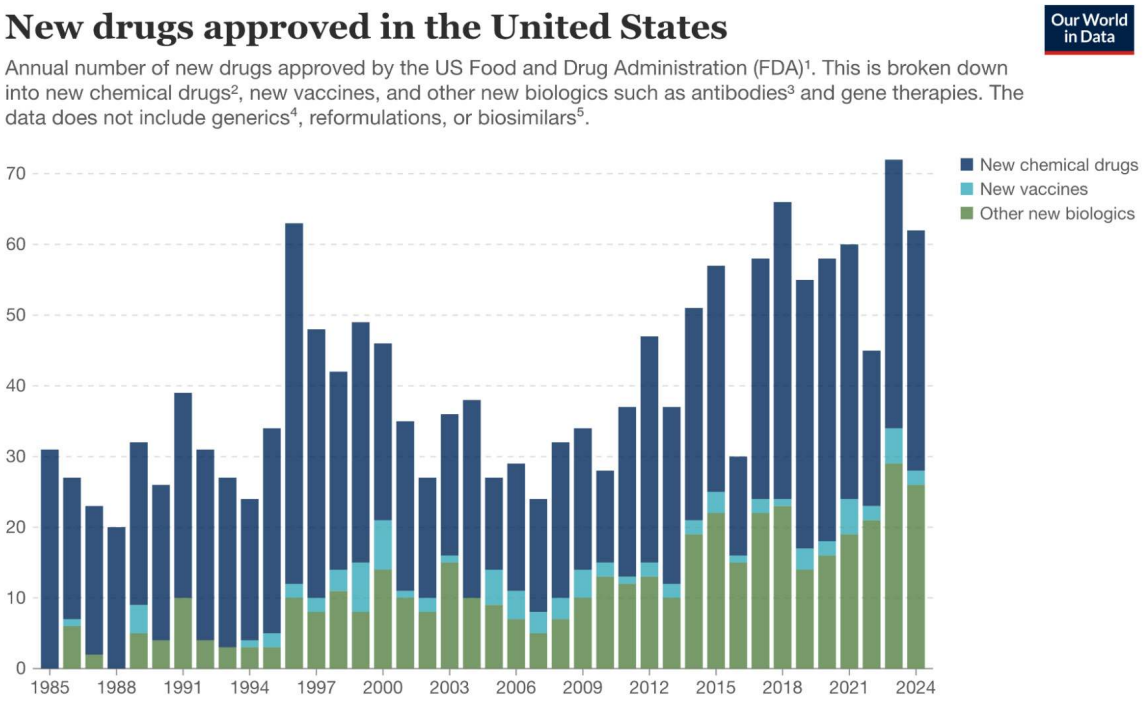

Fundamentally, superintelligence does not equal super-solutions. Grand societal challenges are rarely intelligence-limited problems but systems-limited ones. The uncomfortable reality is that the primary bottlenecks to curing cancer are systemic problems and misalignments within our current power to change. Even if superintelligent AI genies existed today, their wish-granting ability in medicine would be severely constrained by factors no amount of compute can overcome.

Reckoning with Past Failures

Tech's repeated attempts to cure healthcare reveal an abiding misunderstanding of medicine's complexity. From Watson Health to Calico, bold promises have collapsed under data gaps, translational barriers, and flawed assumptions about intelligence.

Silicon Valley's Healthcare Graveyard

The current promises of cancer cures must be contextualized within a long history of failed efforts by technology companies to transform healthcare. **These failures reveal a pattern: consistently overestimating the value of intelligence and technology while greatly underestimating the complexity, data challenges and institutional complexities of medicine.**

IBM's Watson Health serves as perhaps the most cautionary tale. At a 2017 health IT conference, IBM CEO Ginni Rometty told attendees that AI "is real, it's mainstream, it's here, and it can change almost everything about health care," promising to "usher in a medical golden age." After $5 billion invested, Watson Health was sold for parts in 2022 with minimal real-world impact. As Professor Gary Marcus observed, "Watson was a great system for Jeopardy, but it didn't mean you could expect the same technology to solve cancer too." The fundamental problem was not insufficient AI capability but a misunderstanding of healthcare's data landscape and clinical workflows.

Microsoft's HealthVault, launched with fanfare as a platform to revolutionize health records, was shuttered in 2019 after failing to gain meaningful market adoption. Google's life sciences spinoff Verily, despite $2.5 billion in investment over six years, was already being called "a case study in the power, and the limitations, of tech-industry hubris" by 2019. The company's ambitious efforts to develop comprehensive health monitoring and intervention tools floundered on the complexity of translating sensor data into clinically actionable insights.

Google's longevity spinoff Calico represents perhaps the most striking example of the intelligence illusion. Calico was founded in 2013 with the explicit mission to solve aging and extend human lifespan. Radical life extension is another ASI promise. ASI will not only cure cancer, but extend our lifespans by 1,400 years or even longer through some yet-to-be-defined magic. Calico attracted top scientific talent and significant funding. The company operated in stealth for years, the

assumption being that Google-scale intelligence and resources applied to biology would yield breakthroughs. In 2025, the company publicly revealed that its lead drug candidate for ALS had failed in clinical trials, twelve years and billions of dollars later, with little to show for the investment. The failure illustrates a fundamental misunderstanding: that intelligence and capital alone can overcome the biological complexity and translational challenges inherent in therapeutic development.

Haven, the joint healthcare venture from Amazon, Berkshire Hathaway, and JP Morgan, represented unprecedented corporate muscle applied to healthcare reform. If anyone could break through healthcare's ossified structures, surely it would be three of the world's most powerful companies working in concert. Yet Haven collapsed in 2021, falling victim to the very perverse incentives and fragmented local structures it sought to overcome. As Professor Scott Snyder observed: "Given the tremendous training, process discipline, and specialization needed to deliver high-quality care, and the natural risk aversion with patient lives at stake, attempts to reinvent or circumvent the traditional care delivery system have typically failed."

Even Apple's healthcare ambitions, while not outright failures, have fallen far short of CEO Tim Cook's 2019 proclamation that "if you zoom out into the future, and you look back, and you ask the question, 'What was Apple's greatest contribution to mankind?' It will be about health." While the Apple Watch includes health sensors and has found some clinical applications, it has not transformed healthcare delivery or clinical practice in the revolutionary way Cook promised.

These failures share common threads: Silicon Valley has repeatedly stormed into healthcare with the hubris of outsiders attempting to reinvent a system they do not fully understand, failing to learn from past mistakes. The repeated pattern suggests not mere execution challenges but systemic misunderstanding of what's actually limiting medical progress.

The AI Drug Discovery Reality Check

If established tech giants with unlimited resources have failed in healthcare, perhaps specialized AI companies would succeed? The story of AI drug discovery from 2013 to present suggests otherwise. It's largely been a story of disappointment relative to the hype. The term "TechBio" emerged to rebrand biotech, asserting the importance of the "tech" by putting it first. As the head of the UK Bioindustry Association acknowledged in 2024, "I see TechBio as a state of mind—an approach

emerging in companies large and small at the interface of data, AI, life science and innovation." Companies leading the techbio movement were primarily AI drug discovery startups.

Industry pioneer Recursion Pharmaceuticals, founded in 2013, centered its mission on realizing a "techbio" future. It promised to use AI to discover drugs faster and more efficiently than traditional approaches. Exscientia and BenevolentAI were other early leaders, attracting significant investment and generating substantial media attention with their promises of AI-revolutionized drug discovery. The pitch was compelling: AI would identify novel drug targets, design optimal molecules, and dramatically accelerate the path to clinical success.

A decade later, the results tell a more sobering story. Following significant challenges, Exscientia was folded into Recursion in 2024 with the hope of creating an AI drug discovery superpower, representing an implicit admission that neither company had achieved the promised breakthroughs independently. By June 2025, Recursion had laid off 20% of its workforce. In January 2026, Recursion was trading around $4.20, far below its mid-2021 high of $40. BenevolentAI was delisted from Euronext Amsterdam in March 2025, merging with Osaka Holdings. Other long-running AI drug discovery companies including Atomwise and Relay have experienced significant challenges.

> *Thirteen years into the AI drug discovery movement, we still lack a single FDA-approved drug that cleared the full bar of regulatory approval, reimbursement, and clinical adoption.*

A retrospective analysis revealed that AI-discovered molecules show substantially higher success rates in Phase I clinical trials compared to historical norms. This validates AI's superior ability to find novel molecules with drug-like properties. This is genuine progress and should not be dismissed. But this advantage disappears by Phase II trials, when drugs are tested for efficacy to see if they actually treat disease. The bottleneck is not finding promising molecules, it's predicting which will actually work in humans. As one venture capitalist estimated in 2023: "If you take the hype and PR at face value over the last 10 years, you would think it [AI drug discovery] goes from 5% to 90%, but if you know how these models work, it goes from 5% to maybe 6% or 7%."

One recent ray of hope came from Insilico Medicine's success in advancing an AI-discovered candidate for idiopathic pulmonary fibrosis, a lung disease with few good treatment options. Notably, Insilico developed AI tools to target specific bottlenecks at each phase of drug development, from target identification through optimization of clinical trial design, not only at the discovery phase. **Leveraging AI as a tool to manage complexity and reduce friction throughout the drug development process has emerged as more promising than the narrow focus on molecular discovery alone.** While encouraging, this success is still preliminary, as true success in drug development is defined by FDA approval, insurance coverage, and physician confidence in favorable benefit-to-risk ratios for patients.

Thirteen years into the AI drug discovery movement, we still lack a single FDA-approved drug that cleared the full bar of regulatory approval, reimbursement, and clinical adoption. The failure point is not intelligence but the complexity of human biology, the limitations of preclinical models, the friction of drug development, and the challenges of clinical translation.

Market Failures for Life Saving Medicines

The story of AI discovered antibiotics deserves special attention because it represents a controlled experiment. The science was strong enough that there should have been a bidding war to develop the drug, yet it faced significant market headwinds. In 2020, researchers published bespoke AI tools that successfully identified novel antibiotic candidates. The result was an exciting new antibiotic candidate the scientists named Halicin, an homage to the original AI to go mainstream, HAL 9000 from 2001: A Space Odyssey. Unlike HAL, the AI worked. The chemistry worked. The mouse studies worked. Further, compared to most drugs, antibiotics that work in mice have a high predictive value to work in humans. The clinical need is desperate, with antibiotic resistance killing an estimated 1.27 million people globally each year.

But, five years later, where are these antibiotics? The problem wasn't with the science, it was with the market. New antibiotics are inherently unprofitable, as they must be used sparingly to avoid promoting resistance, limiting revenue potential. Major pharmaceutical companies have largely abandoned antibiotic development programs despite the pressing public health need. The antibiotic startup graveyard is instructive: Achaogen went bankrupt in 2019 despite FDA approval for Plazomicin. Melinta Therapeutics filed for bankruptcy. Aradigm abandoned antibiotic development entirely. **The pattern is clear: the market structure can make life-**

saving medicines commercially unviable. The market failures are so significant, researchers had to start a social venture, Phare Bio, with philanthropic and government funding to generate a route to move AI discovered antibiotics forward. While AI can, and already is, reducing friction and costs in drug development, the biggest expenditures, clinical trials and market launch, remain controlled by regulators and payors with no incentive to compromise.

Market forces may block antibiotics, but surely not promising cancer therapies? Tanespimycin, which showed positive results in late-stage trials for multiple myeloma and breast cancer, tells a different story. Very few promising molecules reach late-stage trials, and combined with encouraging data, the drug seemed primed for approval. Instead, the study was "closed prematurely for resource-based reasons." Bristol Myers Squibb offered no official explanation, but a lead investigator suspected "drug most likely was halted because Tanespimycin is difficult to produce and there is only limited time remaining before the drug's patent expires. Combined, these factors would make further investment in the drug difficult to justify financially."

Killing the program was not an act of corporate malice, but a necessary balancing act, weighing the value of a portfolio of promising drugs against market forces. Without extended patent exclusivity, companies cannot justify the enormous costs of late-stage trials, FDA approval, and market launch. Especially when other patented drugs compete internally for investment. **Tanespimycin isn't alone: by some estimates,** 268 off-patent drugs **and** 1,574 natural compounds **have shown safety and potential anti-cancer effects, yet regulatory and market constraints prevent the testing needed to get these drugs to patients.**

Current systemic constraints are already curtailing the potential of AI drug discovery capabilities, and that of biology's progress more generally. The intelligence exists, the science works, the medical need is dire, and yet the drugs don't reach patients. If we cannot solve this problem for current AI-discovered antibiotics or promising cancer drugs, what reason do we have to believe more compute is the solution?

Misguiding Myths and Errors

The belief that cancer can be cured with more intelligence alone ignores decades of evidence showing its staggering complexity. This section explores the myths that distort expectations, why breakthroughs resist simplification, and how oversold narratives about AI obscure the realities of cancer biology.

1. Cancer is the Most Complex of All Diseases - Inherent Barriers to a Universal Cure

"When scientists underestimate complexity, they fall prey to the perils of unintended consequences," wrote Siddhartha Mukherjee in his Pulitzer Prize-winning history of cancer, The Emperor of All Maladies. "Cancer cure" sits atop the list of ASI panaceas, yet it may be the complex disease least likely to yield to any universal curative breakthrough, superintelligent or otherwise.

First, most of us already have potential cancers in our bodies. Our bodies continuously produce cells with mutations that could become cancerous, yet our innate mechanisms of cell removal and immune surveillance eliminate these threats. Cancer is not a foreign invader but a breakdown in our own regulatory systems. **What we call "cancer" is not a single disease but thousands of different diseases, each a product of distinct molecular drivers, cell types of origin, evolutionary pressures, and microenvironments. As Mukherjee eloquently put it: "**Cancer's life is a recapitulation of the body's life, its existence a pathological mirror of our own.**"** In fact, even in a single tumor there can be significant biological heterogeneity, meaning cells within the tumor have different pathways gone awry.

Early research into long-lived, cancer resilient mammals such as elephants, naked mole rats, and the 200-year old bowhead whales, sparked hope of finding a singular mechanism for cancer resistance. However, again nature proved more complex. Rather than a shared compound or a universal protein target, evolution discovered entirely distinct genetic pathways to achieve resilient phenotypes.

> *AI tools are already accelerating the era of precision oncology*

As our understanding of biology advances, cancer continues revealing greater complexity rather than simplification. The landmark 2000 paper "Hallmarks of Cancer" by Douglas Hanahan and Robert Weinberg identified six key capabilities that drive cancer development. By 2011, the authors expanded this to include the

roles of metabolism, immune evasion, and the tumor microenvironment. Cancer was no longer just about rogue cells but about their relationship with surrounding tissues and the immune system. The 2022 update added new dimensions including cellular plasticity and the microbiome. The framework had grown substantially over two decades, not because the original understanding was wrong, but because cancer kept revealing new complexity as our ability to study it improved.

This trajectory represents a shift from understanding cancer as a single cell gone rogue to recognizing it as an individualized state of continuous evolution, presenting an inherent barrier to any universal AI-discovered panacea. **The disease adapts, develops resistance, finds alternative pathways, and fundamentally is an adversarial co-evolutionary process, not a static problem.** Progress in oncology has come from graduating from a futile search for universal solutions towards leaning into complexity and individual variation.

Through the 20th century, cancer care focused on surgical removal of tumors and nonspecific elimination of rapidly dividing cells with chemotherapy. The genomic revolution revealed that specific genetic mutations of cancer cells drive outcomes, regardless of the organ in which a tumor develops. This led to targeted therapy, Gleevec, approved in 2001, was the first drug targeting a specific genetic mutation. Recognition of the immune system's vital role unlocked immunotherapy drugs. More recently, CAR-T cell therapy, personalized cancer vaccines, epigenetic reprogramming, and individualized drug development have all emerged as our understanding of cancer has advanced. Finally, given the time for an idea in the lab to become an approved therapeutic, the science in the clinic today actually reflects discoveries from over a decade ago.

AI tools are already accelerating the era of precision oncology, a fact that gets overshadowed by Big Tech's promise of an absolute cure, as the work is often decentralized across academia, small startups, and internal development programs at pharmaceutical companies. AI is helping to discover new drug targets, biomarkers, predict toxicity, and has helped expand the design of potential new biologics beyond the traditional set of druggable targets. AI is identifying which treatments a patient is likely to respond to, predict resistance in advance, and minimize toxicity. AI is even helping surgeons in the operating room better identify tumor margins to make sure they get all of the cancer and don't need an additional operation.

Examples of Entities Developing, Using & Advocating for AI Tools to Cure Cancer:

AI Drug Discovery & Target Identification: *Developing AI tools to find disease targets and design drugs to block them, faster than traditional lab methods*

Merck*, Novartis*, Pfizer*, Sanofi*, Roche*, AstraZeneca*, Isomorphic Labs, In Silico Medicine, Recursion Pharmaceuticals, Schrodinger, Xaira Therapeutics, Insitro, Relay Therapeutics, BigHat Biosciences, Atomwise, Nimbus Therapeutics, ATOM Consortium

AI Genomics: *Developing AI tools to read and analyze genetic data at massive scale to uncover drivers of cancer and disease*

UKB Genomics (Amgen*, AstraZeneca*, GSK*, Johnson & Johnson*), Memorial Sloan Kettering-IMPACT, AACR Project GENIE, Broad Institute, Wellcome Sanger Institute, Sophia Genetics, Tempus AI (Hardware: Illumina, 10x Genomics, Oxford Nanopore Technologies)

AI Proteomics & Biomarkers: *Develop AI tools to map the proteins the body produces to uncover early or hidden signals of disease presence, progression, and treatment response*

UKB Proteomic Project (Alden Scientific, Amgen*, AstraZeneca*, Biogen, Bristol Myers Squibb*, Calico Life Sciences, Roche*, GSK*, Isomorphic Labs, Johnson & Johnson*, MSD, Novo Nordisk*, Pfizer*, Regeneron*, Takeda*), Human Tumor Atlas Network (Hardware: Olink, SomaLogic, Alamar Biosciences, Nautilus Biotechnology, NanoString)

AI Toxicity Prediction: *Developing AI tools trained on compounds already known to be safe or harmful, predict whether a new compound will be toxic in humans in advance*

University of Vienna ToxCoder, Axiom Bio, Simulations Plus, Instem, Tox21 & ToxCast, Lhasa Limited, Chemaxon

AI In Silico Modeling & Virtual Cells: *Building AI-powered digital replicas of cells and biological systems to simulate experiments and better study basic biology.*

Microsoft Research (BioEmu), Genetech/Roche, Google DeepMind, CZ Biohub, Arc Institute, Broad Institute, D.E. Shaw Research, Allen Institute for Cell Science, Human Cell Atlas, NVIDIA (BioNeMo)

⚠ **AI Drug Repurposing:** *Using AI tools to mine existing approved drugs and natural compounds for hidden potential against diseases*

Broad Institute/PRISM Program, EveryCure, NuMedii, Lantern Pharma, Healx, DreamBio/ OpenTargets

AI Clinical Trials & Regulatory Affairs: *Using AI to accelerate, improve and redesign how trials are run, analyzed, and approved by replacing slower conventional methods.*

NCI Digital Twins Consortium, Friends of Cancer Research (ai.RECIST), Core

Cancer Europe, IQVIA, Medidata, ConcertAI, Unlearn.AI, QuantHealth, Saama Technologies, Flatiron Health, Certara, PAREXEL,TriNetX, Komodo Health

AI Early Detection, Liquid Biopsies & Imaging: *Using AI to process blood, scans, and tissue samples to detect cancer earlier than current methods.*

Foundation Medicine, GRAIL, Veracyte, Delfi Diagnostics, Guardant Health, Exact Sciences, Freenome, C2i Genomics, Volition, OHSU Knight Cancer Institute, Owkin, PathAI, PaigeAI, Proscia, Ibex Medical Analytics, Mindpeak, Caris Life Sciences, Personalis

AI Clinical and Surgical: *Bringing AI into the clinic and operating room to guide decisions, improve precision, and personalize care delivery*

Cancer AI Alliance, MD Anderson Lyda Hill Department of Bioinformatics,

Mayo Clinic Platform, Intuitive Surgical, Activ Surgical, Caresyntax,, Champions Oncology, RaySearch Laboratories, Varian, Elekta, Accuray RefleXion, Cancer Research UK, Dana-Farber Cancer Institute Profile Program, Parker Institute for Cancer Immunotherapy, BostonGene

Advocacy for AI Tools in Oncology: *Leading cancer organizations pushing for AI-driven advances to reach patients through policy, funding, and clinical adoption*

American Association of Cancer Research, American Society of Clinical Oncology,

Stand Up To Cancer, Friends of Cancer Research, Cancer Research Institute

**Major pharmaceutical companies are integrating AI throughout discovery, drug development, and internal processes. The specifics of these programs have limited visibility due to corporate considerations.*

But all of this progress raises the question: why aren't cancer survival rates dramatically increasing? One of the strongest prognostic factors is early detection, and outside of mammograms and colonoscopies, there has been little material progress that has gone mainstream. This is partly because new diagnostics struggle to find viable business models, but also because early detection itself is medically complex. Not all early-detected cancers are destined to become life-threatening, creating risks of overdiagnosis and overtreatment. The South Korean screening program for thyroid cancer, resulted in a 15-fold increase in thyroid cancer over 2 decades, but mortality remained stable, as most of the cancers detected were ultimately small tumors unlikely to cause problems, putting patients through the risks of treatment without benefits. Conversely, broad screening can save lives, with mammography being a prime example, and one recently improved by AI tools. In fact, physician Eric Topol has argued that the largest clinical trial ever conducted on AI in medicine now supports mandatory AI-assisted mammograms to improve early detection of breast cancer.

FDA requirements mandate cancer therapies are studied in populations, not individuals, imposing limits on leaning into the complexity of disease. The narrower you define a patient population, the longer the study takes, burning precious patent time. Worse are the economic incentives driving pharmaceutical firms towards late-stage trials. Bluntly, advanced-stage cancer patients die faster, reaching survival endpoints sooner, meaning faster, cheaper clinical trials. Consequently, new therapies most often meet patients in advanced stages of disease, when the body is worn down and the cancer is aggressive. Medical, regulatory and financial institutions continue to grapple with financial and ethical frameworks for deploying such therapies earlier than studied in trials, but given the life-or-death stakes and high costs there are few clear answers.

These dynamics help to explain why we have so few treatments for the prevention and early treatment of some of the most pressing unsolved diseases - the clinical trials are long and often require a large study population to detect a significant effect. Together these factors equate to "prohibitive development costs."

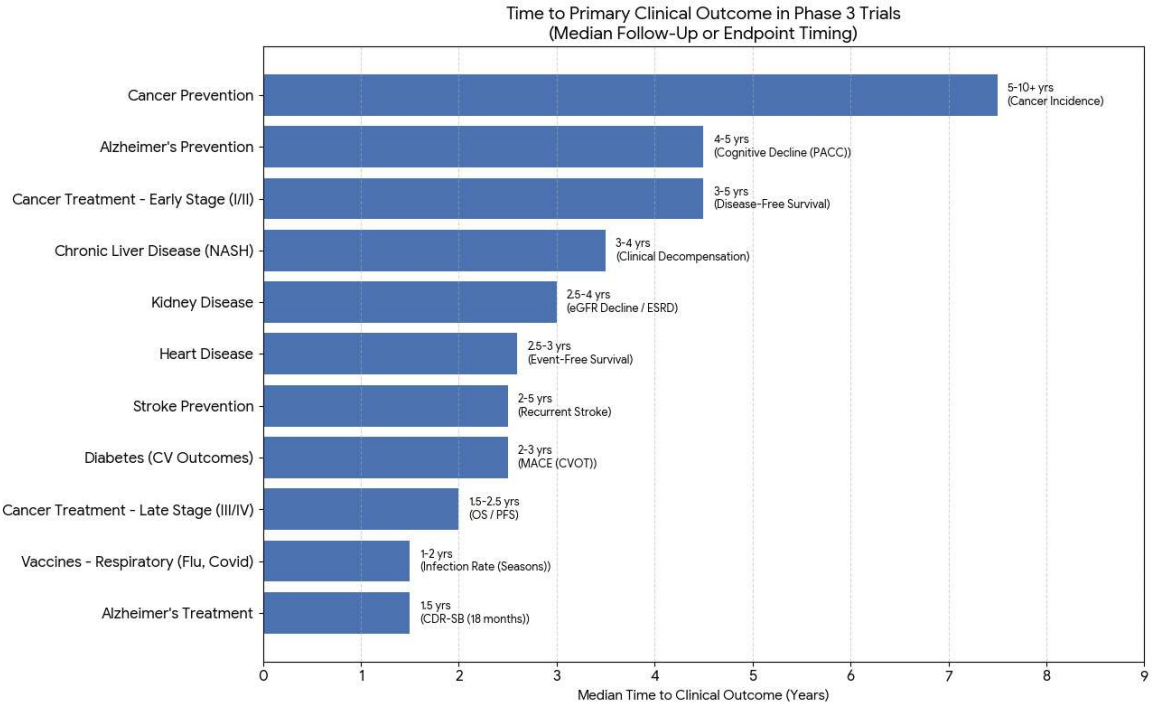

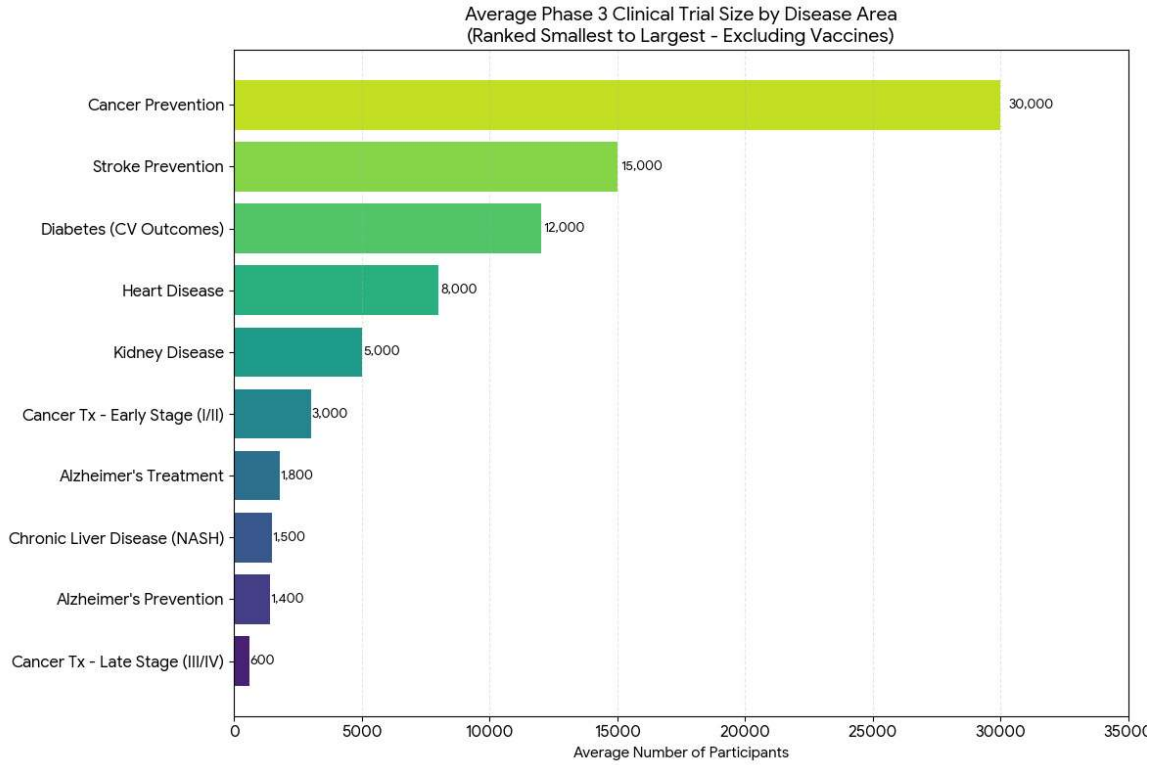

Fundamentally, the more individualized a disease and the earlier you seek to intervene, the more difficult diagnostic and therapeutic development becomes.

2. Myth of the "Eureka" Moment - Discovery is Not Impact

Silicon Valley often misses the bigger picture in therapeutic development by placing an erroneous equal sign between novel insight and clinical impact. The "Eureka myth" refers to the narrative that invention derives from a moment of brilliant insight. As Sir Harold Evans wrote in the Harvard Business Review in 2005, we romanticize the "epiphany of total clarity," creating neat origin stories that obscure the messy reality of innovation.

> *"Eureka," Greek for "I have found it," remains the motto of Silicon Valley's home state of California, which fittingly appropriates a moment of scientific discovery to commemorate the Gold Rush.*

Like the "Great Founder" narrative, the Eureka myth spotlights a lone genius and packages innovation into a single divine moment. This romanticism resonates particularly strongly in tech, where founding a software company can genuinely occur through such moments: a brilliant mind, a laptop, and venture capital ready to invest in scalable solutions.

But therapeutic development bears little resemblance to software engineering. Discovery is not the endpoint but the beginning of what more closely resembles Homer's Odyssey: a decade-long journey tossed at sea, facing insurmountable challenges, gatekeepers, and beating improbable odds to arrive at the just the right moment.

Drug development requires precise intervention in the human body, a highly complex, chaotic, cooperative ecosystem, all while avoiding unintended consequences. Success requires coordinated advances in clinical data generation, rigorous measurement, enabling technologies, relevant experimental models, substantial capital, evolution of scientific paradigms, regulatory policy alignment, social dynamics, and, invariably, quite a bit of luck.

Evans noted that while we obsess over the invention moment, "the act of inventing and improving is far more often a long, hard slog. And the act of capitalizing on invention and of managing the transition from a brain wave to the bustle of the marketplace—is the really hard part." Derek Thompson captures this sentiment in his essay on American progress, observing that if there were a Nobel Prize for

deploying and adopting technology at scale, "our legacy wouldn't be so sterling." In biology, where discovery launches rather than concludes a story, this deployment challenge is exponentially more severe.

Even if a potential cure were discovered in a laboratory today, there is no guarantee we would recognize it as such, nor that it would successfully navigate translational research, regulatory approval, manufacturing scale-up, reimbursement negotiations, and clinical adoption to reach patients who need it, at the moment they need it. Like we saw earlier in the case of Tanespimycin, even drugs who successfully navigated most of the Odyssean voyage can be shelved steps from Ithaca. We likely already possess undiscovered life extending therapies, but they're stuck behind roadblocks and knowledge gaps in our broken system.

3. Staking Hope on New Technologies as Rosetta Stones

Overhyped promises of transformative new technology unlocking therapeutic revolutions predate Big Tech's arrival and the current fanfare around ASI. When President Clinton announced the completion of the Human Genome Project in 2000, he proclaimed that "humankind is on the verge of gaining immense new power to heal" and that the technology "will revolutionize the diagnosis, prevention and treatment of most, if not all, human diseases." The statement could have been written by an AI CEO today, yet it wasn't about AI but genomics.

The project's lead, Francis Collins, speculated that "perhaps in another 15 or 20 years, you will see a complete transformation in therapeutic medicine." The assumption was straightforward: mapping the human genome would provide a Rosetta Stone between disease cause and cure. If we knew each person's fundamental code, we could simply identify and debug the errors causing disease. **It was software thinking applied to biology: find the bug, fix the code, ship the cure. Twenty-five years later, we confront the reality that this assumption was profoundly naive.**

Nearly all diseases proved vastly more complex than single bugs in source code. As a Scientific American review of the Human Genome Project's overstated promises noted, biologist Kenneth Weiss presciently paraphrased Tolstoy in 1993: "All healthy families resemble each other; each unhealthy family is unhealthy in its own way." The diversity and individualization of disease mechanisms has only become more apparent as new methods of measuring biology have come online.

The hope that new technologies will simplify our understanding of disease, only to find that new measurement techniques make things more complex has been the recurring story of modern medicine. At face value, this is obvious. As one zooms into sand on the beach or human tissues, things get more complex at each increasing level of resolution. Each technological advance that promised to unlock biology's secrets instead revealed additional layers of complexity: junk DNA turned out to be functional. Epigenetics showed that genes are regulated by more than sequence. Proteomics revealed that RNA translation to protein is far from deterministic. Metabolomics highlighted the importance of small molecules. The microbiome demonstrated that human biology cannot be understood in isolation from our bacterial passengers. The virome added another layer, and single-cell sequencing showed that even "the same" tissue contains remarkable cellular diversity. It turns out the whole deal of biology is more complicated than the knee bone being connected to the thigh bone. For tech, where "keep it simple stupid" remains a guiding principle of software development, such complexity is quite often the inconvenient truth.

> *It was software thinking applied to biology: find the bug, fix the code, ship the cure.*

With each advance, disease has become more complex and individualized, not less. We may discover that we are all healthy in unique ways as well. **Medicine, humbled by this ever-growing complexity through lived experience of failed predictions, tends to approach new technologies with cautious optimism. Silicon Valley, by contrast, operates within its geographic and ideological bubble of software, where problems do have clean solutions and intelligence can be directly translated into value.** Lacking medicine's hard-earned experience with trials, failures, and entrenched inefficiencies, tech often holds idealized notions of healthcare's challenges and opportunities.

AI Solutions Derive from First Principles or Data

AI thrives when grounded in clear rules or robust datasets, conditions biology rarely grants. With limited first principles and a fragmented data landscape, AI faces inherent constraints in modelling human disease. This chapter outlines why better measurement, curated data, and scientific humility matter far more than raw computational power.

No First Principles for Biology

On the surface, AI is ideal for complex problems. The giant asterisk lurking out of sight is that AI is ideal provided you have first principles and/or robust, high-fidelity data to work with. Equating increasing AI capabilities with scientific breakthroughs is correct in domains with well understood rules of the road, like math and physics. Human biology has neither ingredient.

> *The history of medicine tends to validate H.L. Mencken's axiom:* "*For every complex problem there is an answer that is clear, simple, and wrong.*"

From fundamental biology to clinical care, every stage of scientific inquiry involves substantial information loss and compression, limited by what is measurable, captured, and communicated.

Human biology is neither an engineered system nor fully deterministic. It is the product of billions of years of evolution. Human biology was not designed, but learned through cycles of death and survival. Silicon Valley often sees biology as a complicated engineering problem rather than an incredibly complex, evolved system. In engineering, if you've an incredibly challenging problem like getting a rocket to Mars, you turn to physics. Once you've got the basic rules and building blocks, you have everything you need to build at any scale of ambition. Hence, it is understandable that AI's potential to model biology's basic building blocks, molecules and cells, may be mistaken as providing tools to effectively solve biology at any scale. How wrong this is.

Human bodies are composed of approximately 30 trillion cells, each following local rules that give rise to emergent properties at higher levels of organization. Further, as one travels across layers - from physics to chemistry to biology and then from cells to tissues to organs to organ systems to entire bodies - emergent properties arise that are not fully predictable from lower layers. Biologist Michael Levin argues for reframing biology as collective intelligence, given that functionality of higher layers (i.e. organs) is contingent on the ability of lower levels (i.e. cells, tissues) to cooperate, with each layer evolved to solve different problems that are inaccessible to layers above or below it.

Physicist Stephen Wolfram introduced the concept of "computational irreducibility" to highlight that in some complex systems, the only way to reach fundamental answers is through simulating each step. For some problems like biology, there are no shortcuts and no compression algorithms to capture the essence of the system.

Trying to solve human diseases like cancer through pure computation is akin to trying to solve climate change by simulating every molecule of air. Even if biology were governed by physics alone, simulating human biology from first principles is physically impossible. **Simulating just one week of a single full human (10^28 atoms) using classical physics, with GPUs covering the entire Earth running at their theoretical thermodynamic limit, would take the age of the universe.** Even simulating a single second would be impossible on a timescale relevant to humans. Worse, classical physics is unlikely to be up to the task. And simulating the quantum physics correctly is exponentially harder, requiring either an impossible amount of classical computation, or a quantum computer orders of magnitude larger than we can presently build. While advances in AI's simulation capabilities may unlock exciting insights into molecular and cellular biology, there remains an enormous gap between molecular simulation and human simulation. We cannot compute the fundamental truth of biology, we can only measure it. However, this does not mean that we abandon the role AI can play in improving our ability to simulate biology, improve our measurement abilities, and discover potential "rules of the game," a task scientists are working on. But it does require us to acknowledge that biology is noisy and stochastic. Even in a perfect simulation, one random molecular collision or mutation could be the difference between health and disease given the inter-connectedness of biological systems, posing a fundamental limit on our ability to predict clinically relevant, individual biology from physics.

Data Desert & Biomedical Epistemics

If you don't have first principles to work from, then you need data. Where intelligence-limited bottlenecks do exist in biomedical research, they are likely solvable with current AI capabilities, provided sufficient high-quality, relevant, curated data, which is itself often contingent on innovation in measurement techniques. **These critical measurement and data caveats are often glossed over in discussions of AI's potential.**

Google DeepMind's successes with AlphaFold and GenCast illustrate this point perfectly. Both solved longstanding problems that had stumped researchers for decades, validating AI's capability to accelerate scientific progress in intelligence-limited domains. The secret ingredient in both successes: decades of curated, high-quality datasets spanning global scientific collaboration. AlphaFold was trained on the Protein Data Bank, a database of protein structures maintained by a global consortium of scientists since it was first announced in 1971, representing over five decades of careful data curation. GenCast relied on 40 years of weather data from the ERA5 dataset. These weren't datasets scraped from the internet or assembled ad hoc, they were the product of sustained institutional commitment to data generation, standardization, and sharing in a particular domain.

In science, the potential of AI tools is likewise fundamentally contingent on having relevant, representative, high-quality data from both the bench and the bedside. It is reasonable to infer that DeepMind chose to tackle protein folding and weather prediction precisely because such data existed. Fields where foundational data is broadly available and well-curated, like chemistry and materials science, are primed for AI acceleration. By contrast, biology and medicine, where data is siloed, outdated, incomplete, and often locked behind corporate walls, face far greater challenges to realise the benefits of AI intervention. Further, medicine is full of tacit knowledge that instructs clinicians when to look beyond the data, with being able to walk into a room and determine if a patient is "sick/not sick" being a key tenet of medical training.

Patient Data

In applying AI to medical data, patient records would seem the most logical place to start. Most medical data represents highly compressed, coarse representations of underlying biology: retrospective symptom descriptions, a few hundred standardized blood tests, and medical imaging, all functioning as inaccurate proxies for the biological processes being studied. Electronic health records (EHRs), the

primary source for much medical AI, were not even designed to collect health data but rather to extract and structure information from providers to increase billing revenue, optimizing for sickness not biological reality. As ambient AI scribe systems marketed to restore the human side of medicine by automatically taking notes, are in reality being used to upcode visits and maximize revenue extraction, the clinical accuracy of EHR data is only likely to worsen. **Treating messy, incomplete data that systematically misrepresents clinical reality as ground truth will not yield reliable insights, and may even teach systems that chronic disease maximization is the goal.**

Like AlphaFold's ability to leverage the Protein Data Bank, we recently had a successful example in medicine of how a large, high quality clinical dataset can change the game. The UK Biobank tracked genetic, lifestyle and environmental factors in over half a million volunteers with the goal of collecting 30 years of longitudinal data. Since the release of the first batch of data, applying AI has led to breakthroughs in early diagnosis, clinical care, drug discovery and precision medicine. **Major medical insights historically have come from large scale, structured, longitudinal studies of healthy patients,** such as the Framingham Heart Study, started in 1948, which revealed that reducing high blood pressure and lipids was key to reducing cardiovascular mortality. Similarly, the Nurses' Health Study, started in 1976, taught us about the dangers of trans fats and the role diet plays in driving Type II diabetes.

Beyond scale, time and accuracy, **the UK Biobank is arguably the first longitudinal population scale dataset incorporating modern measurement techniques**. Data collection went far beyond the few coarse measurements routinely collected in clinical care, to include state of the art biology measurement techniques, ranging from whole body imaging to proteomics and metabolomics. The project leaned into the arsenal of modern measurement technologies, adding in new tools as they came on line, to generate torrents of data about human biology at resolutions unimaginable even a decade ago. Mapping precise measurement of biological states to clinical states will be key to unlocking the potential of AI in an otherwise computationally irreducible system. **If the AI build out is to cure disease, shouldn't a fraction of that be invested in the data generation required to solve the problem?**

Such measurement is vital to move beyond unrefined descriptions of symptoms and precisely understand what is health, what is disease, and the variance between individuals. Yet outside of the structured research setting, we've barely begun to integrate genomics, proteomics or metabolomics into routine practice.

The standard annual clinical panel includes roughly 135 routine blood tests, a laughably limited snapshot. Even worse, while some tests are reliably stable with little variance, many have wide variation making it difficult to interpret a single result in a medical record or compare between tests taken at different time points.

In one case study, a healthy young man had blood work done after intense exercise and weightlifting. Initial results suggested he had fatty liver disease, yet when repeating those same tests at rest, they were normal. Even novel modalities like organ-age estimation longevity clocks vary significantly across the day and almost 30 years of variation depending on what tissue type is studied. AI's medical impact is constrained by our failure to measure biology better and refine methods to distinguish signal from noise. These are, ultimately, prerequisites for defining and diagnosing disease with precision.

While routine implementation of next generation measurement into clinical care is not feasible given the current state of the United States healthcare system, spinning up more data moonshots like the UK Biobank is within reach. Such data moonshots are expensive, but are a drop in the bucket relative to expenditures on the race for ASI. The UK Biobank is transforming medicine with the 303 million pounds ($413mn USD) invested so far, yet the projected spend on the ASI race for 2026 alone is estimated at $540 Billion. In more tangible terms, if the total AI spend was $1,000, the entire UK Biobank would amount to a mere 77 cents.

> *If the AI build out is to cure disease, shouldn't a fraction of that be invested in the data generation required to solve the problem?*

As medicine becomes more personalized and precise, the data requirements scale faster than our ability to collect them. No amount of AI capabilities can extract signal from data that was never captured. Intelligence cannot substitute for the measurement of real human beings.

The Literature

The insufficient data problem extends beyond capturing quality clinical data. **Much of what we think we know from biomedical research is likely wrong, we just don't know which parts.** The replication crisis represents more than a quality control problem, it's an epistemic catastrophe that poisons the entire knowledge base on which medical AI would be trained. When 70% of researchers report failing to replicate colleagues' findings, and 68% of papers lack adequate raw data for

reproduction, we're not dealing with occasional errors but systematic knowledge corruption. AI systems trained on this corpus will confidently learn the patterns in published research, which include failures like publication bias, career incentives, and institutional politics mixed in among the biological truths.

In medicine, a basic science finding can be true and replicable in cells and mice, yet clinically insignificant in humans. A common myth is that cells and mice reliably predict human response and that's why they are used as models. The truth is they're simply the best tools we have, but they're not great. When asked whether he could cure cancer, pioneering cancer scientist Dr. Judah Folkman famously replied, "Yes, in mice." **We can cure cancer readily in cells and mice, yet** 97% of **those mouse-cures fail to deliver benefit to humans**, with the main failure mode being lack of efficacy.

Given the low predictive value, biological insights about human disease derived from or engineered into these models are often problematic. The classic example is Alzheimer's research, where researchers spent decades developing and testing drugs to clear amyloid from mouse models, based on the "Amyloid Hypothesis" that build-up of this protein is what causes disease. The mouse models worked, yet 99.6% of the drugs failed to meaningfully improve the disease in humans.

The emergence of the erroneous amyloid hypothesis was not surprising, but its dominance over decades was. It highlighted how **scientific inquiry can be as** limited by institutional gatekeepers **governing an Overton Window as it is by available preclinical models.** The questions allowed to be investigated, and which assumptions can be challenged, are controlled by a peer review publication and NIH grant process that largely relies on legacy incumbent researchers to award funding and decide which papers get published.

The result has been an institutionally low-ambition process that funds fewer out-of-the-box ideas in favor of incremental science, making science less disruptive over time. AI has supercharged this misalignment: advancing scientists' careers at the expense of exploring novel and breakthrough ideas. Dissent is also quashed. There were scientists who questioned the causal assumptions behind the amyloid hypothesis; but they were largely sidelined by what journalists dubbed the "amyloid mafia," institutional incumbents defending established the dogma they helped to create in the first place.

Admitting when we're wrong is vital to the health of science and medicine. Medicine has repeatedly discarded paradigms once considered definitive, from bloodletting to spinning chairs, medical history is peppered with countless therapies abandoned when evidence proved them ineffective. Yet rather than breeding humility, institutional incentives have pushed medicine toward a culture of authority: doctor knows best. The result is heightened barriers to questioning assumptions and pressure to project certainty about what remains nuanced and incomplete. One of the sharpest disconnects between Silicon Valley and the medical establishment is epistemic confidence, as both tech and the public believe medicine knows far more than it does. Medicine's failure to communicate the limitations of existing knowledge (informed doubt) has eroded trust and created a vacuum filled by anyone with an opinion (uninformed doubt). That vacuum fills with questions and theories, much of which is warranted truth-seeking, but some of which is opportunistic proliferation of unscientific content. While speculation may not end up in the literature, it inevitably ends up polluting AI models scraping the internet .

Complicating matters, scientific literature also suffers from severe systemic biases. Firstly, scientific publications only reward the publication of impactful positive results, leaving the negative results and the many experimental iterations required to obtain a significant signal in biology absent from the literature base. **In fact, the value of most of the science being done is lost.** While the NIH has transitioned to mandating the use of electronic lab notebooks internally, externally in academia scientists are still using paper-based data capture. We desperately need massive amounts of comparable, raw biological data, yet most experiments are recorded in formats and silos that render them lost to the scientific commons.

When AI is trained on a literature base riddled with missing data, irreproducible findings, uncertain clinical relevance, and bias toward incumbent thinking, what confidence can we have in its outputs? The problem isn't the algorithm, it's the foundation it's built on. The fix requires mandated data capture, standardized reporting and sharing requirements, tighter links between preclinical science and clinical outcomes, and sustained investment in strategic data generation rather than research that merely validates what institutions already believe. ARPA-H, the DARPA for health, offers a promising model that isn't beholden to incumbent gatekeepers. The agency aims to fund high-risk, high-reward science that breakthrough biomedical progress requires.

Hard Limits to Data Generation

If data is mission critical, beyond improving capture and availability of the data already being collected, how can we scale up our capacity to generate more? The current basic biology equivalent of the UK Biobank is the Human Cell Atlas, a project to map all human cells in the body. The NIH is already funding the BRAIN Initiative, a large-scale coordinated effort to map the human brain. This is exactly the kind of work NIH should pursue, less incremental grants to legacy academics and more investment in the high-quality, large-scale datasets that build the scientific commons. Following in the tradition of the Protein Data Bank, the project begins with identification of what high-quality, curated dataset is missing to advance science.

AI is often put forth as the answer to scale data generation in biology. Robotic automation is already improving standardized processes such as protein engineering, DNA synthesis, CAR-T manufacturing, high-throughput screening and quality control. However, discovery-oriented experiments are often artisanal because iterating and inventing new methods is part of the breakthrough. This requires tacit knowledge, technical mastery, and real-time problem-solving that resists automation, as scientists cannot codify what they're observing and discovering for the first time.

In biology, data generation is not brain-power limited, it's resource limited. There is already a significant oversupply of highly trained scientists relative to available research funding. In the biomedical sciences, there is a single professor position for each 6.3 PhD graduates, an oversupply that cannot be accommodated by a biotech industry in the doldrums. Unlike software engineers who can build with a laptop, biology is a resource-intensive endeavor, with experiments needing to be conducted on a physical lab bench with expensive materials, cutting-edge capital equipment, and a finite amount of hard assets available for experimentation.

Improving the scalability and precision of preclinical work is vital, such as developing virtual cells to better simulate cell biology. But the core challenge remains: most cell biology has limited predictive value for human clinical outcomes. The best preclinical models dramatically improve prediction by using human samples: organs-on-chips with human primary cells, patient-derived xenografts that implant tumor samples onto immunocompromised mice, and Phase 0 trials that test microdoses in patients to assess safety and target modulation.

However, **all of these approaches face a hard constraint: finite human resources. Patient cells, tumor samples, and clinical trial volunteers willing to participate without therapeutic benefit cannot be scaled, they can only be allocated wisely.** While AI can generate thousands of potential targets, we need robust mechanisms to identify which ones justify the allocation of scarce patient resources.

Realizing the potential of AI to accelerate biomedical research requires radically scaling and improving allocation of our data generation infrastructure, precisely at a time when both public and private investment are declining. We must identify data gaps and fund coordinated moonshot efforts to fill them. This means resisting the temptation to ignore what's costly and time-consuming but essential for progress: rich, longitudinal human datasets and novel methods that push the boundaries of biological interrogation.

AI won't solve the intensifying competition for dwindling grant funding and venture capital among academic institutions, nonprofits, and startups. On the contrary, it's actually pulling money from all other markets like a black hole. **Until we generate the experimental and clinical data needed to better model a computationally irreducible system, even the most powerful AI will have limited impact. Filling these gaps takes time, decades for clinical data, years for experimental data, and we haven't yet begun to resource this work in earnest.**

Systemic Bottlenecks & Misalignments

Even with strong science and advancing AI, breakthroughs for cancer stall inside systems built on misaligned incentives, regulatory inertia, and narrow optimization. This section explores how structural constraints limit progress: from reimbursement models and risk-averse regulation to market failures that discourage cures.

The Cost of Narrow Optimization

Silicon Valley's exceptional success has been built on an obsessive commitment to optimization of narrow metrics. Identify a key metric correlated with profit, such as user engagement, search queries, or conversion rates, and align all company activities around maximizing it. The techbio hubris around discovering a "cancer cure" is arguably a byproduct of this philosophy. **Warning lights should be blinking red in applying this culture to biomedical science, as narrow optimization in complex adaptive systems trades short-term gains for long-term cascades of unforeseen consequences.**

The cautionary tale is social media. In optimizing for user engagement, narrow AI algorithms successfully drove profits but also left behind increased rates of depression, impaired cognitive development in youth, erosion of social trust, and the spread of misinformation. From Big Tech's perspective, this approach proved extraordinarily profitable and the negative externalities were borne by users and society, not the companies. Such positive feedback creates a powerful reinforcement function for companies to apply the same process to the next domain of disruption.

> *The narrow optimization problem has already triggered a death spiral in the United States healthcare system, long before Silicon Valley showed up.*

Medicine and human biology represent far more complex systems than social media platforms. **The human body exists in a delicate homeostatic balance: zero inflammation is as pathological as excessive inflammation, zero lipids are as dangerous as too many, driving any biological parameter to an extreme typically causes harm.** Yet the tech industry's muscle memory is to pick an optimization func-

tion and drive it to maximum. This paradigm, so successful in software, becomes dangerous when applied to evolved biological systems characterized by dynamic balance, redundancy, and context-dependent regulation.

The narrow optimization problem has already triggered a death spiral in the United States healthcare system, long before Silicon Valley showed up. Insurers optimize for denying claims to maximize profits. Hospitals and physicians optimize for delivering more care, not better outcomes, because they're paid fee-for-service. Pharmaceutical companies optimize for drugs people take chronically rather than cures, since recurring revenue is more valuable than one-time payments. Regulators optimize for avoiding high-profile failures rather than enabling beneficial innovation, since punishment for approving a harmful drug far exceeds any reward for accelerating a life-saving therapy. Each stakeholder rationally optimizes their narrow objective function and the cumulative result is a system that fails patients: survival rates remain stubborn, patients struggle to access care, and those who can risk financial bankruptcy.

As long as the players benefiting from misalignments are different from those experiencing the negative externalities, incentives for change remain weak. Applying AI, particularly opaque systems, biological or systemic, and optimizing for unclear objective functions, risks supercharging existing misalignments rather than fixing them. AI tools have already been shown to increase billing through more aggressive coding rather than improving care, which risks driving an already struggling system off a fiscal cliff.

In considering the concept of aligning to desired outcomes, it's worth noting that medical AI faces uniquely difficult challenges. Whose definition of health should AI optimize for, a patient or a hospital administrator? When a treatment extends life by three months at the cost of severe side effects, who decides if that tradeoff is worthwhile? These aren't edge cases but the everyday reality of medical decision-making, where values, preferences, and context determine what "care" actually means.

The complexity of biological systems, the stakes involved in healthcare decisions, and the existing misalignments in the system all argue for extreme caution in deploying a culture of narrow optimization-focused AI. We need AI designed with an understanding that biology seeks dynamic balance, not maximum values. We need systems with transparency and interpretability so we can understand what's being optimized and to catch perverse outcomes before they compound. A**nd most**

critically, we need to fix the underlying incentive structures before deploying powerful optimization technologies that will amplify whatever objectives the system currently rewards, whether those objectives serve patients or not.

Regulatory and Institutional Barriers

Having spent the duration of their life cycles largely free from liability under Section 230, and as they are currently advocating for federal AI amnesty, Silicon Valley is uniquely ill-prepared to confront the significant regulatory constraints governing the transition of successful science into approved therapy. **Even when intelligence** could **accelerate discovery, the path to clinical impact requires navigating formidable regulatory and institutional structures.** These structures exist for important reasons, protecting patient safety and ensuring efficacy, but they also create bottlenecks that computation alone cannot overcome.

The FDA is thought of as an agency designed to make sure new drugs are safe and effective prior to being sold in the United States, but as with all things, the devil is in the details. **Close examination reveals a 20th century agency ill-equipped to manage accelerating scientific understanding. The result is an increasing gap between what we could do to help patients based on the state of the science, and what is legally available and accepted by medicine to help patients.**

The FDA's requirements have evolved through layers of well-intentioned administrative expansion. The modern FDA, established in 1938, initially required only that drugs prove their safety. The 1962 Kefauver-Harris Amendments added efficacy requirements to prove that drugs actually work. Our understanding of what is a "disease," what is "effective" and what is "safe" has evolved dramatically since 1962, and lacks the objectivity desired by the regulations.

What constitutes a "disease?" Cancer classification increasingly incorporates genetic mutations alongside traditional organ-based categories; what were once grouped as rheumatic illnesses are now recognized as distinct autoimmune conditions with different mechanisms; dementia has fractured into multiple distinct conditions. Yet regulators often still treat these as monolithic categories. A drug for "Alzheimer's disease" is tested as though Alzheimer's is one thing, when emerging evidence suggests it may be several distinct pathological processes requiring different interventions. Conversely, the FDA refuses to recognize aging as a disease despite growing evidence for shared molecular mechanisms underlying numerous

age-related conditions. Modern biology reveals that disease categories are both too broad (lumping together distinct pathologies) and too narrow (separating conditions that share root causes), yet the regulatory apparatus cannot adapt to either.

The concepts of "effective" and "safe" are equally problematic when applied uniformly. Efficacy is currently defined by endpoints selected through negotiation between the FDA and pharmaceutical industry, not by what degree of improvement actually matters to patients' lives. A drug might meet its primary endpoint of extending progression-free survival by weeks while failing to improve quality of life, or vice versa. Beyond this, efficacy is determined by mean effects, ignoring massive heterogeneity in treatment responses. A drug that fails its primary endpoint may work remarkably well for a subset of patients who will never gain access to it. Safety is similarly treated as objective when it's inherently personal. Patients with the same diagnosis face identical regulatory risk-benefit determinations despite varying personal circumstances, values, and risk tolerances.

The 2006 Physician Labeling Rule took things even further, requiring a detailed mechanism of action, where companies must explain how their drug helps patients, not just if it helps. Each layer was intended to protect patients, but together they have created an environment where rational drug design has become limited to crafting compelling regulatory narratives more than objective biological truth.

***The* current regulatory paradigm demands answers to narrow questions: What is *the* target? What is *the* molecule? What is *the* mechanism? What is *the* outcome? Given *the* complexity of human disease and our incomplete understanding of biology, *the* fact that we've developed any effective treatments at all within *the*se constraints is remarkable.** The entire possibility space of multiple targets, synergistic molecules working in concert, multiple potential mechanisms, or acknowledging incomplete mechanistic understanding while still delivering clinical benefit is functionally off limits from a regulatory perspective.

Yet many highly effective drugs still in use today were approved before mechanism-of-action requirements and we still do not fully understand how they work: valproic acid for seizures, lithium for bipolar disorder, guaifenesin for cold and flu symptoms. Some of our most successful drugs were discovered through "irrational" observation or serendipity during development rather than rational design. Sildenafil (Viagra) was designed for hypertension but proved better at treating erectile dysfunction through mechanisms not originally anticipated. Minoxidil (Rogaine)

had a similar story for hair growth. GLP-1 drugs like Ozempic were thought to work on the stomach but primarily act on the brain, a complete inversion of the rationalized mechanism during development.

What breakthrough therapies are we missing by demanding premature mechanistic certainty? How many potential treatments languish because they don't fit into neat regulatory boxes? What is the unrealized potential of current science to save lives and reduce suffering? AI's ability to identify novel multi-target approaches, synergistic drug combinations, personalized therapeutics, or treatments that work through incompletely understood mechanisms is already constrained by regulatory frameworks designed for simpler paradigms. Again, Big Tech is likely making the erroneous assumption that the lack of biomedical progress is due to scientific shortcomings, naive to the amount of science already bottlenecked by an outdated, sclerotic regulatory system.

Rethinking Paternalism

Ultimately, the potential of contemporary science, and an AI-supercharged future, is contingent on a paradigm shift. **The FDA embodies a mid-20th century worldview of biomedicine as producing static, objective, uniform truths, where knowledge is stable enough to encode in regulations and apply universally. Modern biomedical science reveals the opposite: our understanding is provisional and rapidly evolving, and the frontier of therapeutics is moving toward personalization rather than universal treatments.** In the FDA's defense, it is trying to modernize, but the pace of change is woefully insufficient, relative to the pace of science. A regulatory system designed for slow-moving, population-level knowledge cannot accommodate a science that regularly discovers new disease subtypes and increasingly attempts to predict individual treatment responses. The mismatch is not just about who decides, but about whether centralized decision-making can possibly keep pace with the speed and granularity of current discovery.

Policy already acknowledges this reality in extreme cases. Right to Try laws recognize that terminally ill patients can access experimental treatments without FDA approval after exhausting standard options. **If we trust dying patients to evaluate experimental therapies when they have nothing to lose, why not trust patients with adequate information to make treatment decisions earlier, when interventions might actually work?**

The answer lies in mid-20th-century paternalism, when the presumption was that doctors know everything and always know what's best for you. At the time, such paternalism was defensible when science was stable, information asymmetries were insurmountable, and patients had accessible primary care doctors. Benchmarking to an old, idealized standard of care most patients never experience is intellectually dishonest. In reality, millions of Americans lack regular physician access, let alone specialists offering personalized recommendations. The relevant comparator isn't a world-class specialist, it is more often no care, an overwhelmed doctor with 15 minute appointments, or insurance company algorithms.

Further, the information gap that justified regulatory paternalism is narrowing with AI's democratization of medical knowledge to patients**, and** democratization of accelerating scientific **knowledge to providers.** Medical progress requires a combination of regulatory updates and a rethinking of what decisions should be made by centralized regulators, and what decisions should be made by providers and empowered patients. Such a rethinking is going to require more than tweaking FDA rules. It will require the development of high quality, trustworthy AI tools to empower patients. Physicians will require retraining as scientific interpreters rather than protocol-followers. Equally critical is liability reform, so providers do not bear malpractice risk for individualized decisions that deviate from population-averaged standards of care.

Limits to Compressing Time

There are non-bureaucratic factors limiting the speed in which science can be adopted, which is the time needed to see if something works. Clinical trials themselves present timescales that are difficult to compress. On average, it takes 10.5 years for a drug to move from Phase I through regulatory approval. For drugs entering Phase I, 90% will fail somewhere along the pathway, with lack of clinical efficacy representing 40-50% of failures and safety concerns another 30%. This is not a problem of insufficient intelligence in trial design, it's the inherent challenge of safely testing interventions in humans on biological timescales.

Some companies report compressing preclinical workflows from 3-6 months to 18 months. But this creates a misleading impression: **while AI's role in accelerating drug discovery sounds like a 90% improvement to the public, the reality is more modest, perhaps 10-20% time savings because you're only radically condensing the initial pre-clinical phase of drug development.** More importantly, any increase in drug candidates inevitably collides with a fixed bottleneck: the limited pool of

patients available for clinical trials, already a major challenge with existing pipelines. Clinical trials are harder to compress because they operate on the timescales of human biology and regulatory requirements. **AI can eliminate friction in the process and tighten the feedback loop between scientific discovery and clinical application, but compression is ultimately constrained by the pace of human health and disease itself.**

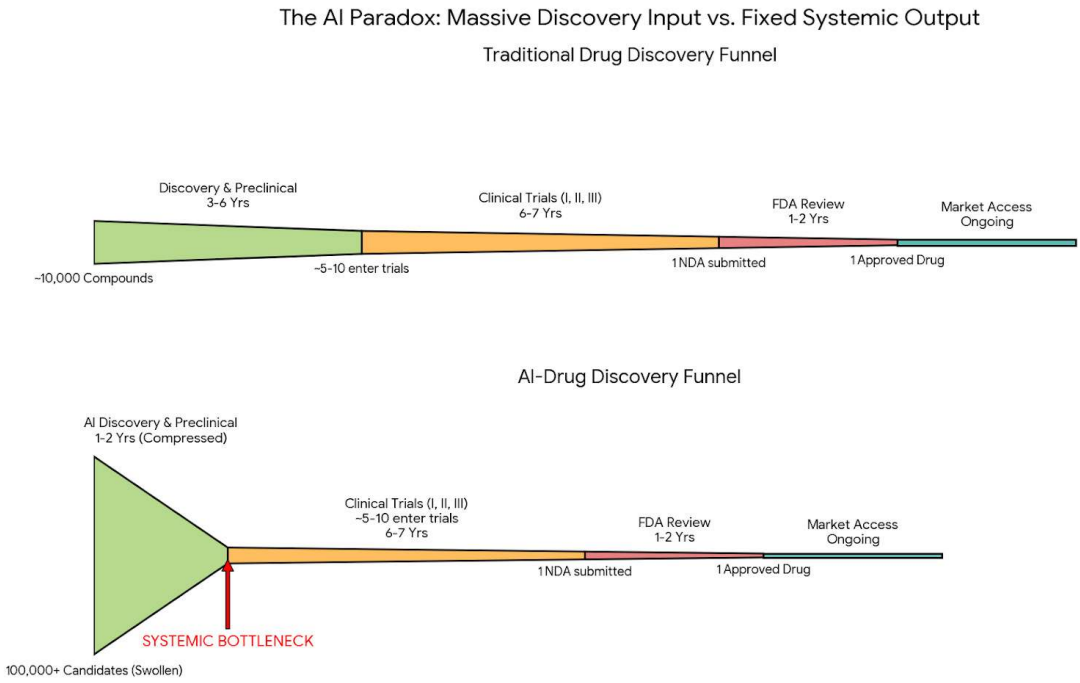

Yet, some ASI proponents suggest that their genies will cure most diseases in 5-10 years. The evidence cited for this assertion is the speed of the COVID vaccine approval. This was approved in less than a year, and it's suggested that this demonstrates long timelines are due to FDA bureaucracy more than biologic limitations. This comparison ignores crucial context. Moderna was founded in 2010 and submitted its first Investigational New Drug application for an mRNA flu vaccine in 2016. By the time COVID struck, the company had nearly a decade of foundational research, safety data, and manufacturing infrastructure already in place. The nine-month timeline counted only from COVID emergence, not from when mRNA technology development began.

> *The logistics of biological timescales, long-term follow-up and patient retention over years present challenges that no AI can simply optimize away*

COVID also presented unique trial advantages that chronic diseases fundamentally lack. With a novel virus spreading throughout the general population where almost everyone is due to be infected, trial participants were easy to find and recruit at scale. Outcomes were clear and rapid: did people get sick, require hospitalization, or die within weeks or months? Some researchers even proposed human challenge trials, in which healthy volunteers would be deliberately exposed to the virus, an ethically defensible approach when community transmission meant most people would likely be infected anyway. Such trial designs can produce results within weeks.

By contrast, chronic diseases like cancer or Alzheimer's require far longer and more complex trials to know if something is working or not, and for whom. Outcomes are slower to manifest. Trial readouts for cancer treatments often require years of follow-up to assess survival benefits. Recruitment is dramatically harder since researchers must identify and enroll patients with specific disease subtypes rather than drawing from the general population. Disease heterogeneity means that subgroup analyses are essential but require larger sample sizes. The logistics of biological timescales, long-term follow-up and patient retention over years present challenges that no AI can simply optimize away.

The speed of COVID vaccine approval also reflected unprecedented political will to find a cure. Operation Warp Speed mobilized resources not only to save lives but also to stabilize the healthcare system and reopen the economy. Government funding de-risked investment, regulatory agencies provided real-time guidance and accelerated review, and manufacturing scale-up began at financial risk before approval. That sense of urgency and resource mobilization does not exist for chronic diseases, despite their enormous cumulative burden. Cancer kills more than 600,000 Americans annually, yet commands nothing close to the focused governmental response COVID received.

There are inherent limits to acceleration of FDA review. The FDA modernization Act 2.0 has already enabled new methods of preclinical testing, such as microdosing human trials and organs-on-chips, to compress timelines between drug discovery and safety testing in humans. AI tools are improving clinical trial patient selection and creating synthetic control groups to try to increase the productivity of clinical efficacy trials. AI is also enabling the creation of digital twins, a key enabling feature for conducting individualized trials (n=1 clinical trials). Even if the FDA were to revert to approvals based on the demonstration of safety, the question as to whether a new drug actually moves the needle for an individual patient remains bounded by the pathophysiology of the disease.

Similar to moonshot investments in clinical data collection, the critical question for new drugs is: how to know reliably if an intervention is working before a clinical outcome? In essence, how can we compress clinical timescales? For some conditions like infectious disease, time cycles are short, but many diseases play out over decades. The ability to know whether something is working is especially salient as medicine transitions to preventing disease in high risk individuals, rather than treating sickness. There are promising methods in this area, like complex biomarker development. However, it takes time to see if these surrogate endpoints match clinical outcomes. You need at least one full cycle of disease pathophysiology to get results.Further, in a system that values treatments far more than diagnostics, there is again a resource constraint on both the capital and patience required to develop the tools ultimately needed to compress biological timescales.

Follow the Money: Who Pays for What?

In the absence of understanding the incentives of the United States healthcare system, one would assume that cures to diseases, especially diseases such as cancer, would be invaluable, and certainly worth the cost. But, in a system that profits from treating diseases, a healthy patient represents lost revenue.

> *As a Goldman Sachs analyst famously asked: "Is curing patients a sustainable business model?" Gilead Sciences developed a cure for Hepatitis C, but analysts watched drug revenues plummet from their peak of $12.5 billion in revenue as the number of patients was exhausted.*

Consider CAR-T therapy for cancer, priced at $475,000 for a one-time treatment. Or Zolgensma for spinal muscular atrophy at $2.1 million. These aren't arbitrary prices, they're essentially estimates of the equivalent lifetime value of chronic treatment alternatives. The pricing explicitly assumes cures must capture the same total revenue as ongoing disease management. The system has encoded a preference for sickness over health.

Similarly, when doctors and hospitals are paid based on how much care they deliver (fee-for-service) rather than health outcomes (value-based-care), the healthcare system thrives economically when people remain sick. **Early diagnosis and prevention eliminate the procedures, tests, and prescriptions that generate revenue, creating a fundamental misalignment between profit and patient welfare.**

Financial rewards in pharmaceutical development are usually decoupled from the eventual clinical outcome, as the typical exit path for small pharmaceutical companies is acquisition by a large pharmaceutical company. This is often before regulatory approval or even large-scale clinical trials. This creates systemic pressure to hype early findings and operate in domains with high acquisition appetite, resulting in "herding effects" where innovation tends to cluster around popular targets and therapeutic areas.

Further, when incumbents purchase promising candidates, they integrate them into existing pipelines subject to existing incentives, or simply shelve them if they threaten profitable drug franchises. The science behind the current blockbuster GLP-1 drugs was already known by the late 1980s, and a decision by Pfizer to end exploration with the startup exploring the science in 1990, set the field back decades. This path dependency ensures that even genuinely revolutionary AI discoveries get filtered through the same misaligned incentive structures the technology was supposed to bypass.

The pharmaceutical industry's relationship with the FDA is clouded by the appearance of sophisticated gaming as opposed to neutral scientific evaluation. Companies employ former FDA officials who understand agency priorities and informal requirements. They design trials to meet established regulatory endpoints rather than answer meaningful clinical questions. This isn't necessarily malicious, it's a rational adaptation to incentive structures. But adding AI to this system risks enabling more sophisticated gaming rather than genuine innovation.

An ASI optimist may argue that the technology can also solve such systemic problems. This reveals the fundamental faith underlying ASI promises: that sufficient intelligence equals power. **Market and regulatory structures don't yield to smart arguments**. An AI might propose perfect healthcare policy, but implementation requires navigating Congress, state legislatures, corporate lobbying, and entrenched interests.

While these incentive misalignments are unlikely to stand in the way of a cancer cure, they do shape the probable paths of capital allocation by companies whose fundamental goal is profit. This prompts reflection as to whether the companies promising AI genies have a history of transcending incentives to advance the betterment of humanity, or whether they are likely to follow the money like anyone else.

Conclusion

The seductive promise of ASI curing cancer allows us to avoid confronting uncomfortable truths. Namely, that our scientific institutions have structural problems limiting innovation, that our regulatory frameworks lag behind biological understanding, that incentives are fundamentally misaligned with patient welfare, and that we lack the data infrastructure to fully leverage AI's potential. Waiting for Team Techbio's intelligent genies absolves us of the collective responsibility to build, reform, and innovate today. When patients are suffering and dying each day, bankrupting families to access basic care, this abdication is unconscionable.

We have examples of what medical progress actually looks like: the UK investing in the Biobank and making the data freely available to researchers globally. The FDA approving the first gene therapies for sickle cell disease after 30 years of foundational research. Operation Warp Speed demonstrating that removing financial risk and regulatory uncertainty can compress timelines when urgency is real. The Framingham Heart Study, started in 1948, teaches us the fundamentals of cardiovascular disease prevention. **The pattern with all of these examples is sustained institutional commitment, coordinated data infrastructure, strategic de-risking of valuable research, and long time horizons. These successes came from systems-building, not intelligence explosions.**

The prosaic truth is that the path forward is not through intelligence explosion but through hard work. Edison had it right: we need 1% inspiration and 99% perspiration. Specifically, that means generating better data through sustained institutional investment, reforming regulatory frameworks to enable innovation while maintaining safety, realigning commercial incentives with medical objectives, and empowering domain experts to build AI tools matched to specific problems they understand deeply. This requires confronting the reality that grand challenges are rarely intelligence-limited. Grand challenges are data-limited, regulation-limited, incentive-limited, and coordination-limited.

Current AI capabilities, thoughtfully applied by people who understand biomedical problems deeply, can help. They can accelerate therapeutic development, reduce trial timelines, improve preclinical models, and help restructure broken incentive systems. We see this potential being realized in distributed fashion across universities, startups, and pharmaceutical companies using current AI as a tool

for specific applications. But realizing this potential at scale requires abandoning the intelligence illusion and embracing the complex, unglamorous work of fixing systems that actually stand between discovery and cure.

The gap between AI's advertised potential and its limited real-world impact on medical progress is not evidence that we need more intelligent AI. It's evidence that intelligence is not the binding constraint. We mistake incentive problems for intelligence problems, data problems for computational problems, and coordination failures for capability limitations. The result is a dangerous complacency: rather than undertaking difficult institutional reforms, we wait for the techbio bros to save us.

Superintelligence may eventually arrive. But cancer patients cannot wait for genies when the real barriers are human-made systems we could reform today if we had the will and the permission to do so. The question is not whether AI can help cure cancer, current capabilities already can, serving as one tool among many in researchers' arsenals. The question is whether we will undertake the systemic changes needed , or just continue to mistake the promise of future intelligence as a substitute for present action.

Acknowledgements

Further to Anthony Aguirre's Standard of Levels of AI Involvement in Creative Works, this would score a 3.5/10. All of the prose was originally human generated and used AI (Gemini, Claude and Grok) to cut down the content. AI was also used as a research assistant for finding citations, feedback, and red-teaming of weak points in some arguments. Gratitude to Anthony for his cosmology consultation on modeling human biology from physics. I'd also like to acknowledge: Eric Topol for his nuanced coverage of AI in medicine in Superagers; Siddhartha Mukherjee for bringing the foe to life in The Emperor of All Maladies; Derek Thompson, co-author of Abundance, for charting a path toward better science; Stephen Wolfram and Michael Levin for their work on why first principles fail in biology and alternatives to pursue; Gary Marcus for his commitment to calling out the AI hype; Charles Piller for his investigative journalism on the history of Alzheimer's research; the Institute for Progress, Friends of Cancer Research, and the many kindred organizations working to unlock medical progress; and all of my wonderful colleagues at The Future of Life Institute, whose conversations and wisdom have shaped my thinking on building a positive future with AI.

The Roadmap Forward

AI vs Cancer explores why artificial superintelligence (ASI) won't deliver on the cancer curing hype, and details the many data, economic, systemic, and institutional challenges that bottleneck progress. But if ASI isn't the answer, where do we go from here to actually have a chance at defeating cancer and how can AI help? This roadmap outlines the work that is already being done and towards this problem, but also highlights the real issues that we need to address if we are going to cure cancer.

This roadmap is also an invitation - to everyone who reads it to help us to refine the ideas and to collectively bring momentum to the issue. Cancer has proven an intractable enemy for over 600 million years since multicellular life appeared. If we want to change that then it will require input from every sector of humanity to bring ideas, skills and the demand for political and societal changes.

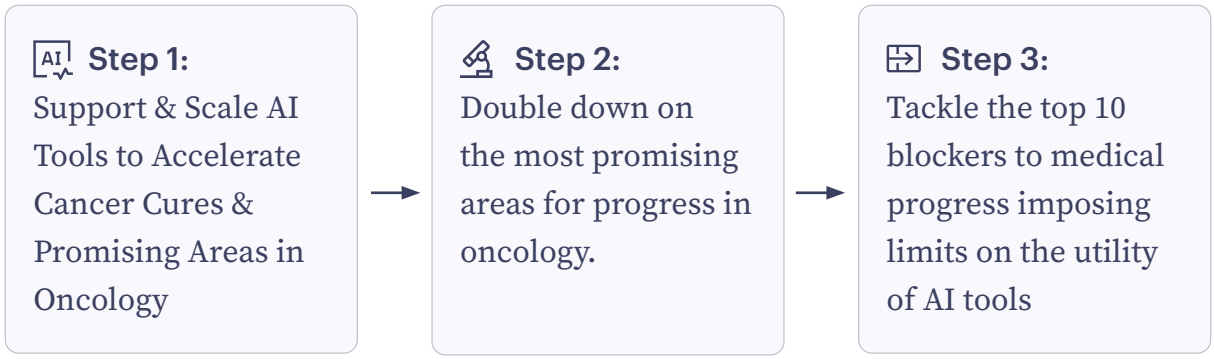

Why Now is the Right Time

Beyond the urgency to tackle cancer given the human cost of patients dying every day, we sit at a unique point in modern history that makes going on the offense against disease more promising than ever. U.S. healthcare has long been trapped in what Scott Alexander calls a Moloch problem: a coordination failure where every stakeholder sees the system collapsing yet no one can escape the dynamics driving it down. Hospitals are closing, insurers are in death spirals, pharmaceutical companies face a $236 billion patent cliff by 2030, science funding is waning, doctors are burnt out, and patients can't access care, all while healthcare consumes 27% of federal spending and threatens the nation's fiscal and geopolitical footing.

It's clear that tinkering with reform at the edges is no longer viable economically or politically. The same resignation haunts cancer specifically: after Nixon's 1972 moonshot and Obama's in 2016, mortality rates have barely budged, breeding a fatalism that curing disease is simply unsolvable. But there has been no better time to shake off fatalism and go on the offense. Ironically, a system in late-stage collapse also creates the conditions to rebuild. Culturally, within US politics, there is an appetite for radical institutional disruption, with polling during elections showing 83% wanting substantial change or complete upheaval.

Further, existing AI capabilities are already set to radically disrupt the workforce and incumbent businesses, whether for the better or worse remains unclear. What is clear is the combination of collapsing healthcare, political appetite for disruption, and the availability of AI's tools to reduce friction, align incentives, and manage coordination at scale, creates a once in a generation window to explore new avenues that may at last enable us to develop a concrete path to better healthcare and a genuine plan to cure cancer.

Step 1: Support & Scale AI Tools to Accelerate Cancer Cures

The AI capabilities that we have right now contain immense promise to accelerate science. There are countless companies, non-profits and academics hard at work developing AI tools often coupled with novel methods of measuring human biology to tackle key bottlenecks in oncology.

From the Cancer AI Alliance's development of a federated AI framework to empower researchers to learn from deidentified patient data, to the National Cancer Institute's Digital Twin development, non-profits, large pharmaceutical companies, small startups, government institutions and academics are all hard at work on AI tool development.

While BigTech has tried to claim the mantle of building, these are the true builders we should be celebrating, resourcing and scaling. Below is a sample of some of the key actors working at various points from bench to bedside to unlock the power of AI to deliver real benefits to patients.

Examples of Entities Developing, Using & Advocating for AI Tools to Cure Cancer:

AI Drug Discovery & Target Identification: *Developing AI tools to find disease targets and design drugs to block them, faster than traditional lab methods*

Merck*, Novartis*, Pfizer*, Sanofi*, Roche*, AstraZeneca*, Isomorphic Labs, In Silico Medicine, Recursion Pharmaceuticals, Schrodinger, Xaira Therapeutics, Insitro, Relay Therapeutics, BigHat Biosciences, Atomwise, Nimbus Therapeutics, ATOM Consortium

AI Genomics: *Developing AI tools to read and analyze genetic data at massive scale to uncover drivers of cancer and disease*

UKB Genomics (Amgen*, AstraZeneca*, GSK*, Johnson & Johnson*), Memorial Sloan Kettering-IMPACT, AACR Project GENIE, Broad Institute, Wellcome Sanger Institute, Sophia Genetics, Tempus AI (Hardware: Illumina, 10x Genomics, Oxford Nanopore Technologies)

AI Proteomics & Biomarkers: *Develop AI tools to map the proteins the body produces to uncover early or hidden signals of disease presence, progression, and treatment response*

UKB Proteomic Project (Alden Scientific, Amgen*, AstraZeneca*, Biogen, Bristol Myers Squibb*, Calico Life Sciences, Roche*, GSK*, Isomorphic Labs, Johnson & Johnson*, MSD, Novo Nordisk*, Pfizer*, Regeneron*, Takeda*), Human Tumor Atlas Network (Hardware: Olink, SomaLogic, Alamar Biosciences, Nautilus Biotechnology, NanoString)

AI Toxicity Prediction: *Developing AI tools trained on compounds already known to be safe or harmful, predict whether a new compound will be toxic in humans in advance*

University of Vienna ToxCoder, Axiom Bio, Simulations Plus, Instem, Tox21 & ToxCast, Lhasa Limited, Chemaxon

AI In Silico Modeling & Virtual Cells: *Building AI-powered digital replicas of cells and biological systems to simulate experiments and better study basic biology.*

Microsoft Research (BioEmu), Genetech/Roche, Google DeepMind, CZ Biohub, Arc Institute, Broad Institute, D.E. Shaw Research, Allen Institute for Cell Science, Human Cell Atlas, NVIDIA (BioNeMo)

AI Drug Repurposing: *Using AI tools to mine existing approved drugs and natural compounds for hidden potential against diseases*

Broad Institute/PRISM Program, EveryCure, NuMedii, Lantern Pharma, Healx, DreamBio/ OpenTargets

AI Clinical Trials & Regulatory Affairs: *Using AI to accelerate, improve and redesign how trials are run, analyzed, and approved by replacing slower conventional methods.*

NCI Digital Twins Consortium, Friends of Cancer Research (ai.RECIST), Core Cancer Europe, IQVIA, Medidata, ConcertAI, Unlearn.AI, QuantHealth, Saama Technologies, Flatiron Health, Certara, PAREXEL,TriNetX, Komodo Health

AI Early Detection, Liquid Biopsies & Imaging: *Using AI to process blood, scans, and tissue samples to detect cancer earlier than current methods.*

Foundation Medicine, GRAIL, Veracyte, Delfi Diagnostics, Guardant Health, Exact Sciences, Freenome, C2i Genomics, Volition, OHSU Knight Cancer Institute, Owkin, PathAI, PaigeAI, Proscia, Ibex Medical Analytics, Mindpeak, Caris Life Sciences, Personalis

AI Clinical and Surgical: *Bringing AI into the clinic and operating room to guide decisions, improve precision, and personalize care delivery*

Cancer AI Alliance, MD Anderson Lyda Hill Department of Bioinformatics, Mayo Clinic Platform, Intuitive Surgical, Activ Surgical, Caresyntax,, Champions Oncology, RaySearch Laboratories, Varian, Elekta, Accuray RefleXion, Cancer Research UK, Dana-Farber Cancer Institute Profile Program, Parker Institute for Cancer Immunotherapy, BostonGene

Advocacy for AI Tools in Oncology: *Leading cancer organizations pushing for AI-driven advances to reach patients through policy, funding, and clinical adoption*

American Association of Cancer Research, American Society of Clinical Oncology, Stand Up To Cancer, Friends of Cancer Research, Cancer Research Institute

**Major pharmaceutical companies are integrating AI throughout discovery, drug development, and internal processes. The specifics of these programs have limited visibility due to corporate considerations.*

Step 2: Double down on the most promising areas for progress in oncology

In parallel, it is also important to double down on the most promising areas for progress in oncology.

Figure 2: Concrete areas of promise in oncology

Early Detection & Prevention

- Resource reduction efforts for known drivers of cancer such as smoking, obesity, and alcohol consumption
- Increase access to proven screening (mammography, colonoscopy/stool tests, low-dose CT for high-risk smokers) and prevention measures (HPV & Hepatitis B vaccination)
- Research and address environmental carcinogens (PFAS contamination, benzene, occupational exposures, and air particulate matter) through policy interventions
- Scale and insurance coverage of multicancer early detection blood tests. Also resource large scale studies to determine if they deliver mortality benefits
- Research promising drugs to prevent cancer in high risk groups defined by genetics or new multi-omic risk analyses coming online
- Invest in methods, such as blood tests and real time monitoring, to detect minimal residual disease or cancer resistance as soon as emerges/ in real time to enable rapid intervention

Data & Clinical Trials

- Build a national cancer data commons linking genomic, imaging, treatment, and outcomes data to end the fragmentation blocking AI-scale discovery
- Develop novel incentives to generate large, longitudinal datasets for example buying individual naming rights to datasets, creation of philanthropic consortia, public-private partnerships, government investment in data moonshots
- Replace single-drug trials with newer, adaptive platform trials (basket/umbrella/ master protocol) testing multiple drugs simultaneously in biomarker-selected populations
- Establish a non-profit pharma accelerator to push scientifically credible but commercially unprofitable drugs (rare cancers, pediatric, off-patent combinations) through FDA approval

Improving Research & Treatment

- Continue innovation in radiation oncology (proton beam therapy, MR-guided adaptive radiotherapy, FLASH radiotherapy, etc) & surgical oncology (real-time margin assessment, fluorescence-guided imaging, robotics) while increasing supply of services and democratizing access to treatments
- Develop a robust supply chain for existing generic chemotherapy drugs and invest in novel methods of manufacturing cutting edge treatments to bring down costs (Allogenic CAR-T to bring cost from $500k to $10-20k, personalized mRNA cancer vaccines)
- Resource innovation in promising areas of drug development (PROTACs, ADCs, bispecific T cell engagers, next generation checkpoint inhibitors, tumor microenvironment targeting, radiopharmaceuticals, etc)
- Increase access to functional profiling and testing (patient derived organoids and tumor xenografts)
- Launch initiatives like Human Genome Projects to study the proteome, virome, microbiome, metabolome and immunome in oncology to help unlock precision medicine
- Resource survivorship medicine (study of cancer survivors) and also convergence between cancer biology and aging biology

Metascience & Institutional Reform

- Restructure NIH toward an ARPA-H model with empowered program officers funding high-risk, unconventional bets that peer review currently filters out & create high-novelty funding tracks
- Address the preclinical replication crisis through mandatory electronic laboratory notebook use for federally funded research, pre-registration, funded replication studies, and incentives for publishing null results
- Implement outcome-linked reimbursement for cancer therapies: link drug payments to confirmed long-term remission and quality-of-life outcomes rather than per-dose volume. Align pharmaceutical incentives with patient benefit rather than treatment duration
- Dedicated FDA Oncology Approval Track with conditional approvals, mandatory confirmatory trials, and fast-tracked withdrawal when those trials fail
- Address financial toxicity directly: Oncology accounts for only 7% of total healthcare spending. Pursue cost-effectiveness frameworks and drug price negotiation, as well as broader systemic reform to ensure citizens do not go bankrupt for cancer treatment.

Step 3: Tackle the main blockers to medical progress that are limiting AI tools

Despite there being promising areas within oncology and AI tools being built - both will fail to meet their full potential when operating in current structural constraints that quietly determine their ceiling. Success in these domains is contingent on solving deeper data and structural problems. In addition to asking what we should build, we need to ask what is stopping us from building it. An attempt to highlight the top blockers limiting. AI's utility in advancing cures for disease and medical progress more generally are:

Data Creation

1. **How do we create data to build a real-time map of what healthy actually looks like?**

Reliable, high-resolution definition of disease, requires reliable high-resolution definition of health. Until we know what deviations in biology are normal and which drive disease, we cannot tease the signal of disease from that of health. We need funding and coordination for more population scale, longitudinal, multiomic baselines for human health across age, sex, ethnicities, geographies, and environments. In parallel, we also need to push the frontier of clinical measurement capabilities and integration of established methods into clinical practice.

2. **How do we create data to compress biological timescales across health and diseases? How do we know something is working in an individual before the disease plays out?**

This is the highest-leverage and most valuable technical problem in medicine. If we can validate surrogate biomarkers and real-time biological readouts, we could both collapse decade-long trials into years, making drug development for prevention and early treatment financially viable. It would also enable individualized medicine, at scale. This will require at least a single cycle of human disease to validate, but will position us well for the future. Further, even without full validation, construction of well-founded surrogate biomarkers is better than the vacuum of information guiding treatment today. Confidence in a surrogate can be built incrementally, based on biological plausibility and retrospective correlation with hard endpoints, even before prospective validation is complete.

3. **How do we generate data to build preclinical models that actually predict human outcomes?**

Cells and mice fail us. The hard limit on experimentation is our inability to study human biology reliably without directly studying humans or their primary cells and tissues, but we are still far short of that limit. Progress starts with optimizing experimental capabilities for human-derived organoids, patient-derived xenografts, and Phase 0 microdosing trials, all of which dramatically improve predictive value over standard cell and mouse models. Ultimately, innovating in and scaling our capacity to measure human biology is the deeper solution, explored in Question 1.

4. **How do we build AI models trained on comparable data rather than publication bias? How can we capture and integrate negative results, failed experiments, and unpublished data?**

Essentially this is solving the garbage in, garbage out problem in distinct ways. The first is capturing what's missing: negative results, failed experiments, and unpublished data that currently disappear into filing cabinets. Mandated electronic laboratory notebooks with enforcement for all federally-funded research is one concrete mechanism. The second is making existing data comparable across labs and experimental iterations, which requires standardization of reagents, protocols, batch correction, and reporting. The field has pursued solutions for decades with limited success, suggesting both technical and incentive structure difficulties. Tackling both begins with rewarding meaningful data generation, not just positive results.

Economics and Incentives

5. **How do we make preventing or curing disease more profitable than treating it chronically?**

Models such as value-based care and outcome-based pricing take on new meaning in the age of AI, which can help to democratize data collection, and improve both transparency and analysis of patient outcomes. Further, political will for payment reform and elimination of administrative waste with AI could generate savings that can be reinvested in experimenting with different reimbursement models. As discussed explored in the Why Now is the Right Time section, AI-enabled outcome tracking, the fiscal crisis forcing change, insurance death spirals, hospital bankruptcies, and an imminent patent cliff create a unique window for experimenting with new models.

6. **How do we fund the therapeutic development that markets won't fund but need to succeed?**

Markets reliably fail to deliver potentially life-saving treatments to patients from antibiotics, to rare diseases. The antibiotic crisis showed that philanthropy can step in when markets fail, but philanthropy alone isn't sustainable. The more durable answer lies in changing the economics. As AI and next-generation measurement techniques like multiomics bring down the cost of clinical trials (and potentially enable n=1 studies) the financial barrier to researching commercially unviable treatments falls significantly. This opens the door to a broader funding ecosystem: government grants, ARPA-H style high-risk bets, single-payer systems with different incentive structures, and philanthropic endowment models designed to generate returns that sustain the work beyond the initial gift. Decentralized science, community-governed research funding coordinated through digital infrastructure (such as Decentralized Autonomous Organisations), is an emerging addition to this ecosystem and improving the economics of clinical research may finally unlock its ability to scale.

7. **How do we incentivize genuinely disruptive science rather than incrementalism?**

Metascience, the application of scientific methods to improve science itself, has established clearly that peer review and grant funding, systematically reward incremental work by established institutions over high-variance bets by unconventional thinkers. ARPA-H represents the most promising institutional attempt to correct this, modeling the risk tolerance of venture capital with the rigor of science. Beyond funding, AI offers a genuinely novel lever: tools that systematically surface underexplored hypotheses, flag assumptions baked into dominant paradigms, and identify promising contrarian findings buried in the literature. While the default development of AI is automation of creativity, alternative development pathways exist to augment it. AI tools could help counteract the institutional gravity that pulls science toward safe, fundable, publishable questions. The remaining challenge is downstream: disruptive findings still face the same peer review gatekeepers that filtered them out in the first place, suggesting that funding reform and publication reform need to advance together.

Institution, Systems and Coordination

8. How do we modernize the FDA to build a regulatory framework that is personalized, not population-averaged?

The 20th century FDA was designed for a simple world of one disease, one drug, one mechanism, one target, one population-averaged result. That world no longer reflects the state of biology. As diseases fracture into molecularly distinct subtypes, a drug that works remarkably well for 30% of patients may fail its population-level endpoint, or pass it while delivering negligible benefit to the majority. Either outcome highlights a regulatory framework mismatched to science. FDA modernization requires at least three shifts:

1) from binary (yes/no) to conditional approval with mandatory real-world evidence collection 2) from population-averaged endpoints to adaptive enrichment designs that identify responder subgroups during trials rather than after 3) from generic disease categories to molecularly defined indications

None of these are radical, and in fact, versions of each exist within current FDA authorities but are chronically underused. The harder question, addressed below, is evidentiary: personalized approval requires personalized evidence, and generating that without massive trial sizes remains an unsolved problem modernization must confront honestly rather than ignore.

9. How can we accelerate development without compromising safety?

The tension between acceleration and safety can be partially resolved by sequencing. The current model demands certainty before approval, but then learns relatively little afterward. A better model lowers the bar for initial approval based on early biological signal, but raises the bar for mandatory confirmatory trials with hard endpoints, real-world evidence collection, and rapid market withdrawal when confirmatory data fails. Accelerated FDA approval pathways already embody this logic, but weak confirmatory trial enforcement has undermined and eroded public trust. Regulators using the carrot more than the stick gives the appearance that "acceleration" may simply mean a faster way for drug-makers to profit. Public skepticism is reinforced by the revolving door between FDA reviewers and industry. There are real, unsolved challenges in the balancing act between development and safety. Shorter trials mean less power to detect rare adverse events which can only be observed at scale. No approval model fully resolves this. Further, ultimately safety is a deeply personal question. Different people have different risk toler-

ances, yet are constrained by a single risk tolerance (FDAs). Respecting individual autonomy while protecting against corporate exploitation of patient desperation is the genuine design challenge.

10. **How do we prevent AI from supercharging the existing misalignments rather than fixing them?**

AI accelerates existing incentive structures. As we saw with AI scribes upcoding under the guise of improving the patient experience, even good intentions are insufficient protection when the underlying system rewards the wrong outcomes. AI's turbocharging ability accelerates the speed asymmetry where misalignments compound far faster than institutional reforms. Early governance and prophylactic reforms are vital wherever possible. As explored above, rapid technological disruption hitting a system in late-stage collapse creates genuine political and economic conditions for reform that have not existed before. Seizing this window requires us to get specific about what gets rebuilt and in whose interest, and what objective functions we choose to optimize a new system for. These decisions will determine if US healthcare accelerates off of a cliff, or if it begins a redemptive journey of repair.

An Invitation To Change

As stated in the introduction, this roadmap is an invitation. The conclusions reached here are deliberately incomplete. It is meant as a stimulus for discussion and feedback, not a comprehensive plan. Much of it may be misguided or wrong. What's missing from current discourse is not diagnosis of the problem, which is abundant, but solution-oriented generative thinking. Ideas are too often killed by criticism before they're properly examined, when iteration and collaborative troubleshooting would serve us better. Even if AI progress stopped today, the disruption already underway will force us to reimagine how we work, live, and organize our institutions.

Meeting the moment will require us to move beyond listing grievances into a mode of generative criticism, where the question is not why an idea is wrong rather what are the positive elements to amplify and identifying how it could be different or better. In that spirit, I welcome feedback on what warrants further development in this plan or how plans could be improved, better prioritized, or more complete. A follow-up piece will incorporate this input in the months ahead.

Sign up to receive updates at curecancer.ai/roadmap

AI vs Cancer

How AI Can, and Can't, Cure Cancer

View this document online:
curecancer.ai