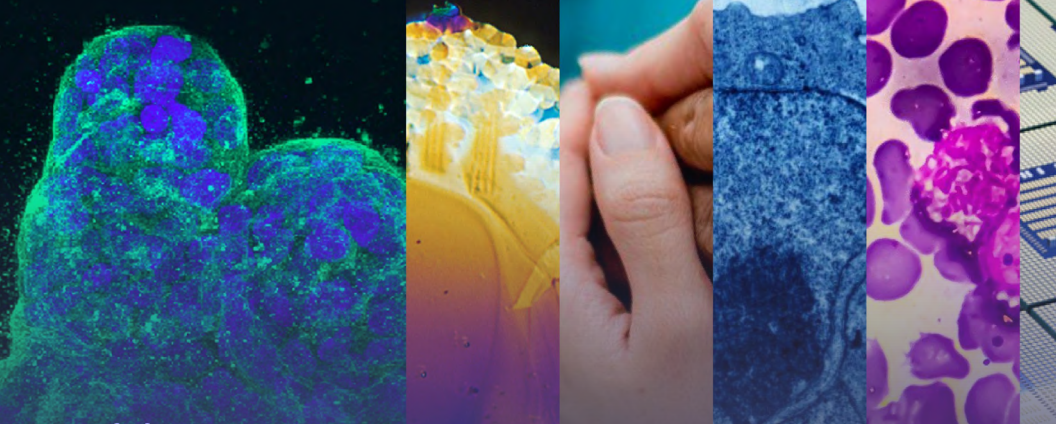

AI vs Cancer

How AI Can, and Can't, Cure Cancer

by **Emilia Javorsky** MD, MPH | curecancer.ai

The Problem: Intelligence is not the Bottleneck

Major AI companies promise that Superintelligence will cure cancer. Yet despite exponentially accelerating biomedical knowledge; cancer mortality and FDA drug approvals have remained largely flat. Big Tech's bold healthcare promises are nothing new: from IBM Watson to Google's Calico to 13 years of AI drug discovery startups, the pattern is consistent: transformative claims followed by underwhelming results, rooted in a fundamental misunderstanding of what actually blocks progress. Grand challenges are data-limited, regulation-limited, and incentive-limited — not intelligence-limited.

Why Now: A Once-in-a-Generation Reform Window

U.S. healthcare is 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, a $236 billion patent cliff looms by 2030, and healthcare already consumes 27% of federal spending. But a system in late-stage collapse also creates the conditions to rebuild. Polling shows 83% of Americans want substantial change to existing institutions. Collapsing healthcare economics, genuine political appetite for disruption, and AI tools capable of reducing friction and aligning incentives at scale create a once-in-a-generation window for biomedical innovation.

Roadmap: How to Cure Cancer

Step 1: Resource and scale AI tools already making progress

Current AI — applied by domain experts to specific, well-defined problems with quality data — already delivers real value in oncology. The builders in each category, not Big Tech's general-purpose genies, deserve our resources:

AI Drug Discovery & Target ID
Finding disease targets and designing drugs faster than traditional lab methods

AI In Silico Modeling & Virtual Cells
Building digital replicas of cells and biological systems to simulate experiments

AI Genomics
Reading and analyzing genetic data at scale to uncover cancer drivers

AI Early Detection & Liquid Biopsies
Processing blood, scans, and tissue to detect cancer earlier

AI Proteomics & Biomarkers
Mapping proteins to surface early or hidden signals of disease and treatment response

Clinical Trials & Regulatory AI
Accelerating & redesigning how trials are run, analyzed, and approved

AI Drug Repurposing
Mining existing approved drugs & natural compounds for potential against cancer

Clinical & Surgical AI
Guiding decisions in the clinic and OR to improve precision and personalize care

AI Toxicity Prediction
Predicting whether a new compound will be toxic in humans in advance

AI Policy & Advocacy
Pushing AI-driven advances into policy, funding, and clinical adoption

Step 2: Double down on promising areas of oncology

- **Early Detection & Prevention:** Scale multicancer early detection blood tests; expand proven screening (mammography, colonoscopy, low-dose CT); invest in real-time minimal residual disease monitoring.
- **Data & Clinical Trials:** Build a national cancer data commons linking genomic, imaging, and outcomes data; replace single-drug trials with adaptive platform trials (basket/umbrella protocols) in biomarker-selected populations; establish a non-profit pharma accelerator for commercially unprofitable but scientifically credible drugs.
- **Research & Treatment:** Drive down costs through allogeneic CAR-T (~$500K → ~$10–20K) and personalized mRNA cancer vaccines; resource promising modalities (PROTACs, ADCs, bispecific T-cell engagers, radiopharmaceuticals); launch "-omics moonshots" to map the proteome, microbiome, virome, and immunome for precision oncology.
- **Metascience & Institutional Reform:** Scale ARPA-H model funding high-risk, unconventional bets; implement outcome-linked reimbursement tying drug payments to remission rather than dose volume; create a dedicated FDA oncology approval track with conditional approvals and mandatory confirmatory trials; address financial toxicity so cancer treatment doesn't bankrupt patients.

Step 3: Tackle top 10 blockers to medical progress

1. Map Health Baselines
Fund longitudinal multi-omic population datasets across age, sex, ethnicity & geography. Can't define disease without a high-resolution definition of health.

2. Compress Timescales
Validate surrogate biomarkers for real-time biological readouts: the highest-leverage technical problem in medicine. Collapse decade-long trials into years.

3. Better Preclinical Models
97% of mouse cures fail in humans. Scale organoids, patient-derived xenografts & Phase 0 microdosing. We are far short of the limit of human-derived models.

4. Fix Training Data
Mandate e-lab notebooks; reward negative results; standardize across labs. AI trained on today's literature learns corruption as confidently as truth.

5. Make Cures Profitable
Outcome-based pricing + value-based care. Use the political disruption window and fiscal crisis to pilot new models before the system fully collapses.

6. Fund What Markets Won't
NIH funded datasets, ARPA-H bets, philanthropic engagement, data naming rights, DAOs, and AI-lowered trial costs to make off-patent drugs viable to develop.

7. Reward Disruptive Science
ARPA-H program officers + AI tools to surface contrarian hypotheses. Funding and publication reform must advance together to break the "amyloid mafia" pattern.

8. Modernize the FDA
Binary yes/no → conditional approval. Generic categories → molecular indications. Population endpoints → adaptive enrichment. Versions of all three exist but are chronically underused.

9. Accelerate Safely
Lower bar for initial approval; raise bar for confirmatory trials with fast-withdrawal. Fix weak enforcement and the revolving door that erodes public trust.

10. Prevent AI Amplifying Harm
Early governance is essential. AI accelerates whatever incentives exist. The window for prophylactic reform is now, before misalignments compound faster than institutions can reform.

Help Build the Roadmap

If you're a researcher, clinician, technologist, policymaker, patient advocacy organization, philanthropist, or other stakeholder with a perspective to share, we'd love to hear from you. Your input will help us develop a practical roadmap for AI in cancer research.

Sign up to access a private feedback survey and receive invitations to curated workshops where we'll bring the roadmap to life.

curecancer.ai/roadmap

Why Superintelligence Faces Hard Limits in Biology

No First Principles

Unlike chess or physics, biology has no clean ruleset. Human biology is evolved, not engineered; with emergent properties at each layer of biological complexity. Simulating a single human for a week with Earth-covering GPUs would take the age of the universe. Biological truth cannot be computed. It can only be measured. Right now, we aren't measuring human biology at resolution or scale.

Data Deserts & Replication Crisis

If there are no first principles, AI needs high-quality data. Electronic Health Records were designed to maximize billing, not capture biology. Up to 70% of biomedical findings fail to replicate. An AI trained on this corpus doesn't learn biology: it learns publication bias and institutional politics, expressed with complete confidence.

Cannot Compress Time

You cannot speed up a pregnancy with more engineers. Questions are answered on the timescale of disease, not compute. The COVID vaccine's nine-month timeline is routinely cited as proof otherwise, but COVID takes two weeks, chronic diseases take decades.

> *Cancer patients cannot wait for genies when the real barriers are human-made systems we could reform today. Current AI already helps — the question is whether we will do the hard systemic work required.*

Broken Incentives & Market Failures

A Goldman Sachs analyst asked: "Is curing patients a sustainable business model?" In the current system, it isn't. $1mil+ gene therapies try to match the lifetime revenue of chronic treatments it replaces. Halicin, an AI-discovered antibiotic validated in 2020, still isn't available to patients not because the science failed, but because antibiotics are structurally unprofitable. More compute cannot fix a market that profits from sickness.

Cancer is Individual

The "Hallmarks of Cancer" grew from 6 to 14+ between 2000 and 2022, not because the original science was wrong, but because cancer kept revealing new complexity as measurement improved. Even within a single tumor, cells have different pathways gone awry. As our understanding deepens, cancer gets more complex, not less.

Regulatory Barriers

The FDA's efficacy framework was set in 1962 designed for one disease, one drug, one population-averaged result. It cannot accommodate multi-target therapies, molecularly defined subtypes, or treatments that work through incompletely understood mechanisms. A drug that works remarkably well in 30% of patients may fail its population-level endpoint and never reach the people it would cure.