Techstrong
SPECIAL REPORT
A FUTURUM COMPANY

SEPTEMBER 2026

MEDICINE / DISCOVERY / WELLNESS

AI and the Promise of Better Health

From Discovery to the Patient

BY ALAN SHIMEL

The possibilities are extraordinary.
The measure is better lives.

Inside the promise

A journey from scientific possibility to evidence, access and better health.

ALSO IN THIS EDITION

01 / THE REPORT

Why we are hearing so much about AI and health

There seems to be a lot of coverage lately about the role AI could play in healthcare. Stories about helping doctors diagnose disease appear alongside reports of biological discoveries, new approaches to drug development and predictions about extending human life. Read enough of them and you begin to wonder whether we are approaching a different era of medicine or simply experiencing the next phase of the AI sales pitch.

I think that is worth a deeper look, because those possibilities are not mutually exclusive. Companies can have something to sell and still be doing important work. Scientists can make a meaningful discovery without having a treatment ready for patients. And a prediction can be worth considering without being a promise anyone knows how to keep.

The stories themselves cover very different territory. Ynet's reporting on NVIDIA's Israeli research effort describes an investment in scientific discovery and collaboration. Anthropic's announcement about Claude and a newly characterized enzyme system concerns an early biological finding. Neither is an announcement that patients can now receive a new cure. [1,2]

Yet both contribute to a larger narrative: AI may do more than help us manage the healthcare system we have. It might help us develop medicine we do not have.

We recently examined what AI could mean for doctors, including why making a physician more productive does not automatically make care better or more affordable. That distinction still matters. But concentrating entirely on the consultation room leaves out the much bigger ambition now attracting attention.

Could AI help scientists understand diseases that remain poorly understood? Could it make the search for treatments more productive? Could it make individualized medicine practical for more people, or help us spend more of our lives healthy and independent?

These are powerful propositions. Most of us do not need a market forecast to understand their appeal. We have people we worry about, conditions we live with and treatments we wish existed. An industry that attaches itself to those hopes has found a compelling argument for investment and public support.

That is a reason to examine the claims, not dismiss them. Rejecting an important discovery because its sponsor sells AI would be as foolish as accepting a medical forecast because a successful technology executive delivered it.

For this report, the question is what connects the promise to the patient. That requires looking at the future being proposed, the research that might get us there and the decisions that determine who benefits. The number of announcements is not the measure. What medicine becomes able to do, and what people become able to obtain, is.

1. Ynet / Calcalist, NVIDIA research group. https://www.ynetnews.com/tech-and-digital/article/bjfsyrl9fx

2. Anthropic, ART announcement. https://www.anthropic.com/news/claude-discovers-novel-enzyme-system

02 / THE REPORT

Amodei's big bet on biology

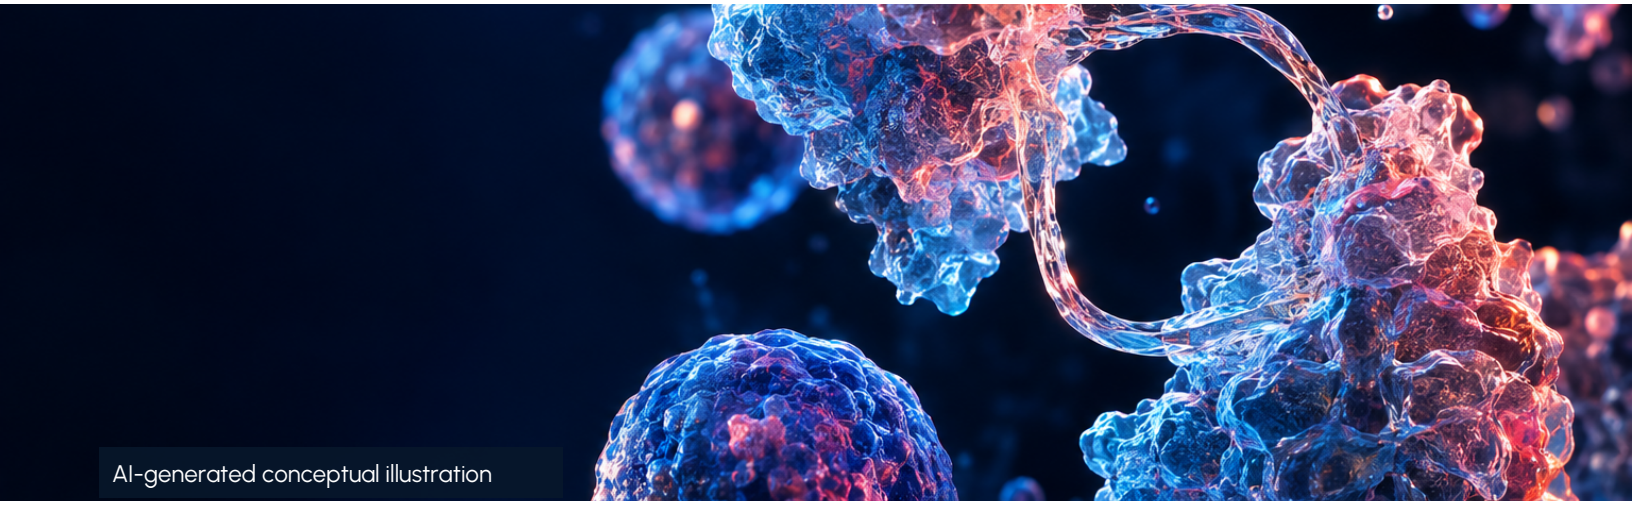

AI-generated conceptual illustration

Anthropic CEO Dario Amodei offers one of the most ambitious statements of the case in Machines of Loving Grace. His October 2024 essay imagines powerful AI helping compress 50 to 100 years of biological progress into five to ten years. That clock begins after systems with the capabilities he describes become available, not when the essay was published. [3]

His vision extends beyond software that answers research questions. It includes systems that help devise experiments, develop tools and carry out sustained scientific work, with many efforts proceeding in parallel. The proposed benefits reach into major diseases, mental health and longer healthy lives. He also acknowledges constraints, including experimental time, missing data and biological complexity. These are explicitly uncertain projections, not clinical findings.

There is a serious argument here, even if you do not accept the schedule. The value of intelligence in medicine need not stop at interpreting information after someone else has collected it. It could influence which questions get asked and which evidence gets collected next.

That is the part I would focus on. A system that summarizes a hundred papers can save a researcher time. A system that helps identify the experiment most likely to resolve an important uncertainty could change the direction of a research program. Both would be useful, but they are different contributions.

3. Dario Amodei, Machines of Loving Grace. https://darioamodei.com/essay/machines-of-loving-grace

02 / THE REPORT

Amodei's big bet on biology

My test for this vision would start there, with decisions rather than descriptions. Does the assistance lead scientists to better experiments? Does it help them abandon weak ideas earlier? Does it expose assumptions that would otherwise survive until an expensive failure? Can another team reproduce the advantage without depending on the original developers to make everything work?

Those questions provide a way to evaluate progress without having to settle an argument about superintelligence first. They also make room for significant gains that fall short of the most ambitious forecast. We should not define success so narrowly that only an enormous leap counts. Preventing years of wasted effort on a poor target could matter, even if nobody announces the end of disease.

There is another distinction I would make. A better research process can increase the opportunity to discover something important without guaranteeing what the discovery will be. We should be careful about sliding from confidence in a tool to confidence in a particular medical outcome. Funding a more capable scientific effort is not the same as placing an order for a cure.

The business questions are equally relevant. Who can afford to run these systems? Who owns the results? Can academic and public-interest researchers use the same capabilities as well-funded pharmaceutical companies? If a model's most consequential contribution is choosing what to investigate, the priorities built into that choice deserve attention.

I would also want to know how much useful disagreement remains in the process. Generating many candidate ideas is valuable only if the process can discriminate among them. An impressive volume of mutually reinforcing explanations could create confidence without creating knowledge. The experimental evidence has to retain the power to overturn the system's preferred answer.

None of that requires us to start from pessimism. It means giving a large claim an appropriately demanding examination. The question is not whether the forecast sounds exciting. It is whether the individual capabilities needed to support it are emerging, whether they work together and whether their benefits survive outside carefully selected demonstrations.

For readers trying to make sense of the coverage, that is a more useful place to stand than choosing between believing every prediction and declaring the whole exercise hype.

"Funding a more capable scientific effort is not the same as placing an order for a cure."

03 / THE REPORT

What the promise could mean for our lives

AI-generated conceptual illustration

The easiest way to lose this discussion is to organize it around the companies. People do not experience health as a competition between AI platforms. They experience it as something they can do, something they can no longer do or something they fear losing.

Start with prevention. The attractive possibility is not simply that a system detects more abnormalities. It is that useful information arrives while there is still something effective to do about it. A warning without a reliable interpretation or an available intervention can leave a person with more anxiety and no better options.

That gives us a practical standard for evaluating predictive medicine. What decision changes because of the prediction? Does making that decision improve the person's health? What happens to people who receive an incorrect warning, and what happens to those falsely reassured? Earlier information is valuable when it supports better action, not merely because it arrives earlier.

Then consider treatment. A more productive search for medicines could help researchers find an intervention where existing approaches have failed. It could also help them understand why a treatment works for one group and not another. We should leave room for both kinds of progress. A more reliable way to select an existing treatment could be deeply consequential without producing a new blockbuster drug.

Personalized medicine takes that ambition further. In 2025, Penn Medicine and Children's Hospital of Philadelphia described a personalized gene-editing treatment for an infant with CPS1 deficiency, a rare metabolic disorder. Their account reported encouraging early observations following treatment, with continued monitoring necessary. The work built on a much longer foundation of scientific development. [4]

4. CHOP, personalized gene-editing treatment.
https://www.chop.edu/news/worlds-first-patient-treated-personalized-crispr-gene-editing-therapy-childrens-hospital

03 / THE REPORT

What the promise could mean for our lives

That was not an AI-designed cure, and it should not be enlisted as one. Its relevance is the question it raises about the future: could tools that improve design, testing and interpretation help make highly individualized treatments more repeatable? Could what takes an exceptional effort become possible for more patients?

Those are questions, not results from that case. They are nevertheless worth asking. If the public benefit of a technology depends partly on making difficult work practical, then exceptional medical achievements can help identify the work we should try to make less exceptional.

Longevity brings an even larger ambition, along with a considerable risk of confusing a measurement with an outcome. There is a difference between a person living longer, a person remaining healthy longer and a laboratory model assigning that person a younger biological age.

A 2026 Nature Biotechnology paper illustrates the distinction. Researchers applied six protein-based aging clocks to samples from 42 participants in a short rentosertib trial and reported shifts toward younger predicted profiles. They also acknowledged that those clocks could not, on their own, distinguish effects on aging from effects on disease. The study did not establish that participants would live longer or that healthy people would benefit. [5,6]

There can still be scientific value in such measurements. The problem comes when the implication grows larger than the observation. For me, the most useful longevity question is whether an intervention preserves function, independence and well-being. A changed score might help investigate that question. It cannot answer the whole thing by itself.

Finally, there is the promise of reaching people who are not well served now. Better tools might make research on some overlooked conditions more feasible or help scarce clinical expertise support more communities. But we should ask who chooses those applications. A reduction in the cost of research does not automatically determine which diseases receive attention.

The future worth pursuing includes more than unusually sophisticated medicine for unusually well-resourced patients. It includes making useful knowledge and effective treatment available where they are missing.

That is a demanding vision. It is also a better account of the opportunity than a future in which everyone has a more articulate health chatbot but the same obstacles to receiving care. Answering questions is useful. Changing what can be done about the answers is the larger prize.

5. Nature Biotechnology, proteomic aging clocks. https://www.nature.com/articles/s41587-026-03286-y

6. NAD.com, aging-clock analysis and limitations. https://www.nad.com/news/aging-clocks-applied-to-human-trial-data-support-ai-designed-lung-drug-lowering-age

04 / THE REPORT

How AI could change the process of discovery

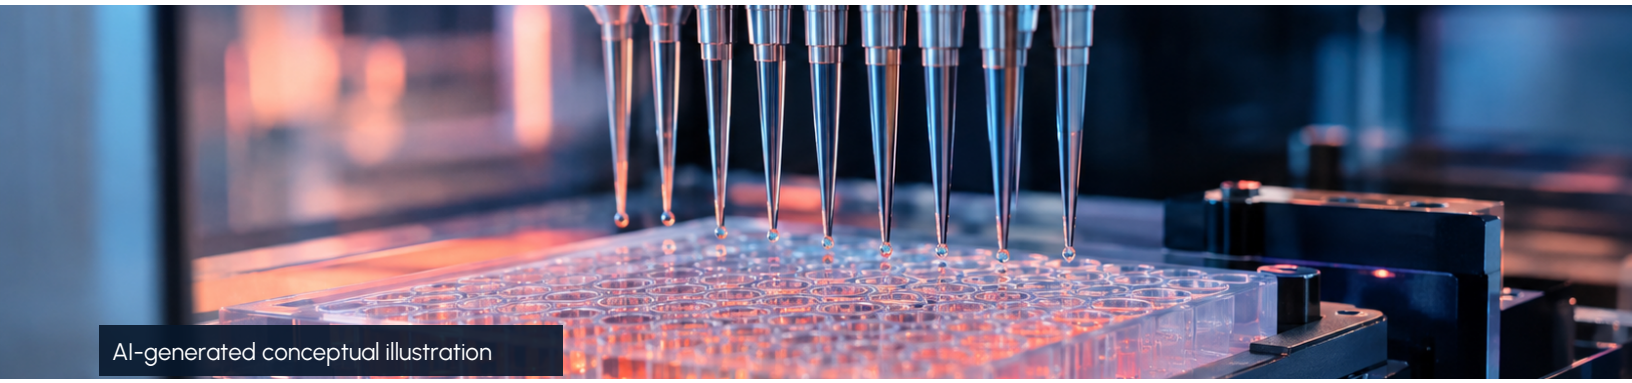

AI-generated conceptual illustration

To understand whether that vision has substance, we need to look inside the research process. A scientific project does not begin with a fully defined problem and end when a computer produces an answer. Researchers decide what to investigate, work out how to test it, interpret imperfect results and decide what to do next.

The potential contribution of AI is different at each stage. Searching literature is not the same as designing a useful experiment. Designing an experiment is not the same as conducting it reliably. And a result that agrees with a hypothesis is not necessarily enough to establish that the hypothesis is correct.

There is now published work addressing more of that cycle. A Nature paper on Google's AI co-scientist describes a multi-agent approach to generating, evaluating and refining hypotheses. Selected drug-repurposing and combination proposals related to acute myeloid leukemia were tested in laboratory experiments. Those tests were not evidence of successful leukemia treatment in patients, but they moved beyond judging the quality of generated text. [7]

Another Nature paper describes FutureHouse's Robin, a semi-autonomous research system connecting literature work, experimental proposals and analysis. Its research included laboratory investigation of candidates relevant to dry age-related macular degeneration. This was not a demonstration that AI had cured blindness. It was an attempt to evaluate assistance across connected research activities. [8]

The connection is important. An assistant that performs one step well may still create work elsewhere. Researchers need to know whether the whole process becomes more useful, not just whether an isolated task becomes faster. If reviewing the output consumes the time saved in generating it, a favorable benchmark may tell us little about research productivity.

7. Nature, Co-Scientist. https://www.nature.com/articles/s41586-026-10644-y

8. Nature Portfolio, Co-Scientist and Robin. https://www.natureasia.com/en/info/press-releases/detail/9330

04 / THE REPORT

How AI could change the process of discovery

Physical experimentation adds another dimension. A Nature Communications study of automated enzyme engineering combined computational tools with a biofoundry, an automated experimental facility. The researchers reported improved enzyme activity through four rounds of work over four weeks. The process began with defined proteins and measurable experimental objectives, rather than an open-ended instruction to solve biology. [9]

That boundary is part of what makes the example useful. We can ask what was supplied, what the system changed and how success was measured. A constrained loop can be a real achievement without demonstrating general scientific autonomy.

Claude's enzyme-system announcement belongs in this part of the story, too. Its importance, if the work stands up to further scrutiny, lies in AI helping identify something worth investigating. The sidebar explains the reported finding and its limits. An unresolved biological function should make us curious, not tempt us to fill in the blank with a medical application.

My interpretation of these efforts is that the most consequential opportunity may be a shorter, more productive loop between an idea and evidence about that idea. That is a narrower claim than an AI scientist capable of everything, but it is a substantial one.

It also changes what I would ask a developer to demonstrate. Show the unsuccessful candidates as well as the successful one. Explain the human work required. Compare the process with a serious alternative, not an artificially weak baseline. Tell us whether a researcher learned something useful sooner, including when the answer was that an attractive idea did not work.

Those requirements are not obstacles to recognizing progress. They are how we distinguish a reusable capability from an impressive story. A result dependent on exceptional human intervention might still be valuable, but the intervention belongs in the account of how it was achieved.

A recent Nature Reviews Drug Discovery perspective emphasizes the importance of context, problem definition and clinical relevance in applying AI to drug discovery. That is a useful counterweight to evaluating tools primarily by prediction scores. The decision being improved matters as much as the score attached to it. [10]

For the technology industry, this should be familiar territory. A component's performance is not the same as the performance of the system in which it operates. In medical research, the system includes experiments, people and biological uncertainty. We should want AI to improve that entire process, even when doing so produces a less spectacular headline than announcing another model milestone.

9. OSTI, automated enzyme engineering. https://www.osti.gov/biblio/2574786

10. Andreas Bender, drug-discovery perspective. https://www.drugdiscovery.net/2026/08/07/artificial-intelligence-in-drug-discovery-what-does-it-mean-and-where-do-we-really-stand/

FIELD NOTES / DISCOVERY

What Did Claude Actually Discover?

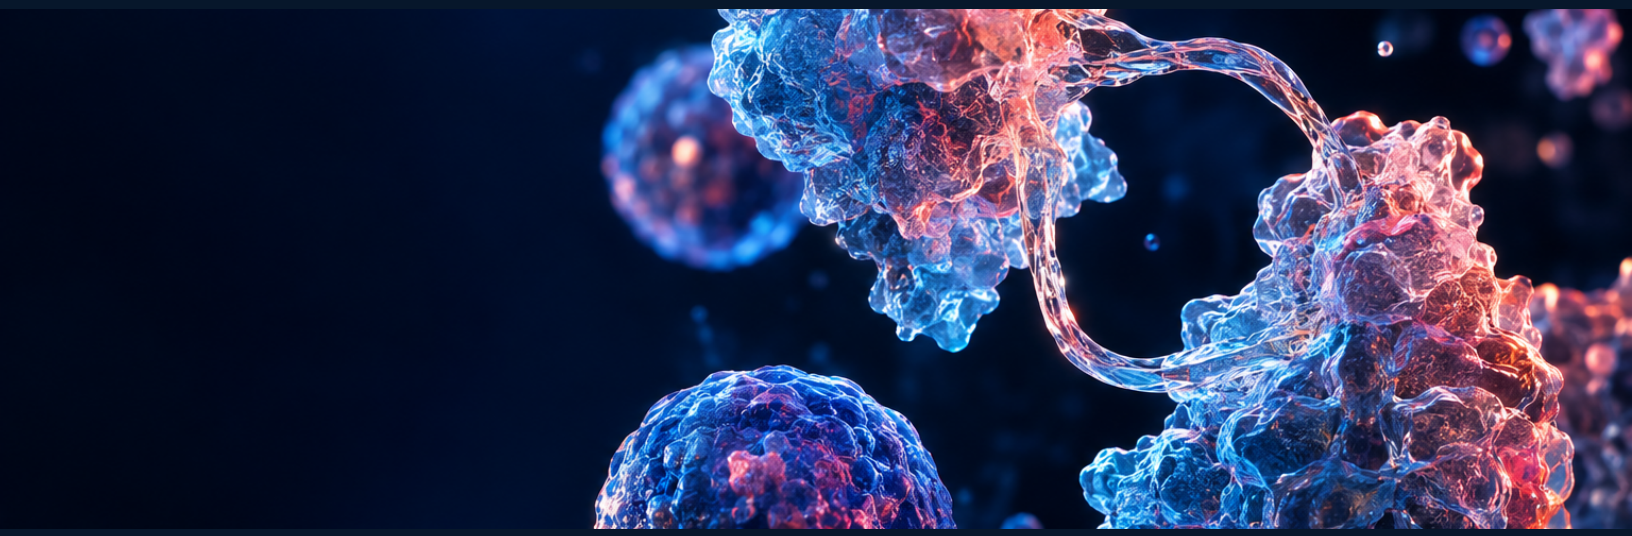

Conceptual illustration, not a depiction of the ART structure.

Anthropic's September 23 announcement describes array-associated reverse transcriptases, or ART, in bacteriophages, viruses that infect bacteria. The reverse transcriptase itself had appeared in previous research. The claimed advance concerns the associated system, including repeat sequences and an accessory protein.[2]

According to Anthropic, Claude contributed to identifying the system, and human scientists performed laboratory experiments. Early work found short RNAs produced by the array. Its biological function remains unknown.[2]

That makes this an early discovery story, not a demonstration of a programmable gene editor or treatment. Features that invite comparison with CRISPR do not establish equivalent capabilities.

The announcement and accompanying preprint should be read as a basis for further investigation, not evidence that Claude has produced a CRISPR replacement.

For readers, the useful next questions concern confirmation, function and practical capability. Can other researchers reproduce the findings? What does the system do? Can it be controlled for a useful purpose? Does a proposed application work outside the conditions in which it was first observed?

Those questions leave plenty of room for the discovery to become important. They also prevent a familiar leap in technology coverage: moving from a suggestive resemblance to an assumed capability before the necessary experiments have been done.

2. Anthropic, ART announcement. https://www.anthropic.com/news/claude-discovers-novel-enzyme-system

05 / THE REPORT

The three clocks of medical progress

I find it useful to think about three clocks running between a promising idea and better health: the discovery clock, the evidence clock and the delivery clock. This is a way to organize the questions, not a scientific law about how medicine develops.

DISCOVERY

Identify something worth pursuing.

Better searches, experiments and analysis.

EVIDENCE

Establish what works, for whom and at what risk.

Meaningful comparisons and sufficient follow-up.

DELIVERY

Make an effective intervention reachable.

Manufacturing, clinical capacity, payment and access.

The discovery clock measures how long it takes to identify something worth pursuing. Better searches, experiments and analysis could shorten that time. The evidence clock measures how long it takes to establish what an intervention does, for whom and at what risk. The delivery clock measures the distance between an effective intervention and someone actually receiving it.

Conceptual framework by Alan Shimel. Not to scale; no durations implied.

05 / THE REPORT

The three clocks of medical progress

The clocks overlap. Manufacturing work can proceed while trials continue, and information from practice can lead back to new research. But faster movement in one part does not guarantee faster movement everywhere.

Rentosertib provides a useful example. The candidate for idiopathic pulmonary fibrosis was developed with AI contributing to target identification and molecule design. Its published Phase 2a trial involved 71 patients over 12 weeks. The highest-dose group had a mean increase in a lung-function measure, forced vital capacity, while the placebo group had a mean decline. The small groups and short follow-up limited what could be concluded about durable benefit. [11]

On September 9, 2026, Insilico Medicine announced the first patient had been dosed in its Phase III trial. That is an important development milestone. It is a company-reported start to a larger test, not an announcement that the drug has passed it. [12]

A reader should be able to recognize the progress without being asked to assume the outcome. Entering a more demanding trial is exactly what a promising development program should do. The next evidence matters because the earlier evidence cannot settle the question.

Could AI help improve trial design, recruitment or the interpretation of data? Those are reasonable applications to investigate. But a model's confidence about a future outcome is not a substitute for observing the outcome. If we need to know whether a benefit lasts, the passage of time can be part of the evidence rather than merely an administrative delay.

Conversely, we should not treat every delay as medically necessary. Repetitive paperwork, poor coordination and avoidable recruitment problems deserve examination. The useful distinction is between making evaluation more efficient and weakening what the evaluation must establish.

Regulators belong in that conversation as institutions responsible for evidence and safety, not as stock villains standing between innovators and patients. The FDA's generative-AI discussion paper raises questions about evaluation and monitoring. It is a discussion document, not final guidance or a new set of binding rules. [13]

The delivery clock demands different questions. Who can manufacture the intervention consistently? Which facilities can provide it? What training is needed? Who pays, and what does the patient have to do to qualify? None of those questions disappears because discovery involved AI.

This is why I would be wary of a single claim about how much faster medicine will move. We need to know which clock is being accelerated. A faster discovery process is worth having. Its public value grows when evidence and delivery receive comparable attention.

11. Insilico Medicine, Phase IIa results. https://insilico.com/news/tnrecuxsc1-insilico-announces-nature-medicine-publi

12. Insilico Medicine, Phase III announcement. https://www.prnewswire.com/news-releases/insilico-medicine-doses-first-patient-in-genesis-ipf-3-the-worlds-first-phase-iii-trial-of-a-generative-ai-driven-innovative-drug-302873749.html

13. FDA, GenAI medical-device discussion paper. https://www.fda.gov/news-events/press-announcements/fda-seeks-public-feedback-inform-regulatory-approach-generative-ai-enabled-medical-devices

FIELD NOTES / TRANSLATION

From a Promising Molecule to a Medicine

An AI-assisted drug candidate still needs evidence that it is a useful medicine. The main stages answer different questions, and none guarantees success at the next.

1

Discovery and laboratory work

Researchers identify a target, develop candidate molecules and investigate whether they have the intended effects. A favorable prediction must survive experimental testing.

2

Early human studies

Testing investigates safety, dosing and how the candidate behaves in people. Initial evidence of benefit can help determine whether larger studies are justified.

3

Confirmatory trials

More demanding comparisons evaluate benefit and risk in the intended population. The endpoints and duration determine which claims the results can support.

4

Regulatory evaluation and delivery

Evidence must support the proposed use. Manufacturing, clinical capacity, payment and access then determine whether patients can receive the intervention. Monitoring continues to matter after deployment.

Rentosertib illustrates the distinction between stages. Its small Phase 2a trial produced encouraging findings in patients with idiopathic pulmonary fibrosis. Insilico subsequently announced the first patient dosed in Phase III. Beginning that larger study is progress toward an answer, not the answer itself. [11,12]

When a headline says a drug is AI-designed, ask what AI contributed and which evidence stage the candidate has reached. The label describes part of its origin. It does not establish its clinical value.

11. Insilico Medicine, Phase IIa results. https://insilico.com/news/tnrecuxsc1-insilico-announces-nature-medicine-publi

12. Insilico Medicine, Phase III announcement. https://www.prnewswire.com/news-releases/insilico-medicine-doses-first-patient-in-genesis-ipf-3-the-worlds-first-phase-iii-trial-of-a-generative-ai-driven-innovative-drug-302873749.html

06 / THE REPORT

What has already crossed into patient care?

The future-facing argument should not obscure evidence from actual care. We have examples that are more informative than a model's score on medical questions, precisely because they test a defined intervention in a defined setting.

The Swedish MASAI breast-screening trial compared AI-supported screening with standard double reading. Its interval-cancer analysis examined cancers diagnosed between screening rounds. Rates were 1.55 versus 1.76 per 1,000 participants, respectively. The study established noninferiority for this endpoint: the AI-supported approach met the trial's specified standard for not being unacceptably worse. [14]

That is different from proving it reduced interval cancers by 12%. The uncertainty around the estimate included no difference. It also was not proof of reduced breast-cancer mortality. Those qualifications do not erase the result. They tell us what the result means.

The TAILORED-AF trial examined a cardiac-ablation strategy incorporating AI-identified treatment targets. Among 370 patients who underwent the procedures, the reported one-year freedom from documented atrial fibrillation was 88% for the tailored strategy versus 70% for conventional treatment. The broader endpoint covering any atrial arrhythmia after one procedure did not significantly differ, and procedure and ablation times were roughly doubled. [15]

Here, the intervention was a treatment strategy carried out by clinicians, not software operating independently. The primary result and the additional procedure time both belong in an account of its value. Neither should be hidden to make the story simpler.

These studies illustrate why asking whether AI works in medicine is too broad. We need to ask which intervention, for which patients, compared with what and measured how. A system can be useful for a specific purpose without earning a general license to perform other medical tasks.

The person using it is also part of the intervention. In Oxford researchers' study of 1,298 participants using hypothetical medical scenarios, access to the tested language models did not improve performance over usual information sources. The models did better when given the scenarios directly than when people interacted with them. These were simulated situations, not a measurement of clinical injury, but the distinction is important. [16]

It is tempting to assume that expertise demonstrated by a model becomes expertise available to its user. The study gives us a reason to test that assumption rather than build a service around it. Asking questions, providing relevant context and acting on an answer are all part of using a tool.

Taken together, the evidence supports a more precise discussion than either enthusiasm or dismissal supplies. Some applications can be evaluated in real workflows and show useful results. Other claims remain much further from patient benefit. We should be able to recognize both without demanding that every study decide the future of AI in medicine.

14. The Lancet / PubMed, MASAI trial. https://pubmed.ncbi.nlm.nih.gov/41620232/

15. Nature Medicine, TAILORED-AF trial. https://www.nature.com/articles/s41591-025-03517-w

16. Oxford research team, public-use study. https://oxrml.com/llm-medical-assistants/

EVIDENCE IN VIEW / 01

Breast screening: read the uncertainty

MASAI compared AI-supported screening with standard double reading. The endpoint was interval cancers, diagnosed between screening rounds.[14]

INTERVAL CANCERS PER 1,000 PARTICIPANTS

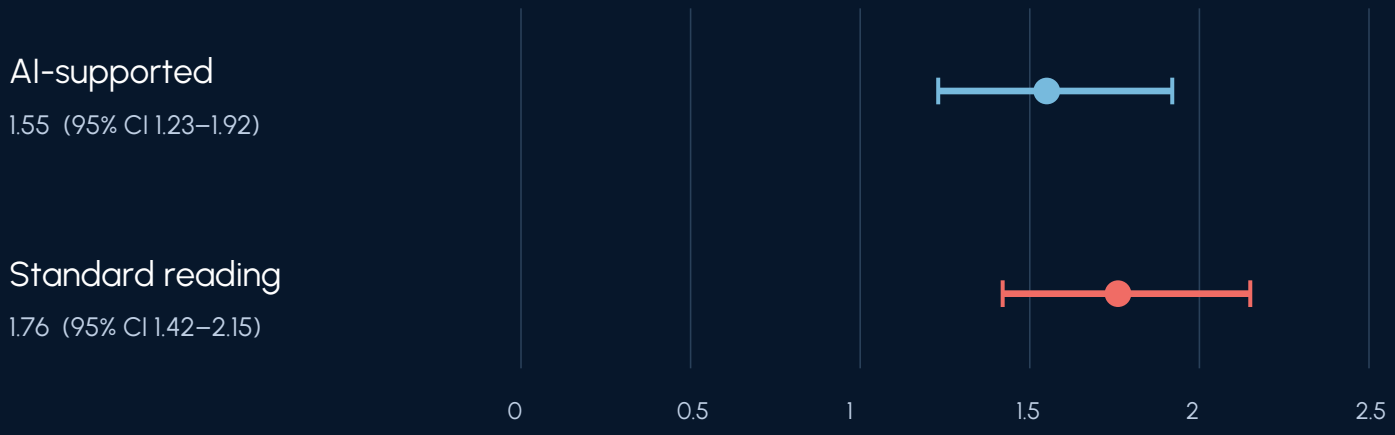

Dots show rates; whiskers show 95% confidence intervals.

Noninferior is not the same as proven superior.

Proportion ratio: 0.88 (95% CI 0.65–1.18). The uncertainty includes no difference. The trial met its noninferiority criterion; it did not establish a reduction in breast-cancer mortality.[14]

14. The Lancet / PubMed, MASAI trial. https://pubmed.ncbi.nlm.nih.gov/41620232/

EVIDENCE IN VIEW / 02

A better result. A longer procedure.

TAILORED-AF tested a clinician-delivered ablation strategy with AI-identified treatment targets, not autonomous software.[15]

FREEDOM FROM DOCUMENTED AF AT ONE YEAR

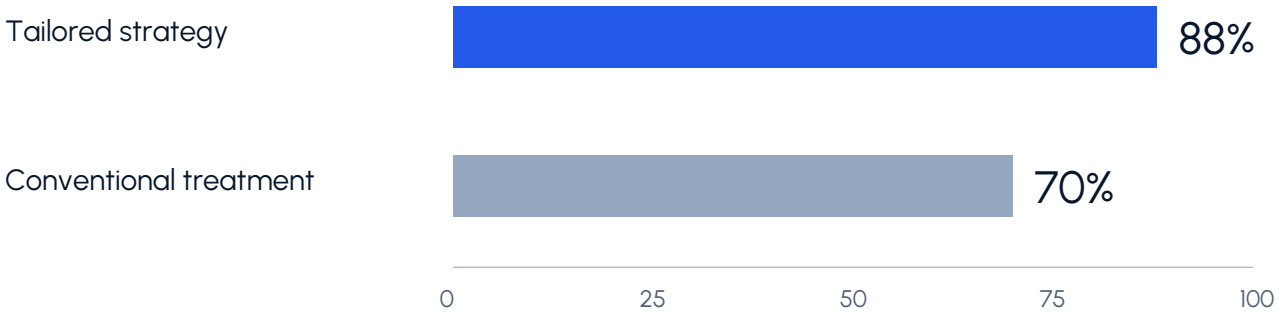

Estimated primary-endpoint probabilities; modified intention-to-treat analysis, n = 357. A total of 370 patients underwent ablation.[15]

MEAN TOTAL PROCEDURE TIME

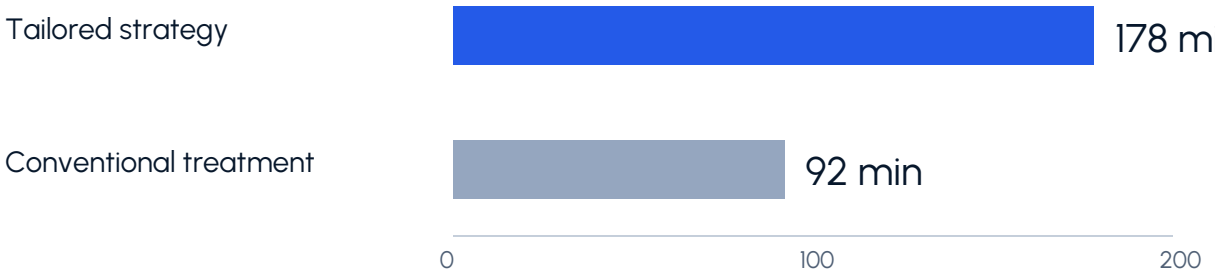

The broader endpoint of freedom from any atrial arrhythmia after one procedure did not significantly differ. Procedure time and radiofrequency delivery time were roughly doubled.[15]

15. Nature Medicine, TAILORED-AF trial. https://www.nature.com/articles/s41591-025-03517-w

07 / THE REPORT

Who gets the healthier future?

AI-generated conceptual illustration

Suppose the scientific promise develops substantially. Researchers discover useful treatments more efficiently. Clinical tools become more reliable. Some difficult tasks become easier and less expensive. The next question is whether the rest of us experience those gains as better health and more accessible care.

That is not a detail to address after the technology succeeds. It is part of what success should mean.

Consider the word cost. It can refer to the expense of developing a medicine, the expense of providing a service, the amount an insurer pays or the bill a patient receives. Those are connected, but they are not the same number. Lowering one does not automatically lower all the others.

If AI makes development less expensive, what happens to the price of the resulting treatment? If a practice saves time, does it offer appointments sooner, give clinicians a more manageable workload or simply raise expectations for throughput? If an insurer automates decisions, does the patient get a faster resolution or a faster denial? These are questions about incentives and management as well as technical performance.

A current argument over coding illustrates the problem. A Blue Cross Blue Shield Association analysis links increased coding intensity amid AI adoption to additional spending. The American Hospital Association disputes the broader criticism of provider coding practices, emphasizing documentation obligations, patient acuity and insurer practices. These are interested parties with different financial perspectives, not a settled accounting of AI's overall effect. [17,18]

17. Becker's, BCBSA coding-intensity analysis.
https://www.beckerspayer.com/payer/bcbs-study-hospital-ai-billing-tools-may-be-driving-up-healthcare-costs-by-billions/

18. AHA, coding-intensity response.
https://www.aha.org/fact-sheets/2026-07-31-fact-sheet-artificial-intelligence-and-coding-intensity

07 / THE REPORT

Who gets the healthier future?

For a patient, the relevant question is not which institution won the automation contest. It is whether the care was appropriate and the resulting cost reasonable. A system that helps one party collect or retain more money might accomplish its owner's objective perfectly without improving either outcome.

The same distinction applies to access. An AI service might offer useful assistance to someone who previously had little. That possibility should not be dismissed because the service is not equivalent to a full clinical team. The comparison must include what was actually available before.

But we should also ask whether an organization is expanding assistance or substituting it for care people should be able to obtain. An optional tool with a reliable route to a clinician is different from a tool that becomes the only door open to patients who cannot pay more.

West Health–Gallup survey findings released in April 2026 make the access issue concrete. In a survey conducted in late 2025, 25% of U.S. adults reported using AI for health information or advice. Among recent users, some cited inability to afford a physician visit or access a provider. The findings describe self-reported behavior, not proof that the resulting advice improved health. [19]

We should not mistake people adapting to unmet need for evidence that the need has been met. Nor should we ignore the possibility that a well-designed service could help. The question is whether assistance leads people toward appropriate care, with a workable route to human help when required.

"We should not mistake people adapting to unmet need for evidence that the need has been met."

19. West Health–Gallup, AI health survey. https://news.gallup.com/poll/707789/americans-turning-supplement-healthcare-visits.aspx

07 / THE REPORT

Who gets the healthier future?

The global-health opportunity deserves the same seriousness. The Gates Foundation and Anthropic announced a partnership in May 2026 involving $200 million over four years in grants, API credits and technical support across health, education and agriculture. That is not $200 million entirely in health spending or cash, and it is a commitment rather than a demonstrated health outcome. [20]

Such initiatives give us something concrete to evaluate. Do local institutions help define the problems? Can the resulting tools operate in the languages and settings where they are needed? Are there clinicians, supplies and referral services available when a tool identifies a problem? Who supports the service after the initial project ends?

Privacy and responsibility are part of that delivery work. Patients should be able to understand who receives their information, how it is used and who can correct an error. They should also know who is accountable for the decisions a service influences. An assurance that a human remains involved means little if that human lacks the time, authority or information to intervene.

The HHS guidance on health information and personal devices is a reminder that information entered into a consumer app is not automatically protected in the same way as records held within a covered healthcare relationship. That does not mean consumer services are wholly unregulated. It means patients cannot infer the protections from the subject matter alone. [21]

This is where the earlier doctors article and the larger scientific promise come together. More capability is valuable. But the people paying for, deploying and governing it make choices about whose needs it serves. We should evaluate those choices as closely as the models themselves.

“More capability is valuable. But the people paying for, deploying and governing it make choices about whose needs it serves.”

20. Gates Foundation, Anthropic partnership.
https://www.gatesfoundation.org/ideas/media-center/press-releases/2026/05/ai-anthropic-partnership

21. HHS, health information on personal devices.
https://www.hhs.gov/hipaa/for-professionals/privacy/guidance/cell-phone-hipaa/index.html

EVIDENCE IN VIEW / 03

AI is already part of the health search.

25% of U.S. adults reported having used AI for health information or advice.[19]

WHY RECENT USERS TURN TO AI

Among those using AI for health in the past 30 days.[19]

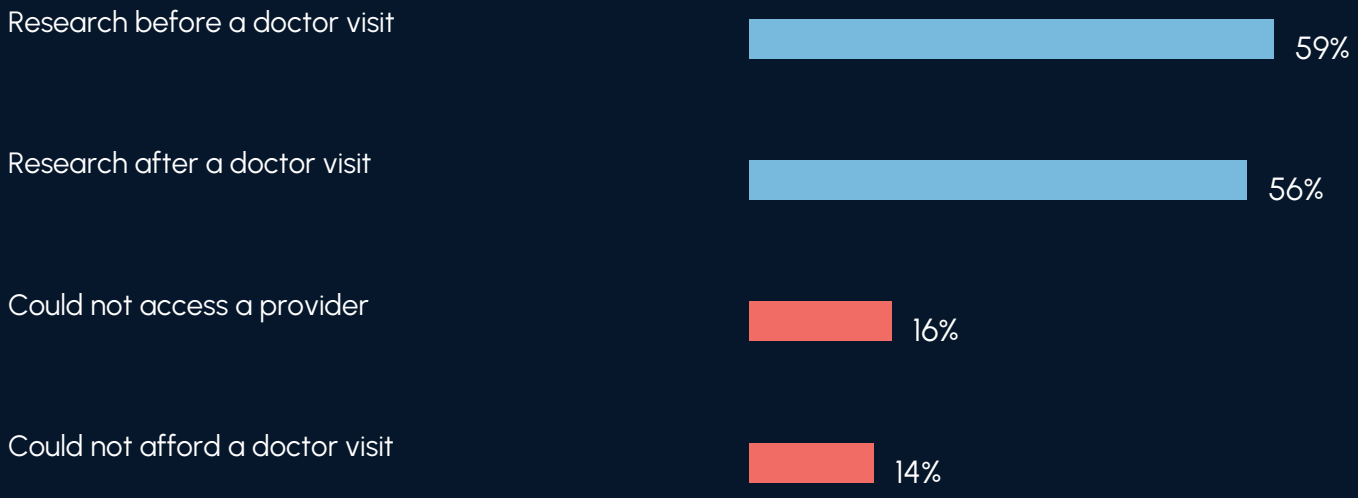

Use is not an outcome.

Self-reported behavior does not show that advice improved health. Survey: 5,660 U.S. adults, Oct. 27–Dec. 22, 2025; findings released April 2026. Reasons are not mutually exclusive.[19]

19. West Health–Gallup, AI health survey. https://news.gallup.com/poll/707789/americans-turning-supplement-healthcare-visits.aspx

08 / THE REPORT

A promise worth taking seriously

The recent coverage gives us a good reason to pay attention. It does not give us a single conclusion to accept. Research investments, laboratory findings, clinical trials and longevity forecasts are different pieces of a story that is still developing.

My view is that the promise is substantial enough to deserve serious effort. Helping researchers choose better experiments, develop useful interventions and understand why treatments succeed or fail could be an important contribution to medicine. We do not need to accept every forecast to recognize the value of pursuing those capabilities.

We also should not make skepticism a competition to find the least impressive interpretation of every result. A bounded finding can matter. A narrow clinical success can be worth building on. An early discovery can justify further work without justifying a claim about a future cure.

The corresponding obligation falls on the people making the promises. Tell us what was tested. Separate what the system did from what people supplied. Explain what remains unknown. Give independent researchers enough information to examine the claim. When the evidence changes, change the claim with it.

I would apply the same discipline to the economics. If an organization says AI makes care less expensive, ask less expensive for whom. If it says access improves, ask what patients can now obtain. If it says a clinician remains accountable, ask whether that clinician can actually exercise control.

There is no need to choose between wanting a remarkable future and insisting on these answers. A future involving powerful medical tools makes the answers more important, not less. The ambition should include building the institutions and services that allow people to benefit safely and fairly.

As this story develops, I will be looking for a progression that readers can follow: findings other researchers can confirm, interventions tested against meaningful alternatives, benefits that endure and care that people can reach and afford. Not every project will complete that journey. Reporting should make it clear how far each has traveled.

What would persuade me most is not another executive's confidence about the year everything changes. It would be accumulating evidence that more people are avoiding serious illness, receiving effective treatment or retaining abilities they otherwise would have lost.

That is the future behind the appealing headlines. It is worth pursuing with ambition and examining without sentimentality. The eventual measure will not be how much AI we put into healthcare. It will be what people are able to do with their lives because their health is better.

THE MEASURE THAT MATTERS

“It will be what people are able to do with their lives because their health is better.”

ALAN SHIMEL

About this report

EDITORIAL

AI and the Promise of Better Health:
From Discovery to the Patient
By Alan Shimel

A Techstrong Special Report
A Futurum Company
September 2026

PRODUCTION NOTES

This report connects scientific discovery with clinical evidence, access and patient benefit. Source notes, explanatory graphics and sidebars distinguish research findings from forecasts and potential applications.

Original cover and feature imagery were generated with AI. They are conceptual illustrations, not photographs of actual study participants, clinicians, facilities or molecular structures. Quantitative charts were drawn programmatically from the cited data.

HOW TO READ THE EVIDENCE

Source links are clickable and printed in full beside the relevant pages. Company announcements are identified as such. Forecasts, laboratory findings, clinical endpoints and access measures should not be read as interchangeable evidence of patient benefit.

The clinical graphics retain the uncertainty and endpoint distinctions in the manuscript. The three-clock diagram is an authorial framework, not a forecast of development time. The survey graphic distinguishes all-adult prevalence from reasons reported by recent users.

Editorial information only. This report is not personal medical advice.