The Future of AI-Facilitated Medicine

Eric J. Topol

AI's most significant contribution to medicine to date is its ability to produce more accurate and comprehensive interpretations of medical images, as validated through randomized clinical trials. Next up is the opportunity to address the steady erosion of the patient-doctor relationship over several decades and a global burnout crisis among clinicians. By reducing the data clerical work of clinicians and giving patients more autonomy, AI has the potential to restore the humanity in medicine and care. Freeing physicians from constant screen time could bring back the physician's presence during clinic visits and help foster empathy and trust. Furthermore, AI's ability to aggregate and contextualize a "full stack" of patient data can enable new opportunities for prevention of major age-related diseases, significantly extending human health span. Despite this vast potential, the implementation of medical AI faces substantial challenges, including mitigating bias, ensuring data privacy and security, and establishing appropriate regulatory and reimbursement frameworks.

The pioneers and leaders in the field of artificial intelligence have consistently singled out advances in medicine and health care as its most substantial and unequivocally positive impact. As Geoffrey Hinton put it: "I always pivot to medicine as an example of all the good it can do because almost everything it's going to do there is going to be good."[1] Similarly, speaking about the end of disease, Demis Hassabis said: "I think that's within reach. Maybe within the next decade or so, I don't see why not."[2] Then there's Dario Amodei, who believes AI will cure most cancers, prevent Alzheimer's, and double the human lifespan, and Sam Altman, who predicted that "we will see diseases cured at an unprecedented rate thanks to AI."[3] And there's Mustafa Suleyman, who said: "One of the applications that I'm most excited about is this quest for medical superintelligence. I think it's pretty clear in the next couple of years we are going to build models that are better than the vast majority of expert clinicians in any discipline at diagnosis."[4]

Encouraging as these proclamations are, the specifics of how AI can and will ultimately transform medicine remain vague. This essay intends to address the gaps and provide a futuristic perspective of how AI can achieve profound improvements in the practice of medicine and in health outcomes for patients – by improving diagnostic accuracy, restoring humanity to the doctor-patient relationship, and preventing diseases.

https://doi.org/10.1162/DAED.a.986

The medical community tends to minimize its errors. However, an important report by researchers at Johns Hopkins estimates that nearly eight hundred thousand Americans per year suffer permanent disability or death due to significant diagnostic medical errors.[5] And that doesn't include errors in treatment, which represent another serious problem. With the common use of the term "precision medicine," there is no acknowledgment that if you make the same mistake repeatedly, it is, in a perverse way, very precise. We desperately need to promote diagnostic accuracy in medicine.

Today we're making strides, particularly with respect to interpretation of medical imaging. The largest randomized trial in medical AI compared radiologists using AI against radiologists without AI in a sample of more than one hundred thousand women undergoing mammography.[6] The radiologists working with AI support detected approximately 30 percent more instances of breast cancer and 25 percent more invasive tumors than radiologists without AI. Similarly, across forty randomized trials in gastroenterology, machine vision–supported colonoscopy (compared with colonoscopy without machine vision) increased the overall detection of adenomas and polyps by more than 20 percent.[7]

Beyond these randomized trials, there is supportive evidence for improved accuracy of all types of medical images, including X-rays, ultrasound, positron emission tomography (PET) scans, CT scans, and MRI. The next step is for the initial image review process to be fully automated, with AI systems capable of not just screening specific traits but of synthesizing information to produce comprehensive reports that can be reviewed by supervising physicians. Researchers recently did just that in one of the most complex areas of medical imaging: echocardiography. The motion of the heart and the need to acquire and interpret multiple views of not just static images but of ultrasound video create a data-rich and complex environment in which to develop a foundation AI model. Trained using over twelve million echocardiographic videos and paired text reports across twenty-three benchmarks of heart structure and function, the EchoPrime model achieved state-of-the-art performance in interpretation of cardiac features and diagnoses.[8] There is work in progress to do the same for all medical scans, with the potential to improve diagnosis accuracy and streamline workflows.

That brings us to the concept of *superhuman vison*. The conventional interpretation of medical images – answering a specific query – may become obsolete. For example: Does a chest X-ray show pneumonia? With computer vision, we now can detect information embedded in chest X-rays about the presence of type 2 diabetes, the risk of atherosclerosis in the next ten years, the coronary artery calcium score, ejection fraction, and osteoporosis. The same is true for mammograms that can accurately detect heart disease risk, or chest CT scans finding pancreatic cancer. This "opportunistic" medical scan interpretation is a reflection of the power of digital eyes to "see" things that expert clinicians cannot.[9]

Superhuman vision is not limited to radiologic scans. Perhaps the most extraordinary example is the retina image, whether it be a photo or optical coherence tomography (OCT). Various studies, and more recently AI foundation models, which use vast and diverse datasets for a wide range of tasks, have indicated how remarkably informative the retina can be.[10] AI interpretation of retinal imaging can tell us about a patient's blood pressure, glucose control, and risk of heart, kidney, liver, and gall bladder disease. The images are a window into the future risk of Alzheimer's and Parkinson's disease in people without any symptoms. In the retina, AI can also detect the likelihood of subclinical (no symptoms, no narrowings that limit blood flow) atherosclerosis of the coronary arteries and silent brain injury, as well as the risk of stroke.

The aforementioned 12-lead electrocardiogram likewise lends itself to AI insights that were not previously obtainable through imaging alone. They include the determination of the age and sex of the patient; the presence of anemia, diabetes, or kidney disease; ejection fraction of the left ventricle; filling pressure of the heart; prediction of atrial fibrillation; and the risk of stroke. Or, with AI, a pathology slide of a questionable cancer tissue can determine the genomic driver mutations of the tumor, its malignancy status, the site of origin, and patient prognosis. The possibilities of what all is encoded in different types of medical images seem limitless, and were certainly not envisioned before the advent of deep learning and generative AI. Whether and when all this information will move from research publications to medical practice is uncertain, and will require further validation, but it is fair to say we won't be leaving so much data on the table thanks to the power of AI superhuman vision.

AI is making great progress toward more accurate and comprehensive diagnostic imaging, but the contributions of AI for promoting accuracy go well beyond that. Specially adapted large language models have demonstrated the ability to concisely and accurately synthesize all of a patient's data, which can be quite extensive and from multiple sources.[11] Improved AI performance extends to better patient communication and management plans. In one study involving 150 case scenarios and patient actors, AI demonstrated superior performance to specialist clinicians in thirty of thirty-two tasks.[12] While that study is not representative of the real world of medical practice, there has since been a randomized trial in real-world medical centers from four continents that confirms the advantages of using AI for ophthalmology care with respect to diagnosis of seven major eye diseases and patient management.[13] The clinical benefits of AI are also not limited to high-income countries: A randomized trial of primary care at fifteen clinics in Nairobi involving nearly forty thousand patient visits showed that clinician use of an AI consult tool reduced diagnosis and treatment errors. Similar results have followed for many other medical tasks, such as diagnosis of tuberculosis, childhood cancer, diabetic retinopathy, and rheumatic heart disease in low- and middle-income countries.[14]

For several decades, there has been a steady erosion of the patient-doctor relationship, in part because of the very limited time allotted to clinic visits and bedside hospital rounds. The business of medicine and the drive for clinician productivity have left physicians with the sense they cannot provide adequate care for patients – the principal reason they entered the profession. A typical doctor-patient encounter today lasts only a few minutes without much, if any, eye-to-eye contact – the legal and ethical requirements of progress notes push doctors from the bedside to the keyboard as they fulfill data-clerk responsibilities. The patients perceive the doctor's lack of presence, which in turn engenders a lack of trust, care, and concern. What was an intimate human-human bond decades ago has slipped to an estranged encounter, devoid of compassion, often with very limited physical examination and laying of hands. In my 2019 book *Deep Medicine*, I wrote that there was potential for AI to foster humanity in medicine.[15] Many find that idea to be counterintuitive: How could technology make medicine more human again?

Liberation from the keyboard could markedly improve physicians' visits with a patient. With the use of AI scribes that synthesize notes with many downstream functions, doctors can redirect their attention to the patient, bringing back face-to-face communication. At the same time, ambient conversations can produce excellent notes with audio links for the patient to return to after the appointment. On top of that, the AI assistant can schedule labs, follow-up appointments, and scans; run preauthorizations with insurance companies and submit prescriptions; do coding and billing; and even nudge patients after the visit to see if they are following up on planned actions, such as measuring their blood pressure or going for walks. In this way, administrative use of AI is humanistic. The automation is a major step toward relieving clinicians of their data-clerk burden and allowing them to focus more fully on their patients. Such systems have already been shown to save clinician time and reduce burnout, a global crisis among health care workers.[16] The uptake of this technology has been remarkably fast; before long, automated note-taking and AI support for time-consuming administrative tasks will be the norm.

And patients are gaining empowerment. Even prior to generative AI, algorithmic diagnostics were changing how patients monitor their health, enabling them to measure abnormal heart rhythm by smartwatch, screen for diabetic retinopathy at grocery stores, detect urinary tract infections via an over-the-counter kit, and classify skin rashes or lesions by app.[17] Now patients are using AI chatbots to make diagnoses (some of which had eluded doctors and workups), get first and second opinions, enter their lab data, obtain advice for optimal treatment, or prepare for a doctor visit.[18] This trend was reflected in a recent *New York Times* article entitled "Empathetic, Available, Cheap: When AI Offers What Doctors Don't."[19]

I had not considered that empathy could be facilitated by machines. But in a systematic review of fifteen studies comparing patient perception of empathy, all but two of the studies showed that patients reported higher levels of empathy in their text conversations with AI-based chatbots than in their text communications with clinicians.[20] While AI isn't empathetic, its ability to foster a sense of empathy through appropriate language is significant inasmuch as the perception of empathy promotes better patient outcomes. In the future, this ability will likely lead to AI coaching of clinicians to be better, more empathetic communicators. Going further, think of how AI review of the physician's automated note could lead to productive questioning: "Why did you interrupt Mrs. Jones after nine seconds? Why didn't you give her the chance to express her real concerns?" While this was certainly not envisioned as a role for AI and many clinicians will bristle at the thought, it would not be surprising for such coaching to become part of an annual licensing requirement for the practice of medicine. For clinicians to actualize the benefits of improved AI accuracy, recognize its drawbacks, and promote trust with patients, specific training with these new systems will be necessary.

Improving the patient-doctor relationship requires the gift of time; it is the starting point for deeper physician-patient discussion and careful physical examination. But AI also introduces the potential for health care administrators to use the technology to enhance productivity and profitability: instead of creating more time for meaningful doctor-patient interactions, these tools could simply allow clinicians to see more patients in less time. Such an approach to these technologies would be destructive and obviate the chance for AI to help restore the essential humanity in health care. To counter this risk, I have called for solidarity among clinicians to stand up to administrators.[21] Clinicians should use their bargaining power to increase the time allotted for in-person outpatient visits, reserving this precious encounter for important matters while taking advantage of alternative patient services including AI chatbots, telehealth, and virtual coaches, when appropriate.[22]

Cancer, cardiovascular disease, and neurodegenerative disorders are the three main age-related diseases. Cumulatively, and with onset usually after age sixty, they are the primary reasons why the human health span – the years lived in good health – is cut short. Besides their increasing incidence with age, they share two important features: 1) a long incubation period of approximately two decades before symptoms manifest and 2) an immune system dysregulation and loss of protection, with untoward inflammation. The terms "immunosenescence" and "inflammaging" are used to characterize the latter.

Primary prevention is a treatment approach to take preemptive steps to anticipate and address a person's risk and stop a disease before it occurs. Although frequently discussed, primary prevention of the three major age-related diseases is rarefied in medical practice. Mass cancer screening is done with mammography, colonoscopy, and prostate-specific antigen, but these tests, at best, are used for detection, not prevention of cancer. Once cancer is detected by a scan there are already billions of cancer cells present and an increased chance it has spread within the body. But there are now ways through genomics to know who is at risk of cancer, including polygenic risk scores and whole genome sequencing. Further, our intact immune system is essential to prevent the occurrence of cancer and we have only recently found ways to determine the health and functionality of our immune responsiveness. Until now, surprisingly, there has not been any test in the clinic to gauge the health of our immune system. But the recent advance of organ clocks – a panel of blood proteins that tracks the pace of aging for each organ of the body and our immune system – offers a way forward.[23]

While there has been a marked reduction of heart attacks in the past two decades, there are still more than eight hundred thousand per year in the United States that are not prevented. Focus has recently shifted to identifying vulnerable patients with inflamed atherosclerotic plaque, well before a potentially catastrophic event occurs.[24]

For Alzheimer's disease, the breakthrough blood test known as p-tau217 can now detect the risk of this condition many years before any symptoms of mild cognitive impairment.[25] The risk of Alzheimer's is also increased in people with a family history, one or two copies of the APOE4 allele, and a high polygenic risk score. Brain biomarkers like p-tau217, similar to LDL cholesterol for heart disease, serve a dual purpose of defining risk and measuring changes from lifestyle modifications and medications.

With a long runway of being able to identify high-risk individuals for each of these age-related diseases, the opportunity to achieve primary prevention has never been more realistic and attainable. But it requires having the "full stack" of patients' data, including electronic health records, lifestyle factors, labs, images, retina scan, genomics, proteins (organ clocks), relevant blood biomarkers, environmental exposures, and wearable sensor data. We've learned that while doctors and patients typically dismiss lab tests in the normal range as unimportant, AI can pick up significant trends and contextualize these results with the medical records.[26] AI can also draw important data from the notes themselves, both structured and unstructured components, such as for detecting a high risk for pancreatic cancer many years before this difficult and deadly diagnosis could otherwise be made.[27] I have reviewed the opportunistic capabilities of medical scans and retinal imaging for predicting risk of cardiovascular and neurodegenerative diseases. In addition, AI models can now differentiate between "normal" mammograms, reclassifying some as showing a high risk of breast cancer in the next five years.[28]

In aggregate, these layers of data set up multimodal AI to determine the risk status of each individual for one or more of the major age-related diseases and to track whether the risk has been decreased via interventions. The interventions include not only healthy lifestyle factors – such as an anti-inflammatory diet, regular exercise, and good sleep – but also drug treatments. The field of AI drug discovery has garnered much attention, with dozens of companies active in this space, some partnered with the major pharmaceutical firms. We tend to think of a drug as solely for treatment of symptoms and disease, but in the years ahead, there will be a shift to preventive medications. Perhaps the pluripotent action of glucagon-like peptide (GLP-1) family of drugs, with benefits that have extended from treating diabetes and obesity to preventing heart, kidney, liver, and addiction disorders (to some extent independent of weight loss), will be viewed as the turning point for preventive drugs. While optimizing lifestyle will always be regarded as fundamental and low-cost, it is clear more will be needed to achieve prevention of the big three age-related diseases.

Adding to newfound capabilities to predict disease is the publication of the first large health model, known as Delphi-2M.[29] Researchers built and validated Delphi-2M from electronic health records and lifestyle data of four hundred thousand UK Biobank participants and nearly two million people in Denmark. By tokening each individual's past data, the model, based on GPT-2, a relatively low-parameter and early generative AI precursor, was able to predict disease over the next twenty years – both the *disease or health event* and *when* – with a surprising level of accuracy. The prediction performance increased when polygenic risk scores were introduced as an additional layer of data. In the future, more advanced large health models will be validated with the full stack of data described above, taking the predictive capability to higher levels of accuracy.

There is also considerable potential for a digital-twin resource to predict diseases.[30] Previously, computer scientist and investor Kai-Fu Lee and I envisioned a global data resource with billions of individuals and their full stacks of medical and health data, along with all data captured during an extended follow-up period. With such a dataset – whose privacy and security must be protected by federated learning and encryption – we would have the capability of using nearest-neighbor analysis to find digital twins for any individual. Such data analysis would complement individuals' multimodal AI and large health models as another means of forecasting health events and disease risk to the highest level of accuracy, enabling primary prevention. No such resource yet exists, but the potential to promote human health span is too great to ignore, and it is likely an inevitability. The enormous economic upside of reducing the burden of these age-related diseases, concurrent with the marked aging of the U.S. population, makes investment in these AI tools and resources and the ideal of primary prevention exceptionally alluring.

While there is vast potential for AI to promote medical accuracy, humanism, and prevention in the future, there is no shortage of challenges. At this juncture, large language models lack readiness for broad medical use, with notable deficiencies determined by stress testing and representative medical benchmarks.[31] I have mentioned the importance of data privacy and security, something we do not have a good track record with to date. But a shift to individuals owning their own medical data on a secure platform, which could also incorporate data from sensors and genomics that are not currently part of our electronic health records, could lead to improvement. As for intrinsic cultural biases of these and other inputs, we must fully interrogate them to avoid their propagation. I have argued that AI can reduce health inequities, but the fact is that without deliberate initiatives that prioritize fairness, AI systems are likely to exacerbate them. Of course, AI-induced medical errors are also a concern, especially where there is direct connection between generative AI systems and patients with insufficient clinician oversight. We must lead a dedicated effort to interrogate and refine any patient-facing models before making them publicly available. To date, we have no objective data for how direct patient access is faring with respect to benefit and harm; that needs to be captured and monitored on an ongoing basis. There are also significant risks and issues with an approach of permissive regulatory oversight that has chiefly relied on low-bar clearance instead of formal approvals for applications of AI algorithms. Reimbursement standards for AI in health care markets have not adapted to the progress being made: stakeholder incentives remain misaligned, which contributes to stasis and delayed adoption of useful technologies. We must act on and build upon the recommendations we have for how to best proceed.[32] And most important of all, we need to establish trust in the medical community and the public, which stems from transparency and compelling medical research to identify benefits and harms. So much of the medical AI research today is left compartmentalized, not leading to implementation in medical practice. That gaping chasm between research and practice needs to be bridged.

Medical AI's future impact will be transformative, especially if it attends to and eventually overcomes the many obstacles and challenges ahead. I believe that preventing age-related diseases, thereby significantly extending human health span, will be the single most important contribution of AI in medicine in the decades ahead.

ABOUT THE AUTHOR

Eric J. Topol is Chair of the Department of Translational Medicine at Scripps Research Institute, Director and Founder of the Scripps Research Translational Institute, Senior Consultant for the Division of Cardiovascular Diseases at the Scripps Clinic, and the Gary and Mary West Chair of Innovative Medicine in the Department of Translational Medicine at Scripps Research. He is the author of *Super Agers: An Evidence-Based Approach to Longevity* (2025), *Deep Medicine: How Artificial Intelligence Can Make Healthcare Human Again* (2019), and *The Patient Will See You Now: The Future of Medicine is in Your Hands* (2014).

ENDNOTES

[1] Eric Topol, "Geoffrey Hinton: Large Language Models in Medicine. They Understand and Have Empathy," *Ground Truths* (podcast), December 8, 2023, https://erictopol.substack.com/p/geoffrey-hinton-large-language-models.

[2] Scott Pelley, "Artificial Intelligence Could End Disease, Lead to 'Radical Abundance,' Google DeepMind CEO Demis Hassabis Says," CBS News, August 3, 2025, https://www.cbsnews.com/news/artificial-intelligence-google-deepmind-ceo-demis-hassabis-60-minutes-transcript.

[3] Anderson Cooper and Nichole Marks, "Why Anthropic CEO Dario Amodei Spends So Much Time Warning of AI's Potential Dangers," CBS News, November 16, 2025, https://www.cbsnews.com/news/anthropic-ceo-dario-amodei-warning-of-ai-potential-dangers-60-minutes-transcript; and Matteo Wong, "AI Executives Promise Cancer Cures. Here's the Reality," *The Atlantic*, April 25, 2025, https://www.theatlantic.com/technology/archive/2025/04/how-ai-will-actually-contribute-cancer-cure/682607.

[4] Dan Fox and Richard Hodson, "The Future of AI," *Nature*, November 14, 2025, https://www.nature.com/immersive/d41586-025-03701-5/index.html#section-mT0JSji4Ai.

[5] David E. Newman-Toker, Najlla Nassery, Adam C. Schaffer, et al., "Burden of Serious Harms from Diagnostic Error in the USA," *BMJ Quality & Safety* 33 (2) (2024): 109–120, https://doi.org/10.1136/bmjqs-2021-014130.

[6] Veronica Hernstrom, Viktoria Josefsson, Hanna Sartor, et al., "Screening Performance and Characteristics of Breast Cancer Detected in the Mammography Screening with Artificial Intelligence Trial (Masai): A Randomised, Controlled, Parallel-Group, Non-Inferiority, Single-Blinded, Screening Accuracy Study," *Lancet Digital Health* 7 (3) (2025): https://doi.org/10.1016/S2589-7500(24)00267-X.

[7] Saeed Soleymanjahi, Jack Huebner, Lina Elmansy, et al., "Artificial Intelligence-Assisted Colonoscopy for Polyp Detection: A Systematic Review and Meta-Analysis," *Annals of Internal Medicine* 177 (12) (2024): 1652–1663, https://doi.org/10.7326/ANNALS-24-00981.

[8] Milos Vukadinovic, I-Min Chiu, Xiu Tang, et al., "Comprehensive Echocardiogram Evaluation with View Primed Vision Language AI," *Nature* 650 (8103) (2026), https://doi.org/10.1038/s41586-025-09850-x.

[9] Eric J. Topol, "AI-Enabled Opportunistic Medical Scan Interpretation," *Lancet* 403 (10439) (2024): 1842, https://doi.org/10.1016/S0140-6736(24)00924-3.

[10] Eric J. Topol, "As Artificial Intelligence Goes Multimodal, Medical Applications Multiply," *Science* 381 (6663) (2023): adk6139, https://doi.org/10.1126/science.adk6139.

[11] Dave Van Veen, Cara Van Uden, Louis Blankemeier, et al., "Adapted Large Language Models Can Outperform Medical Experts in Clinical Text Summarization," *Nature Medicine* 30 (4) (2024): 1134–1142, https://doi.org/10.1038/s41591-024-02855-5.

[12] Tao Tu, Mike Schaekermann, Anil Palepu, et al., "Towards Conversational Diagnostic Artificial Intelligence," *Nature* 642 (8067) (2025): 442–450, https://doi.org/10.1038/s41586-025-08866-7.

[13] Yilan Wu, Bo Qian, Tingyao Li, et al., "An Eyecare Foundation Model for Clinical Assistance: A Randomized Controlled Trial," *Nature Medicine* 31 (10) (2025): 3404–3413, https://doi.org/10.1038/s41591-025-03900-7.

[14] Robert Korom, Sarah Kiptinness, Najib Adan, et al., "AI-Based Clinical Decision Support for Primary Care: A Real-World Study," arXiv (2025), https://doi.org/10.48550/arXiv.2507.16947.

[15] Eric Topol, *Deep Medicine: How Artificial Intelligence Can Make Healthcare Human Again* (Basic Books, 2019).

[16] Kristine D. Olson, Daniella Meeker, Matt Troup, et al., "Use of Ambient AI Scribes to Reduce Administrative Burden and Professional Burnout," *JAMA Network Open* 8 (10) (2025): e2534976, https://doi.org/10.1001/jamanetworkopen.2025.34976; and Majid Afshar, Mary Ryan Baumann, Felice Resnik, et al., "A Pragmatic Randomized Controlled Trial of Ambient Artificial Intelligence to Improve Health Practitioner Well-Being," *NEJM AI* 2 (12) (2025): AIoa2500945, https://doi.org/doi:10.1056/AIoa2500945.

[17] Nelson Barrera, Maria Solorzano, Yomary Jimenez, et al., "Accuracy of Smartwatches in the Detection of Atrial Fibrillation," *JACC: Advances* 4 (11.1) (2025): 102133, https://doi.org/doi:10.1016/j.jacadv.2025.102133.

[18] Teddy Rosenbluth and Maggy Astor, "Empathetic, Available, Cheap: When A.I. Offers What Doctors Don't," *The New York Times*, November 17, 2025, https://www.nytimes.com/2025/11/16/well/ai-chatbot-doctors-health-care-advice.html; and Brittany Trang, "Patients Are Already Using AI. How Can the Health Care System Help?" *STAT*, November 26, 2025, https://www.statnews.com/2025/11/26/patients-artificial-intelligence-health-care-system-ai-prognosis.

[19] Rosenbluth and Astor, "Empathetic, Available, Cheap."

[20] Alastair Howcroft, Amber Bennett-Weston, Ahmad Khan, et al., "AI Chatbots Versus Human Healthcare Professionals: A Systematic Review and Meta-Analysis of Empathy in Patient Care," *British Medical Bulletin* 156 (1) (2025), https://doi.org/10.1093/bmb/ldaf017.

[21] Eric Topol, "Why Doctors Should Organize," *The New Yorker*, August 5, 2019, https://www.newyorker.com/culture/annals-of-inquiry/why-doctors-should-organize.

[22] A. Ali Heydari, Ken Gu, Vidya Srinivas, et al., "The Anatomy of a Personal Health Agent," arXiv (2025), https://arxiv.org/abs/2508.20148.

[23] Eric Topol, "The Emergence of Protein Organ Clocks: Predicting and Tracking Organ-Specific Aging, Health and Disease," Ground Truths, July 13, 2024, https://erictopol.substack.com/p/the-emergence-of-protein-organ-clocks.

[24] Eric Topol, "The Big Shift in Cardiology to Atheroma and Inflammation," Ground Truths, December 3, 2023, https://erictopol.substack.com/p/the-big-shift-in-cardiology-to-atheroma.

[25] Eric Topol, "Predicting and Preventing Alzheimer's Disease," Ground Truths, May 31, 2025, https://erictopol.substack.com/p/predicting-and-preventing-alzheimers; and Eric Topol, "The Breakthrough Blood Test for Alzheimer's Disease," Ground Truths, April 13, 2025, https://erictopol.substack.com/p/the-breakthrough-blood-test-for-alzheimers.

[26] Eric Topol, "Your Lab Tests: What Is Being Missed by Relying on Population-Based Reference Values," Ground Truths, December 14, 2024, https://erictopol.substack.com/p/your-lab-tests.

[27] Davide Placido, Bo Yuan, Jessica X. Hjaltelin, et al., "A Deep Learning Algorithm to Predict Risk of Pancreatic Cancer from Disease Trajectories," *Nature Medicine* 29 (5) (2023): 1113–1122, https://doi.org/10.1038/s41591-023-02332-5.

[28] Sadia H. Zapp, "A Breakthrough in Breast Cancer Prevention: FDA Clears First AI Tool to Predict Risk Using a Mammogram," Breast Cancer Research Foundation, June 2, 2025, https://www.bcrf.org/blog/clairity-breast-ai-artificial-intelligence-mammogram-approved.

[29] Artem Shmatko, Alexander Wolfgang Jung, Kumar Gaurav, et al., "Learning the Natural History of Human Disease with Generative Transformers," *Nature* 647 (8088) (2025): 248–256, https://doi.org/10.1038/s41586-025-09529-3.

[30] Eric Topol and Kai-Fu Lee, "It Takes a Planet," *Nature Biotechnology* 37 (8) (2019): 858–861, https://doi.org/10.1038/s41587-019-0214-z.

[31] Yu Gu, Jingjing Fu, Xiaodong Liu, et al., "The Illusion of Readiness: Stress Testing Large Frontier Models on Multimodal Medical Benchmarks," arXiv (2025), https://arxiv.org/abs/2509.18234.

[32] Narges Razavian, Prem Batchu-Green, Vikas Chowdhry, et al., "Catalyzing Health AI by Fixing Payment Systems," *NEJM AI* 2 (12) (2025): AIpc2500871, https://doi.org/doi:10.1056/AIpc2500871.