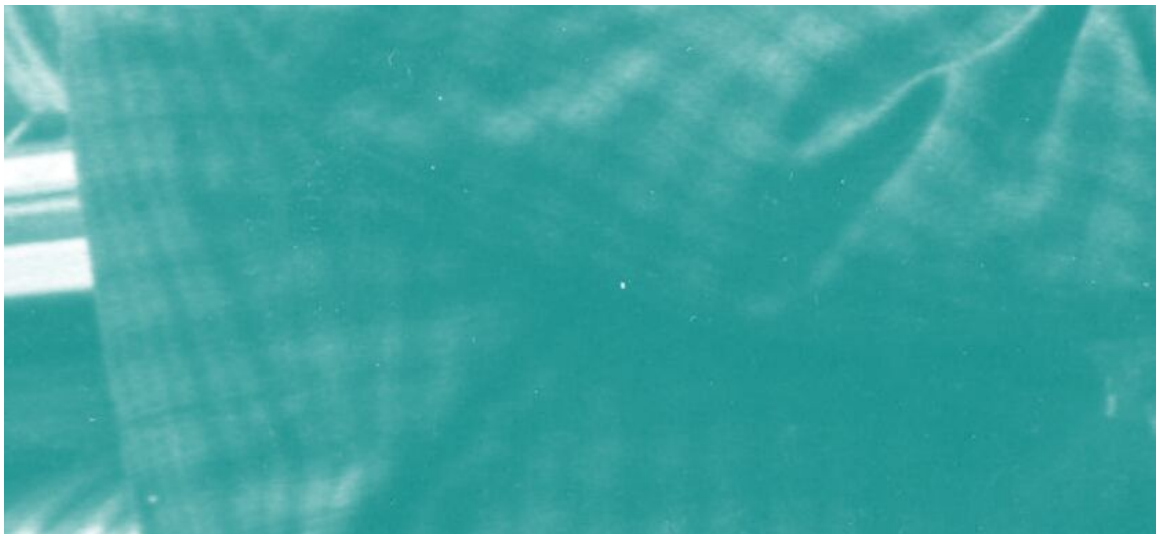
06

PART I · THE ATOMIC WORLD

Build Medicine's Holy Grail: A Simulated Human Body

John-Clark Levin



Computer specialist John Smith arranges canisters of magnetic tape used in processing medical data, National Library of Medicine, 1960s.

National Library of Medicine.

SUMMARY

Clinical trials are too statistically underpowered to capture many biological variations that can determine whether a treatment cures or kills a given patient. Yet AGI would enable a scientific megaproject culminating in full-body simulations with high enough fidelity to replace *in vivo* experiments with *in silico* trials that can predict treatment effects more accurately, rapidly, safely, and cheaply. Brute-force molecular dynamics approaches are computationally intractable, but the macro-scale behavior of biological systems can be predicted with many orders of magnitude less compute—once AGI discovers the correct laws for accurate multiscale abstraction. This will require not only enormous computing power, but an aggressive program of wet-lab science to resolve open biostructure questions so these simulations can be populated properly. The article concludes with an overview of how U.S. government action can accelerate these efforts, and an assessment of how full-body biosimulation can unlock much broader medical breakthroughs.

EVERY YEAR, SOMEWHERE AROUND 100,000 AMERICANS DIE FROM ADVERSE DRUG reactions. That's not counting meth overdoses or opioid abuse, nor doctor errors or pharmacy mistakes. That's just cases where a study said the drug would help the patient, the physician prescribed a proper dose in good faith, and instead it killed the patient. Why?

The key reason is that clinical trials are grossly underpowered statistically. How a drug affects a given patient—cures, kills, or has little effect either way—depends on numerous variables such as genetics, demographics, lifestyle, comorbidities, drug interactions, and disease variation. Even in a large phase III trial with 3,000 patients, there is probably not a single one who matches *you* in every clinically relevant way. But AGI can unlock a profound solution shimmering in the distance: personalized *in silico* trials that simulate each patient digitally.

Yet if your brow isn't wrinkling right now with deep skepticism, you don't fully understand the challenge. Although today's biosimulation projects like the Virtual Physiological Human[1] and various digital twin[2] platforms can be narrowly useful, they capture nowhere near enough detail to match the predictive power of flesh-and-blood human trials. Our biology is so complex that even extremely small systematic errors or omissions in simulations can lead via the butterfly effect to fatally wrong predictions. So to safely replace *in vivo* trials with AI modeling, we will need to build on a foundation of approximately single-atom precision. AlphaFold-family models have already achieved that for many protein structures, and new versions will likely soon achieve similar accuracy for modeling many protein complexes—followed by simulations of their ongoing interactions. From there, the roadmap climbs the complexity ladder from high school biology: organelles, cells, tissues, organs, and eventually the full body. Yet this quickly outstrips our ability to scale more computing hardware. A true atom-by-atom simulation of one entire human would take a data center almost 1,000 times larger than the sun—or, if it's not too urgent, the El Capitan supercomputer running 100 million *billion* times longer than the current age of the universe. So how can this be anything other than science fiction?

The answer lies in effective complexity—how much micro-scale detail about a system is relevant to predicting its macro-scale behavior. A party balloon contains around 10^{23} gas molecules, but if you are trying to predict whether it sinks or rises, you don't have to model them all individually. Instead, we have discovered the ideal gas law, which summarizes all the chaotic pinballing of hundreds of sextillions of nitrogens and oxygens into vastly simpler quantities of temperature, volume, and pressure. In a similar way, gene expression does not depend on which of a cell's billions of identical adenine molecules fits into a given base pair, and tissue oxygenation does not depend on which of billions of nearby red blood cells enters a capillary. The problem is, we do not

1 Marco Viceconti and Peter Hunter, "The Virtual Physiological Human: Ten Years After," *Annual Review of Biomedical Engineering* 18 (2016): 103–123.

2 Evangelia Katsoulakis et al., "Digital Twins for Health: A Scoping Review," *npj Digital Medicine* 7, no. 1 (2024): 77.

yet know the "ideal gas laws" to bridge scales accurately enough and make full-molecular-fidelity biosimulation computationally tractable. That is where AGI comes in.

As we build these biosimulations at larger and larger scales, we will use AGI to test and validate which details can be pruned and abstracted into simplified laws that are much cheaper computationally. Next, AGI will integrate those simplified models into even larger-scale simulations and repeat the validation-and-abstraction process—exhaustively checking for unexpected causal influences. For example, to test a lung cancer drug, we know we will need to model pneumocytes in detail, but a keratinocyte in the whorl of a fingerprint is almost certainly irrelevant. Yet it is not clear how much detail we'd need for hepatocytes in the liver, and pure theory cannot give us a reliable answer—we will need AGI to actually run the simulations and find the patterns. This does not require perfection. The project's guiding principle should be to demonstrate better predictive accuracy than today's clinical trials.

Attaining these capabilities will entail developing safe AGI and building out the enormous computing power needed to run *in silico* trials. But, contrary to the overconfident "AI solves everything" handwaves of some Silicon Valley moguls, this will not be sufficient for full-body biosimulation. Even if we had AGI and unlimited compute tomorrow, we still would not know many details necessary to populate the simulations accurately. For example, scientists do not fully understand the transport mechanism of the Golgi apparatus within cells,[3] or have a complete proteomic map of lysosomes.[4] Realistically, it will take a few hundred billion dollars more of wet-lab science—the painstaking work of real-world experimental biology—to answer all these open questions.

This is not out of reach. It would only entail around a 10 percent increase in medical R&D spending over the next decade,[5] and is dwarfed by the potential benefits. Around 1.9 million Americans die every year from chronic disease, and annual U.S. healthcare spending is projected to reach $8.6 trillion by 2033.[6] If cheaper drug development and breakthrough cures can shave even 1 percent off this, biosimulation would pay for itself several times over. The main challenge is coordination. Current market incentives and public-sector grantmaking do not efficiently push wet-lab research toward resolving the "shopping list" of unknowns required by a biosimulation megaproject. While increased government funding is urgently needed, just as important is the role bodies like the National Institutes of Health (NIH) and Advanced Research Projects Agency for Health (ARPA-H) can play in defining, publicizing, and updating the shopping list, har-

3 Hieng Chiong Tie et al., "Quantitative Intra-Golgi Transport and Organization Data Suggest the Stable Compartment Nature of the Golgi," *eLife* 13 (2025): RP98582.

4 Ali Ghoochani et al., "Cell-Type Resolved Protein Atlas of Brain Lysosomes Identifies SLC45A1-Associated Disease as a Lysosomal Disorder," *Cell* 189, no. 3 (2026): 765–782.e31.

5 Research!America and TEConomy Partners, LLC, *U.S. Investments in Medical and Health Research and Development, 2016–2020* (Arlington, VA: Research!America, 2022), 1.

6 Matt McGough et al., "How Much Is Health Spending Expected to Grow?" *Peterson-KFF Health System Tracker*, August 4, 2025.

monizing American science to fill it exhaustively while minimizing wasteful duplication.

The White House's Genesis Mission aims to marshal American AI technology to accelerate scientific research. First, policymakers should expand those authorities[7] into the AGI era, building roboticized "self-driving labs" that greatly accelerate wet-lab science across the U.S. research infrastructure. Second, the National Institute of Standards and Technology (NIST) and the Food and Drug Administration (FDA) should incrementally validate components of the biosimulation stack—regularly determining which applications physical experimentation is still superior for, and where it should be replaced by cheaper and faster simulations. If regulators do not start this work until the technology fully matures, we will lose years of life-saving cures while agencies figure out from scratch whether *in silico* trials are safe enough. Third, since private-sector compute scaling is driven by short-term incentives, there is too little investment in strategies to unlock much more abundant compute in 10–15 years—a classic market failure. For example, federal funding to accelerate fusion power development will likely be needed to meet the electricity demands of affordable biosimulation.

Achieving full-fidelity, full-body biosimulation will give us a platform to replace *in vivo* clinical trials with digital trials that are enormously faster (hours or days, instead of years), safer (no human subjects endangered), and cheaper (millions of dollars, compared to billions between preclinical R&D and actual trials) than the current paradigm. Following initial drug validation, affordable individual trials could be personalized to all known data about a given patient.

Building such a tool would not just supercharge pharmaceutical discovery but allow us to understand human biology much more fundamentally—not as messy clinical approximations but in something like the exhaustive way that aerospace engineers understand the workings of airplanes. This would enable not just breakthrough curative medicines for major diseases, but broad-spectrum defenses against pandemic pathogens, and methods of slowing and reversing damage from biological aging itself.

In this way, a biosimulation megaproject is our realistic path to far longer and healthier lives. We just need the vision to make it a reality.

ABOUT THE AUTHOR

John-Clark Levin is the head of research in Ray Kurzweil's lab, where his work focuses on timelines and roadmaps to AGI, strategies for keeping AGI safely aligned with human wellbeing, and applications of AGI to STEM research, especially in life sciences. As a Yorktown Institute fellow and senior AI advisor at Greenmantle, he regularly advises governments, corporations, and NGOs on AI and its impacts.



From *Futurecraft: The Day After AGI*, Foundation for American Innovation, 2026.
Read or listen online at futurecraft.thefai.org/essays/build-medicines-holy-grail-a-simulated-human-body.

© 2026 by the Foundation for American Innovation. All rights reserved.

7 U.S. Department of Energy, "The Genesis Mission."