

The Next Fifty Years: Drug Discovery in the Age of the Computational Self

Kosi Gramatikoff^{1*}, Ed Judge², Klaus Brusgaard^{3,4}

Abstract

For most of biomedical history, drug discovery has been organized around populations: diseases defined through shared symptoms, therapeutics evaluated in statistical cohorts, approval granted for averaged responses. The paradigm produced extraordinary successes. It also embedded a foundational assumption now under pressure from several independent directions at once. AI-native drug discovery, multi-scale digital twins, pharmacogenomics, patient-derived organoids, adaptive trial design, closed-loop biosensing, decentralized science — these are not isolated innovations. They are convergent expressions of a single directional tendency: the progressive migration of therapeutic authority from population-scale structures toward the individual. This is the biomedical expression of a broader civilizational vector, in which the organizing unit of decision-making resolves continuously toward smaller, more personalized scales. Twelve convergent indicators of this transition are introduced, along with a fifty-year roadmap projecting five successive phases of drug discovery, from AI-assisted precision medicine through full therapeutic sovereignty. The Asymmetric Epistemic Burden (AEB) model formalizes the growing divergence between AI-accelerated hypothesis generation and the bottlenecked experimental and regulatory validation that now constitutes the primary constraint on progress. The defining transformation of twenty-first century medicine may not be the discovery of a better drug class. It may be the conceptual and institutional replacement of the average patient as the fundamental unit of therapeutic reasoning — and the emergence of a pharmacotherapy organized not around populations, but around continuously modeled, computationally instantiated, biologically sovereign individuals.

Keywords: precision medicine, AI drug discovery, digital twins, therapeutic sovereignty, N-of-1 therapeutics, personalized pharmacotherapy, asymmetric epistemic burden, closed-loop therapeutics, decentralized science, pharmacogenomics

Introduction

1. The End of the Average Patient

The concept of the ‘average patient’ has been the organizing fiction of modern pharmacotherapy. From trial design to dosing guidelines, from disease categorization to drug approval, medicine has systematically assumed that interventions optimized for populations would produce acceptable outcomes for the individuals within them. The assumption was not unreasonable. For much of the twentieth century it was not merely defensible but necessary: the tools for characterizing individual biological variation did not exist, and population-level reasoning was the only epistemically tractable approach.

That constraint is dissolving. The simultaneous maturation of genomics, computational biology, AI, and digital sensing is making individual-level characterization increasingly tractable. The biological uniqueness of each patient — genomic profile, transcriptomic state, metabolomic signature, microbiome composition, immunological history, environmental exposure — is becoming computationally accessible in ways inconceivable fifteen years ago [5,6].

These developments are not simply making existing medicine more accurate. They are initiating a structural reorganization of drug discovery itself — a transition from a paradigm organized around populations toward one organized around individuals. The end-point of this trajectory is what we call **therapeutic sovereignty**: the condition in which the individual becomes the primary organizing center of their own therapeutic ecosystem, with institutional structures playing consultative rather than authoritative roles.

The argument develops as follows. Section 2 situates this transition within a broader civ-

¹Villa Veritas, Yanko Slavchev 9, 9000 Varna, Bulgaria, kosi.gramatikoff@gmail.com

²EPAM Systems Ltd, 9 Devonshire Square, 5th Floor, London, EC2M 4YE, UK

³Klinisk Genetik, Vejle Sygehus, Department of Clinical Genetics, Lillebaelt Hospital, Beriderbakken 4, 7100 Vejle

⁴Institute for Regional Health Research, University of Southern Denmark, Campusvej 55, 5230 Odense, Denmark, klaus.brusgaard@rsyd.dk

*Corresponding author:

Kosi Gramatikoff

E-mail: kosi.gramatikoff@gmail.com

DOI: 10.2478/ebtj-2026-0018

© 2026 Authors. This work was licensed under the Creative Commons Attribution-NonCommercial-NoDerivs 3.0 License.

Table 1. Authority Migration Across Civilizational Domains

Domain	Historical Authority Locus	Emerging Authority Locus	Mechanism
Technology	States / Corporations	Individual (smartphone, wearable)	Miniaturization + personalization
AI / Knowledge	Universities / Experts	Personalized AI interface	Algorithmic mediation
Finance / Trust	Banks / States	Protocols / Distributed networks	Blockchain, DeFi
Medicine / Health	Hospital / Population protocols	Individual genomic + digital twin	Precision medicine, AI diagnostics
Energy	National grids / Utilities	Household solar + storage	Distributed generation
Drug Discovery	Big Pharma / Population trials	AI-native N-of-1 design	AI + patient-derived biology
Clinical Trials	Large randomized cohort	Virtual patient / Adaptive N-of-1	Digital twins, organ-on-chip
Identity / Self	Family / Community / Nation	Self-authored identity	Cultural de-institutionalization

ilizational tendency toward individualized authority. Section 3 identifies twelve convergent indicators of the therapeutic transition. Sections 4–15 develop each indicator in detail. Section 16 applies the Asymmetric Epistemic Burden (AEB) model to pharmaceutical R&D. Section 17 presents a fifty-year roadmap. Section 18 addresses governance, equity, and regulatory challenges. Section 19 draws conclusions.

2. A Civilizational Tendency: Authority Migration and Its Biomedical Expression

This transition is not unique to medicine. Across computing, energy, finance, communication, and education, the same directional pattern is visible: authority migrating from centralized, institutional, population-scale structures toward distributed, individualized, computationally mediated systems. The smartphone consolidates capabilities that once required state infrastructure. Blockchain relocates trust from institutions to protocols. Solar energy makes the household a producer. AI transfers expertise from credentialed intermediaries to personalized interfaces [1](Table 2).

Medicine is undergoing the same movement. Therapeutic authority resided historically with the physician, the hospital, the pharmaceutical manufacturer, and the population cohort as the unit of evidence. Each of these structures is now under pressure: genomic testing returns biological information directly to patients; digital twins create individualized computational models; AI-guided diagnostics routinely outperform clinical judgment on specific tasks; decentralized science platforms create alternative pathways for research funding and data governance [7,8].

The convergence is what demands explanation. Technologies with entirely different origins — molecular biology, computer science, materials science, synthetic biology — are all

producing outputs that increase biological resolution, compress the therapeutic scale from populations to individuals, and shift decision-making authority toward the patient.

This convergence is not accidental. It reflects the same structural dynamic visible across other domains: computation becomes cheaper, individualized data becomes abundant, algorithmic mediation becomes pervasive, and the organizing unit of decision-making resolves toward smaller, more personalized scales. Drug discovery is the latest — and potentially the most consequential — domain in which this resolution is occurring.

The 2026 Stanford AI Index documented that AI is now involved in the majority of new drug approvals and that AI-enabled biological hypothesis generation is growing at a rate that outpaces institutional evaluation capacity, a direct empirical manifestation of this civilizational dynamic [4].

3. Twelve Convergent Indicators of the Therapeutic Transition

Twelve convergent developments — arising from different disciplines and different institutional contexts — point toward the same structural destination. We call these the Therapeutic Authority Vector Indicators (TAVIs). Each represents a genuine innovation. Together, they constitute evidence of a directional reorganization of drug discovery (Table 2).

Figure 1 maps the network architecture of the 12 TAVIs across four thematic quadrants — Individual Biology & Simulation; Computational Governance & Authority; Continuous, Autonomous Intervention; New Systems & Modalities — illustrating structural interdependence rather than mere parallelism. The color-coded vectors trace causality chains, feedback loops, shared technological foundations, and evolutionary successions that collectively displace centralized, reactive medical authority toward a distributed, continuous, autonomous ther-

Table 2. The Twelve Therapeutic Authority Vector Indicators (TAVIs)

#	Indicator	Core Transformation	Authority Vector Direction
1	N-of-1 Drug Discovery	Discovery begins from individual biology, not population disease categories	Population → Patient as origin point of therapeutic design
2	Multi-Scale Digital Twins	Therapeutic decisions are pre-evaluated in continuously updated individual computational models	Population evidence → Individual simulation
3	AI-Native Therapeutic Design	AI systems generate, evaluate, and iterate drug hypotheses with minimal human involvement	Institutional expertise → Computational mediation
4	Self-Driving Clinical Trials	Virtual patient models and organ-on-chip replace or precede human cohort trials	Cohort → Simulated individual validation
5	Therapeutic Operating Systems	Drug + sensor + model + adaptation replaces the static pill	Episodic intervention → Continuous computational regulation
6	Closed-Loop Biology	Sense → Predict → Intervene → Update replaces Diagnose → Treat → Wait	Reactive treatment → Continuous autonomous optimization
7	Personal Cell Factories	Patient's biology becomes the substrate of therapeutic production	Centralized manufacturing → Patient-origin production
8	Computational Authority	The model predicts displaces the expert recommends as the locus of therapeutic trust	Institutional authority → Algorithmic authority
9	Search-Driven Discovery	Drug discovery becomes navigation of biological possibility spaces rather than mechanistic inference	Hypothesis-driven → Search-driven exploration
10	Therapeutic Sovereignty	Individuals own their genomic data, digital twins, and health records; patient as principal organizer	Institutional authority → Individual sovereignty
11	Decentralized Science (DeSci)	Blockchain, DAOs, IP-NFTs, and community funding distribute research governance	Centralized pharma → Distributed therapeutic communities
12	Recursive Self-Optimization	Medicine becomes continuous self-improvement; health becomes a dynamic engineering discipline	Disease treatment → Capability enhancement

apeutic ecosystem.

Figure 2 shows how the structural weight of these four clusters transforms over time. The present state is roughly symmetric. The projected future is not: Computational Governance & Authority expands to a dominant position, exerting directed influence over the other three, while Continuous, Autonomous Intervention becomes the operational engine of the system.

4. Indicator 1 — N-of-1 Drug Discovery

Conventional drug development begins with a disease category: Alzheimer's, metastatic lung cancer, type 2 diabetes. Within that category, cohort studies define the target population, preclinical models approximate population-level biology, and clinical trials seek effects that are statistically significant across groups. The individual patient is an entry in a register, a member of a class. Success is measured against a population mean. N-of-1 drug discovery inverts this logic. Discovery begins with the individual: their specific mutational landscape, transcriptomic state, protein expression profile, physiological context. The disease category becomes a provisional label, not a discovery target. The therapeutic candidate is evaluated for predicted

efficacy within a specific biological context rather than a statistical average [9].

Evidence for the viability of this approach is accumulating. In rare disease therapeutics, antisense oligonucleotides have been designed in weeks for individual patients based on their specific mutations [10]. CAR-T cell therapies are inherently individualized — the therapeutic molecule is the patient's own cells, engineered *ex vivo*. Personalized cancer vaccines designed against individual neoantigen profiles are progressing through clinical development. A 2026 review in *Nature Medicine* articulated five fundamental tenets for advancing evidence-based precision medicine, emphasizing precision trial designs and strategies that address heterogeneity in therapeutic response [9,11].

N-of-1 drug discovery is not an incremental improvement on precision medicine. It is a categorical shift in the unit of discovery: from disease to individual, from population to person. The economic consequences run deep. The blockbuster model — one drug for millions of patients — is directly challenged by a model in which therapeutic value derives from biological specificity rather than scale. Existing regulatory frameworks,




Figure 1. Topological Map of the Twelve Therapeutic Authority Vector Indicators (TAVIs) and Their Interconnected Structural Relationships. This schematic diagram maps the network architecture of the 12 TAVIs, organized into a self-organizing map layout that positions conceptually proximal indicators within four distinct thematic quadrants to illustrate the shifting paradigms of modern medicine. The upper-left quadrant anchors Individual Biology & Simulation, capturing the transition from population-level averages to patient-specific modeling. The upper-right quadrant defines Computational Governance & Authority, illustrating how epistemic trust shifts from institutional experts toward algorithmic systems and individual sovereignty. The lower-left quadrant represents Continuous, Autonomous Intervention, mapping the evolution of treatment from episodic, reactive events into real-time computational engineering loops. The lower-right quadrant aggregates New Systems & Modalities, covering the decentralization of manufacturing, knowledge discovery, and research governance. Color-coded vectors represent the structural dynamics between nodes, revealing interdependencies that go beyond sequential progression: solid blue arrows denote direct dependencies and causality, where one indicator is a necessary prerequisite for another; bi-directional orange arrows denote feedback loops and reciprocal optimization between adjacent paradigms; dotted purple lines denote shared technological foundations and common computational infrastructure; solid green and dashed teal arrows trace evolutionary successions and co-dependent relationships across operational domains. Where a node-pair is connected by more than one vector type — as between Closed-Loop Biology (6) and Recursive Self-Optimization (12), which are linked by a solid blue dependency arrow, a bi-directional orange feedback arrow, and a dashed teal succession arrow — each line represents a distinct and independently meaningful relationship rather than a redundant restatement of the same connection: continuous biological regulation is a structural precondition for recursive self-optimization, the two paradigms refine each other's operational targets through ongoing feedback, and self-optimization represents the evolutionary extension of closed-loop regulation beyond disease management. Together, the network visualizes how the twelve indicators converge to displace centralized, reactive medical authority with a distributed, continuous, autonomous therapeutic ecosystem.

designed around population-level evidence, require fundamental adaptation to accommodate this [12].

5. Indicator 2 — Multi-Scale Digital Twins

The digital twin concept — a continuously updated computational replica of a physical system — originated in aerospace and manufacturing. Its application to human biology may be among the most consequential transformations in the history of medicine [13,14]. A medical digital twin integrates multiple

layers of individual data: genomic, transcriptomic, proteomic, metabolomic, physiological, behavioral, environmental. Not a static model. A dynamic one, continuously updated by data streams from wearable sensors, electronic health records, and laboratory measurements.

The therapeutic significance of digital twins lies in their potential to shift the primary site of drug evaluation from biological systems to computational ones. Instead of asking whether a drug works for a population, the digital twin asks whether

Simplification of <TAVI> TAVIs into Thematic Clusters




Figure 2. Evolutionary Dynamics and Structural Reconfiguration of TAVI Clusters from the Present to a Future Sovereign Ecosystem. This two-panel comparative diagram illustrates the macro-level transformation of the four primary TAVI clusters, tracking how their relative structural impact, governing authority, and internal dependencies evolve over time. The left panel depicts the baseline Present State: a symmetric, decentralized configuration in which Individual Biology & Simulation, Computational Governance & Authority, Continuous, Autonomous Intervention, and New Systems & Modalities operate as distinct, largely parallel pillars connected by bidirectional, uniform exchanges (gray arrows), with traditional institutional structures still mediating the interaction between individual patient data and therapeutic development. The right panel models the anticipated Future State, in which the ecosystem undergoes a shift in both scale and directional governance: Computational Governance & Authority expands and elevates to a dominant, top-tier position, exerting directed regulatory and design influence (thick blue down-arrow) over the other domains; Continuous, Autonomous Intervention expands to become the operational engine of the system, fueled by heavy feedback loops (orange vectors) that continuously cycle data from Individual Biology & Simulation into real-time optimization; green and teal successional arrows track how these computational frameworks mature into decentralized, community-driven New Systems & Modalities, shifting the locus of control away from centralized pharma. All arrow colors and styles in both panels follow the same legend used in Figure 1; no additional line types are introduced. The transition depicted moves from static, episodic medical intervention toward a highly integrated, algorithmically governed, self-optimizing therapeutic paradigm.

it works — and at what dose, timing, and co-administration context — for a specific modeled individual. Drug interactions, side effects, pharmacokinetic profiles, and disease progression trajectories can be explored computationally before any biological intervention occurs [15].

Clinical relevance is accelerating. A 2026 *Frontiers in Digital Health* review documented multi-scale digital twin architectures integrating molecular, cellular, tissue, organ, behavioral, and environmental data layers for individualized clinical decision-making [15]. A 2026 review in *Exploration of Digital Health Technology* described the shift from provider-centered to patient-centered medicine mediated through digital twin infrastructure [14]. A review in *Innovation in Health Sciences* identified data interoperability, validation standards, and regulatory acceptance as the primary constraints on digital twin deployment at scale [16]. The *Journal of Medical Internet*

Research published a 2025 perspective on the mathematical foundations of digital twins — differential equations, Bayesian networks, Markov models, reinforcement learning — for individual health trajectory prediction and treatment optimization [17].

The FDA's December 2025 qualification of the first AI-based tool for clinical trial drug development signals that regulatory frameworks are beginning to incorporate computationally mediated evidence [18]. As digital twin fidelity improves and validation methods mature, the twin may become the primary site at which first-in-human therapeutic decisions are effectively made — biological administration confirming computational prediction rather than initiating discovery.

6. Indicator 3 — AI-Native Therapeutic Design

The May 2026 *Nature* papers — Robin, Co-Scientist, and ERA

— mark the emergence of AI-native science: a paradigm in which artificial intelligence functions not as an instrument of inquiry but as an integrated epistemic participant [5,6,7]. Applied to drug discovery, this transition transforms the discovery process from hypothesis-driven experimentation into large-scale computational search. A companion Nature editorial declared that AI scientists are fundamentally changing research and that institutions, funders, and publishers must urgently adapt [19].

Foundation models for biology — large-scale pre-trained architectures learning from genomic, proteomic, transcriptomic, and chemical data — are the technical substrate of AI-native therapeutic design. AlphaFold3, developed by Google DeepMind and Isomorphic Labs, predicts the structures and interactions of virtually all biomolecules with unprecedented accuracy, achieving 76.4% accuracy in protein-ligand docking, a 1.8-fold improvement over prior methods [20,21,22]. RFDiffusion generates novel protein scaffolds using diffusion models. MAMMAL, a molecular multi-modal foundation model, enables cross-modal drug discovery across chemical, protein, and genomic modalities simultaneously [23].

The practical consequences are concrete. Insilico Medicine identified a novel target for idiopathic pulmonary fibrosis and advanced a drug candidate to preclinical trials in 18 months at a cost of approximately \$150,000 — a process that previously required four to six years. Exscientia designed the first AI-designed molecule to enter human clinical trials in under 12 months. A 2025 review in *Patterns* documented the transformation of multiple stages of the drug development pipeline by large language models, including target identification, lead optimization, and clinical trial design [24]. The application of LLMs to drug target discovery is one of the most active current research fronts [25]. A comprehensive 2025 review catalogued AI-driven drug discovery platforms in clinical or preclinical use [26], and a 2024 review in *Pharmaceutics* documented AI applications spanning target identification through clinical optimization for personalized medicine [27].

The structural shift is decisive. When drug discovery becomes search — when the binding constraint is no longer generating hypotheses but evaluating them — the entire organizational logic of pharmaceutical R&D changes. Evaluation here means something more demanding than checking a candidate against known biology, because for most biological processes the underlying mechanism is itself poorly understood. The evaluator is not confirming a result against a known ground truth; they are conducting primary research into systems whose causal structure remains partially opaque even after the experiment is run. That asymmetry — a generator producing candidates at near-zero marginal cost, against an evaluator who must do original biological discovery just to assess each one — is what the Asymmetric Epistemic Burden model formalizes in Section 16. Autonomous agent frameworks are extending the generative side further still: a 2025 review in *Nature Machine Intelligence* documented integrated discovery pipelines orchestrating scientists, language models, code, and physical

simulation with minimal human oversight at individual experimental steps [28].

7. Indicator 4 — Self-Driving Clinical Trials and Virtual Validation

Clinical trials are the most expensive, time-consuming, and ethically consequential stage of drug development. A Phase III trial costs over \$300 million and requires years of patient recruitment and follow-up. The failure rate is extraordinary: roughly 90% of candidates entering clinical development never reach market [29]. AI and digital technologies are beginning to transform all stages of this process.

Virtual patient populations — computational models derived from real-world patient data — can substitute for or precede human cohorts in early trial phases. Digital twin representations of individual patients allow adaptive trial designs that continuously update dosing, endpoint definitions, and patient stratification in real time. Organ-on-chip systems provide human-relevant microphysiological models that improve on animal models for specific, isolated questions — tissue-level toxicity, local drug-receptor interactions — at a fraction of the cost and time [30]. They do not yet capture organismic-level pharmacology: systemic drug metabolism, adverse effects arising from whole-body feedback, and complex multi-organ drug interactions remain outside their scope. Closing that gap is itself an active research problem, not a solved one.

A landmark development came in April 2026, when the FDA formally launched an AI-Enabled Optimization of Early-Phase Clinical Trials Pilot Program, requesting industry input on AI applications including adaptive designs, biomarker-based patient stratification, endpoint validation, and earlier Phase 1 to Phase 2 decision support [18]. In December 2025, the FDA qualified the first AI-based tool for use in drug development clinical trials — a cloud-based platform for liver biopsy scoring in NASH/MASH trials — representing formal regulatory acceptance of AI within the clinical evidence framework [31]. The convergence of virtual patient models, AI-driven trial optimization, and novel trial designs is transforming clinical development from a sequential, cohort-based process into a continuously adaptive, partially simulated one. The fully self-driving trial — in which adaptive AI manages all stages of validation — remains ahead. But the architecture is already being assembled.

8. Indicator 5 — Therapeutic Operating Systems

Pharmacotherapy has historically operated in open loop. A drug is prescribed at a dose, taken on a schedule, and its effects are assessed at the next clinical encounter. Between prescription and reassessment — days, weeks, sometimes months — the patient's physiology changes, the therapeutic response changes, and the intervention remains fixed.

The therapeutic operating system concept challenges this architecture. A therapeutic OS integrates continuous sensing, real-time monitoring, predictive modeling, and adaptive drug delivery into a unified feedback architecture. Within this system, a drug is not a static product but a component of a larger

computational-biological process. Physiological data continuously informs dosing decisions; therapeutic regimens adapt dynamically to real-time biomarker levels; algorithms optimize intervention strategies over time. A 2026 review on precision medicine in medical technology documented the emergence of personalized devices, implantable sensors, adaptive stimulators, and closed-loop drug-delivery platforms as the hardware substrate of this model [32].

The prototype already exists in diabetes management. Closed-loop insulin delivery systems — commonly called the artificial pancreas, though the name is a convenience, not a biological description — integrate continuous glucose monitoring, insulin pump delivery, and control algorithms into a system that adapts dosing continuously without clinical intervention. It replicates one function of the organ, insulin release in response to glucose, through engineering rather than tissue; it is not a synthetic reconstruction of the pancreas itself. A 2025 review in *Nature Reviews Bioengineering* examined smart closed-loop systems integrating real-time biosensing with automated drug delivery, AI-enhanced control algorithms, and synthetic biology components for personalized adaptive therapies across chronic disease domains [33]. Wearable electrochemical biosensors are achieving continuous real-time tracking of multiple drug plasma levels — tobramycin, vancomycin, levodopa — enabling precision pharmacokinetic monitoring outside the clinic [34].

The extensions of this model across oncology, cardiovascular disease, neurodegeneration, and psychiatric illness represent one of the most consequential trajectories in pharmacotherapy. Consider a cancer patient whose tumor microenvironment is continuously monitored, whose treatment regimen adapts continuously to emerging resistance, whose adverse effect profile is continuously managed. Each component already exists. Integration is the challenge.

9. Indicator 6 — Closed-Loop Biology

The therapeutic operating system is a specific expression of a more general principle: the replacement of episodic, open-loop intervention with continuous, closed-loop biological regulation. In an open-loop system, the physician observes, diagnoses, prescribes, and waits — feedback slow, mediated by clinic visits, dependent on patient self-report. In a closed-loop system, sensing, prediction, intervention, and monitoring occur continuously and automatically, with human oversight operating at a higher level of abstraction.

Diabetes care demonstrates this transition most clearly. The progression from fingerstick glucose measurement (open loop) to continuous glucose monitoring (real-time sensing) to the closed-loop artificial pancreas (autonomous intervention) exemplifies the trajectory. The same progression is underway across multiple domains: cardiac monitoring moving toward predictive arrhythmia intervention; wearable neurology enabling closed-loop neuromodulation for epilepsy and Parkinson's disease; psychiatric medicine exploring continuous biomarker monitoring for mood disorder management.

A 2025 review in *Micromachines* documented the transformation of healthcare through wearable and implantable BioMEMS devices, enabling continuous biomarker tracking, closed-loop diagnostic systems, and adaptive therapeutic delivery across metabolic, cardiovascular, and neurological domains [35]. Miniaturized electronics, AI-driven signal processing, advanced materials, and synthetic biology are progressively extending closed-loop regulation beyond glucose toward the full spectrum of physiological control.

10. Indicator 7 — Personal Cell Factories

The industrial pharmacotherapy model was organized around centralized manufacturing: a drug synthesized in a factory, approved by a regulator, distributed through a supply chain, and administered to a patient. The patient was the terminal endpoint — present at consumption, absent from creation. Emerging technologies are beginning to dissolve this boundary.

Autologous cell therapies — CAR-T cells, TIL (tumor-infiltrating lymphocyte) therapies, engineered immune cells — are inherently patient-derived: the therapeutic molecule is produced from the patient's own biology. Patient-derived tumor organoids serve as personalized drug screening platforms, allowing ex vivo testing of treatment regimens against the patient's actual tumor before in vivo administration [36]. Personalized mRNA vaccines against tumor neoantigens are designed de novo from each patient's mutational landscape. Induced pluripotent stem cell (iPSC) technology enables patient-derived production of virtually any cell type for therapeutic application. Patient-derived zebrafish xenograft avatars provide a complementary platform for rapid in vivo therapeutic decision support, though the value of any cross-species avatar is bounded by genomic divergence from the host species — zebrafish, like other animal models, differ substantially from humans in non-coding genome content, limiting what such systems can reveal about regulatory biology even when the implanted tumor cells are themselves patient-derived [37].

The integration of AI into patient-derived therapeutic pipelines is accelerating. A 2026 study in *npj Drug Discovery* demonstrated that AI-guided small molecule design targeting individualized tumor immunomodulation profiles can substantially compress the time from patient biopsy to therapeutic candidate identification [38]. Comprehensive reviews have documented the convergence of AI design, patient-derived biology, and personalized delivery into an emerging unified pipeline [39].

The implications reach beyond oncology. When the patient's biology — cells, proteins, nucleic acids — becomes the primary substrate for therapeutic production, the supply chain of medicine runs through the individual rather than around them. Drug production migrates toward the point of care. Therapeutic design becomes inseparable from individual biology.

11. Indicator 8 — Computational Authority

One of the least examined but most consequential shifts in medicine is the substitution of algorithmic authority for expert

authority. The physician's recommendation has historically been the primary interface between medical knowledge and individual treatment — authority derived from training, experience, credentialing, community recognition. Constitutively human.

The emergence of AI decision support systems, and their progressive improvement relative to human clinical judgment across multiple domains, introduces a new form of authority: computational authority. The trusted entity becomes *the model predicts* rather than *the expert recommends*. Neither institutional authority nor individual authority in the traditional sense. Algorithmic authority — mediated by systems whose outputs are often more accurate than human judgment but whose reasoning processes are frequently opaque.

A 2026 Frontiers review documented multi-modal AI integration of genomics, imaging, and EHR data for clinical insights, demonstrating AI assistance in disease risk assessment, early diagnosis, drug discovery, and therapeutic optimization [40]. A 2025 review of AI and multi-omics in pharmacogenomics in Mayo Clinic Proceedings documented the progression from foundational genomics through regulatory adoption to real-world implementation of FDA-cleared AI decision-support panels [41]. AI-powered genomic analysis is increasingly assisting variant classification in oncology settings, improving consistency and throughput where human interpretation remains the primary, though now partially augmented, constraint [42].

As computational authority becomes more accurate and more pervasive, the physician's role evolves from primary therapeutic authority toward expert oversight, quality control, and ethical stewardship of algorithmic outputs. Not a diminishment. A reorganization. The governance challenge this creates — establishing accountability frameworks for AI clinical decisions — is addressed in Section 18.

12. Indicator 9 — Search-Driven Drug Discovery

The classical model of drug discovery was hypothesis-driven. A researcher formed a mechanistic theory — this protein is involved in this pathway and is abnormally regulated in this disease — and designed experiments to test it. Hypothesis preceded experiment. The research program was organized around theoretical propositions subsequently confronted with empirical reality.

AI-native drug discovery fundamentally alters this logic. When large-scale generative models can propose millions of candidate molecules against a given target, when virtual screening can rapidly evaluate candidates across multiple parameters, when self-driving laboratories can physically test compounds at a rate of thousands per day, drug discovery increasingly resembles search — the systematic navigation of vast biological possibility spaces — rather than classical hypothesis-driven inquiry [5,26]. Comprehensive analyses of AI drug discovery platforms document this structural shift across the full pipeline, from target identification to lead optimization [39].

This is not merely a change in technique. It is a change in

the epistemological structure of discovery. In search-driven discovery, the causal story is often reconstructed post hoc, after a candidate has been identified by screening rather than predicted by theory. The relationship between mechanism and discovery is inverted. Understanding follows finding rather than preceding it. A 2025 review in Nature Machine Intelligence documented how autonomous agent frameworks — orchestrating language models, code execution, and physical simulation — are transforming scientific discovery into a recursive computational search process [28]. What these systems do well is navigate a space of possibilities already implicit in existing data and known chemistry. What they do not yet do is originate the kind of conceptual reframing — a new hypothesis nobody thought to test, an analogy drawn from an unrelated field — that has historically driven the largest leaps in biology. The reduction in human involvement is real at the level of individual experimental decisions. It is not yet a reduction in the need for human imagination at the level of framing the problem.

The consequences for pharmaceutical R&D are structural. As the cost of generating drug candidates collapses — AI reducing the marginal cost of hypothesis generation to near zero — the binding constraint on innovation shifts from discovery to validation. This is the central claim of the Asymmetric Epistemic Burden model, developed in Section 16.

13. Indicator 10 — Therapeutic Sovereignty

Therapeutic sovereignty describes the emerging condition in which individuals acquire the capacity to organize, direct, and optimize their own therapeutic ecosystems. Not a technological development in the conventional sense. A reorganization of the authority structure of medical decision-making.

Multiple technological threads converge toward this condition. Direct-to-consumer genomic testing provides individuals with detailed information about their pharmacogenomic profile — variants in CYP450 enzymes, drug transporter genes, disease-associated loci — without clinical mediation. Wearable sensors generate continuous physiological data that belongs to the wearer. Digital health records, mandated to be patient-accessible in many jurisdictions, return individuals' medical histories to their control. These three developments are already in place. What is not yet in place, but is emerging, is a fourth: AI-guided decision support tools capable of translating this information directly into actionable therapeutic guidance, independent of clinical interpretation. Today's tools largely assist clinicians rather than patients directly; the shift toward patient-facing, autonomous decision support remains a near-term trajectory rather than a current capability.

At the institutional level, decentralized science is creating alternative architectures for research funding, data governance, and drug development that distribute authority away from centralized pharmaceutical companies and regulatory agencies. VitaDAO, AthenaDAO, Molecule Protocol, and Hippocrat — patient-owned health data networks on blockchain — represent early instantiations of distributed therapeutic governance [43,44].

The endpoint — a patient who possesses their complete multi-omic profile, maintains a continuously updated digital twin, participates in decentralized research communities, and uses AI decision support to direct their own therapeutic strategy — represents a qualitative shift in the authority structure of medicine that no single technology alone produces. It is what all twelve indicators, acting simultaneously, converge toward.

14. Indicator 11 — Decentralized Science (DeSci)

The pharmaceutical industry is organized around centralized authority. Large companies fund research, patent discoveries, and control production. Academic institutions operate within funding structures controlled by government agencies whose paylines have tightened sharply. Regulatory agencies maintain gatekeeping authority over what may be administered to patients.

Decentralized science challenges this architecture through blockchain-enabled mechanisms that distribute research funding, data governance, intellectual property management, and publication. Decentralized Autonomous Organizations (DAOs) pool community funds for research projects through governance voting. IP-NFTs tokenize intellectual property, allowing researchers to prove ownership and share revenue. Patient-owned data networks give individuals control over their health data and the ability to selectively share it for research purposes [43,44].

As of early 2026, approximately 50 active DeSci initiatives span funding DAOs, publishing platforms, healthcare data systems, and IP management. The convergence of collapsed federal research paylines — NIH success rates for early-stage investigators falling from 26% in 2024 to 19% in 2025 — with the maturation of DeSci infrastructure is accelerating the migration of research governance away from centralized institutions [43].

The pharmaceutical implications are potentially transformative. A DeSci ecosystem in which patients own their biological data, contribute it to research DAOs, and participate in the governance of therapeutic pipelines would represent a drug discovery model organized around patient sovereignty rather than corporate control. The distance between this vision and current reality remains substantial. It also conceals a democratic problem the optimistic framing tends to skip: meaningful participation in DAO governance, token economics, and data-sharing decisions demands financial literacy, technical fluency, and time that most patients do not have in surplus. A sovereignty model that requires patients to become part-time governance experts in their own biology risks reproducing the same resource asymmetry it claims to dissolve — substituting technical and administrative capital for the institutional capital it displaces. Whether DeSci infrastructure can be designed to lower this burden, rather than simply relocate it, is an open question addressed further in Section 18.2.

15. Indicator 12 — Recursive Self-Optimization

The final indicator is the most radical. The first eleven describe

a trajectory whose endpoint is the replacement of disease treatment with continuous self-optimization. If each individual possesses a comprehensive multi-omic profile, a continuously updated digital twin, a closed-loop sensing and intervention system, and AI decision support, the object of medicine shifts from eliminating pathology to continuously enhancing capability.

This is already visible in certain domains. The closed-loop insulin system does not merely treat diabetes; it optimizes metabolic function continuously. Performance optimization medicine, precision nutrition, pharmacological cognitive enhancement, longevity-oriented interventions — each extends this logic beyond managing disease toward engineering wellbeing. This extension creates a problem the closed-loop insulin case does not: diabetes management operates within a tightly regulated evidentiary framework, while much of the wellness-optimization space does not. As medicine's boundary expands into domains historically occupied by unregulated supplements, biohacking communities, and direct-to-consumer health claims, the non-expert is left to distinguish validated intervention from marketing without a reliable arbiter. No institution currently occupies the role of objective guide across this expanded territory — physicians are not trained for it, regulators do not yet cover it, and platforms hosting health claims have no obligation to adjudicate them. This is not a peripheral concern. It is a precondition for the legitimacy of recursive self-optimization as a medical category rather than a consumer one.

If the line between disease management and capability optimization dissolves, the social compact of medicine — organized historically around the treatment of pathology and the relief of suffering — requires fundamental renegotiation. Questions of equity, enhancement justice, and the social meaning of health move from the margins to the center. The pharmaceutical industry and the regulatory apparatus are not equipped to answer them [45,46].

16. The Asymmetric Epistemic Burden in Pharmaceutical R&D

The transformation described in Sections 4–15 introduces a structural imbalance in the pharmaceutical research enterprise, formalized here as the Asymmetric Epistemic Burden (AEB) model [2,3] (Table 3). Three functional roles are identified in any knowledge-production system:

- **The Generator (G):** produces candidate drug molecules, mechanisms, and therapeutic hypotheses
- **The Explorer (E):** navigates biological possibility spaces and identifies promising candidates
- **The Evaluator (V):** validates candidates against biological reality through experimental and clinical testing

The AEB is formally defined as the differential burden imposed on the Evaluator relative to the Generator as generative capacity scales. If generative throughput $G(t)$ grows at approximately the rate of AI capability improvement — exponential since 2020 — while evaluative capacity $V(t)$ grows at the pace of experimental infrastructure expansion and expert judgment

Table 3. Asymmetric Epistemic Burden (AEB) Model Applied to Pharmaceutical R&D.

Role	Function	Pre-AI Scarcity	Post-AI Scarcity	AEB Status
Generator (G)	Molecular hypothesis, target ID, lead generation	High — bottlenecked by human creativity, labor, and chemical library size	Very low — AI reduces marginal cost of candidate generation toward zero	Collapsing
Explorer (E)	Virtual screening, property prediction, ADMET optimization	High — requires domain expertise, experimental resources, iterative cycles	Medium — AI automates much; target context and strategy still matter	Narrowing
Evaluator (V)	Biological assay, IND-enabling studies, clinical trials, regulatory review	Medium — experimental validation was also limiting	Critical — now the primary bottleneck across the entire pipeline	CRITICAL

development — approximately linear — then:

$AEB(t) \propto G(t) / V(t)$ The AEB widens over time. This is not a transitional state but a structural feature of any system where generation is dominated by statistical pattern-matching while evaluation requires causal understanding, experimental validation, and regulatory judgment. The 2026 Stanford AI Index confirms this dynamic empirically: AI-generated scientific outputs are now growing at a rate that substantially exceeds the growth of peer review, experimental validation, and regulatory review capacity [4].

The pharmaceutical implications are direct. The historical attrition rate — approximately 1 in 5,000 preclinical compounds reaching market [29] — does not improve through accelerating generation alone if validation infrastructure does not scale proportionally. AI systems that generate ten times more candidates per year increase the absolute volume of false positives requiring expensive experimental work, not merely the number of true positives reaching patients.

The organizations and nations that define the next era of pharmaceutical innovation will not be those that generate the most AI-designed molecules. They will be those that solve the validation bottleneck — investing in AI-assisted experimental validation, regulatory science, digital trial infrastructure, and institutional frameworks for evaluating AI-native therapeutics. The self-driving laboratory paradigm, integrating AI hypothesis generation with robotic experimental execution and automated data analysis, is the most promising current direction for narrowing the AEB on the evaluative side [47,48,49,50].

17. A Fifty-Year Roadmap: From Population Medicine to Therapeutic Sovereignty

The twelve indicators and the AEB model converge on a directional trajectory whose stages can be projected across the coming fifty years. This roadmap does not predict specific technologies. It identifies the organizational logic and structural prerequisites for each successive phase, based on trends visible in 2026 (Table 4). Figure 3 illustrates projected temporal trajectories and shifting systemic influence of the four TAVI clusters. Figure 4 tracks the micro-trajectories of all 12 individual TAVIs

across the same timeline.

A note on the roadmap figures. Figures 3, 4, and 5 are conceptual projections, not empirical forecasts. The phase boundaries in Table 4, and the curve shapes, inflection points, and intercepts shown in Figures 3 through 5, are derived from the structural logic of the TAVI framework and the AEB model developed in this review — not from time-series data, simulation, or quantitative forecasting models. Where curve shapes are sigmoidal, this reflects a modeling choice consistent with standard technology-adoption dynamics: gradual S-curves for infrastructure requiring institutional and regulatory uptake, and steeper, earlier-inflecting curves for capabilities that scale with computational cost reduction, as formalized in Section 16. These figures should be read as visual arguments for the directional claims made throughout this review, not as quantitative predictions of timing, magnitude, or rate. The fifty-year horizon is organized around structural prerequisites — what must be institutionally and technically true for each phase to begin — rather than calendar forecasting, and the specific year ranges in Table 4 should be treated as illustrative rather than precise.

18. Governance, Equity, and Regulatory Challenges

18.1 The Regulatory Lag

Regulatory agencies designed to evaluate population-level evidence now confront a therapeutic paradigm generating individualized, computationally mediated, rapidly adapting innovations. The FDA's 2025 draft guidance on AI in regulatory decision-making is a step. The gap between regulatory adaptation and technological development remains substantial [12,18].

The central regulatory challenge is not technical. It is epistemological. Existing frameworks ask whether a therapy works for a population and is safe on average. They are poorly equipped to evaluate whether a therapy works for a specific individual modeled by a digital twin, whether an AI-generated hypothesis meets evidentiary standards, or whether a continuously adaptive system has maintained an acceptable safety profile over time. New paradigms — adaptive approval, real-world

Table 4. The Fifty-Year Roadmap of Drug Discovery and Pharmacotherapy

Phase	Defining Transformation	Key Technologies and Indicators	Dominant Constraint
Phase 1 2025–2035 AI-Assisted Precision Medicine	AI augments but does not replace human therapeutic decision-making. Multi-omics profiling becomes standard clinical practice. Pharmacogenomics guides most drug prescribing in high-income settings. AI drug discovery pipelines enter early clinical development.	Foundation models for biology; AlphaFold3; multi-omics integration; FDA AI guidance; pharmacogenomics panels	Regulatory adaptation; clinical implementation infrastructure
Phase 2 2035–2045 Digital Twin Medicine	Individual digital twins become the standard site of first-pass therapeutic evaluation. Clinical trials incorporate virtual patient components. Simulation precedes most in vivo human testing. Closed-loop systems manage chronic disease across major therapeutic areas.	Multi-scale digital twins; organ-on-chip; AI clinical trials; closed-loop biosensing; wearable pharmacokinetics	Twin validation methods; regulatory acceptance of simulated evidence
Phase 3 2045–2055 N-of-1 Therapeutic Design	Drug discovery routinely originates from individual biological contexts. Patient-derived biology becomes a standard manufacturing substrate. N-of-1 adaptive trials replace traditional cohort designs in most therapeutic areas. Personal cell factories are clinic-adjacent.	N-of-1 discovery platforms; patient-derived organoids; personalized vaccines; autologous cell therapies; iPSC manufacturing	Manufacturing scale; cost of individual production; intellectual property
Phase 4 2055–2065 Adaptive Therapeutic Ecosystems	Therapy becomes a continuous process rather than an episodic event. Therapeutic operating systems integrate sensing, modeling, intervention, and adaptation into unified ecosystems managed by AI. Computational authority largely replaces physician authority for routine decisions.	Therapeutic OS; recursive self-optimization; AI physicians; real-time pharmacogenomics; ambient biosensing	Algorithmic accountability; human oversight architecture; equity of access
Phase 5 2065–2075 Therapeutic Sovereignty	Individuals organize, direct, and own their therapeutic ecosystems. DeSci provides alternative governance for therapeutic development. The distinction between disease management and capability optimization dissolves. The average patient ceases to exist as a meaningful scientific construct.	Full DeSci; patient data sovereignty; therapeutic DAOs; recursive self-optimization; AI-native governance	Normative frameworks; enhancement justice; societal adaptation

evidence integration, algorithmic credibility assessment — are beginning to emerge. They require substantial development before they can support the therapeutic sovereignty model [46].

18.2 Equity and the Personalization Dividend

Personalized medicine creates a fundamental equity challenge. If therapeutic optimization requires extensive individual biological data, expensive computational infrastructure, and AI-guided decision support, the benefits of this transition

will accrue disproportionately to those with the resources to access them. The history of medical innovation provides little comfort: precision oncology has substantially widened survival gaps between high- and low-income patients despite its aggregate population benefit.

Addressing this requires deliberate intervention: public investment in data infrastructure that gives all patients access to their own biological information; regulatory frameworks ensuring AI-guided decision support is available across health-




Figure 3. Projected Temporal Trajectories and Relative Systemic Impact of TAVI Clusters Over a 50-Year Horizon. Individual Biology & Simulation (light blue) initiates from a strong baseline before plateauing as digital modeling becomes standardized infrastructure. Computational Governance & Authority (purple) exhibits a steep sigmoidal acceleration, rapidly overtaking other domains to establish itself as the primary locus of systemic control. Continuous, Autonomous Intervention (orange) rises steadily in lockstep with governing computational frameworks. New Systems & Modalities (teal) undergoes a protracted incubation period before accelerating mid-horizon, signaling the systemic shift toward distributed, patient-centric therapeutic ecosystems.




Figure 4. Micro-Trajectory Projections and Architectural Shift of the Twelve Individual TAVIs Across a 50-Year Timeline. This diagram tracks the individual evolutionary paths of all 12 TAVIs from the present baseline to a highly integrated 50-year future state, using a unified four-cluster color-coding scheme. Near-term, N-of-1 Drug Discovery, Multi-Scale Digital Twins, and AI-Native Therapeutic Design dominate as baseline drivers. Over the mid-to-long term, Closed-Loop Biology, Computational Authority, and Recursive Self-Optimization experience steep sigmoidal acceleration to form a highly optimized regulatory core. The DeSci and Therapeutic Sovereignty trajectories ascend rapidly in the final decades, delivering systemic control into distributed, patient-centric networks.




Figure 5. Macro-Evolutionary Convergence and Long-Term Systemic Stabilization of Consolidated TAVI Functional Layers (2026–2076). This long-term macro-projection illustrates the consolidation of the twelve TAVIs into two functional layers, grouped according to the “Authority Vector Direction” classification in Table 2: indicators describing a shift in the origin of therapeutic authority (data and biological substrate ownership — Indicators 1, 2, 7, 10, 11) constitute Group A, the Decentralized Biocomputational Substrate; indicators describing a shift in the execution of therapeutic authority (computational control and autonomous intervention — Indicators 3, 5, 6, 8, 9, 12) constitute Group B, Autonomous Algorithmic Orchestration. The horizontal axis represents time from a 2026 baseline; the vertical axis is a qualitative index of relative systemic impact, not a measured or empirically calibrated quantity. Group A's trajectory is modeled as a gradual sigmoid, consistent with historical technology-adoption patterns for infrastructure requiring regulatory and institutional uptake (e.g., electronic health records, genomic testing). Group B's trajectory is modeled as a steeper, earlier-inflecting sigmoid, consistent with the Generator-side growth dynamics described in the Asymmetric Epistemic Burden model (Section 16), where algorithmic capability scales faster than institutional adoption. The curves' intercept represents the point at which algorithmic execution capacity reaches parity with the maturity of the underlying data substrate; the subsequent flattening of both curves reflects saturation of the addressable institutional transition, consistent with standard S-curve technology diffusion dynamics. This figure is a conceptual, illustrative model derived from the TAVI framework and the AEB model developed elsewhere in this review. It is not based on empirical time-series data, simulation, or quantitative forecasting, and the specific curve shapes, inflection timing, and intercept point are heuristic rather than measured.

care settings; international governance preventing the concentration of therapeutic computing capacity in a small number of nations or corporations.

18.3 Authorship, Attribution, and the AI Therapeutic

When a drug candidate is identified by an AI system trained on proprietary biological data, designed by a generative model, and validated by a self-driving laboratory, who owns the resulting intellectual property? The question is not merely legal but ethical. Existing IP frameworks assume human inventors and institutional assignees. They were not designed for AI-native discovery environments in which the line between human and machine contribution to therapeutic innovation is systematically unclear [45,47].

A second and more immediate problem precedes the question of ownership: the question of incentive. As the population of rare and ultra-rare disorders requiring variant-specific compounds continues to grow, the per-patient economics of development become harder to justify under any conventional return model. A compound designed for one patient, or a handful, may never recoup its development cost regardless of who is credited as inventor. AI may lower the marginal cost of designing such a compound, but it does not resolve who bears the cost of validating, manufacturing, and delivering it when the addressable market is a single person. Without new financing models — public subsidy, pooled risk across rare-disease portfolios, or some restructuring of how value is captured from N-of-1 therapeutics — the IP question may end up moot for a

large share of cases: there will be no compound to own.

The emergence of AI as an autonomous contributor to scientific discovery raises parallel questions about authorship in academic publication — though not the question the phrase “AI authorship” suggests. Most journals, following ICMJE guidance, explicitly exclude AI systems from authorship on the grounds that authorship requires accountability for the work, which a model cannot hold. The actual point of contention is narrower and more practical: how to disclose and attribute substantial AI contribution to a paper's methods, analysis, or even hypothesis generation, without either overstating the tool's role or concealing it. Wang and Gong's 2026 cross-disciplinary analysis documents growing inconsistency across journals, funders, and institutions in how this disclosure is handled in practice [46]. Nature's 2026 editorial on AI scientists frames the resulting attribution problem as a defining institutional challenge for the research ecosystem — one the pharmaceutical sector will confront in its most acute form as AI-native therapeutics, designed and validated with minimal human intervention at individual steps, enter regulatory review, where the question is not who is the author but who is accountable for the result [19].

Conclusion

The history of pharmacotherapy can be understood as a progressive reduction in therapeutic scale, from populations to subpopulations, from subpopulations to individuals, and ultimately from individuals to continuously updated computational-biological representations of individuals. The most important discovery of the next fifty years may therefore not be a molecule. It may be the replacement of the average patient as the fundamental unit of therapeutic reasoning.

The transition from population therapeutics to therapeutic sovereignty is not a prediction about a possible future. It is a description of a directional tendency already visible in twelve convergent developments across pharmaceutical science, computational biology, digital health, and decentralized governance.

The drugs of the future will not be merely more selective or more potent. They will be components of continuous therapeutic processes — computational-biological ecosystems organized around the individual, updated continuously rather than prescribed episodically, governed increasingly by the patient rather than by external authorities.

The institutional challenge is substantial. Regulatory frameworks, clinical training, pharmaceutical business models, patent systems, equity frameworks — all were designed for a medicine organized around populations. Adapting them for a medicine organized around individuals who are computationally instantiated, biologically sovereign, and therapeutically autonomous is the central institutional task of twenty-first century pharmacotherapy.

The Asymmetric Epistemic Burden model identifies where the binding constraint now lies: not in generating therapeutic

ideas — which AI has made cheap and abundant — but in validating those ideas against biological reality [2,3]. The organizations, institutions, and nations that invest in evaluation as deliberately as they have invested in generation will define the next era of pharmaceutical innovation. Figure 5 offers a conceptual map of how these forces might converge into two stabilizing layers: a decentralized biocomputational substrate and an autonomous algorithmic orchestration layer — the long-term architecture of therapeutic sovereignty.

The average patient is ending. The sovereign patient is beginning. What comes next will be unlike anything medicine has seen before, and yet, in its deepest logic, it is the continuation of a process that has been underway for centuries: the progressive resolution of authority toward the individual.

Declaration / Disclaimer:

The authors (including Ed Judge) declare that the research, analysis, and opinions presented in this article are entirely their own.

The content of this publication does not reflect the views, opinions, or official stance of EPAM Systems, Inc. (EPAM).

Ethical Approval

Not applicable.

Acknowledgements

The author thanks the editors of Academia Drug Development and Pharmacotherapy for their engagement with this framework. The theoretical foundations of this review draw on the Authority Vector framework developed in the author's prior work [1].

Conflict of Interest

The author declares no competing financial interests. All ideas and frameworks presented are original.

Funding

The authors received no financial support for the research.

References

1. Gramatikoff K. The Authority Vector: A Unified Theory of Civilizational Direction from Transcendence to the Self. PhilPapers. 2026. Available from: <https://philpapers.org/rec/GRATAV-2>
2. Gramatikoff K. AI and the Coming Evaluation Crisis in R&D: How AI Restructures the Economics of Knowledge in Science and Pharma R&D [Asymmetric Epistemic Burden]. Medium. 2026. Available from: <https://medium.com/p/e127cf62a9b3>

3. Gramatikoff K. Asymmetric Epistemic Burden in Academic Research. Medium. 2025. Available from: <https://medium.com/p/7f02694e58ab>
4. Gil Y, Perrault R, et al. Artificial Intelligence Index Report 2026. Stanford Human-Centered Artificial Intelligence (HAI). 2026. Available from: <https://hai.stanford.edu/ai-index/2026-ai-index-report>
5. Ghareeb AE, Chang B, Mitchener L, et al. A multi-agent system for automating scientific discovery. *Nature*. 2026. Available from: <https://doi.org/10.1038/s41586-026-10652-y>
6. Gottweis J, Weng W-H, Daryin A, et al. Accelerating scientific discovery with an AI co-scientist. *Nature*. 2026. Available from: <https://doi.org/10.1038/s41586-026-10644-y>
7. Aygün E, Belyaeva A, Comanici G, et al. An AI system to help scientists write expert-level empirical software. *Nature*. 2026. Available from: <https://doi.org/10.1038/s41586-026-10658-6>
8. Lu C, et al. Towards end-to-end automation of AI research (The AI Scientist). *Nature*. 2026. Available from: <https://doi.org/10.1038/s41586-026-10265-5>
9. Coral DE, Sargent JL, Carrasco-Zanini J, et al. Five tenets for advancing evidence-based precision medicine. *Nat Med*. 2026. Available from: <https://doi.org/10.1038/s41591-026-04309-6>
10. Kim J, Hu C, Moufawad El Achkar C, Black LE, Douville J, Larson A, et al. Patient-customized oligonucleotide therapy for a rare genetic disease. *N Engl J Med*. 2019. Available from: <https://doi.org/10.1056/NEJMoa1813279>
11. Adam D. The AI co-scientist is here. *Nat Med*. 2026. Available from: <https://doi.org/10.1038/s41591-026-04275-z>
12. U.S. Food and Drug Administration. Considerations for the Use of Artificial Intelligence to Support Regulatory Decision-Making for Drug and Biological Products. Draft Guidance. FDA. 2025. Available from: <https://www.fda.gov/regulatory-information/search-fda-guidance-documents/considerations-use-artificial-intelligence-support-regulatory-decision-making-drug-and-biological>
13. Saratkar SY, Langote M, Kumar P, et al. Digital twin for personalized medicine development. *Front Digit Health*. 2025. Available from: <https://doi.org/10.3389/fdgth.2025.1583466>
14. Mazuz K, Biswas S. Towards the future of personalized medicine: digital twin technology. *Explor Digit Health Technol*. 2026. Available from: <https://doi.org/10.37349/edht.2026.101181>
15. Vallee A. Multi-scale digital twins for personalized medicine. *Front Digit Health*. 2026. Available from: <https://doi.org/10.3389/fdgth.2026.1753906>
16. Menon G, Malave B, Mhaske M, et al. Digital twin technologies in medicine: innovations, barriers, and future directions. *Innov Health Sci*. 2026. Available from: <https://doi.org/10.1016/j.inhs.2025.100043>
17. Vallée A. Digital twins for personalized medicine require epidemiological data and mathematical modeling. *J Med Internet Res*. 2025. Available from: <https://doi.org/10.2196/72411>
18. U.S. Food and Drug Administration. AI-Enabled Optimization of Early-Phase Clinical Trials Pilot Program; Request for Information. *Fed Regist*. 2026. Available from: <https://www.federalregister.gov/documents/2026/04/29/2026-08281>
19. Nature Editorial Board. AI scientists are changing research: institutions, funders and publishers must respond. *Nature*. 2026. Available from: <https://www.nature.com/articles/d41586-026-00934-w>
20. Abramson J, Adler J, Dunger J, Evans R, Green T, Pritzel A, et al. Accurate structure prediction of biomolecular interactions with AlphaFold 3. *Nature*. 2024;630(8016):493-500. Available from: <https://doi.org/10.1038/s41586-024-07487-w>
21. Jumper J, Evans R, Pritzel A, Green T, Figurnov M, Ronneberger O, et al. Highly accurate protein structure prediction with AlphaFold. *Nature*. 2021;596(7873):583-589. Available from: <https://doi.org/10.1038/s41586-021-03819-2>
22. Fang Z, Ran H, Zhang Y, Chen C, Lin P, Zhang X, Wu M. AlphaFold 3: an unprecedented opportunity for fundamental research and drug development. *Precis Clin Med*. 2025. Available from: <https://doi.org/10.1093/pcmedi/pbaf015>
23. Shoshan Y, Raboh M, Ozery-Flato M, et al. MAMMAL, Molecular Aligned Multi-Modal Architecture and Language for biomedical discovery. *npj Drug Discov*. 2026. Available from: <https://doi.org/10.1038/s44386-026-00047-4>
24. Zheng Y, Koh HY, Ju J, et al. Large language models for drug discovery and development. *Patterns*. 2025. Available from: <https://doi.org/10.1016/j.patter.2025.101346>
25. Zhao X, et al. Application of artificial intelligence large language models in drug target discovery. *Front Pharmacol*. 2025. Available from: <https://doi.org/10.3389/fphar.2025.1597351>
26. Ferreira FJN, Carneiro AS. AI-driven drug discovery: a comprehensive review. *ACS Omega*. 2025. Available from: <https://doi.org/10.1021/acsomega.5c00549>
27. Serrano DR, Luciano FC, Anaya BJ, et al. Artificial intelligence applications in drug discovery and drug delivery: revolutionizing personalized medicine. *Pharmaceutics*. 2024;16(10):1328. Available from: <https://doi.org/10.3390/pharmaceutics16101328>
28. Zhou L, Ling H, Fu C, et al. Autonomous agents for scientific discovery: orchestrating scientists, language, code, and physics. *Nat Mach Intell*. 2025. Available from: <https://doi.org/10.48550/arXiv.2510.09901>
29. Kola I, Landis J. Can the pharmaceutical industry reduce attrition rates? *Nat Rev Drug Discov*. 2004;3(8):711-715. Available from: <https://doi.org/10.1038/nrd1470>
30. Ma C, Zhang H, Rao Y, et al. AI-driven virtual cell models in preclinical research: technical pathways, validation

- mechanisms, and clinical translation potential. *npj Digit Med*. 2025. Available from: <https://doi.org/10.1038/s41746-025-02198-6>
31. Sergi CM, Sesso HD. Artificial intelligence and the future of clinical trials. *Contemp Clin Trials Commun*. 2025. Available from: <https://doi.org/10.1016/j.conctc.2025.101545>
 32. Qoronfleh MW. Precision medicine in medical technology (medtech): the future of healthcare through personalized devices. *Health Technol (Berl)*. 2026. Available from: <https://doi.org/10.1007/s12553-026-01065-y>
 33. Paci MM, Saha T, Djassemi O, et al. Smart closed-loop drug delivery systems. *Nat Rev Bioeng*. 2025. Available from: <https://doi.org/10.1038/s44222-025-00328-z>
 34. Duan H, Peng S, He S, et al. Wearable electrochemical biosensors for advanced healthcare monitoring. *Adv Sci*. 2024. Available from: <https://doi.org/10.1002/advs.202411433>
 35. Abhinav V, Basu P, Verma SS, et al. Advancements in wearable and implantable BioMEMS devices: transforming healthcare through technology. *Micromachines (Basel)*. 2025;16(5):522. Available from: <https://doi.org/10.3390/mi16050522>
 36. Xiao Y, Li Y, Jing X, et al. Organoid models in oncology: advancing precision cancer therapy and vaccine development. *Cancer Biol Med*. 2025. Available from: <https://doi.org/10.20892/j.issn.2095-3941.2025.0127>
 37. Fontes M, Garcez D, Ruivo J, et al. Glioblastoma zebrafish avatars guide therapeutic decisions in a patient with gliosarcoma: a case report. *Commun Med (Lond)*. 2025;6:53. Available from: <https://www.nature.com/articles/s43856-025-01314-7>
 38. Sutanto H, Fetaryani D. Integrating artificial intelligence into small molecule development for precision cancer immunomodulation therapy. *npj Drug Discov*. 2026. Available from: <https://doi.org/10.1038/s44386-025-00029-y>
 39. Lopes AB, Rodrigues CF, Rodrigues FAM. From algorithm to medicine: AI in the discovery and development of new drugs. *AI (Basel)*. 2026;7(1):26. Available from: <https://doi.org/10.3390/ai7010026>
 40. Khan SN, Danishuddin, Khan MWA, Guarnera L, Akhtar SMF. Multi-modal AI in precision medicine: integrating genomics, imaging, and EHR data for clinical insights. *Front Artif Intell*. 2026. Available from: <https://doi.org/10.3389/frai.2025.1743921>
 41. Zack M, Stupichev DN, Moore AJ, et al. Artificial intelligence and multi-omics in pharmacogenomics: a new era of precision medicine. *Mayo Clin Proc Digit Health*. 2025. Available from: <https://doi.org/10.1016/j.mcp-dig.2025.100246>
 42. Srivastava R. Advancing precision oncology with AI-powered genomic analysis. *Front Pharmacol*. 2025. Available from: <https://doi.org/10.3389/fphar.2025.1591696>
 43. Hussain A, Dershem M, Nasr J, et al. Why pharmacoepidemiology must embrace patient data sovereignty and decentralized technologies. *Blockchain Healthc Today*. 2026. Available from: <https://blockchainhealthcareday.com/index.php/journal/article/view/495>
 44. Kumar HA, Venkatram CP, Saran N, et al. Decentralised digital health ecosystems: a unified architecture for AI-enhanced medical record management. *Front Digit Health*. 2025. Available from: <https://doi.org/10.3389/fdgth.2025.1685628>
 45. Levy N. Responsibility is not required for authorship. *J Med Ethics*. 2024;51(4):e109912. Available from: <https://doi.org/10.1136/jme-2024-109912>
 46. Wang Z, Gong M. Cross-disciplinary analysis of AI policies in academic peer review. *Learn Publ*. 2026. Available from: <https://doi.org/10.1002/leap.2035>
 47. Gramatikoff K. The Emergence of AI-Native Science: From Tools to Autonomous Research Systems. *PhilPapers*. 2026. Available from: <https://philpapers.org/rec/GRATEO-83>
 48. Tobias AV, Wahab A. Autonomous self-driving laboratories: a review of technology and policy implications. *R Soc Open Sci*. 2025. Available from: <https://doi.org/10.1098/rsos.250646>
 49. Hartung T. AI, agentic models and lab automation for scientific discovery: the beginning of scAInce. *Front Artif Intell*. 2025. Available from: <https://doi.org/10.3389/frai.2025.1649155>
 50. Dai Y, Chan H, Vriza A, et al. Adaptive AI decision interface for autonomous electronic material discovery. *Nat Chem Eng*. 2025. Available from: <https://doi.org/10.1038/s44286-025-00318-3>