Hypothesis

World Models for Biomedicine: A Steerability Framework

Jianghui Xiong *

Posted Date: 7 May 2026

doi: 10.20944/preprints202605.0366.v1

Keywords: biomedical world model; biomedical steerability; Capomics; mIC vector; constraint checkpoints; closed-loop architecture; counterfactual reasoning; intervention-response semantics; quality-control feedback; state-transition modeling

 doi:10.20944/preprints202605.0366.v1

Hypothesis

World Models for Biomedicine: A Steerability Framework

Jianghui Xiong

DeepoMe Limited, Beijing, China; jianghui@deepome.com or xiongjh77@163.com

Abstract

Recent medical world-model rubrics have mainly described a linear progression from representation and forecasting to action-conditioned simulation, counterfactual evaluation, and planning/control. This Perspective starts from a different goal: biomedical world models should not merely predict likely trajectories, but help make biological trajectories steerable. Steerability requires five linked functions: defining state, measuring state, specifying intervention-induced state movement, simulating alternative transitions, and inspecting deviations. We therefore propose the **Deductively Constrained Capomics World Model**, a closed-loop architecture organized around five corresponding constraint checkpoints: **CP1 state representation**, **CP2 intrinsic-capability quantification**, **CP3 intervention-response semantics**, **CP4 counterfactual transition**, and **CP5 quality-control feedback**. The framework shifts biomedical world modeling from a "what-if" simulator toward a quality-controlled "why-not" steering system, in which failed or unexpected transitions can be traced to state measurement, intervention specification, module response, state transition, or downstream phenotypic propagation. Within this architecture, **module-level intrinsic capability (mIC)** provides the proposed state variable, and **Capomics** provides its measurement framework. In the current prototype, DNA methylation is used to estimate module-level mIC values and assemble them into an **mIC vector**, while other omics and physiological readouts may be incorporated in future implementations. The accompanying depression case study illustrates how the cycle can be instantiated as a thought experiment for state-matched intervention reasoning and deviation inspection. The framework does not claim validated treatment planning or guaranteed efficacy; it is intended as a hypothesis-generating scaffold for biomedical world models, longitudinal intervention studies, and future biomedical applications.

Keywords: biomedical world model; biomedical steerability; Capomics; mIC vector; constraint checkpoints; closed-loop architecture; counterfactual reasoning; intervention-response semantics; quality-control feedback; state-transition modeling

1. Introduction

1.1. From Predictive Medicine to Simulative Medicine

Clinical care is interventional by nature. A physician choosing between two treatments, a clinical trial designer selecting a dosing schedule, a drug developer prioritizing a candidate molecule, or a person deciding whether to adopt a lifestyle change is asking a counterfactual question: how would the trajectory differ under one intervention route rather than another?

Yet most AI systems deployed or developed for medicine are designed to answer factual questions. Imaging classifiers, risk scorers, screening algorithms, molecular predictors, and literature-synthesis systems can describe what is present, classify what category a case belongs to, or estimate what is likely under the status quo. They cannot reliably answer what would happen if the system were intervened upon. The gap between medicine's inherently interventional questions and AI's predominantly factual answers is the foundational motivation for biomedical world models.

 doi:10.20944/preprints202605.0366.v1

This is why biomedical AI needs world models. A biomedical world model should represent the current biological or patient state, model how that state evolves over time, accept interventions as model inputs, and roll forward alternative future trajectories under competing interventions. In embodied intelligence and reinforcement learning, a world model is an internal simulator of environmental dynamics [1]. In medicine, the analogous system must become an internal simulator of biological and clinical trajectories under intervention.

This transition changes the central abstraction of medical AI. The goal is no longer only to answer "What is the current state?" or "What is the most likely risk under the status quo?" The goal is to answer:

- If intervention A is given instead of intervention B, how will the disease trajectory change?
- If inflammation-related state change is attempted before metabolic-state remodeling, will that be superior to the reverse sequence?
- If an aging module has already entered a low-mIC state, which intervention can induce a meaningful module response?
- If an expected transition does not occur, did the deviation arise at the level of state measurement, intervention specification, module response pattern, mIC state change, or downstream phenotypic propagation?

These questions require medical AI to move from predictors to simulators, from passive forecasting to action-conditioned rollout, and from correlational pattern recognition to counterfactual reasoning. A world model is therefore not merely a more powerful predictor. It is a system for asking what would happen if one intervention, sequence, or state-dependent response route were taken rather than another.

Recent medical world models, virtual cell models, and tumor-treatment simulation systems indicate that this direction is rapidly forming [2–7]. Medical world models have been used to simulate tumor evolution in hepatocellular carcinoma under different TACE treatment schemes [4]. VCWorld predicts cell-level drug perturbation responses by integrating biological knowledge with large language model reasoning [5], while recent virtual-cell evaluation work frames virtual cells as causal world models that should be assessed by intervention validity, counterfactual consistency, trajectory faithfulness, and mechanistic alignment [6]. Spatial-omics-driven virtual-cell foundation models are also entering pharmaceutical R&D workflows through the GSK–Noetik oncology partnership [7]. These developments show that biomedical and pharmaceutical AI is moving from "recognizing patterns" to "simulating worlds."

1.2. Common Challenges of Existing Routes

Existing biomedical world models can be broadly divided into three routes [2–6].

The first route is data-driven world models. These models learn state transitions directly from imaging, omics, medical records, or longitudinal follow-up data [2–4]. Their advantages are scalability, trainability, and high performance in specific tasks. However, their causal structures mainly come from statistical regularities in data. When such models are asked to simulate interventions outside the training distribution, their reliability depends on generalization capacity rather than an explicit causal scaffold.

The second route is knowledge- or LLM-driven world models. These models use literature, knowledge graphs, pathway databases, and large language model reasoning to organize biological knowledge into executable simulation structures [2,3,5,6]. Their advantages are stronger interpretability and broader knowledge coverage, but the knowledge sources remain collections of human-published inductive findings. Gaps, contradictions, biases, and outdated mechanisms in the literature may enter the model itself.

The third route is statistical causal or virtual cell models. These models attempt to predict intervention effects through causal graphs, perturbation data, multi-omics data, and mechanistic modeling [5,6,8,9]. Their advantage is that they directly address causal questions. However, in

 doi:10.20944/preprints202605.0366.v1

complex human biological systems, whether variables are complete, whether causal graphs are correct, and whether unobserved confounding can be ignored remain core challenges [8,9].

These routes are not wrong. On the contrary, they constitute necessary foundations for contemporary biomedical world models [2–6]. The claim of this article is not to replace them, but to point out an insufficiently solved problem: **before data modeling begins, can first principles be used to constrain the state space, intervention-response semantics, and quality-control feedback structure of medical world models?**

In other words, existing models usually learn data first and then attempt causal explanation. The route proposed here defines the causal scaffold first and then uses data to calibrate and refute it.

1.3. Contributions of This Article

In the formative stage of biomedical world models, complementary frameworks are emerging. Recent work has proposed stage-based rubrics that define what a medical world model should be able to do, from state representation to planning [2,3]. This article addresses a complementary question of equal importance: what essential components must a qualified biomedical world model possess? We propose a five-constraint-checkpoint architecture that defines the irreducible structural requirements for representing state, quantifying intrinsic capability, specifying intervention-response semantics, performing counterfactual transition, and converting deviations into quality-control feedback.

This article proposes the **Deductively Constrained Capomics World Model** as a candidate architecture for biomedical world modeling. We argue that a qualified biomedical world model requires more than predictive accuracy. It requires a biologically interpretable state space, explicit intervention-response semantics, a causal transition scaffold, counterfactual rollout, planning logic, and a quality-control loop that makes deviations diagnosable and correction-oriented. Because the field is still in its formative stage, the central task is not only to build larger models, but to define the necessary conditions that biomedical world models should satisfy.

Here, "deductive constraint" does not mean that the model is fully derived by logic, nor does it mean that data are unnecessary. It means that the high-level causal architecture of the model is first constrained by explicit first principles; data no longer bear the entire burden of discovering the causal world from scratch, but are used to calibrate, test, and revise a pre-declared causal scaffold.

Deductive constraint is therefore not a claim of epistemic certainty. It is a method for making high-level assumptions explicit before modeling, so that they can be calibrated, challenged, and falsified rather than hidden inside learned representations.

Layer	Role	Corresponding element in this article
Theory layer	Defines the first-principle assumptions	Life as adaptive capacities; exposures and interventions mapped to module response patterns
Architecture layer	Defines the minimum world-model control structure	Five constraint checkpoints
State layer	Defines the common state representation	mIC vector
Prototype layer	Makes the state representation computable	Capomics as the measurement framework for mIC; DNA methylation as the first computable prototype
Ecosystem layer	Implements the framework in longitudinal intervention modeling	CAPOVIME and related engineering infrastructure

 doi:10.20944/preprints202605.0366.v1

The framework can therefore be understood as a set of **constraint checkpoints** rather than a set of independent pillars. These checkpoints define how a biomedical world model should represent state, quantify intrinsic capability, specify intervention-induced module response patterns, perform counterfactual transition, and inspect deviations. Their purpose is not ever-finer point prediction, but **biomedical steerability**: guiding biological trajectories toward desired directions while keeping every state transition quality-controlled, auditable, and revisable.

This article makes six contributions:

1. **It proposes a new framework for biomedical and pharmaceutical world models.** This framework defines life at the system level as an ensemble of adaptive capacities, defines module-level state as **module-level intrinsic capability (mIC)**, and defines environmental exposures or interventions by the module response patterns and mIC-state transitions they induce.
2. **It proposes a minimal causal scaffold.** Disease and aging phenotypes are abductively traced to a limited set of upstream **mIC-level causal determinants**, whose complete description requires at least a root-cause layer, a functional layer, and a phenotypic layer.
3. **It proposes the mIC vector as a candidate interoperable state representation format.** Different biomedical world models may estimate module states using different internal methods, but output mIC-format vectors for cross-model comparison, composition, and validation.
4. **It proposes Capomics as the measurement framework for mIC.** Capomics is proposed here as the omics and measurement science of module-level intrinsic capability, intended to measure and compare intrinsic capability across biological modules and across scales of the human body. DNA methylation is presented as the first computable prototype, not as the only possible substrate.
5. **It proposes a quality-control loop for intervention success and deviation.** The framework does not promise unconditional efficacy, but makes unexpected or failed transitions decomposable into explicit premises through five-gate inspection.
6. **It aligns with and extends emerging stage-based biomedical world-model rubrics.** The framework maps onto recent requirements for representation, forecasting, action-conditioned simulation, counterfactual evaluation, and planning/control. Beyond these requirements, it argues that biomedical world models also need an additional structural requirement: **quality-control feedback and five-gate inspection**. This extends a model from a "what-if" simulator into a "why-not" diagnostic and steering system.

The aim of this paper is not to report a new empirical result, but to make the causal assumptions of a biomedical world model explicit, inspectable, and open to refutation. Counterfactual reasoning deserves not only a statistical foundation, but also an explicit causal scaffold.

 doi:10.20944/preprints202605.0366.v1

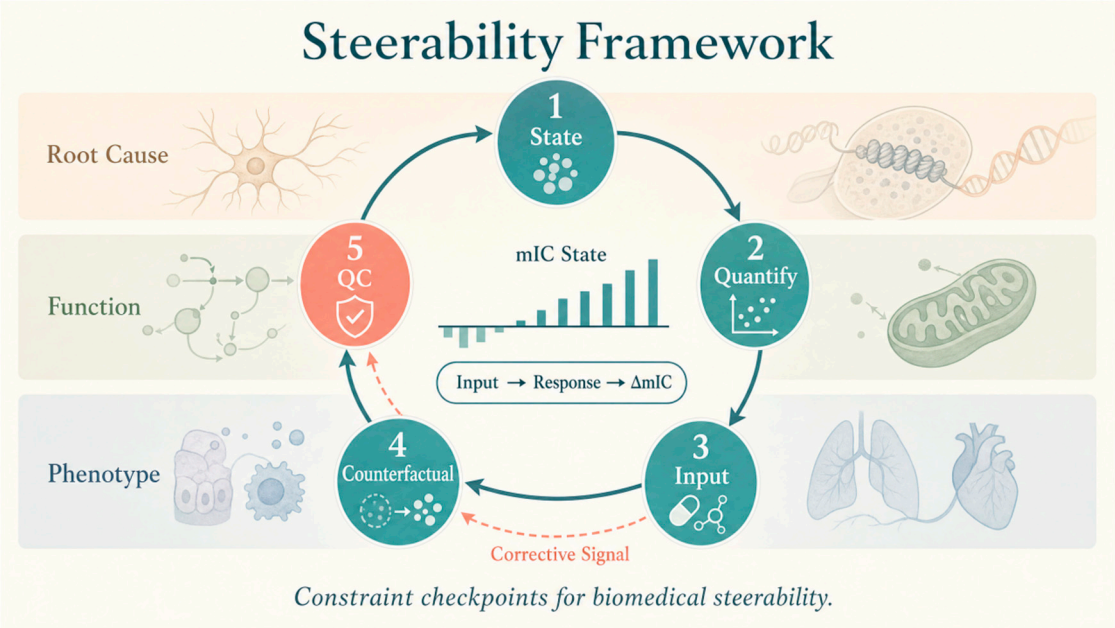

Figure 1. Master architecture of the biomedical steerability framework. The framework is organized as a closed-loop set of constraint checkpoints: state representation, intrinsic-capability quantification, intervention-response semantics, counterfactual transition, and quality-control feedback. The mIC vector provides the state representation, module response patterns provide computable intervention-response semantics, and QC feedback turns unexpected deviations into corrective signals.

1.4. Core Concepts

Concept	Definition	Role in this article
World model	A model that represents a system's current state and simulates how that state changes over time under different interventions	Overall AI framework
Deductive constraint	Constraining high-level causal structure using first principles	Provides the pre-declared causal scaffold that constrains state representation, intervention-response semantics, transition modeling, and feedback inspection
Biomedical steerability	A general systems property: the capacity to keep state change directionally guided, quality-controlled, feedback-correctable, and non-runaway when a system is perturbed toward an intended state	Defines the control objective for biomedical world-model construction: simulated state transitions should remain directionally guided, quality-controlled, auditable, feedback-correctable, and non-runaway
Constraint checkpoint	An indispensable control element for maintaining a system's steerability: a constraint that keeps state change guided, inspectable, correctable, and resistant to uncontrolled deviation	In biomedical world-model construction, specifies the required checkpoints for state representation, intrinsic-capability quantification, intervention-response semantics, counterfactual transition, and quality-control feedback

 doi:10.20944/preprints202605.0366.v1

Adaptive capacity	A system-level property of a living system: the capacity to maintain sustainability and functional continuity under changing environmental conditions	First-principle starting point: living systems are products of environmental adaptation, and this principle motivates the later derivation of mIC, Environmental Information, and constraint checkpoints for biomedical world-model construction
module-level intrinsic capability (mIC)	A module-level concept proposed here by analogy to WHO Intrinsic Capacity (IC): whereas WHO IC measures whole-person functional capacity, mIC defines an IC-like property for modular biological structures and functions, including pathways, gene networks, cellular programs, tissue modules, and organ-level subsystems	Provides the multi-scale state variable for biomedical world-model construction: by defining mIC from molecular and cellular modules to tissue, organ, and whole-body scales, the model can measure, guide, and finely regulate system-state evolution at different granularities
Capomics	The omics and measurement framework for studying mIC across biological modules and body scales; DNA methylation is one computable prototype	Measurement framework for making mIC computable
CAPOVIME	Capomics Virtual Intervention Model Ecosystem	Engineering ecosystem for implementing Capomics-based virtual intervention modeling and world-model-based medicine
Environmental Information (EI)	Information from the environment that a living system or cell can sense and process; intervention information is a special, intentionally designed form of environmental information, such as a drug regimen or exercise prescription, that is introduced to guide biological response	Provides the broad information-theoretic input concept; in the world model, EI becomes computable only through the module response pattern and mIC-state change it induces
Act	An intentional action taken to actively change a system state, such as exercise, drug treatment, behavioral intervention, nutritional intervention, or other deliberate intervention	External action applied to the human body or biological module and the driver of state change; its computational representation depends on the induced module response pattern
Mechanism of action (MOA)	The biochemical or physiological process by which a drug or intervention produces its biological or therapeutic effect	In the world model, explains how an exposure or intervention is sensed, which modules respond, how mIC changes, and whether the simulated trajectory is mechanistically plausible

 doi:10.20944/preprints202605.0366.v1

2. Three Current Routes of Biomedical World Models

2.1. The Data-Driven Route

Data-driven world models learn state transitions from observational data [2–4]. Medical imaging world models can generate post-treatment images from pre-treatment images; longitudinal electronic health record models can predict disease progression; multi-omics models can predict drug response from molecular states. Their core task can be summarized as learning or simulating the future state given the current state and a specified intervention.

Their advantages are clear: the larger the data, the stronger the model; the more specific the task, the higher the performance. They are particularly suitable for high-dimensional inputs such as imaging, spatial omics, and single-cell expression maps.

However, this route faces three limitations:

1. **Training-distribution limitation.** Unobserved intervention combinations, dosages, sequences, and populations often require extrapolation.
2. **Insufficient intervention-response semantics.** The model may know the label "drug A was given," but may not understand what response pattern that intervention induces inside the living system.
3. **Difficulty of deviation diagnosis.** When prediction fails, it is difficult to determine whether the deviation arises from insufficient data, model architecture, distribution shift, or incorrect causal structure.

2.2. The Knowledge- and LLM-Driven Route

Knowledge-driven or LLM-driven models attempt to use literature, pathway databases, protein-interaction networks, and large language model reasoning to organize life systems into simulatable structures [2,3,5,6]. Their advantage is stronger interpretability, especially for low-data scenarios and mechanism-hypothesis generation.

But this route also faces limitations. The literature itself is a collection of inductive knowledge. It is incomplete, uneven, and constrained by research trends, publication bias, and experimental conditions. LLMs can organize knowledge, but cannot guarantee that the knowledge structure itself is correct. They can generate mechanistic explanations, but explanation is not equivalent to causal validity.

2.3. The Statistical Causal and Virtual Cell Route

Statistical causal models, virtual cell models, and perturbation omics models attempt to model causality more directly [5,6,8,9]. Through gene perturbations, drug treatments, multimodal omics, and spatial context, they predict cellular or tissue responses to interventions. This is one of the most important directions in contemporary AI for biology.

However, in complex human systems, the fundamental challenges of causal inference remain: have enough variables been observed? Is the causal graph correct? Are unmeasured confounders negligible? Do perturbation data cover real clinical interventions? Can in vitro cellular perturbations transfer to tissues, organs, and whole human populations? These questions do not disappear simply because models become larger.

2.4. The Framework Proposed Here: Deductive Constraint, Not Pure Deduction

The field of biomedical world models is still in its formative stage [2–6]. Its terminology, architecture, state representation, intervention-response semantics, validation criteria, and quality-control requirements remain unsettled. This article therefore proposes a candidate framework for what a biomedical world model should contain: an explicit biological state space, a computable

 doi:10.20944/preprints202605.0366.v1

intervention-response layer, a causal scaffold for state transition, a counterfactual rollout structure, and a mechanism for diagnosing why a predicted or desired trajectory does not occur.

We call this a **deductively constrained biomedical world model**:

> *A deductively constrained biomedical world model is not a model in which all biological facts are derived purely from logic. Rather, it is a world model whose high-level causal architecture is first constrained by first principles and then calibrated, tested, and revised by experimental and clinical data.*

Its core is not "no data," but "not allowing data alone to bear the entire burden of discovering causal structure." In this framework, data are not asked to discover the causal structure of the living system from scratch. Instead, data are used to calibrate, test, and revise a causal scaffold declared in advance.

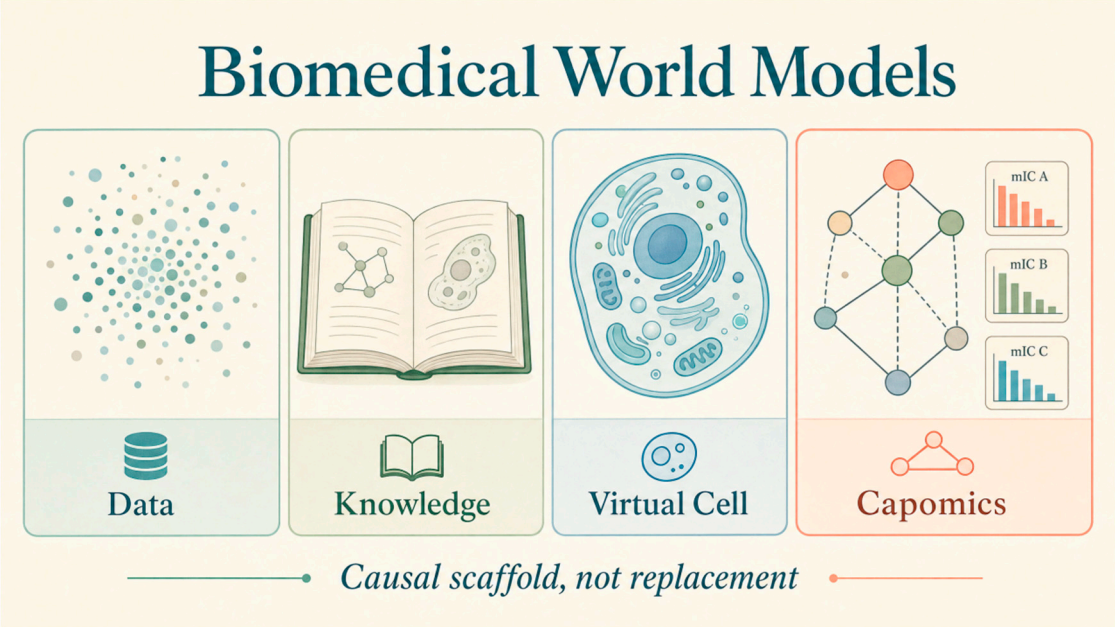

Figure 2. Landscape of biomedical world-model routes and the role of deductive constraint. Existing biomedical world models can be broadly grouped into data-driven, knowledge- or LLM-driven, and statistical causal or virtual-cell routes. The Capomics framework is proposed as a causal scaffold that constrains state, intervention-response semantics, transition, and feedback rather than replacing these routes.

3. Two Working Postulates

This article uses the term "working postulates" rather than "absolute axioms." A working postulate is not an irrefutable truth, but a minimal theoretical starting point for model construction. Its value is evaluated by explanatory power, internal consistency, computability, and falsifiable predictions.

3.1. Postulate One: Life as an Ensemble of Adaptive Capacities

A living system can be operationally viewed as an ensemble of adaptive capacities.

Here, "adaptive capacity" is not an abstract slogan. It refers to a system-level property of a living system: the capacity to maintain sustainability and functional continuity under changing environmental conditions, including the ability to maintain function, restore homeostasis, update state, and anticipate future challenges. It includes multi-scale capacities from molecules to cells, tissues, organs, and whole-body behavior.

 doi:10.20944/preprints202605.0366.v1

Definition 1: Adaptive capacity. Adaptive capacity is the system-level ability of a living organism to maintain sustainability, restore functional state, and preserve functional continuity under changing environmental conditions.

Definition 2: module-level intrinsic capability (mIC). The WHO concept of **Intrinsic Capacity (IC)** refers to the overall functional capacity of the whole person [10]. Such whole-person capacity is supported by lower-level biological components. Structure and function are therefore essential biological concepts, and living systems are organized through structural modules and functional modules. To measure an IC-like property at the module level, this article defines **module-level intrinsic capability (mIC)** by analogy to WHO IC. mIC is proposed here as an operational, model-facing construct for representing the biological potential of a pathway, gene network, cellular program, tissue module, or organ-level subsystem to be mobilized when called upon. It is not a directly observed biological entity, nor is it identical to current function, such as ATP output or inflammatory cytokine level, or to structure, such as mitochondrial abundance or synaptic morphology. It refers to the estimated latent capability of a module: what the module can still do under challenge.

Definition 3: Aging. Aging is the gradual decline, mismatch, or dysregulation of mICs across multiple biological modules, leading to reduced system-level adaptive capacity, a narrowed range of tolerable perturbations, slower recovery, increased cost of homeostatic maintenance, and higher probability of systemic failure [11,12].

To make mIC measurable rather than only conceptual, this article introduces **Capomics**: the omics and measurement framework for studying mIC across biological modules and body scales. Capomics asks how the intrinsic capability of pathways, gene networks, cellular programs, tissues, organs, and whole-body systems can be represented as measurable multi-scale readouts.

3.2. Postulate Two: Module Response Patterns as the Common Encoding Space

Environmental exposure and intentional intervention become computable in a biomedical world model only when they are mapped to module response patterns.

The current form of the living system is the result of long-term human–environment interaction. Representative environmental exposures such as nutrition, pathogens, temperature, and social stress have shaped receptors, metabolic sensors, neuroendocrine loops, and epigenetic programs [10–15]. These systems do not respond to abstract labels such as “stress,” “exercise,” “drug,” or “meditation.” They respond to signals that biological modules can sense and process.

At the module level, Environmental Information (EI) should be defined from the perspective of the receiving biological unit. A cell, for example, does not directly receive “psychosocial stress” or “mindfulness meditation” as semantic labels. Instead, it receives changes in its molecular and physical environment, such as stress hormones and inflammatory cytokines. If the cell is treated as a biological computing unit, these signals are its inputs. The cell processes them according to its genetic program, epigenetic state, receptor repertoire, and pathway wiring. The output is a module response pattern, such as altered pathway activity, a shifted transcriptional state, or a changed secretion profile.

This module-level view also applies to intentional interventions. A drug is not computable merely as a treatment name; it becomes biologically meaningful through concentration, exposure time, receptor engagement, and downstream response. Exercise, nutrition, behavioral intervention, and cell therapy likewise acquire computable meaning only through the response patterns they induce in specific modules.

Therefore, natural environmental exposures and designed interventions should be encoded in the same response-pattern space. In the Capomics implementation, this response pattern is represented by the module’s mIC state and its state change: an exposure or intervention is evaluated by how it shifts the mIC profile of the receiving module and how that shift propagates downstream. This makes the transition from “module response pattern” to “mIC change” explicit. A shared mIC-based response space allows heterogeneous exposures and interventions to be compared, combined,

 doi:10.20944/preprints202605.0366.v1

sequenced, and simulated through their induced module responses rather than through surface categories. The same behavioral intervention may shift mIC differently in different starting states; the same molecular intervention may be beneficial, neutral, or harmful depending on current mIC and the response program of the receiving module.

Mechanism of action (MOA) is the interpretable path through this response-pattern space: how an exposure or intervention is sensed, which modules respond, how the response propagates across the root-cause, functional, and phenotypic layers, and whether the resulting transition changes mIC in the intended direction. In this sense, MOA is not merely a pharmacological label. It is the explanatory structure connecting signal encoding, module response, mIC transition, and phenotypic propagation.

This formulation does not exclude drugs, gene therapy, cell therapy, organ transplantation, or engineered interventions. It places them in the same computable space as lifestyle, behavioral, environmental, and physiological exposures or interventions. What matters for a biomedical world model is not whether an intervention is called a drug, a behavior, a stressor, or an environmental exposure, but whether its module response pattern can be specified, measured, compared, and used to predict state transition.

4. The Minimal Causal Scaffold

4.1. Proposition One: Adaptive Capacity Is Modularly Organized

Environmental challenges are highly heterogeneous. Oxidative stress, nutrient deprivation, pathogen invasion, and psychological stress require responses at different timescales, spatial scales, and molecular mechanisms.

No single mechanism can efficiently process all challenges. Therefore, adaptive capacities in living systems are modularly organized. Here, a "module" is not a completely isolated device, but a relatively stable and identifiable functional unit with specific input–output characteristics.

Modules may include:

- DNA damage repair modules;
- mitochondrial energy adaptation modules;
- inflammation resolution modules;
- metabolic flexibility modules;
- neural plasticity modules.

These modules are coupled, share nodes, and contain feedback and compensation [13–15]. Therefore, this article does not claim that the living system is a mechanically assembled set of isolated modules. Rather, it argues that modularity is a necessary abstraction for state representation and intervention planning in a computable world model.

4.2. Proposition Two: Disease and Aging Phenotypes Can Be Abductively Traced to Upstream Intrinsic-Capability Determinants

If system-level adaptive capacity is organized through biological modules, disease and aging phenotypes should not be viewed only as terminal manifestations. They should be abductively traced back to a limited set of upstream intrinsic-capability causal determinants.

Therefore, this article avoids the overly strong expression that "every disease must have a single root-cause module" and adopts a more robust definition:

> *Any disease or aging phenotype should be localizable within a finite network of biological modules as an abductively inferable combination of one or more upstream intrinsic-capability causal determinants.*

This definition is compatible with complex diseases. Cancer may involve genomic stability, immune escape, metabolic reprogramming, and microenvironmental adaptation [16,17]. Alzheimer's

 doi:10.20944/preprints202605.0366.v1

disease may involve proteostasis, neuroinflammation, metabolic dysfunction, vascular function, and synaptic plasticity. Frailty may involve coupled decline in muscle, immunity, endocrine function, neural function, and nutritional state.

Thus, "root cause" is not synonymous with a single target. It refers to an intrinsic-capability causal determinant with upstream influence in the causal chain.

4.3. Proposition Three: A Three-Layer Minimal Causal Scaffold

Disease is usually encountered first as a visible phenomenon. A clinical diagnosis, imaging abnormality, symptom cluster, functional decline, or disease endpoint is what brings a biological state into clinical view. However, the level at which disease is recognized is not necessarily the level at which it is causally explained. To support abductive tracing from phenotype back to causal explanation, this article therefore distinguishes three layers: the phenotypic layer, the root-cause layer, and the functional layer.

The order below follows the order of clinical recognition: from visible phenotype, back to upstream root-cause determinants, and then to the functional modules that connect them. Causally, however, the propagation runs in the opposite direction: root-cause constraints are expressed through functional modules and finally appear as phenotypic outcomes.

1. **Phenotypic layer.** The layer that expresses capacity outcomes, including clinical biomarkers, imaging manifestations, functional tests, symptoms, disease endpoints, and lifespan or healthspan outcomes. This is the layer at which disease is usually named and recognized. It captures the observable phenomenon, but it does not by itself specify which upstream causal determinant produced it.
2. **Root-cause, or upstream constraint, layer.** The layer that stores or constrains intrinsic capability, including epigenetic state, genomic stability, stem-cell reserve, mitochondrial genetic integrity, long-term immune memory, and tissue-structural foundation. In this scaffold, the hallmarks of aging can be interpreted as recurrent common-feature summaries of upstream root-cause determinant families for aging-related diseases [11,12]. Likewise, the hallmarks of cancer summarize recurrent root-cause patterns that sustain cancer-related phenotypes, including genomic instability, immune evasion, proliferative signaling, metabolic reprogramming, and microenvironmental adaptation [16,17]. These hallmarks should not be treated as single causes for every case, but as conserved upstream abnormality families that repeatedly generate disease vulnerability.
3. **Functional layer.** The intermediate dependency layer that invokes capacity and transmits root-cause constraints into phenotypic outcomes. It includes metabolic flux, immune response, organ-level function, and tissue-function dynamics. Different functional modules depend on different root-cause capacities to different degrees. This differential dependency explains why impairment of the same upstream hallmark may preferentially produce disease in some organs, tissues, or physiological systems rather than others.

These three layers constitute a minimal causal scaffold:

> *Disease is recognized at the phenotypic layer, abductively traced back to conserved root-cause determinant families, and expressed through functional modules whose dependency on those determinants differs across organs and tissues.*

This dependency structure is crucial. Root-cause hallmarks are common-feature summaries, but they do not map one-to-one onto diseases. A given hallmark impairment becomes clinically meaningful only when it passes through a functional module that is highly dependent on that

 doi:10.20944/preprints202605.0366.v1

capacity. The same upstream causal determinant can therefore generate different disease spectra depending on which functional systems are most vulnerable.

For example, mitochondrial integrity and energy adaptation are especially important for high-energy-demand systems such as the brain, heart, skeletal muscle, liver, renal tubules, and retina. If mitochondrial mIC is depleted, the resulting phenotypic spectrum may preferentially include neurodegenerative vulnerability, cognitive or mood-related dysfunction, cardiomyopathy or heart-failure risk, myopathy and fatigue, metabolic liver dysfunction, renal vulnerability, or retinal degeneration. The same root-cause determinant does not produce one universal disease; it is filtered through the dependency structure of functional modules.

Similarly, stem-cell reserve and tissue-repair capacity are root-cause-layer causal determinants whose phenotypic expression depends on which tissues most require renewal and repair. Hematopoietic stem-cell decline may appear as immune aging, anemia, or reduced recovery from stress. Intestinal stem-cell dysfunction may affect epithelial renewal and barrier integrity. Muscle satellite-cell decline may contribute to sarcopenia and impaired repair. Skin and connective-tissue repair deficits may appear as poor wound healing or tissue fragility. Thus, a hallmark such as stem-cell exhaustion becomes clinically meaningful only when mapped through tissue-specific functional modules to a disease spectrum.

This scaffold also clarifies the distinction between **structure**, **function**, and **intrinsic capability**. Structure refers to the material or morphological basis of a module, such as mitochondrial abundance, synaptic morphology, receptor density, stem-cell niche architecture, or tissue architecture. Function refers to the module's current performance, such as ATP output, inflammatory cytokine production, insulin response, tissue repair, or pathway activity. Intrinsic capability refers to the module's latent biological potential beyond structure and current function: whether mitochondrial programs can increase energy adaptation under demand, whether an insulin-signaling pathway can respond efficiently to a small signal, whether a stem-cell compartment can regenerate tissue after injury, or whether an immune module can mount and resolve a response without chronic inflammation.

This scaffold is not a complete ontology of the living system, nor does it claim to exhaust all layers. Environmental exposure, spatial organization, cell type, behavior, social factors, developmental history, and timescale may all require additional modeling. But as a minimal structure for biomedical world models, the three-layer scaffold is sufficient to connect clinical phenomena to upstream causal determinants, explain organ- and tissue-specific vulnerability, and support state representation, intervention-response semantics, and counterfactual rollout.

4.4. Structural Basis of Counterfactual Reasoning

In a pure predictive model, the question “what would happen if intervention A rather than B were applied?” usually depends on statistical extrapolation. In a deductively constrained framework, the foundation of counterfactual reasoning is not simply how many similar examples the model has seen, but whether the model has an explicit causal scaffold [8,9].

Once the root-cause, functional, and phenotypic layers are explicitly distinguished, an intervention is no longer merely an external treatment label. It becomes a biologically specified intervention whose module response pattern can be represented as an mIC-state transition and compared with an intended or desired state change.

Counterfactual reasoning can therefore be decomposed into five questions:

1. What is the current mIC state?
2. What intervention is introduced, and what module response pattern is expected?
3. What state change is desired?
4. Does the observed mIC transition move toward the desired state change and propagate to the phenotypic layer?
5. If it fails, at which premise did the failure occur?

 doi:10.20944/preprints202605.0366.v1

This is the core difference between a deductively constrained model and a black-box predictive model.

An inductive model mainly asks: "Given existing data, which outcome is more likely?" A deductively constrained model further asks: "Given the causal scaffold, which premise must be changed for the outcome to change?"

5. Capomics: A Computable Prototype

5.1. Why a Computable State Representation Is Needed

A world model must have a state space. If state cannot be defined, the model cannot simulate state transitions. If intervention-response semantics cannot be defined, the model cannot evaluate interventions. If the target cannot be defined, the model cannot plan.

In this framework, state is not a single age, a single risk score, or a single disease label, but the intrinsic-capability state of multiple biological modules. An ideal state representation should satisfy five conditions:

1. measurable from human samples;
2. longitudinally traceable;
3. mappable to functional modules;
4. associated with disease and aging phenotypes;
5. responsive to intervention.

5.2. Definition of Capomics

Capomics is derived from "capability" and "omics."

Capomics is the proposed omics and measurement science for quantifying ***module-level intrinsic capability (mIC)*** *across biological modules and across scales of the human body.*

In this terminology, a module has structure, current function, and mIC. mIC is the object being measured: the latent potential of a biological module that may not be fully visible in its current output but can determine whether the module can respond when challenged.

Capomics is not, in principle, restricted to DNA methylation. Future implementations may integrate transcriptomics, proteomics, metabolomics, single-cell data, spatial omics, clinical physiology, wearable signals, and behavioral readouts. The reason DNA methylation is emphasized here is more specific: it provides a practical first prototype for representing relatively durable biological state.

An intuitive computational analogy is useful before listing the biological rationale. Transcriptomic, proteomic, and metabolomic states often resemble information in a computer's working memory: they reflect what the cell is currently reading, executing, producing, and consuming. DNA methylation, by contrast, is closer to information written into a persistent configuration layer or hard disk. After a document is edited in working memory, it must be saved to become durable; similarly, many transient environmental or cellular responses become long-lasting biological state only when they are encoded into more persistent regulatory layers. This analogy should not be taken literally, because proteins, chromatin states, metabolites, and cell structures can also have different degrees of persistence. Its purpose is to emphasize why DNA methylation is especially attractive for representing durable biological state rather than momentary activity.

In the current prototype, Capomics mainly uses DNA methylation patterns as state readouts. This is not because DNA methylation is assumed to provide a complete causal explanation for aging and disease, nor because other omics lack useful state information. Rather, DNA methylation is selected as the first computable prototype because it sits at the interface of persistent regulation, environmental memory, cell identity, and long-term biological state:

 doi:10.20944/preprints202605.0366.v1

1. **Long-term stability.** Compared with transient transcription or protein phosphorylation, DNA methylation is relatively stable and more suitable for reflecting durable biological state.
2. **Gene-regulatory meaning.** DNA methylation functions as a relatively durable regulatory layer of gene activity, helping record which genes or programs are more open, silenced, or poised in a given cellular context.
3. **Environmental memory with mitotic persistence.** DNA methylation is sensitive to environmental exposure, stress, metabolism, inflammation, and lifestyle-related inputs, and some methylation states can be maintained across cell divisions.
4. **Cell identity and lineage record.** DNA methylation patterns preserve information about cell differentiation, cell lineage, and cell-type identity, making them useful for cell-type deconvolution, annotation, and sorting-related interpretation.
5. **Age and module relevance.** Epigenetic clocks have demonstrated that DNA methylation can strongly predict age and many aging-related phenotypes [18–21], while CpG patterns can be mapped through genes, pathways, gene ontology, and regulatory networks to functional modules.

Therefore, this article does not claim that DNA methylation exhausts the possible molecular substrates of this framework. It is used here as a state-estimation substrate, not as proof that DNA methylation is the primary causal carrier of mIC. It is a computable prototype that makes the theoretical concept of module-level intrinsic capability measurable, comparable, and modelable.

5.3. From a Single Aging Clock to a Modular Intrinsic-Capability Spectrum

Traditional aging clocks usually output a scalar: the biological age of a person. This scalar is useful for risk prediction, but insufficient for intervention planning. Two people with the same biological age may have completely different causally relevant intrinsic-capability profiles:

- one may mainly have reduced inflammation-resolution capability;
- one may mainly have reduced mitochondrial response capability;
- one may mainly have reduced immune-surveillance capability;
- one may mainly have reduced metabolic-flexibility capability.

Therefore, aging should not be only a number. It should be a spectrum.

The goal of Capomics is to extend a single aging clock into a modular intrinsic-capability spectrum:

Aging is not only a scalar. Aging is a vector of intrinsic-capability decline.

5.4. State Space of the Capomics World Model

In the Capomics world model, an individual's state can be represented as a combination of states across multiple intrinsic-capability modules. Each module has an intrinsic-capability readout and can be compared with young healthy baselines, age-matched populations, disease populations, or the individual's own pre-intervention state.

We operationalize the mIC of a module through a **Capability Index (CI)**. For a given pathway or module, its **Pathway Aging Index (PAI)** reflects the accumulated molecular burden or aging-associated deviation of that module. The module's remaining mIC is therefore defined as:

$$CI = 1 - PAI$$

When PAI is close to 0, the module retains a youthful mIC. When PAI approaches 1, the module's mIC is substantially exhausted. This formulation is a provisional normalization convention for module-level comparison rather than a claim that all biological capability is reducible to a single scalar or that capability is universally the linear inverse of aging burden; it provides a practical module-level readout that can be assembled into a vector.

 doi:10.20944/preprints202605.0366.v1

This practical readout allows module-level capability estimates to be assembled into an mIC vector for downstream interpretation and simulation.

5.5. mIC as a Candidate Interoperable State Representation Format

The strongest standardization potential of this framework lies in the **mIC vector** as a candidate interoperable state representation or exchange format for biomedical world models. This does not mean that all systems must use Capomics internally. Rather, different world models may estimate module states through different data modalities, algorithms, causal models, or virtual-cell architectures, while exporting their state estimates in a comparable mIC-vector format.

Recent aging-clock literature is converging toward this direction. Xiong proposed pathway-level aging indices across more than 3000 cellular pathways and introduced the Capability Index as a transformation from pathway aging burden to remaining capability [22]. Ageome independently constructed high-dimensional aging representations across functional modules [23]. PathwayAge proposed another pathway-level epigenetic clock strategy using machine-learning aggregation from CpGs to GO and KEGG pathways [24]. Fuentealba et al. further demonstrated that WHO Intrinsic Capacity can be predicted from blood-based DNA methylation, linking molecular readouts to macroscopic functional capacity [25], while Jia and Liu highlighted the significance of this direction for aging assessment [26].

Together, these works suggest a shared transition: aging is not only a scalar clock, but a high-dimensional pathway- or module-level state vector. Their common limitation is that most remain descriptive, reporting pathway aging scores or clock outputs. The mIC transformation, operationalized here as **CI = 1 − PAI**, reframes the question from "how aged is this pathway?" to "how much mobilizable intrinsic capability remains in this module?" This is the shift from diagnostic biomarker to intervention-relevant state variable.

By analogy, TCP/IP does not require every device to run the same operating system; it specifies a packet format that allows heterogeneous systems to communicate. Similarly, mIC does not require every biomedical world model to adopt the same internal architecture. It proposes a common output format for comparing, composing, validating, and eventually standardizing state representations across heterogeneous biomedical simulators.

In this sense, the mIC vector is not proposed merely as one model's hidden representation, but as a candidate **exchange format** for biomedical world models. Heterogeneous systems may preserve their own internal architectures while exporting comparable state estimates in an mIC-compatible form.

6. Intervention Semantics: From Intervention to Module Response to mIC Transition

6.1. Unified Intervention-Response Semantics: Intervention → Module Response Pattern

A biomedical world model requires a computable definition of intervention [3,4,9,10]. In this framework, an intervention is not treated as an external overwrite of the living system, nor merely as a treatment label. It is treated as a model input whose biological meaning is determined by the module response pattern it induces.

Environmental Information (EI) remains useful as a broad information-theoretic concept: it refers to inputs that a living system or biological module can sense and process. However, EI is not the final modeling object. Natural exposures and intentional interventions become computable only when they are translated into measurable module response patterns, and in the Capomics implementation these response patterns are represented as mIC states and mIC-state changes.

The state of an individual is represented as:

$$S_t = [mIC_t(M_1), mIC_t(M_2), \ldots, mIC_t(M_n)]^\top$$

 doi:10.20944/preprints202605.0366.v1

where each M_i denotes a biological module and $mIC_t(M_i)$ denotes its module-level intrinsic capability at time t. An input is represented by the response pattern it induces in the current state:

$$R_t = \mathcal{R}(I_t, S_t)$$

where I_t denotes the environmental or interventional input, $\mathcal{R}$ denotes the state-dependent response-mapping function, and R_t denotes the induced module response pattern. The world-model transition can then be written as:

$$S_{t+1} = F(S_t, R_t) = F(S_t, \mathcal{R}(I_t, S_t))$$

This formulation avoids treating a molecular target or an intervention label as the primitive unit of action. A target is a planning concept: the desired state or desired state change within the mIC state space. What matters computationally is whether an intervention induces the intended module response pattern and whether that response produces the desired mIC-state transition.

6.2. Desired State Change and Mechanism of Action

Planning requires a desired state or desired state change:

$$\Delta S^* = S^* - S_t$$

where S^* denotes the desired future mIC state. The role of the world model is to estimate whether a given intervention is likely to induce a module response pattern that moves the current state S_t toward ΔS^*.

Mechanism of action (MOA) is not a separate primitive of the intervention. It is the explanatory path through the response-pattern space: how an exposure or intervention is sensed, which modules respond, how the response changes mIC, and how that transition propagates through the root-cause, functional, and phenotypic layers. MOA can be evaluated using perturbation-response signatures, molecular and cellular readouts, changes in mIC, and downstream functional or clinical outcomes. Classical precision medicine is a special case of this formulation: when a drug acts on a well-defined molecular node, the desired state change may be narrow and interpretable as a single-target intervention. In complex chronic disease, aging, psychiatry, metabolic disease, and longevity medicine, however, the desired change is usually a multi-module mIC-state transition.

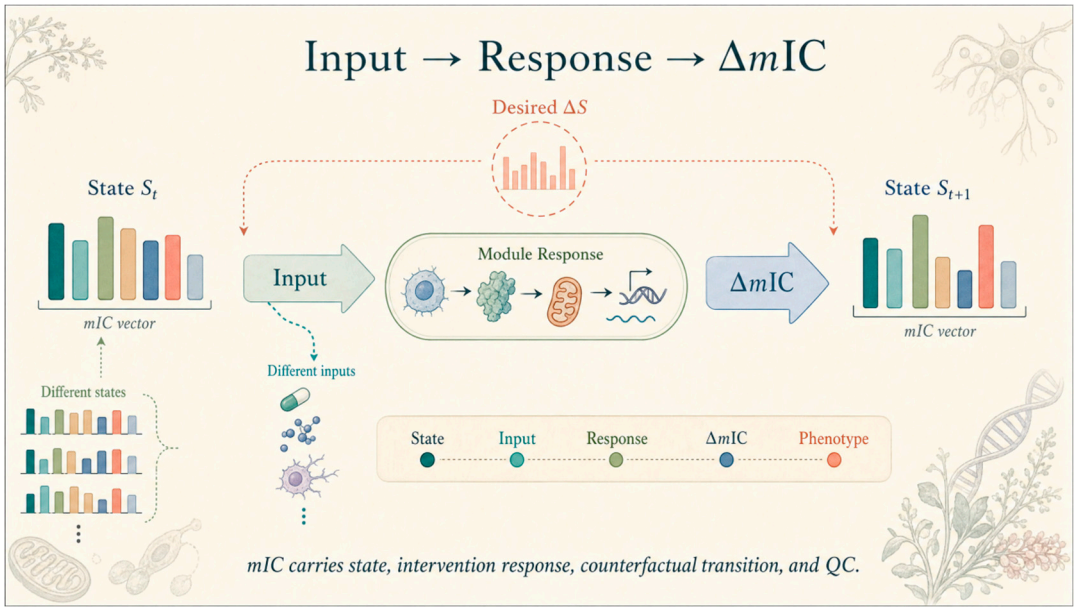

Figure 3. Input-induced module response and mIC-state transition. In the Capomics World Model, an intervention is represented not merely as a treatment label but as an input that induces a module response pattern. This response pattern transforms the current mIC state into a future mIC state, allowing desired ΔmIC movement, counterfactual rollout, and later QC inspection to be represented in a shared state space.

 doi:10.20944/preprints202605.0366.v1

6.3. Quality-Control Loop for Intervention Success and Failure

This framework does not claim that any intervention can produce unconditional efficacy. Its stronger claim is that intervention failure can become diagnosable, correctable, and useful for the next iteration. In an inductive model, prediction failure may arise from insufficient training data, distribution shift, confounding, missing variables, or structural error, but these causes may not be decomposable. In a deductively constrained model, failure can be examined as a premise failure within an explicit causal scaffold.

The quality-control loop decomposes failure into five diagnostic gates:

1. **Gate 1: Was the state measured correctly?** Was the current mIC vector S_t accurately estimated, and were the relevant state patterns interpreted correctly?
2. **Gate 2: Was the intervention specified correctly?** Was the intervention appropriate in modality, dose, strength, timing, frequency, sequence, and context?
3. **Gate 3: Did the expected module response pattern occur?** Did the intervention generate the expected molecular, cellular, physiological, or behavioral response signature?
4. **Gate 4: Did the mIC state move as expected?** Did the observed state change $\Delta S_t = S_{t+1} - S_t$ move toward the desired state change ΔS^*?
5. **Gate 5: Did downstream propagation reach phenotype?** Did the mIC state transition propagate through functional dynamics to phenotypic or clinical readouts?

Thus, therapeutic failure is not merely an opaque negative result; within a deductively constrained world model, it becomes a diagnosable premise failure:

> *When an inductive model fails, the failure may remain opaque; when a deductively constrained model fails, the failure can be traced to explicit premises.*

QC is therefore not an accessory function but the feedback layer that makes biomedical steerability improvable. Each deviation becomes a signal to revise one or more explicit premises: state measurement, intervention design, response-pattern modeling, state-transition assumptions, or downstream propagation logic.

Figure 4. Five-gate quality-control feedback for biomedical steerability. The quality-control loop decomposes an unexpected or failed transition into five diagnostic gates: state measurement, intervention specification, module response pattern, mIC state change, and downstream phenotypic propagation. This converts model or intervention failure into a structured corrective signal.

 doi:10.20944/preprints202605.0366.v1

6.4. Drug Discovery as One Application

For drug discovery as one application, the significance of this framework is not to replace target discovery, but to reframe drug action as module response and mIC-state transition. A target is not only something to be bound or inhibited; in a world-model context, it is part of a desired state-change route. The key question becomes whether a candidate intervention merely changes functional-layer indicators, or whether it induces a response pattern that moves the mIC state toward a desired state change and propagates to phenotype.

The Capomics World Model can therefore serve three tool functions for drug discovery:

1. **State stratification tool.** Use individual or sample mIC spectra to identify populations or disease subtypes whose current states are more likely to respond to a class of interventions.
2. **Mechanism-of-action evaluation tool.** Evaluate whether an intervention produces the expected module response pattern, whether that response changes the mIC vector in the intended direction, and whether the transition propagates to functional or phenotypic readouts.
3. **Combination and sequence planning tool.** When multiple mIC abnormalities coexist, evaluate drug combinations, dosage timing, and intervention sequence by their predicted state transitions rather than by target labels alone.

This shifts drug discovery from "finding druggable targets" toward "finding interventions capable of inducing beneficial module response patterns and remodeling biological state." It also makes combination design and sequence planning a natural use case for biomedical world models: the model can compare not only which candidate works, but in which starting state, through which response pattern, and in what order.

7. Alignment with and Extension of Biomedical World-Model Rubrics

Recent biomedical AI literature has begun to converge on level-based rubrics for world models [2,3]. Qazi et al. proposed a clinical rubric ranging from temporal prediction to action-conditioned prediction, counterfactual rollout, and planning/control [3]. Similarly, the Medical World Model framework organizes the field from representation and forecasting to single-arm projection, comparative treatment evaluation, and planning [4]. These standards are not merely additional references; they are beginning to define the evaluation standards by which medical world models will be assessed. They clarify what a medical world model should be able to do.

The Deductively Constrained Capomics Framework addresses a complementary question: what causal-state architecture is required for these rubric-defined functions to become biologically interpretable, testable, steerable, and quality-control-inspectable? It does not claim to be a fully validated planning system. Rather, it specifies the structural conditions that such a system would require: a biologically interpretable state representation, explicit intervention-response semantics, a causal scaffold for state transition, a structure for counterfactual rollout, and a quality-control loop for unexpected intervention outcomes.

7.1. Alignment: From Level-Based Rubrics to Causal-State Architecture

The two level-based rubrics differ in terminology but converge in their direction of travel: from representing the present, to forecasting the future, to simulating interventions, to comparing alternatives, and finally to planning. The framework proposed here maps naturally onto both rubrics by translating each rubric requirement into a causal-state component.

World-model rubric requirement	Meaning in recent biomedical world-model rubrics	Response in the Deductively Constrained Capomics Framework
Representation / state encoding	Encode the current biological or patient state	Capomics represents the individual as a vector of modular intrinsic capabilities, rather than

 doi:10.20944/preprints202605.0366.v1

		as an uninterpreted latent embedding alone
Temporal prediction / forecasting	Predict natural disease or biological trajectories over time	Module-level intrinsic-capability trajectories provide a biologically interpretable basis for forecasting decline, recovery, or capability depletion
Action-conditioned prediction	Predict future state under a specified intervention or action	Interventions are defined by intervention-induced module response patterns that transform a current mIC state
Single-arm projection	Simulate the trajectory under one intervention	A specified intervention is mapped to an expected module response pattern and mIC-state trajectory
Counterfactual rollout / comparative treatment evaluation	Compare outcomes under alternative interventions	Counterfactual rollout can vary either the intervention while holding state fixed, or the initial mIC state while holding the intervention-response assumption fixed
Planning / control	Select intervention sequence, dosage, timing, or combination to optimize a target trajectory	Planning begins with the current mIC state, desired state change, predicted module response patterns, and feedback correction; actual clinical implementation requires prospective validation

7.2. Extension: Quality-Control Feedback as an L+1 Structural Requirement

Existing stage-based ladders generally culminate in planning or control. We argue that, in safety-critical biomedical settings, planning is not sufficient. A useful biomedical world model should also diagnose why a prediction, simulation, or intervention plan failed and use that diagnosis to improve the next iteration. We refer to this additional requirement as **quality-control feedback**, or an **L+1 structural requirement** beyond planning.

In a purely inductive system, failure may be difficult to localize without post-hoc analysis. In the deductively constrained framework, failure is treated as a premise failure within an explicit causal scaffold. The error can be examined at several levels: state measurement, intervention specification, module response pattern, mIC state change, timing and dosage, or downstream phenotypic propagation.

This structure extends a world model from a "what-if" simulator into a quality-controlled "why-not" steering system: when an expected trajectory does not occur, the model can ask which premise failed and which constraint checkpoint should be revised. Failure itself becomes a source of scientific information rather than an opaque negative result.

7.3. Architectural Advantages over Purely Inductive Systems

Compared with many existing systems, which focus primarily on learning transition functions from data, the proposed framework emphasizes the prior specification of state semantics, intervention-response semantics, explanatory structure, and quality-control feedback. Its potential advantage is therefore not benchmark superiority at the current stage, but architectural completeness for counterfactual biomedical modeling.

Many biomedical world models first learn to simulate and then seek post-hoc explanations. In contrast, the deductively constrained framework specifies the explanatory structure before simulation begins. This is the core theoretical distinction: the model does not merely produce a

 doi:10.20944/preprints202605.0366.v1

simulated trajectory and then ask why it occurred; it defines in advance what counts as state, how interventions become response patterns, how mIC-state transitions should be evaluated, and where failure can be localized if the expected trajectory does not appear.

First, **state is biologically interpretable**. A Capomics state is not merely a latent embedding; it is intended to represent modular intrinsic capabilities. This makes it possible to ask which module has lost capability, which causal determinant is upstream, and which intervention should be considered first.

Second, **intervention has biological semantics through response.** In many action-conditioned models, the action is represented as a treatment code. In this framework, an intervention is represented by the module response pattern it induces in a current biological state.

Third, **counterfactual rollout is scaffolded along two complementary axes**. The first axis is **intervention counterfactual reasoning**: given a current mIC state S_t, the model compares how different interventions may induce different response patterns and future states. This supports intervention comparison and comparative treatment evaluation. The second axis is **state counterfactual reasoning**: given a fixed intervention, the model compares how different initial mIC states may produce different response patterns and trajectories. This supports interpretation of treatment heterogeneity, patient stratification, and quality-control inspection. Purely inductive models can in principle model such heterogeneity when sufficient data are available, but their state representations may remain opaque. The advantage of the mIC vector is that it makes state-dependent counterfactuals biologically interpretable, module-resolved, and empirically testable.

Fourth, **planning begins with state transition.** Instead of treating planning as a generic optimization problem over treatment labels, the framework asks what the current state is, what desired state change is clinically or biologically meaningful, which intervention is most likely to induce the needed response pattern, and how feedback after intervention should update the next decision.

Fifth, **quality-control feedback is treated as a core world-model requirement**. A clinically useful biomedical world model should not only simulate success, but also diagnose failure and convert deviation into a corrective signal. Quality-control feedback may therefore be considered a necessary extension of existing stage-based rubrics.

7.4. Five Constraint Checkpoints for Biomedical Steerability

Section 6 described the quality-control loop operationally, as a way to localize failed or unexpected transitions. Here, the same logic is reinterpreted architecturally: across existing world-model rubrics and their proposed extension, a biomedical world model requires more than a list of functional components. It requires constraint checkpoints: control points that specify what counts as valid state representation, valid measurement, valid intervention-response semantics, valid counterfactual transition, and valid deviation inspection. These checkpoints are not independent pillars. They form a closed-loop structure whose purpose is quality-controlled biomedical steerability.

The five constraint checkpoints are not an arbitrary collection. They map to the irreducible stages of a quality-assured world-model cycle: defining state, measuring state, designing intervention, executing counterfactual transition, checking results, and feeding corrective signals back into the next iteration.

- **CP1 (State representation)** defines what is being modeled. Without it, the model has no ontology of the system it claims to simulate.
- **CP2 (Intrinsic-capability quantification)** makes CP1 measurable. Without it, state remains an abstraction rather than a computable input.
- **CP3 (Intervention-response semantics)** defines how an intervention becomes a computable module response pattern. Without it, the model cannot compare or combine drugs, behavior, nutrition, environmental exposure, and measurement context within a unified biological-response language.

 doi:10.20944/preprints202605.0366.v1

- **CP4 (Counterfactual transition)** projects how state evolves under an intervention-induced response pattern. Without it, the model remains a classifier or passive predictor rather than a world model.
- **CP5 (Quality-control feedback)** closes the loop. Without it, the model has no mechanism to detect, attribute, or learn from its own errors.

This is why the framework is neither a four-point system nor an open-ended list. Four checkpoints would leave the model open-loop: it could simulate with confidence while lacking any mechanism for corrective inspection. Adding a sixth checkpoint, such as separating failure detection from failure correction, would violate functional cohesion. Detection and correction are two sides of the same quality-control responsibility, just as a thermostat integrates sensing, deviation detection, and corrective response within one control function. The five-gate inspection inside CP5 is therefore the internal structure of a single constraint checkpoint, not a set of independent constraints.

In short, CP1 defines "what it is," CP2 quantifies "how much is present," CP3 defines "how it is moved," CP4 projects "what happens if it moves," and CP5 checks "whether it moved correctly" and corrects the next cycle. This mapping exhausts the logical requirements of a biomedical world model that aspires not merely to predict, but to **steer** biological systems within constrained, quality-assured boundaries.

Operationally, these checkpoints function as a compact architectural checklist: **State defined? Intrinsic capability quantified? Intervention-response pattern specified? Counterfactual transition simulated? Deviation inspected and corrected?** A model that cannot answer these five questions may still be useful as a predictor, classifier, or simulator, but it has not yet satisfied the closed-loop requirements for quality-controlled biomedical steerability.

Constraint checkpoint	Common limitation in existing systems	Capomics framework response
CP1: State representation	Patient state is often represented as an image latent, EHR embedding, omics embedding, or general patient representation with limited biological semantics	State is represented as a Capomics intrinsic-capability vector with module-level biological interpretation
CP2: Intrinsic-capability quantification	Biological state variables may be abstract, unmeasured, or difficult to compare across individuals and models	Capomics operationalizes each module's intrinsic capability through measurable molecular readouts, with DNA methylation serving as the current prototype for estimating module-specific CI/PAI-derived values
CP3: Intervention-response semantics	Interventions are often encoded as treatment labels or action codes	Interventions are evaluated by the module response patterns and mIC-state transitions they induce
CP4: Counterfactual transition	Comparison of intervention A versus B, or response differences across baseline states, may depend heavily on statistical extrapolation outside the training distribution	Counterfactuals compare both intervention variation under the same state and state variation under the same intervention, within the root-cause → functional → phenotypic scaffold
CP5: Quality-control feedback	When predictions or interventions fail, the failure may be opaque	Deviations can be inspected across state measurement, intervention specification, module response pattern, mIC state change, or downstream propagation, and then fed back to revise earlier checkpoints

 doi:10.20944/preprints202605.0366.v1

7.5. From Milestone Ladders to an mIC-Centered Steerability Data Flywheel

Existing medical world-model rubrics are primarily milestone-based. They define a sequence of **modeling stages** that a model should pass through, from representation and forecasting to action-conditioned simulation, counterfactual comparison, and planning. This stage ladder is useful for evaluation because it clarifies what a medical world model should eventually be able to do. However, a stage- or milestone-based ladder does not by itself specify a shared biological state space through which these modeling stages can be continuously connected. As a result, each new stage may require additional task-specific data, task-specific transition models, validation criteria, and failure-analysis procedures.

The Capomics framework proposes a different architectural logic. Instead of treating representation, simulation, counterfactual evaluation, planning, and quality control as separated milestones, it organizes them around a shared **mIC state space**. The mIC vector functions as a common carrier of biological state: intervention inputs are interpreted through the module response patterns they induce; response patterns are evaluated as ΔmIC transitions; counterfactuals compare either different inputs under the same mIC state or different mIC states under the same input; and QC feedback localizes failed transitions to explicit premises.

This distinction changes the role of data. In separated stage modeling, later stages often introduce new data requirements and new submodels. In an mIC-centered steerability model, each follow-up measurement, intervention response, and unexpected deviation can be converted into structured feedback for the same state-transition space. The result is a potential **closed-loop data flywheel**: measure mIC state, simulate intervention-induced ΔmIC, observe deviation, assign the deviation to a QC gate, revise the relevant premise, and improve the next rollout. This proposed efficiency advantage remains a hypothesis and should be tested through future benchmarks comparing fragmented milestone-based modeling with mIC-centered closed-loop modeling.

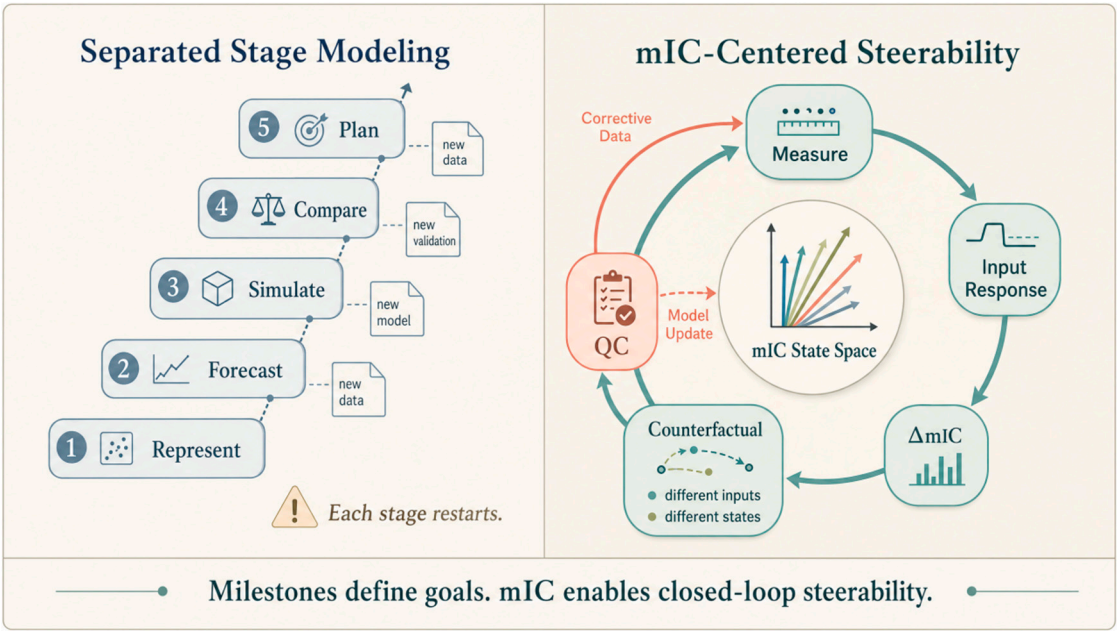

Figure 5. Closed-loop steerability data flywheel versus separated stage modeling. Milestone ladders define progressive modeling goals, but each new stage may require new data, new task-specific models, validation criteria, and failure-analysis procedures. The Capomics framework instead organizes measurement, intervention-response semantics, counterfactual ΔmIC transition, QC inspection, and corrective data collection around a shared mIC state space, forming a closed-loop steerability data flywheel.

 doi:10.20944/preprints202605.0366.v1

7.6. Illustrative Thought Experiment: Counterfactual Rollout Across Three mIC State Patterns

The purpose of the following thought experiment is not to claim clinical validity, but to show what becomes computable once an individual is represented as an mIC vector rather than as a diagnostic label, chronological age, or single risk score. The numerical mIC values below are illustrative normalized values, not clinically validated thresholds.

They are included only to make the direction of reasoning visible; they should not be interpreted as clinical cutoffs, diagnostic thresholds, or treatment-selection criteria.

Consider three hypothetical individuals matched in chronological age at 70 years, but differing in their spectra of module-level intrinsic capability.

Individual A: immune-dominant mIC state pattern. Immune-surveillance mIC is critically low, while mitochondrial-adaptation and metabolic-flexibility mIC are relatively preserved. In illustrative terms, mIC(immune surveillance) might be 0.05, mIC(mitochondrial adaptation) 0.60, mIC(inflammation resolution) 0.30, and mIC(metabolic flexibility) 0.55.

Individual B: mitochondrial-dominant mIC state pattern. Mitochondrial-adaptation mIC is critically low, while immune-surveillance mIC is only moderately depleted. In illustrative terms, mIC(mitochondrial adaptation) might be 0.03, with moderate depletion in inflammation resolution and metabolic flexibility.

Individual C: multi-determinant frailty mIC state pattern. Several modules are simultaneously depleted, with all four illustrative modules below 0.15 and a particularly strong coupling between mitochondrial dysfunction and impaired inflammation resolution.

These three individuals have the same chronological age, but the biological meaning of the same intervention differs across them because their starting mIC states differ. If Individual A receives an immune-directed therapeutic intervention, the deductively constrained model would generate the hypothesis that the current state should move toward improved immune-surveillance mIC, improved immune-cell repertoire diversity or related functional readouts, and eventually reduced infection susceptibility or immune fragility. If the intervention fails, quality-control inspection would ask: Was the starting mIC state measured correctly? Was the intervention specified correctly? Did the expected module response pattern occur? Did the mIC state move toward the desired change? Did downstream propagation occur?

If Individual A instead receives a mitochondria-directed intervention, the framework would expect limited systemic improvement, conditional on the assumed causal scaffold, because mitochondrial adaptation is not the dominant state abnormality in this example. The likely failed premise would be mismatch between the intervention-induced response pattern and the individual's current mIC state. Conversely, in Individual B, a mitochondrial state-remodeling intervention may be more plausible as a first intervention.

Individual C illustrates planning rather than single-action prediction. If the causal scaffold suggests that mitochondrial dysfunction is upstream of impaired inflammation resolution, the model would prioritize interventions predicted to improve mitochondrial-adaptation mIC before interventions predicted to modulate inflammation resolution. A reversed sequence might still affect downstream markers, but would be hypothesized to produce a less durable trajectory because the upstream state abnormality remains unresolved.

This thought experiment illustrates how explanatory structure is specified before simulation begins. A purely inductive world model trained on population data might learn that “intervention X improves outcome Y” on average, but it may struggle to explain why the same intervention produces different trajectories across different mIC states, or why Individual C shows sequence-dependent outcomes. In the deductively constrained framework, treatment heterogeneity is not merely noise around an average effect; it becomes a structured diagnostic signal.

The example generates three falsifiable hypotheses that connect directly to the broader predictions in Section 8: first, individuals should respond more strongly to interventions whose expected response pattern matches their current mIC state than to mismatched interventions; second, in multi-abnormality states, sequencing interventions according to upstream state dependencies

 doi:10.20944/preprints202605.0366.v1

should outperform downstream-first sequencing; third, intervention failure should be traceable to specific premises such as state measurement, intervention design, module response pattern, mIC state change, timing, or downstream propagation.

The abstract archetypes above represent a general reasoning template. As a disease-level illustration, the supplementary case study applies this template to major depressive disorder, showing how mIC profiling could be used to reason about treatment heterogeneity, state-matched versus state-mismatched interventions, sequence dependence, and quality-control feedback: Supplementary Case Study: mIC Profiling of Depression Subtypes.

7.7. What This Framework Does Not Yet Claim

This framework does not claim to have already achieved validated planning/control as defined in existing medical world-model rubrics, nor does it claim to be a clinically deployable treatment-planning system. It also does not claim that Capomics alone is sufficient for all biomedical world modeling tasks. The current contribution is more upstream: a causal-state architecture and intervention-response semantics that can help biomedical world models satisfy and extend the requirements highlighted by existing world-model rubrics.

Actual implementation of high-level planning/control requires longitudinal intervention data, calibrated uncertainty, individualized feedback, prospective validation, and explicit safety constraints. The role of the present framework is to define what must be represented and tested before such systems can become biologically interpretable, quality-controlled, steerable, and clinically useful.

8. Falsifiable Predictions and Research Roadmap

The value of a Perspective / Hypothesis article does not lie in proving all conclusions, but in proposing strong hypotheses that can be tested. This framework generates the following falsifiable predictions.

Before individual predictions are listed, the framework can be translated into three benchmark agendas that make the proposed architecture testable rather than purely conceptual:

1. **State-representation benchmark.** Does an mIC-vector representation explain aging- or disease-related phenotypes better than scalar biological age, diagnostic labels, or uninterpreted latent embeddings?
2. **State-matched intervention-response benchmark.** Does baseline mIC state predict the magnitude and direction of intervention-induced module response patterns?
3. **QC-inspection benchmark.** When an expected transition fails, can the five-gate inspection localize the deviation in a way that improves the next prediction or intervention hypothesis?

Hypothesis	Required data	Empirical test	Falsification criterion
mIC modules age at different rates	Longitudinal methylation or multi-omics data with age and phenotype follow-up	Compare module-specific trajectories within individuals	No reproducible module-specific divergence beyond scalar age
mIC state patterns explain phenotype	mIC spectra, clinical phenotypes, functional measures, and outcomes	Compare single-module, multi-module, and causal-priority models	mIC patterns do not improve explanation or prediction over simpler baselines
State-matched interventions produce stronger response	Baseline mIC, recorded intervention exposure, module response signatures, and follow-up outcomes	Test interaction between baseline mIC state and intervention-induced response pattern	Response is independent of baseline mIC state after appropriate controls

 doi:10.20944/preprints202605.0366.v1

QC inspection improves iteration	Pre/post mIC, intervention specification, response markers, phenotype follow-up, and failed-transition records	Assign deviations to the five gates and test whether gate-specific revision improves subsequent prediction	Deviations cannot be localized above chance or do not improve subsequent modeling

8.1. Prediction One: Modular Aging Rates

Different intrinsic-capability modules within the same individual should show different rates of aging. One person's immune-response capability may show clear aging while their metabolic-flexibility capability remains relatively young; another person may show the opposite.

Testing strategy: Longitudinal DNA methylation cohorts, multi-omics follow-up, functional tests, and clinical phenotype modeling.

8.2. Prediction Two: State Patterns Explain Phenotypes

In individuals with decline across multiple modules, the explanatory unit should often be a state pattern rather than a single module. Upstream or causally prioritized mIC patterns should have greater explanatory power for the overall phenotype than isolated module scores.

Testing strategy: Compare the predictive ability of single-module scores, multi-module state patterns, and causal-priority state representations for clinical outcomes.

8.3. Prediction Three: State-Matched Interventions Produce Stronger Response

The effect of the same intervention should depend on the starting mIC state. Interventions whose expected module response pattern matches an individual's current mIC pattern should produce stronger response signatures, larger ΔmIC shifts toward the desired direction, and more consistent downstream phenotypic change.

Testing strategy: Joint analysis of Capomics state, pathway activity, module response signatures, ΔmIC, and intervention response.

8.4. Prediction Four: Root-Cause-Layer Intervention Is More Durable

Compared with interventions that only alter the functional layer or phenotypic layer, interventions capable of changing upstream mIC should produce more durable trajectory changes.

Testing strategy: Compare short-term biomarker improvement with long-term Capomics trajectory, functional outcome, and disease-risk changes.

8.5. Prediction Five: Sequential Intervention Planning Is Superior to Random Combination

When multiple mIC state abnormalities coexist, intervention sequencing according to upstream state dependency should outperform unordered combination or a single strong intervention.

Testing strategy: N-of-1 longitudinal intervention, real-world data, state-transition simulation, and adaptive clinical trials.

8.6. Research Roadmap

This framework can proceed in four steps:

1. **Construct module maps.** Define mIC modules and establish their gene, pathway, cell-type, and phenotype mappings.
2. **Establish state readouts.** Train or organize module-level Capomics readouts to form mIC spectra and mIC-vector outputs.
3. **Validate causal association.** Use longitudinal cohorts, intervention studies, and perturbation experiments to test relationships between module readouts and functional change.

 doi:10.20944/preprints202605.0366.v1

4. **Develop planning models.** Integrate mIC spectra, intervention specifications, desired state changes, module response readouts, ΔmIC, and feedback data into individualized world models for counterfactual rollout and sequential intervention planning.

8.7. Outlook: Individualized Trial Designs and Mechanism-Based Evidence

An important future application of the proposed framework lies in individualized trial design and mechanism-based evidence generation.

The FDA's 2026 draft guidance on the Plausible Mechanism Framework is specific to individualized therapies for ultra-rare genetic conditions and remains a draft [27]. Nevertheless, it highlights a broader logic for mechanism-based evidence. Three elements are especially relevant here: **defined causal abnormality**, **mechanism engagement**, and **clinically meaningful improvement**.

First, **defined causal abnormality** means that an individualized therapy should be anchored to a specific disease-causing biological abnormality. Second, **mechanism engagement** means that the intervention should plausibly engage the relevant biological mechanism. Third, **clinically meaningful improvement** means that mechanism engagement should be connected to patient-relevant benefit rather than remaining a molecular observation alone.

The architecture proposed here maps naturally onto these elements, while remaining outside any regulatory claim. First, **mIC-vector analysis** may help define an individualized causal state rather than only a diagnostic category. Second, **intervention-induced module response signatures** may provide a measurable layer for mechanism engagement. Third, **downstream phenotypic tracking** corresponds to the final gate of quality-control inspection.

This mapping also explains why individualized evidence cannot be replaced by large foundation models alone. Recent work argues that N-of-1 trials remain essential for individual-level causal inference, while multi-agent N-of-1 AI ecosystems may provide complementary computational infrastructure [28,29]. In this context, the Deductively Constrained Capomics World Model should be viewed as a candidate architecture for future mechanism-aware, quality-control-inspectable, individualized evidence generation.

9. Limitations

9.1. It Is Not a Validated Algorithm

This article does not report new experimental data or provide a validated clinical decision system. It proposes a theoretical framework and computable prototype route. Any clinical application requires independent validation.

9.2. Adaptive Capacity Is Not the Only Possible Definition of Life

This article operationally defines life as an ensemble of adaptive capacities for the purpose of world-model construction, not to exhaust the philosophy of life. Other definitions, including self-organization, metabolism, replication, information processing, and evolutionary selection, can complement this framework.

9.3. Modularity Is an Abstraction, Not Mechanical Segmentation

Living systems exhibit high coupling, feedback, redundancy, and degeneracy. Modularity is an abstraction adopted for computable modeling and intervention planning, and should not be misunderstood as implying that the biological system consists of independent parts.

 doi:10.20944/preprints202605.0366.v1

9.4. Capomics Is a Prototype, Not the Only Substrate

DNA methylation may be a cause, consequence, or common marker. Capomics must be validated through longitudinal data, functional experiments, and intervention studies. Future state representations may integrate transcriptomics, proteomics, metabolomics, spatial omics, wearable devices, and behavioral data.

9.5. The Three-Layer Scaffold Is Not a Complete Ontology

The root-cause, functional, and phenotypic layers constitute a minimal causal scaffold, but real living systems also involve environmental layers, tissue-spatial layers, social-behavioral layers, temporal-scale layers, and developmental history. Different tasks may require more complex architectures.

9.6. Intervention-Response Semantics Cannot Replace Efficacy Validation

"Intervention → module response pattern → mIC-state transition" provides a language for intervention design and hypothesis generation, not a guarantee of efficacy. Whether any specific intervention is effective still requires mechanistic experiments, dose–response studies, safety evaluation, and prospective clinical validation.

9.7. Boundary Conditions for Use

The framework is most applicable when biological state can be measured longitudinally, intervention exposure can be recorded with sufficient detail, response signatures can be tracked, phenotypic propagation can be observed, and safety constraints are explicit. It is less applicable when state variables are unmeasurable, intervention exposure is poorly recorded, outcomes are too sparse, causal layers cannot be mapped, or safety constraints are unavailable. These boundary conditions should be treated as design requirements for future CAPOVIME-like implementations rather than as assumptions that can be ignored.

10. Conclusion

The central contribution of this article is the shift from prediction-centered biomedical AI to **steerability**-centered biomedical world modeling. Prediction remains necessary, but it is not the final objective. A useful biomedical world model should help determine how a biological state can be guided toward a desired direction under explicit causal, measurement, intervention-response, transition, and feedback constraints.

This shift also distinguishes the proposed framework from a purely milestone-based interpretation of medical world-model development. Existing stage ladders define important modeling goals, but those goals may remain separated if each stage requires its own data structure, model family, validation logic, and failure-analysis method. The Capomics framework instead proposes the **mIC vector** as a shared state carrier that connects state measurement, intervention-induced response, counterfactual ΔmIC transition, and QC feedback within one closed-loop data flywheel. Its proposed advantage is not that it has already solved planning/control, but that it defines a unified state-transition space in which planning, deviation diagnosis, and iterative data collection can become mutually reinforcing.

This is why the proposed framework is organized as a closed loop rather than a linear prediction ladder. Steerability requires the model to define state, measure state, specify intervention-induced state movement, simulate alternative transitions, and inspect deviations when the expected transition fails. As the biomedical world-model field consolidates, this article offers these five functions as a candidate set of architectural primitives: state representation, intrinsic-capability quantification, intervention-response semantics, counterfactual transition, and quality-control feedback.

The Deductively Constrained Capomics World Model proposed here attempts to provide a new framework for this problem. It uses adaptive capacity as the system-level operational definition of

 doi:10.20944/preprints202605.0366.v1

life state, **module-level intrinsic capability (mIC)** as the module-level state variable, the root-cause–function–phenotype hierarchy as the minimal causal scaffold, Capomics as the computable state prototype, the **mIC vector** as a candidate interoperable state representation format, and module response patterns as the computable bridge from environmental exposures or interventions to mIC-state transitions.

The framework does not claim to replace data-driven models, but to provide causal constraints for them. It does not claim to provide unconditional efficacy guarantees, but to provide a quality-control inspection structure for unexpected outcomes. It does not claim that Capomics is the only answer, but proposes an immediately actionable computational prototype.

More importantly, the framework specifies a structural requirement that is not yet explicit in many biomedical world-model benchmarks: built-in **quality-control feedback**. While aligning with standard requirements such as representation, forecasting, counterfactual simulation, and planning, it also asks why an expected trajectory did not occur and which constraint checkpoint should be revised. Many biomedical world models first learn to simulate and then seek post-hoc explanations. In contrast, the deductively constrained framework specifies the explanatory structure before simulation begins.

At the implementation layer, this framework can be organized as **CAPOVIME**, the **Capomics Virtual Intervention Model Ecosystem**. CAPOVIME, accessible at capovime.org and currently redirecting to capovime.com, together with its future evolutionary simulation node evosika.org, is proposed as an open engineering ecosystem for world-model-based medicine.

This framework does not compete with existing AI drug discovery engines; it aims to provide a causal coordination layer for them. Future work should validate whether CAPOVIME-like implementations can support safe, quality-controlled biomedical intelligence across longitudinal intervention datasets.

If this framework holds, its main advance will not be a better predictor alone. It will be a steerable reasoning architecture for biomedicine: one that can simulate interventions, inspect unexpected outcomes, and revise the next intervention hypothesis within an explicit causal scaffold. This may provide a new route connecting first principles, molecular measurement, individualized intervention, quality-control feedback, and deviation-correctable reasoning in aging medicine, complex chronic disease management, and drug discovery.

The framework is offered not as a finished theory to be accepted, but as a deductive chain to be inspected, challenged, calibrated, and, where necessary, refuted. Its value will depend on whether its postulates, scaffold, state representation, intervention-response semantics, and quality-control logic can survive empirical testing and support useful counterfactual modeling in real biomedical systems.

Supplementary Materials: The following supporting information can be downloaded at the website of this paper posted on Preprints.org.

Author Contributions: Conceptualization, J.X.; writing, review and editing, J.X. The author has read and agreed to the published version of the manuscript.

Funding: This research received no external funding.

Institutional Review Board Statement: Not applicable.

Informed Consent Statement: Not applicable.

Data Availability Statement: Not applicable.

Conflicts of Interest: The author declares no conflicts of interest.

References

1. Hafner D, Pasukonis J, Ba J, Lillicrap TP. Mastering diverse control tasks through world models. Nature. 2025;640(8059):647–653. doi:10.1038/s41586-025-08744-2.

 doi:10.20944/preprints202605.0366.v1

2. Qazi MA, Nadeem M, Yaqub M. Beyond Generative AI: World Models for Clinical Prediction, Counterfactuals, and Planning. arXiv preprint. 2025; arXiv:2511.16333.
3. Saeed N, Hassan S, Khan S, Qazi MA, Maier-Hein KH, Khan S, Yaqub M. Medical World Model: From Passive Prediction to Active Simulation in Medicine. Preprints. 2026;2026042168. doi:10.20944/preprints202604.2168.v1.
4. Yang Y, Wang ZY, Liu Q, Sun S, Wang K, Chellappa R, Zhou Z, Yuille A, Zhu L, Zhang YD, Chen J. Medical World Model: Generative Simulation of Tumor Evolution for Treatment Planning. Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV). 2025:8319–8329; arXiv:2506.02327.
5. Wei Z, Ma R, Wang Z, Li Z, Song S, Zheng S. VCWorld: A Biological World Model for Virtual Cell Simulation. International Conference on Learning Representations. 2026; arXiv:2512.00306.
6. Callahan TJ, Beckwith Z, Merth T, van der Poel C, Lewis A, Lemos P. Virtual Cells as Causal World Models: A Perspective on Evaluation. NeurIPS 2025 AI4D3 Workshop. 2025.
7. Business Wire. GSK Licenses Noetik's AI Foundation Models in Anchor Partnership to Transform Cancer Therapeutic Research and Development. January 8, 2026. Available from: https://www.businesswire.com/news/home/20260108468293/en/GSK-Licenses-Noetiks-AI-Foundation-Models-in-Anchor-Partnership-to-Transform-Cancer-Therapeutic-Research-and-Development.
8. Pearl J. Causality: Models, Reasoning, and Inference. 2nd ed. Cambridge: Cambridge University Press; 2009.
9. Hernán MA, Robins JM. Causal Inference: What If. Boca Raton: Chapman & Hall/CRC; 2020.
10. World Health Organization. World Report on Ageing and Health. Geneva: World Health Organization; 2015. ISBN:9789241565042.
11. López-Otín C, Blasco MA, Partridge L, Serrano M, Kroemer G. The hallmarks of aging. Cell. 2013;153(6):1194–1217. doi:10.1016/j.cell.2013.05.039.
12. López-Otín C, Blasco MA, Partridge L, Serrano M, Kroemer G. Hallmarks of aging: An expanding universe. Cell. 2023;186(2):243–278. doi:10.1016/j.cell.2022.11.001.
13. Barabási AL, Gulbahce N, Loscalzo J. Network medicine: A network-based approach to human disease. Nature Reviews Genetics. 2011;12(1):56–68. doi:10.1038/nrg2918.
14. Kitano H. Systems biology: a brief overview. Science. 2002;295(5560):1662–1664. doi:10.1126/science.1069492.
15. Edelman GM, Gally JA. Degeneracy and complexity in biological systems. Proceedings of the National Academy of Sciences of the United States of America. 2001;98(24):13763–13768. doi:10.1073/pnas.231499798.
16. Hanahan D, Weinberg RA. Hallmarks of cancer: the next generation. Cell. 2011;144(5):646–674. doi:10.1016/j.cell.2011.02.013.
17. Hanahan D. Hallmarks of cancer: new dimensions. Cancer Discovery. 2022;12(1):31–46. Available from: https://doi.org/10.1158/2159-8290.CD-21-1059.
18. Horvath S. DNA methylation age of human tissues and cell types. Genome Biology. 2013;14(10):R115. doi:10.1186/gb-2013-14-10-r115.
19. Hannum G, Guinney J, Zhao L, Zhang L, Hughes G, Sadda S, et al. Genome-wide methylation profiles reveal quantitative views of human aging rates. Molecular Cell. 2013;49(2):359–367. doi:10.1016/j.molcel.2012.10.016.
20. Levine ME, Lu AT, Quach A, Chen BH, Assimes TL, Bandinelli S, et al. An epigenetic biomarker of aging for lifespan and healthspan. Aging (Albany NY). 2018;10(4):573–591. doi:10.18632/aging.101414.
21. Lu AT, Quach A, Wilson JG, Reiner AP, Aviv A, Raj K, et al. DNA methylation GrimAge strongly predicts lifespan and healthspan. Aging (Albany NY). 2019;11(2):303–327. doi:10.18632/aging.101684.
22. Xiong J. Next Generation Aging Clock: A Novel Approach to Decoding Human Aging Through Over 3000 Cellular Pathways. bioRxiv. 2024. doi:10.1101/2024.06.18.599632.
23. Ying K, Tyshkovskiy A, Chen Q, Latorre-Crespo E, Zhang B, Liu H, et al. High-dimensional Ageome Representations of Biological Aging across Functional Modules. bioRxiv. 2024. doi:10.1101/2024.09.17.613599.

 doi:10.20944/preprints202605.0366.v1

24. Li P, Zhu J, Wang S, Zhuang H, Zhang S, Huang Z, et al. Decoding disease-specific ageing mechanisms through pathway-level epigenetic clock: insights from multi-cohort validation. eBioMedicine. 2025;118:105829. doi:10.1016/j.ebiom.2025.105829.
25. Fuentealba M, Rouch L, Guyonnet S, Lemaitre JM, de Souto Barreto P, Vellas B, et al. A blood-based epigenetic clock for intrinsic capacity predicts mortality and is associated with clinical, immunological and lifestyle factors. Nature Aging. 2025;5:1207–1216. doi:10.1038/s43587-025-00883-5.
26. Jia X, Liu Z. Developing a blood-based epigenetic clock for intrinsic capacity. Nature Aging. 2025;5(7):1188–1190. doi:10.1038/s43587-025-00899-x.
27. U.S. Food and Drug Administration. Considerations for the Use of the Plausible Mechanism Framework to Develop Individualized Therapies That Target Specific Genetic Conditions with Known Biological Cause: Draft Guidance for Industry. Silver Spring, MD: FDA; 2026.
28. Konigorski S, Vedder JE, Owoyele BA, Özkan İ. Personalization of Large Foundation Models for Health Interventions. arXiv preprint. 2026; arXiv:2601.03482.
29. Fard P, Azhir A, Rezaii N, Tian J, Estiri H. An N-of-1 Artificial Intelligence Ecosystem for Precision Medicine. arXiv preprint. 2025; arXiv:2510.24359.