

Control theory of aging and longevity for gerotherapeutic drug discovery: state space, vector fields, modalities, safety, and the minimum safe cost of functional restoration

Alex Zhavoronkov^{1,2,3,4,5}, Bud Mishra⁶

¹Insilico Medicine Hong Kong Ltd., Hong Kong SAR, China

²Insilico Medicine Shanghai Ltd., Pudong 201203, Shanghai, China

³Insilico Medicine AI Ltd., Masdar, Abu Dhabi, UAE

⁴Insilico Medicine US Inc., Cambridge, MA 02138, USA

⁵Buck Institute for Research on Aging, Novato, CA 94945, USA

⁶Courant Institute School of Mathematics, Computer and Data Sciences, New York University, New York, NY 10012, USA

Correspondence to: Alex Zhavoronkov; email: alex@insilicomedicine.com

Keywords: aging, control theory, biological age, drug discovery, gerotherapeutics

Received: May 14, 2026

Accepted: August 6, 2026

Published: September 16, 2026

Copyright: © 2026 Zhavoronkov and Mishra. This is an open access article distributed under the terms of the [Creative Commons Attribution License](https://creativecommons.org/licenses/by/4.0/) (CC BY 4.0), which permits unrestricted use, distribution, and reproduction in any medium, provided the original author and source are credited.

ABSTRACT

Existing aging theories describe what changes with age but do not prescribe how to intervene. We propose a control-theoretic framework that is not merely descriptive but prescriptive: it specifies which intervention, at which dose and sequence, under which safety constraints, will restore a measured biological state to a functional region. Aging is defined as progressive loss of safe controllability; biological age is the minimum safe control cost of functional restoration. Drugs are modeled as vector fields on biological state space whose non-commutativity, quantified by Lie brackets, predicts that intervention order determines outcome. The core differentiation from prior theories is operational: the framework outputs ranked targets, optimal sequences, safety-constrained protocols, and falsifiable predictions directly usable in drug discovery, rather than mechanistic ontologies or correlative biomarkers. We present a five-dimensional ODE model with analytic Lie-bracket derivation, a modality-aware control layer, three translational case studies, an implementation architecture with power analysis, and empirical scoring of aging interventions across five biological epochs. Twenty falsifiable predictions are enumerated. The central claim is that control-value reduction predicts translational success better than Hallmark annotation or biomarker reversal alone. If validated, this provides the missing interventional layer connecting aging biology to rational gerotherapeutic discovery.

INTRODUCTION

From mechanisms to intervention laws

Despite decades of intensive research, aging biology lacks consensus on its most fundamental questions. A recent large-scale survey of leading experts revealed significant disagreement on the definition of aging, its causes, its onset, the meaning of rejuvenation, and

whether aging should be classified as a disease [1], with none of these foundational questions receiving a majority opinion. The disagreement is not merely philosophical: different definitions lead to divergent experimental approaches, therapeutic strategies, and measures of success. The absence of a unifying quantitative framework that can accommodate these perspectives while generating falsifiable, intervention-relevant predictions motivates the present work.

The field has developed complementary frameworks. Evolutionary theories explain why natural selection may permit late-life decline [2–4]. Damage and maintenance theories describe accumulation of molecular and cellular lesions [5, 6]. The Hallmarks of Aging organize diverse mechanisms into a widely used ontology [7, 8]. SENS provides a repair-oriented catalogue of damage classes [9]. Information theories emphasize epigenetic drift and partial reversibility through reprogramming [10, 11]. Geroscience frames aging mechanisms as modifiable drivers of chronic disease [12]. Hyperfunction theory highlights persistent growth signaling as a driver of late-life pathology [13]. These frameworks have transformed the field. As aging biology becomes increasingly connected to drug discovery, however, a different kind of theoretical object is required: it is no longer sufficient to ask only which processes change with age. The therapeutic science of aging must answer a more operational question.

Operational question

Given a measured biological state x , which intervention u , at which dose and duration, in which sequence, under which safety constraints, will move the system into a healthier functional region with acceptable risk and preserved future controllability?

This question is not answered by a list of Hallmarks, a catalogue of damage classes, or a metaphor of information loss. Those frameworks identify relevant biology, but they do not by themselves specify the dynamics of the aged system, the state-dependent response to intervention, the safety-constrained set of admissible actions, or the objective function that defines therapeutic success. The relationship is analogous to that between thermodynamics and statistical mechanics: a phenomenological ontology organizes observations, while a dynamical and interventional layer supplies equations of motion, explicit inputs, response operators, constraints, objective functions, and predictions that can fail (Supplementary Information 1). The purpose of this paper is to propose that layer. We do not argue that existing theories are wrong; the framework depends on their insights. Our contribution is narrower but operationally important: we formulate aging as a problem of safe controllability and define biological age as the minimum safe control cost required to maintain or restore function under a specified intervention library.

Control theory is the natural mathematical language for this problem [14–17]. Its central objects, including state variables, dynamics, inputs, response operators, reachable sets, value functions, constraints, and optimal policies, map directly onto the needs of

interventional biology. A biological system has internal state variables that evolve under endogenous dynamics and stochastic perturbations; drugs, cell therapies, genetic perturbations, nutritional interventions, exercise, and surgery act as inputs whose effects are state-dependent, so the same intervention can be beneficial, neutral, or harmful depending on biological context. Therapeutic success is not movement toward an abstract youthful ideal but restoration or maintenance of function under acceptable risk and preserved future controllability. In this framework a drug is a vector field on biological state space, a state-dependent transformation operator that pushes the system along particular directions. A target is valuable if its modulation reduces the optimal control cost of returning the system to a functional viability set; a combination is valuable if its component vector fields expand the reachable safe set more than either intervention alone; and a sequence matters when the vector fields do not commute.

This last point provides the framework’s central geometric content. Local controllability, in the sense of Kalman [15, 18] and the geometric control theory of Nijmeijer and Schaft [19] and Isidori [20], can be improved by exploiting the noncommutativity of pairs of interventions. Given intervention vector fields g_A and g_B , sequential application along g_A then g_B , and their reverses, generically returns the system to a state displaced from the start by an amount proportional to ε^2 along a new direction, quantified by the Lie bracket $[g_A, g_B]$. The same logic distinguishes local controllability (a small reachable set near the current state) from global controllability (reachability of any youthful state from the aged state). An aged tissue may retain local controllability while having lost global controllability because irreversible coordinates have collapsed the reachable set. The progressive collapse from global to merely local controllability, and eventually to loss of even local controllability in critical dimensions, is the aging process under this framework. A clinical running illustration (type 2 diabetes monitored by continuous glucose monitoring) is developed in Supplementary Information 2 and referenced where relevant below.

$$[g_A, g_B](x) = Dg_B(x) g_A(x) - Dg_A(x) g_B(x)$$

Three developments make this framework timely. First, multi-omics and single-cell technologies increasingly permit high-dimensional measurement of biological state. Second, perturbational datasets, including CRISPR screens, LINCS/CMap profiles [21, 22], drug-response atlases, and animal intervention studies, provide empirical information about how systems respond to inputs. Third, machine-learning methods can in principle estimate state

representations, infer dynamics, rank targets, model response heterogeneity, and integrate heterogeneous data [23–25]. These technologies do not eliminate the need for theory; they make a mathematically explicit theory of interventional aging tractable. This paper presents the framework and a computational pipeline for prioritizing gerotherapeutic interventions; it does not report a trained engine that has learned intervention vector fields from perturbational data. The estimation of state-dependent vector fields from data is developed as a method and flagged as future and companion work (Supplementary Information 5), not as an existing capability. We have previously argued that artificial intelligence provides machinery to integrate heterogeneous aging data into longevity pipelines [24], demonstrated an early pathway-activation instantiation of aging as a multi-dimensional problem [26], and shown that aging clocks can serve as instruments for therapeutic target discovery [27]. The control-theoretic framework supplies the mathematical scaffold connecting these advances: it defines the state space in which clocks operate, the objective function that target discovery optimizes, and the safety constraints under which such interventions must be evaluated. Its value will be measured by whether it improves prediction, experimental design, and therapeutic outcomes.

Background: the landscape of theories of aging

Aging is a hierarchy of phenomena spanning molecular damage, cellular dysregulation, tissue remodeling, organ decline, and rising mortality risk, and the major theoretical families address different levels of it: evolutionary [2–4, 6, 28], damage and repair [5, 9, 29–31], the Hallmarks ontology [7, 8], information and reprogramming [10, 11, 32–34], geroscience [12, 35, 36], hyperfunction [13, 37–41], and reliability and resilience [42–46]. Each contributes essential insight; none alone provides a full control law for intervention (a detailed per-theory review is given in Supplementary Information 1). Despite their differences, most lack eight components required for drug discovery: (1) state variables, the mathematical state of the system; (2) dynamics, how the state evolves without intervention; (3) control inputs, what can be manipulated; (4) response operators, how each intervention affects the state and how that effect depends on context; (5) constraints, which interventions are unsafe, infeasible, or toxic; (6) objective functions, what counts as improvement; (7) optimality conditions, which intervention or sequence minimizes cost while restoring function; and (8) falsifiable predictions, what quantitative outcomes would refute the framework. Their absence does not make previous theories wrong; it indicates they were built to answer different questions. Drug discovery requires a theory of safe control, which the next section formalizes.

RESULTS

The control theory of aging

Biological state space

We represent a biological system as a latent state vector $x(t) \in \mathbb{R}^n$ describing a cell population, tissue, organ, physiological subsystem, or organism. It is latent because no platform measures all relevant variables; observed data are noisy projections,

$$y(t) = h(x(t)) + \varepsilon(t),$$

where $y(t)$ may include transcriptomic, proteomic, metabolomic, epigenetic, imaging, clinical, and functional measurements; h is an observation function; and $\varepsilon(t)$ represents measurement error and unobserved variation. The components of $x(t)$ need not correspond one-to-one to measured biomarkers; in practice $x(t)$ is inferred as a structured representation learned from multi-omics, clinical, and perturbational data, with coordinates spanning epigenetic, transcriptional, metabolic, mitochondrial, proteostatic [47], senescent, immune, regenerative, matrix, vascular, and systemic-functional axes. The full specification, including hierarchical state representations and measurement specifics, is given in Supplementary Information 2.

The functional viability set

Aging research often describes intervention goals as “rejuvenation” or reversal of biological age. These are imprecise: the goal of medicine is not to make every molecular feature identical to youth but to preserve or restore function safely. We define a functional *viability set*

$$\mathcal{V} = \{x: F_i(x) \geq \theta_i, \forall i\}, \quad \text{eq 01}$$

where each $F_i(x)$ is a functional measure and θ_i an acceptable threshold (cardiac reserve, renal filtration, immune competence, wound healing, muscle force, cognition, pulmonary compliance, hematopoietic output, metabolic flexibility, vascular reactivity). Thresholds need not be those of a young adult; they may be age-appropriate, disease-specific, or patient-specific. The viability set avoids treating youth as an attractor: developmental trajectories are not simply reversible aging trajectories, and a function-based objective is more biologically and clinically appropriate.

The separable form in [eq 01] treats each functional measure independently, but physiological constraints are frequently interdependent. The general form of the viability set is

$$\mathcal{V} = \{x: G(x) \geq 0\},$$

where $G: \mathbb{R}^n \rightarrow \mathbb{R}^p$ is a vector of joint constraint functions and the inequality is taken componentwise. The separable threshold form $F_i(x) \geq \theta_i$ is the special case $G_i(x) = F_i(x) - \theta_i$ in which each constraint involves a single functional coordinate. Joint constraints capture the common situation in which an acceptable value of one function depends on other coordinates: an acceptable body weight depends jointly on height, age, and sex; an acceptable cardiac output depends on body size and metabolic demand; an acceptable hematocrit depends on altitude and sex. When such dependencies matter, G is specified directly rather than assembled from independent thresholds, and the constructions below carry over unchanged with $\mathcal{V}$ read in its general form.

Before defining the viability kernel, we specify the admissible safe control set $\mathcal{U}_{\text{safe}}$, the set of control functions $u(\cdot)$ satisfying the safety, feasibility, pharmacologic, dosing, and ethical constraints of the problem, including the hard and soft constraints developed in 4.3; $\mathcal{U}_{\text{safe}}$ is in general state-dependent, so the same intervention may be admissible in one state and forbidden in another. A related concept is then the *viability kernel*, the set of initial states from which some admissible safe control keeps the system in $\mathcal{V}$, or drives it into $\mathcal{V}$ within a finite entry time and holds it there over the horizon,

$$\mathcal{K} = \left\{ x_0: \exists u(\cdot) \in \mathcal{U}_{\text{safe}} \text{ and } t_e \in [0, T] \right. \\ \left. \text{such that } x_t \in \mathcal{V} \text{ for all } t \in [t_e, T] \right\}. \quad \text{eq 02}$$

This definition admits initial states that lie outside $\mathcal{V}$ but can be steered into it and maintained thereafter, so that $\mathcal{V} \subseteq \mathcal{K}$: any state already in $\mathcal{V}$ that can be held there has entry time $t_e = 0$. In the classical viability-theory sense the viability kernel is a subset of the constraint set; the object in [eq 02] is the corresponding *reach-avoid* (capture-basin) generalization that additionally admits states outside $\mathcal{V}$ from which $\mathcal{V}$ can be safely reached and thereafter maintained, and we use it throughout in this sense. States outside the kernel may still be transiently improved but cannot be brought into and held within functional bounds under available safe interventions, providing a formal language for frailty, irreversibility, and late-stage disease (Figure 1).

The viability kernel generalizes the classical basin of attraction. For an autonomous dynamical system $\dot{x} = f(x)$ with a stable attractor $\mathcal{V}$, the basin of attraction is the set of initial states whose uncontrolled trajectories converge to $\mathcal{V}$ as $t \rightarrow \infty$. The viability kernel differs in three

respects. It is control-dependent: membership requires the existence of an admissible control, not spontaneous convergence, so the kernel expands as the intervention library grows and contracts as it shrinks. It is safety-constrained: only controls in $\mathcal{U}_{\text{safe}}$ are allowed, so states reachable only through forbidden trajectories are excluded. It is finite-horizon: entry and maintenance are required over $[0, T]$ rather than asymptotically, which matters when biological deadlines make the asymptotic limit irrelevant. When the control set is empty and $T \rightarrow \infty$, the kernel reduces to the ordinary basin of attraction of $\mathcal{V}$.

Aging dynamics without intervention

In the absence of intervention, the biological state evolves according to stochastic dynamics,

$$dx_t = f(x_t, a, e, g)dt + \Sigma(x_t)dW_t,$$

where f is the endogenous drift; a chronological age; e environment; g genotype; Σ is the state-dependent noise matrix; and W_t is a Wiener process. The inclusion does not imply that age is an independent causal force; it indexes time-dependent exposures, cumulative damage, and changing regulatory regimes. The noise matrix is state-dependent because youthful well-buffered states have small while pre-frail and senescent-cell-rich states show inflated variance; estimation procedures are detailed in Supplementary Information 2. Aging corresponds to several changes: the drift increasingly moves the system away from $\mathcal{V}$ (the functional viability set); stochastic instability rises with transcriptional noise, epigenetic drift, and immune contraction [49–51]; recovery after perturbation slows, a known early-warning signal near critical transitions [45, 46]; the dimensionality of safe control decreases, and the cost of restoration increases until, at advanced stages, no safe path may exist. Gompertz-like increases in hazard [43] may emerge as the probability of crossing viability boundaries rises with accumulating drift, noise, and loss of redundancy.

Intervention as control

Interventions enter the dynamics as control inputs,

$$dx_t = f(x_t)dt + \sum_{j=1}^m g_j(x_t)u_j(t)dt + \Sigma(x_t)dW_t \quad \text{eq 03} \\ = f(x_t)dt + G(x_t)u_t dt + \Sigma(x_t)dW_t,$$

where u_t denotes the intensity, dose, schedule, or exposure of intervention, G is the corresponding intervention vector field, and collects all available

fields. [eq 03] is the core translation from aging biology to drug discovery: a drug is a transformation of biological state rather than a label or pathway inhibitor. The vector field is state-dependent, which is biologically central. Rapamycin in a metabolically overactive, inflammatory, pre-frail state may reduce pathology; in a frail state requiring wound repair it may impair recovery [51–53]. The sign of an intervention is therefore not intrinsic but depends on the value function: locally, intervention is beneficial when and harmful when, formalizing context dependence as a geometric property rather than a rhetorical caveat. $u_j(t)jg_j(x)G(x)$ [eq:controlled-sde-matrix] $g_j(x)jg_j(x)^\top \nabla V(x) < 0, g_j(x)^\top \nabla V(x) > 0$.

Biological age as control cost: the core definition

We define the optimal control value function

$$V(x_0, T) = \min_{u(\cdot) \in \mathcal{U}_{\text{safe}}} \mathbb{E} \left[\int_0^T (\ell(x_t) + \lambda c(u_t) + \rho r(x_t, u_t)) dt + \Phi(x_T) \right]$$

subject to

$$dx_t = f(x_t)dt + G(x_t)u_tdt + \Sigma(x_t)dW_t$$

Here $\ell(x_t)$ penalizes functional loss or distance from $\mathcal{V}$; $c(u_t)$ penalizes intervention burden; $r(x_t, u_t)$ penalizes toxicity and safety risk; $\Phi(x_T)$ penalizes terminal dysfunction; and $\mathcal{U}_{\text{safe}}$ is the admissible safe control set defined in 3.2. A simple choice is $\ell(x) = d(x, \mathcal{V})$ or $\ell(x) = \sum_i w_i \max(0, \theta_i - F_i(x))$. We define *control biological age* as

$$\text{BA}_{\text{control}}(x_0) = \varphi(V(x_0, T)),$$

where φ maps restoration cost to a clinically interpretable scale. Crucially, $\text{BA}_{\text{control}}$ is *intervention-relative*: V depends on the admissible intervention set $\mathcal{U}$, measurement model $\mathcal{M}$, cost function $\mathcal{C}$, horizon T , and safety constraints. A state may be highly controllable using cell therapy and gene editing but poorly controllable using only small molecules.

The Geometry of Aging: Navigating the Biological State Space




Figure 1. The geometry of aging: navigating the biological state space. A 2D projection of the n -dimensional biological state space showing the viability set $\mathcal{V}$ (region of acceptable function), the viability kernel $\mathcal{K}(\mathcal{V})$ (recovery zone), and the absorbing region $\mathcal{R}_{\text{unsafe}}$ (irreversible loss). Consistent with the finite-entry-time definition in [eq:kernel], the kernel contains $\mathcal{V}$ and also states that begin outside $\mathcal{V}$ from which an admissible safe control reaches $\mathcal{V}$ within a finite entry time and remains thereafter, so $\mathcal{V} \subseteq \mathcal{K}(\mathcal{V})$. Trajectory A (rescue) begins outside $\mathcal{V}$ but inside $\mathcal{K}(\mathcal{V})$ and returns to $\mathcal{V}$ under optimal intervention $u^*(\cdot)$; trajectory B begins outside $\mathcal{K}(\mathcal{V})$, crosses the biological deadline, and terminates in $\mathcal{R}_{\text{unsafe}}$.

Core definition

A young system has low and returns to function through endogenous repair or modest intervention. A middle-aged system has higher, requiring stronger intervention. A frail system has high, reflecting narrow safety margins. A terminally damaged system has effectively infinite: no admissible control path reaches. Chronological age measures time lived; biomarker age [54–57] predicts age or outcomes from molecular features; Hallmark burden scores age-associated processes; control biological age measures how difficult it is to restore or maintain function safely. Only the last is directly actionable $VVVV\mathcal{V}$.

Around a state or trajectory, dynamics can be approximated as $\dot{x} = A(a)x + B(a)u$, with $A(a)$ age-dependent endogenous dynamics and $B(a)$ intervention susceptibility. The controllability Gramian over horizon T is

$$W(T, a) = \int_0^T e^{A(a)\tau} B(a) B(a)^\top e^{A(a)^\top \tau} d\tau,$$

and for linear systems the minimum energy to move from x_0 to x_T is

$$E^*(x_0, x_T, a) = (x_T - e^{A(a)T} x_0)^\top W(T, a)^{-1} (x_T - e^{A(a)T} x_0).$$

With age, the smallest eigenvalues of should decline and the minimum restoration energy should increase. The value-function biological age connects to classical survival analysis [58–60]: with hazard and survival, the hazard is a monotone function of the value function,

$$W\lambda(t|x(t))S(t) = \exp\left(-\int_0^t \lambda(s|x(s)) ds\right)$$

$$\lambda(t|x(t)) = \Psi(V(x(t), T_{\text{ref}})),$$

for a non-negative link Ψ and reference horizon T_{ref} . This formalizes biological-age acceleration (dV/dt above the cohort mean), geroprotection (kernel expansion that shifts V downward), and mortality-rate deceleration (survivor selection for low- V states). Horizon T is a problem-specification choice (trial duration for trials; $T_{\text{ref}} = 5$ years for cross-population comparison) that shortens as biological deadlines approach; the full treatment of horizon selection and the Kalman controllability/observability decomposition appear in Supplementary Information 2.

Network controllability, irreversibility, and biological deadlines

The framework is closely related to, but distinct from, the network-controllability program of Liu et al. [61], in which a linearized network is structurally controllable when a minimum set of driver nodes determined by maximum matching is actuated. In aging, nodes represent senescent burden, mitochondrial dysfunction, proteostatic stress, stem-cell exhaustion, epigenetic drift, inflammation, matrix remodeling, nutrient sensing, telomere and genomic state, intercellular communication, and autophagy, with directed edges encoding causal dependencies (Figure 2). This clarifies why single-target geroprotectors have limited restorative power: if unmatched nodes are distributed across hallmark modules, one driver is insufficient, so effective late-life rejuvenation will usually require multiple independent intervention vector fields. Our framework extends this program in four ways, detailed in Supplementary Information 3: we ask whether control is achievable at acceptable biological cost rather than only in principle; we include safety constraints through the viability-constrained reachable set $\dot{x} = Ax + Bu$

$$\mathcal{R}_{\text{safe}}(x_0, T) = \left\{ \begin{array}{l} x_T : \exists u(\cdot) \in \mathcal{U}_{\text{safe}}, x(0) = x_0, x(T) = x_T, \\ x(t) \in \mathcal{V} \forall t \in [0, T] \end{array} \right\};$$

we treat aging as time-varying, since the same drug field can have different effects across ages; and we replace binary controllability with the continuous value function. A plausible minimal intervention basis includes anti-senescent, proteostasis/nutrient-sensing, epigenetic-reprogramming, anti-inflammatory, mitochondrial/metabolic, and matrix/fibrosis fields, though whether it suffices is an empirical network-identifiability problem.

Aging is not composed solely of reversible deviations. We decompose the state into reversible and irreversible coordinates, $x(t) = [x_R(t), x_I(t)]$, where many components of x_I (somatic driver mutations, mitochondrial DNA deletions, neuronal, nephron, and oocyte loss, severe matrix crosslinking, architectural fibrosis, clonal expansion, thymic stromal collapse) are monotonic or nearly monotonic under the admissible control set, reflecting the clinically available controls rather than physical impossibility. An *absorbing boundary* is a region beyond which return to $\mathcal{V}$ is impossible under the specified library and safety constraints: with $\mathcal{R}(x_0, T, \mathcal{U}_{\text{safe}})$ the safe-reachable set, x_0 is controllable to viability when

$\mathcal{R}(x_0, T, \mathcal{U}_{\text{safe}}) \cap \mathcal{V} \neq \emptyset$, and an absorbing boundary is reached when this intersection is empty. Absorbing boundaries imply *biological deadlines*, time points after which the optimal policy changes from restoration to compensation, replacement, or palliation,

$$t^* = \sup \{ t : \Pr[\mathcal{R}(x(t), T, \mathcal{U}_{\text{safe}}) \cap \mathcal{V} \neq \emptyset] \geq 1 - \alpha \}.$$

This matters for preventive geroscience: late rescue is biologically more difficult and clinically more dangerous than early stabilization. When a boundary is crossed, the admissible set can sometimes be expanded by prosthetic controls (insulin pumps, L-DOPA, dialysis, cochlear implants) that compensate for an irreversible coordinate without restoring its biological function, yielding a therapeutic hierarchy: prevent boundary crossing, restore biologically, compensate prosthetically, or manage decline. Extended material on the rare-event (Luria and Delbrück) statistics of irreversible coordinates [62, 63], phase transitions and tipping points, propagation of irreversible losses, and a worked hippocampal example appears in Supplementary Information 3.

Drugs as vector fields

The drug-as-vector-field concept

Drug discovery often describes compounds by their primary target, compressing a high-dimensional intervention into a scalar mechanism. In the control-theoretic framework, a drug induces a vector field on biological state space whose effect depends on current state, dose, duration, tissue context, pharmacokinetics, pharmacodynamics, and endogenous feedback. This provides a natural language for responder heterogeneity (two individuals of the same chronological age but different may experience opposite effects) and for dose-response, since optimal dosing is selection of that minimizes the value function under constraints rather than maximization of target engagement. The drug-as-vector-field is the mean-field projection of the more general drug-as-distribution operator used in single-cell perturbation modeling [64]; the two agree under weak heterogeneity and diverge when subpopulation selection or tail outcomes dominate, in which case the distributional lift is required (Supplementary Information 4). $g_f(x)xu(t)\mu_{t+\Delta t} = T_u\mu_t$.

The Thirteen-Hallmark Aging Network: A Driver/Follower Control Structure




Figure 2. Thirteen-hallmark network of aging with driver and follower classification. Drivers (cellular senescence, loss of proteostasis, deregulated nutrient sensing, epigenetic alterations, chronic inflammation, mitochondrial dysfunction, ECM stiffening and fibrosis) have clinically tractable actuating modality classes; followers are primarily downstream. Edges encode known causal and regulatory dependencies.

Non-commutativity: why order matters

Control theory predicts that sequence can matter independently of the components. Two vector fields commute only if their Lie bracket is zero,

$$[g_A, g_B](x) = \frac{\partial g_B}{\partial x}(x) g_A(x) - \frac{\partial g_A}{\partial x}(x) g_B(x).$$

When the order changes the trajectory: a senolytic may remove cells that constrain tissue remodeling and thereby change the response to reprogramming, whereas reprogramming before senolysis may act on a tissue still dominated by SASP, fibrosis, or immune dysfunction. The bracket has a precise operational meaning: applying, then, then, then, each for time, returns the system to its start to leading order, with a second-order residual displacement of exactly. The bracket is therefore the new direction in state space exposed by exploiting noncommutativity, a degree of freedom inaccessible to either intervention alone. In the geometric control language of Nijmeijer and Schaft [19] and Isidori [20], the controllability Lie algebra at determines the dimension of the locally accessible set; interventions that fail to span this algebra cannot reach the youthful viability set even at unbounded duration and dose. The expected improvement from sequencing is quantified by projected onto the gradient of. These sequence effects are not generally predicted by Hallmark annotation, which can suggest combining senolysis and reprogramming but does not define when one should precede the other. 5 carries out the explicit analytic Lie-bracket calculation for senolysis and reprogramming.

$$[g_A, g_B] \neq 0 \rightarrow A \neq B \rightarrow Ag_A g_B - g_A - g_B \varepsilon \varepsilon^2$$

$$[g_A, g_B](x) + O(\varepsilon^3)x \parallel [g_A, g_B] \parallel V 5$$

Experimental test

For a pair of interventions A and B, estimate g_A , g_B , and their Lie bracket in a relevant aged tissue model. Predict whether $A \rightarrow B$, $B \rightarrow A$, simultaneous treatment, or monotherapy best reduces V. Test predefined functional endpoints, safety markers, and durability. If strong predicted non-commutativity repeatedly fails to correspond to sequence-dependent outcomes, the model is wrong.

Controllability loss, modality classes, and safety as a first-class object

Aging reduces controllability in several ways: intervention susceptibility declines (stem-cell exhaustion,

vascular rarefaction, immune exhaustion); structural variables make the target set harder to reach (fibrosis, crosslinking, neuronal and nephron loss); safety constraints tighten with frailty, shrinking the admissible set; and stochasticity rises. Many interventions fail because they are applied after a controllability threshold has been crossed. The control input $u_j(t)$ is not restricted to small molecules: modern modalities (biologics, mRNA and lipid nanoparticles, viral gene delivery, genome and epigenome editing, engineered immune cells, stem-cell products, synthetic circuits) differ fundamentally in reversibility, latency, durability, targeting, titratability, immunogenicity, and monitoring. Each admissible control should be annotated as

$$u_j = (m_j, d_j, \tau_j, \rho_j, \eta_j, \sigma_j),$$

with modality class, dose, timing, reversibility, tissue-targeting profile, and uncertainty (Table 1). A practical policy follows a *reversibility hierarchy*: in early or uncertain states, reversible interventions are preferred to perturb, observe, and update the state estimate; only after the trajectory is validated should the controller escalate to durable modalities. The escalation rule makes a permanent control u_{perm} admissible only if

$$\Pr[\Delta V(x; u_{\text{perm}}) > \Delta V(x; u_{\text{rev}}) + \varepsilon] > \gamma$$

$$\text{and } \Pr[x(t) \text{ forbidden} | u_{\text{perm}}] < \alpha$$

over a clinically meaningful horizon, preventing premature use of irreversible modalities under model uncertainty (Table 1).

Safety is not a secondary penalty appended after efficacy; it is the central constraint that defines the admissible control set. We define a forbidden region of unacceptable risk (excessive clonal hematopoiesis, loss of epithelial identity, oncogenic proliferation signatures, severe immunosuppression, cytokine-storm risk, impaired wound repair, fibrosis beyond threshold) and require an admissible policy to keep the probability of entering it below a tolerated level α over the horizon, with α much smaller for preventive aging than for life-threatening disease. *Hard constraints* (predicted malignant transformation, expansion of high-risk clones, severe immune collapse, persistent identity loss, organ toxicity above stopping thresholds, germline exposure) exclude a control regardless of predicted benefit; *soft constraints* (mild infection risk, transient wound-healing impairment, reversible cytopenias) enter as context-dependent penalties. The admissible set is state-dependent: the same intervention can be admissible in one patient and forbidden in another. To prevent *reward hacking*, where measured function

Table 1. Modality classes and control properties.

Modality	Reversibility	Onset	Duration	Principal safety concerns	Control implication
Small molecules	High	Hours-weeks	Short-medium	Off-target toxicity, interactions, infection, metabolic	Iterative MPC, dose titration, sequence exploration
Biologics / antibodies	Moderate	Days-weeks	Weeks-months	Immunogenicity, cytokine effects, target depletion	Sustained but adjustable; requires washout modeling
mRNA / LNP	Moderate-high	Hours-days	Days-weeks	Innate immune activation, reactogenicity, tissue accumulation	Pulsed expression; reversible program testing
AAV gene delivery	Low	Weeks	Months-years	Immunogenicity, insertional risk, no redose, persistent expression	High-confidence state selection, long-horizon safety, stopping rules
Base / prime editing	Very low	Days-weeks	Permanent	Off-target, clonal expansion, p53 selection, germline risk	Reserve for irreversible lesions with strong causal evidence
Epigenome editing	Variable	Days-weeks	Weeks-years	Identity loss, dedifferentiation, lineage mis-programming	Identity-preservation constraints, longitudinal monitoring
Engineered immune cells	Low-moderate	Days-weeks	Months-years	CRS, off-target killing, exhaustion, autoimmunity, transformation	Population dynamics, kill switches, antigen escape
Senolytic cell therapy	Low-moderate	Days-weeks	Durable depletion	Loss of beneficial senescent cells, impaired repair, inflammation	Use when senescent burden causal and localized
iPSC-derived cell therapy	Low	Weeks-months	Years	Teratoma, ectopic tissue, rejection, mis-integration	Replacement for crossed boundaries; release + surveillance
Synthetic circuits	Variable	Hours-weeks	Programmable	Circuit escape, mutation, immune recognition	Enables closed-loop control with fail-safes

improves while latent damage accumulates, the terminal cost includes future controllability,

$$\Phi(x(T)) = \Phi_{\text{function}}(x(T)) + \beta_1 D_{\text{irreversible}} + \beta_2 R_{\text{cancer}} + \beta_3 R_{\text{immune}} + \beta_4 C_{\text{future}}^{-1},$$

so that an intervention is successful only if it preserves or expands the future safe reachable set. The relationship to classical PK/PD and systems pharmacology [21, 22, 65–67], and the full treatment of modality reversibility, hard and soft constraints, and reward-hacking prevention, are developed in Supplementary Information 4.

Worked computational example: aged murine liver

To make the framework concrete we use a deliberately low-dimensional, fully specified model of aged murine liver. The purpose is not to claim that five variables capture hepatic aging, but to show how hallmarks, interventions, safety constraints, non-commutativity, and value-function scoring can be placed into a single computable system simple enough to be reproduced, falsified, and extended. Let collect senescent-cell fraction, oxidative damage, epigenetic information integrity, regenerative capacity, and fibrotic burden, initialized at an aged state. The controlled stochastic dynamics, literature-calibrated illustrative parameterization [10, 33, 38, 55, 68, 69], the viability set, and the full equations are specified in Supplementary Information 5. Controls are (senolytic), (reprogramming), (rapamycin), and (TNIK inhibitor),

with senolytic courses and reprogramming pulses normalized to unit integral.

$$x(t) = (s, d, e, r, f) \in [0, 1]^5 \text{ sderfx}_0 \\ = (0.15, 0.40, 0.55, 0.40, 0.30) u_{\text{sen}} u_{\text{rep}} u_{\text{rap}} u_{\text{tnik}}$$

We compare three fixed protocols over $T = 56$ days: Protocol A applies a senolytic course first (days 0 to 7) then reprogramming pulses (days 14 to 42); Protocol B reverses the order; Protocol C administers them simultaneously. Senolysis before reprogramming produces greater restoration of epigenetic integrity ($e_A(T) = 0.72$, $e_B(T) = 0.61$, $e_C(T) = 0.66$ on the mean trajectory; stochastic integration over $n = 200$ runs preserves the ranking, Supplementary Information 5). The reason is mechanistic: reprogramming efficacy scales with $(1-s)r$, so clearing senescent cells first increases the effective gain of the reprogramming pulse (Figure 3). Under a finite-horizon cost with terminal penalty

$$\Phi(x(T)) = w_s s(T)^2 + w_d d(T)^2 + w_e (1 - e(T))^2 \\ + w_r (1 - r(T))^2 + w_f f(T)^2,$$

and the weights specified in Supplementary Information 5, the protocol costs are $J_A = 0.184$, $J_B = 0.257$, $J_C = 0.231$, so that $\hat{V}(x_0, T) = \min\{J_A, J_B, J_C\} = J_A$. This is a finite-protocol approximation rather than the exact Hamilton–Jacobi–Bellman value function, demonstrating how candidate schedules can be scored by the same value-function logic. Loss of controllability

with intervention start age, including a biological deadline beyond which no admissible protocol returns the state to viability, is shown in Figure 4.

The ordering effect is provable analytically. With state ordering $x = (s, d, e, r, f)$, the senolytic field is $g_{\text{sen}}(x) = (-\beta_s s, 0, 0, 0, 0)^\top$ and the reprogramming field is $g_{\text{rep}}(x) = (0, 0, \mu_e(1-s)r, 0, 0)^\top$. Because g_{rep} depends on s while g_{sen} does not, the Lie bracket is

$$[g_{\text{sen}}, g_{\text{rep}}](x) = (0, 0, \beta_s \mu_e s r, 0, 0)^\top \neq 0$$

whenever $s > 0, r > 0$.

At the aged initial state, $\beta_s \mu_e s_0 r_0 = (0.800)(0.150)(0.15)(0.40) = 0.0072$. The bracket is

positive in the epigenetic-integrity coordinate because senolysis increases the effective gain of reprogramming by reducing the inhibitory factor s , which is the mathematical expression of the ordering effect observed in simulation. The full derivation appears in Supplementary Information 5.

Falsifiability

Non-commutativity is not merely a metaphor for “sequence matters.” It is a computable property of controlled biological vector fields. If empirical measurements showed that g_{sen} were independent of s , or that senolysis did not alter the state variable suppressing reprogramming efficacy, the Lie bracket would vanish and the ordering prediction would fail. The claim is therefore falsifiable.




Figure 3. (A) Worked-example terminal-cost trajectory for the intervention protocols in the five-variable aged murine liver model $X = (s, d, e, r, f)$ over the $T = 56$ -day horizon ($n = 200$ stochastic realizations per protocol; Euler–Maruyama SDE integration at $\Delta t = 0.25$ days). Solid curve: no intervention. Dashed: Protocol A (senolytic → reprogramming). Dash-dot: Protocol B (reprogramming → senolytic). Dotted: Protocol C (simultaneous senolytic + reprogramming). Shaded bands are 95% bootstrap confidence intervals. (B) Intervention schedules for the three protocols (dark bars: senolytic course; light bars: reprogramming pulses). Full equations, parameters, and protocol schedules are specified in Supplementary Information 5.

Reprogramming is not a purely beneficial epigenetic field. It also raises an oncogenic-risk and dedifferentiation coordinate whose contribution to the safety cost is state-dependent, larger when senescent or damage burden is high, so that partial reprogramming can reverse sign from beneficial to harmful as baseline damage increases [70, 71]. The Supplement instantiates this negative coordinate explicitly in the worked-example dynamics and cost (Supplementary Information 5). The same Supplement develops an optional state-space extension in which epigenetic erosion derepresses LINE1 retroelements, activating cGAS-STING and type-I interferon signaling in a positive-feedback loop with senescence and SASP, and adds nucleoside reverse-transcriptase inhibitor and SIRT6-activator intervention vectors that yield a natural three-intervention ordering problem [72–75]; this extension is optional and is not required for the results reported here.

The vector-field statement is useful only if the field is empirically calibrated. For intervention, is estimated by combining perturbational datasets across scales (LINCS L1000 and CMap [21, 22], Perturb-seq and pooled CRISPR screens, organoid panels, animal intervention studies [76], and clinical multi-omics [77]) using neural ODEs [78], Gaussian-process state-space models, or Koopman operators with control, together with conditional perturbation autoencoders [64]. Every estimate must carry uncertainty, and a coordinate enters as causal only if perturbing it changes future trajectories as predicted. The full data-source taxonomy, training objective, uncertainty quantification, validation criterion, and an AAV-FOXN1 thymic-regeneration safety example are given in Supplementary Information 5.

$$jg_j(x) = \mathbb{E}[(x(t + \Delta t) - x(t)) / \Delta t \mid x(t) = x, u_j].$$




Figure 4. Conceptual illustration of loss of controllability with intervention start age. (A) Illustrative reachable safe-set fraction. (B) Illustrative restored-function success fraction under a stricter threshold. The curves and shaded envelopes are schematic and are not estimates or confidence intervals from the five-variable liver simulation.

Implementation architecture

The framework can be implemented as a closed-loop computational system that estimates a latent biological state, learns state-dependent intervention vector fields, computes safe reachable sets, selects interventions under uncertainty, and updates the model after repeated measurement. The pipeline consists of six layers (Figure 5): a multimodal measurement layer (plasma proteomics, DNA methylation, untargeted metabolomics, immune profiling, single-cell omics, functional tests, imaging, and clinical chemistry, with rigorous preprocessing and uncertainty propagation); a latent-state estimator implemented as a multimodal VAE, neural ODE [78], Koopman operator, or Gaussian-process state-space model, biologically anchored and trained to predict future trajectories rather than only age; a perturbation-response learner that estimates conditional vector fields; a modality-aware safety model that estimates the probability of entering forbidden regions over short and long horizons; a safety-constrained model-predictive controller.

$$z(t) = E(y(t), m(t), h(t)) G_j(z, d, c)$$

$$\pi^* = \arg \min_{\pi} \mathbb{E} \left[\int_t^{t+T} L(z(s), u(s), s) ds + \Phi(z(t+T)) \right],$$

subject to $u(s) \in \mathcal{U}_{\text{safe}}(z(s), h)$ and $\Pr[z(s) \text{ forbidden}] < \alpha$, which applies only the first action before remeasurement and may recommend no treatment; and a *feedback loop* that updates the latent state via Bayesian filtering after quarterly lightweight and annual deep phenotyping. Robust MPC (conditional value at risk, worst-case objectives) is used when model uncertainty is high, and the problem becomes a POMDP over belief states under partial observation. Full layer-by-layer specifications and a two-year minimal viable

implementation focused on immune aging appear in Supplementary Information 6.

Translational case studies

The framework becomes useful only when instantiated in specific tissues with measurable states, admissible interventions, safety constraints, and go/no-go criteria. Three case studies, developed in full in Supplementary Information 7, illustrate this. In aged thymus and immune rejuvenation, the latent state collects thymic volume, epithelial and stromal integrity, adipose replacement, naïve T-cell output, repertoire diversity, inflammatory tone, and clonal-expansion risk; a staged policy begins with reversible metabolic and inflammatory optimization, proceeds to transient thymopoietic stimulation (IL-7 axis, GH secretagogues with metformin co-administration), and considers durable FOXN1 programs only at a higher evidence threshold, with success defined by improved naïve output, repertoire diversity, and vaccine response without autoimmunity or clonal expansion. In sarcopenia, the state separates mass from quality, mitochondrial function, neuromuscular-junction integrity, fibrosis, adiposity, and functional performance, so that the controller penalizes mass gain without strength, power, or metabolic improvement; reversible interventions (resistance exercise, protein and amino-acid optimization, urolithin A and mitophagy-targeting agents [79], rapamycin schedules) precede biologics (myostatin/activin inhibitors) and durable follistatin gene therapy. In ovarian aging, follicular reserve is a dominant irreversible depletion coordinate (under the standard small-molecule set), making the tissue a stringent test of biological deadlines, with the control goal being maintenance of fertility potential over a defined horizon under strict state-dependent admissibility around conception and pregnancy. Each case specifies measurement layers, candidate interventions, safety




Figure 5. Closed-loop implementation architecture: measurement → latent-state estimation → perturbation-response and safety models → safety-constrained MPC → first action applied to the patient or model → remeasurement. The safety model gates the admissible set at every optimization step.

constraints, and go/no-go criteria, and each test whether control-value reduction predicts functional improvement better than Hallmark labels or clock reversal alone.

Validation and power analysis

A control-theoretic framework should be judged by whether it improves drug-discovery decisions, and AI-driven discovery pipelines [23, 24, 80] provide a natural validation environment. The identification of dual-purpose targets implicated in both aging hallmarks and age-associated diseases has been demonstrated on AI platforms [81], and programs since 2021 against six aging-associated targets (TNIK, PHD1/2, QPCTL, NLRP3, NR3C1, TEAD) constitute a retrospective testbed in which each program is a decision sequence amenable to comparison of target-scoring frameworks. For each candidate one can compute a Hallmark score, age-association score, disease-association score, network-centrality score, druggability score, toxicity-risk score, and control-value score (estimated reduction in restoration cost). The central hypothesis is that control-value score outperforms Hallmark score for predicting translational advancement, functional efficacy, and toxicity-adjusted therapeutic index. An internal pilot is statistically limited: detecting a Spearman correlation near 0.45 with 80% power at two-sided $\alpha = 0.05$ requires roughly 36 to 40 compounds, and advancement is influenced by potency, pharmacokinetics, intellectual property, and strategic decisions beyond biological efficacy. A stronger program expands to public corpora (DrugAge, the Interventions Testing Program, known geroprotectors, and matched negative controls), is pre-registered with frozen compound list, model version, intervention library, cost function, and endpoints, and prefers prospective scoring in aged organoid panels. The full design, the four-quadrant Hallmark-versus-control-value prediction, the pre-registration template, and explicit failure criteria appear in Supplementary Information 8. The framework fails if control-value scoring does not outperform Hallmark or network scoring, if state-dependent predictions do not improve outcomes, or if non-commutativity, irreversibility, or safety-constraint predictions fail when tested; none of these failure modes, and none of the corresponding successes, has yet been demonstrated in pre-registered $V = 0.05$, independent, adequately powered studies.

Twenty falsifiable predictions

The following twenty predictions distinguish the control-theoretic framework from Hallmark enumeration, biomarker optimization, and single-mechanism theories. Each is a testable, falsifiable

hypothesis; full statements, tests, and falsification conditions are given in Supplementary Information 9. They group into state-dependent efficacy (1 through 4), order and combination (5 through 8), controllability and irreversibility (9 through 12), target prioritization (13 through 16), and systems-level behavior (17 through 20).

1. Sign and magnitude of effect depend on baseline state beyond chronological age.
2. Non-responders to rapamycin, senolytics, or partial reprogramming are identifiable from pre-treatment latent state.
3. At least one major intervention has a measurable state region where it is harmful despite being beneficial on average.
4. Optimal dose depends on state, not a single target-engagement curve.
5. In fibrotic/inflamed aged tissues, senolysis before reprogramming restores function better than the reverse order.
6. In fibrotic tissues, antifibrotic remodeling before regenerative stimulation outperforms the reverse order.
7. The magnitude of empirical non-commutativity correlates with the estimated Lie bracket of intervention vector fields.
8. Rationally chosen low-dose combinations match or exceed high-dose monotherapy at lower toxicity.
9. Structurally aged tissues have measurable controllability boundaries beyond which no admissible small molecule restores function even if clocks improve.
10. Recovery dynamics after standardized perturbation predict future decline better than static biomarkers.
11. Ovarian reserve and hippocampal neuronal loss have identifiable deadlines beyond which standard interventions cannot restore function.
12. State-space variance and autocorrelation increase measurably before clinical frailty transitions.
13. Some high-control-value targets lie outside canonical Hallmark annotations.
14. Control-value reduction predicts advancement better than Hallmark membership alone.
15. High network-centrality targets without high control-value reduction underperform in translation.
16. Modality-matched targeting outperforms target-only selection.
17. Closed-loop adaptive MPC outperforms fixed-schedule protocols at matched cumulative exposure.
18. Reversibility-hierarchy escalation reduces serious adverse events without sacrificing efficacy.
19. Controllability structure differs between sexes (ovarian dominance, distinct immune kinetics, different mTOR/IGF-1 sensitivities).

20. Interventions targeting shared aging mechanisms reduce multi-system disease incidence more efficiently than single-disease drugs at matched cost and safety.

Each prediction derives from a specific formal object of the framework rather than from biological intuition alone. Table 2 makes this mapping explicit for the five predictions that are both riskiest and most directly tied to the formalism (Table 2); the full mapping for all twenty predictions appears in Supplementary Information 9. A prediction is riskiest when it is independently testable and is not a built-in modeling assumption, so that its failure would refute the corresponding formal object.

Empirical application: scoring interventions across biological epochs

To show that the framework generates actionable, quantitative predictions, we applied an epoch-stratified scoring rubric to a curated set of interventions drawn from DrugAge (3,423 entries), the NIA Interventions Testing Program (63 entries across ~50 compounds), ClinicalTrials.gov (638 aging/longevity trials), and OpenTargets (2,559 aging-associated targets). Each intervention was mapped onto a 20-dimensional biological state space (mTOR, AMPK, sirtuin, inflammatory tone, senescent burden, epigenetic integrity, proteostasis, mitochondrial function, NAD⁺, stem-cell reserve, telomere length, immune function, vascular integrity, fibrosis, metabolic flexibility, hormonal regulation, circadian amplitude, microbiome diversity, neuroplasticity, genomic stability) with a vector field $g_i \in \mathbb{R}^{20}$, and scored on eight axes (control leverage, safety margin, controllability expansion, durability, reversibility, combinability, evidence tier, sequence position) at each of five biological epochs. The composite score at epoch k is

$$\text{Composite}_k = 100 \times \sum_{i=1}^8 w_{k,i} \cdot \text{Axis}_i \text{ with epoch-specific}$$

weights, and the Lifetime Integrated Score (LIS) is $\sum_{k=1}^5 \text{Composite}_k$. The full rubric, axis definitions, epoch weights, and the complete scored dataset (Supplementary Table 1) are given in Supplementary Information 10; Table 3 reports the top twenty interventions by LIS (Table 3).

Several observations emerge. Lifestyle interventions dominate the top ranks: caloric restriction (LIS 320.2) and aerobic exercise (316.7) outperform all pharmacological agents through broad vector-field coverage, high safety margins, and Tier 1 evidence.

SGLT2 inhibitors emerge as the highest-ranked pharmacological class (dapagliflozin 309.3, empagliflozin 306.0), outranking rapamycin, senolytics, and NAD⁺ precursors through multi-dimensional effects and Tier 1 cardiovascular-outcome evidence, a prediction testable in head-to-head geroscience trials, with GLP-1 receptor agonists ranking third (liraglutide 298.9, semaglutide 286.6). All scores decline monotonically with epoch, formalizing the intuition that prevention is more efficient than late intervention through narrowing safety envelopes and diminishing substrate. Rapamycin ranks outside the top fifteen despite gold-standard ITP lifespan evidence, because its immunosuppressive safety penalty and narrow vector field reduce its composite score relative to multi-dimensional agents; the framework predicts intermittent low-dose deployment at Epoch II to III rather than chronic administration. A central implication is that the current intervention library is incomplete, providing meaningful control over roughly 20% of the state space ($C_0 \approx 0.20$), which motivates an iterative control cycle mapping control gaps, directing discovery toward uncontrollable dimensions, and re-estimating controllability with the expanded library, so that the framework is intended to operate as an iterative pipeline for rational gerotherapeutic prioritization rather than a one-shot solution; the automated learning of vector fields that would close this loop is future and companion work.

Relationship to existing frameworks and formal methods

The control-theoretic framework complements, rather than replaces, existing aging theories; each supplies content the control layer depends on. Hallmark annotations map onto state coordinates; damage inventories inform terminal cost; SENS repair categories map onto intervention vector fields; information-loss theories motivate specific reversible coordinates; geroscience provides clinical context; hyperfunction theory supplies a key vector field whose state-dependent sign is made explicit.

The framework also stands in a specific relation to the resilience and critical-transition literature. Work on early-warning signals for critical transitions [44, 45] and on universal resilience patterns in complex networks [46] shows that systems approaching a tipping point exhibit critical slowing down, rising variance, and reduced recovery rate before overt collapse. These are descriptive, intervention-free diagnostics: they detect that a system is losing dynamical resilience but do not specify which input restores it. The control-cost formulation adds a prescriptive layer on top of this descriptive picture. Loss of resilience appears here as declining conditioning of the controllability Gramian

Table 2. Mapping of the five riskiest predictions to the formal object each derives from.

Pred.	Statement (abbreviated)	Formal object
3	The same intervention is beneficial in some states and harmful in others	Safety-cost sign: $g_j^\top \nabla V$ changes sign
7	Sequenced interventions with nonzero Lie bracket depend on order	Lie bracket $[g_A, g_B] \neq 0$ (ordering)
9	Each tissue has a loss level beyond which no admissible intervention restores function	Viability-kernel boundary / non-responder state
14	Targets ranked by control value outperform Hallmark-count ranking	Value-function gradient identifies high-leverage node
17	Adaptive model-predictive control outperforms a fixed schedule	closed-loop MPC / partially observed control policy

Failure of the prediction refutes the associated formal claim.

Table 3. Top 20 aging interventions by Lifetime Integrated Score (LIS).

Rank	Intervention	Tier	Peak	Best Ep.	LIS
1	Caloric Restriction (30%)	1	72.1	1	320.2
2	Aerobic Exercise (≥ 150 min/wk)	1	70.9	2	316.7
3	Dapagliflozin	1	69.3	1	309.3
4	Empagliflozin	1	68.2	1	306.0
5	Resistance Exercise (2 to 3 $\times$ /wk)	1	67.6	1	301.5
6	Liraglutide	1	67.2	1	298.9
7	Intermittent Fasting (16:8)	2	69.1	1	298.2
8	Telmisartan	1	65.7	1	289.4
9	Semaglutide	1	63.8	2	286.6
10	Sulforaphane	2	65.2	1	285.6
11	Berberine	2	64.5	1	281.7
12	Urolithin A	2	62.9	1	279.3
13	Losartan	1	63.7	1	278.4
14	Methionine Restriction	3	64.5	1	278.4
15	Time-Restricted Eating	2	63.9	1	277.6
16	Pioglitazone	1	62.6	1	276.2
17	Lithium (low dose)	2	63.0	1	275.4
18	Fenofibrate	1	61.4	1	271.3
19	Pterostilbene	3	61.9	1	268.3
20	Magnesium supplementation	1	60.6	1	266.2

Evidence tier: 1 = ITP/Phase 3 validated; 2 = Phase 1 or 2 / strong preclinical; 3 = preclinical only.

and rising restoration cost, and the same slowing of recovery that serves as an early-warning signal is, in control terms, the quantity to be reversed. Beyond detection, the value function specifies an intervention-relative, safety-constrained objective: which admissible control most reduces from the current state, at what risk, and within which reachable safe set. Critical slowing down thus becomes an input to the controller rather than an endpoint of the analysis IV .

The operational difference from the Hallmarks scheme can be stated concretely. The Hallmarks organize mechanisms into categories and support the inference

that a mechanism is a candidate target, but they do not by construction order interventions, identify who will not respond, or predict when a beneficial intervention becomes harmful. The control formulation yields three operationally distinct outputs. First, explicit intervention ordering: when two intervention vector fields have a nonzero Lie bracket, the framework predicts which sequence reduces more, a statement absent from any hallmark annotation. Second, a predicted non-responder state: an initial state lying outside the viability kernel for a given library cannot be restored by any admissible protocol in that library, identifying non-responders from the pre-treatment state rather than post hoc. Third, a

prospectively testable state-dependent sign reversal: the sign of an intervention is set by the projection of its vector field onto the gradient of, so the framework predicts a state region in which a normally beneficial intervention becomes harmful. These three outputs are the content the hallmark ontology does not supply. The recurring gap the control layer fills is the absence of a state vector, dynamics, ordering, an admissible set, and a continuous safe value function in place of binary controllability or annotation-based ranking; a full cross-framework comparison is tabulated in Supplementary Information 10. The formulation also interfaces with three adjacent formal methods, developed in Supplementary Information 10: hybrid automata [82–84] represent discrete clinical stages (robust, pre-frail, frail, disabled, terminal) with mode-dependent flows and enable model-checking of safety properties; asymmetric signaling games [85] model cell-cell communication under selection, where pooling equilibria underlie inflammaging and immune evasion and interventions act on equilibrium structure; and anti-hallmarks and synthetic lethality [86, 87] predict that high-value targets lie outside canonical Hallmark vocabulary in dependencies the aged state acquired through compensatory remodeling, consistent with the high ranking of SGLT2 inhibitors and GLP-1 receptor agonists above *IV*.

Limitations

The framework has substantial limitations. The latent state is not observable, and inferring it from measurements is an identifiability problem compounded when interventions perturb the measurement mapping itself; it should be addressed through multimodal over-determination, perturbational experiments, biological priors, model comparison on held-out perturbational data, and sensitivity analysis. Intervention data are sparse and biased toward cancer cell lines, young mice, and healthy volunteers, so extrapolation to aged, frail, and comorbid populations is risky and prospective aging-specific datasets are essential. The full stochastic control problem is intractable in closed form, and the linearizations, short-horizon MPC, reinforcement learning, and mechanistic submodels used in practice introduce errors that must be quantified. Long-horizon safety prediction is weak because aging drug programs are young and relevant outcomes have long latencies; the framework keeps safety uncertainty in the optimization and prefers reversible modalities when uncertainty is high. Value-function weights reflect patient preferences, clinical judgment, and societal values and must be elicited rather than derived from biology, and control models must be shared with frozen parameters, code, and pre-specified endpoints to be falsifiable. Finally, the framework is a research

program: it does not authorize any specific clinical protocol or off-label use, which must follow regulatory, ethical, and evidentiary standards independent of the framework. A five-phase scalability roadmap, from single-organ animal and organoid models through human pilots, multi-organ combination validation, durable-modality integration, and population-scale geroscience, with defined deliverables and failure modes at each phase, is given in Supplementary Information 6.

DISCUSSION

The core claims are that aging is progressive loss of safe controllability; biological age is the minimum safe control cost of restoring or maintaining function; drugs are state-dependent vector fields; combinations and sequences matter through Lie-bracket structure; safety is a first-class constraint that defines the admissible control set; and modality class determines the reversibility hierarchy of controls. The distinguishing contribution is not the general idea that drugs perturb biological systems, which is standard systems pharmacology, but the casting of gerotherapeutic discovery as a constrained stochastic control problem with explicit state space, safety constraints, irreversibility coordinates, and an intervention-relative definition of biological age. This yields specific, testable consequences: responder stratification by state, order-dependent combinations, controllability boundaries, modality-matched escalation, adaptive MPC, and safety-constraint modeling with forbidden regions and reward-hacking prevention. It addresses the gap between descriptive biology and translational decision-making by asking, at each step, which admissible intervention most reduces the minimum safe cost of functional restoration given the current state estimate and its uncertainty. Several caveats hold: identifiability, data bias, computational complexity, and weak safety modeling are serious; parameters must be empirically calibrated and predictions pre-registered; and value-function weights cannot be derived from biology alone. The central empirical test is whether control-value reduction predicts translational success better than alternative scoring methods, conductible retrospectively on existing portfolios, prospectively in aged organoid and animal panels, and adaptively in clinical cohorts.

MATERIALS AND METHODS

All simulations in the worked example (Figure 3 and Supplementary Information 5) used the specified stochastic differential equation system. Deterministic mean-trajectory integration used explicit Euler stepping with $\Delta t = 0.25$ days over a 56-day horizon, with state

variables clamped to [0,1] after each step; the stochastic reference integration used the same step size with additive Wiener increments scaled by the specified diagonal noise matrix. Control functions were normalized so that a senolytic or TNIK course and a reprogramming pulse each integrate to unity. The value-function approximation is over three hand-specified policies and does not solve the full Hamilton–Jacobi–Bellman equation. The Lie-bracket computation was performed analytically from the vector-field specifications and independently checked in SymPy. All computations are reproducible from the equations, parameters, and initial conditions provided; the reference implementation and validation scripts are provided as Supplementary Code under the MIT License.

In accordance with the ICMJE Recommendations (January 2024 update) and the ICMJE recommendations on chatbots and generative AI [88], the authors disclose that large language models (Anthropic Claude Opus and Sonnet families; OpenAI GPT-4 and GPT-5 families, 2024 to 2026 releases) were used as drafting and editing tools for prose, copy-editing, reference cross-checking against author-provided sources, and LaTeX assistance, and that retrieval-augmented assistants located candidate citations subsequently verified against the primary literature. AI systems did not originate the scientific framework, did not generate claims or predictions not specified by the human authors, did not produce any citation accepted without human verification, and did not carry out the numerical simulations or analytic derivations. AI systems are not listed as authors. The human authors reviewed, edited, and approved every sentence, equation, figure, table, reference, and numerical claim, and accept full responsibility for the accuracy, originality, and integrity of the work. The manuscript was checked for textual overlap against prior publications and indexed sources; all retained passages are original or properly quoted and cited; the submission contains no fabricated or falsified data and no unverified AI-generated citations. No human or animal experiments were performed for this manuscript; translational case studies outline prospective designs requiring institutional review, informed consent, and regulatory oversight. The authors disclose affiliations with Insilico Medicine and collaborating institutions; specific financial disclosures appear in the journal submission record. The framework is proposed as an open scientific program requiring independent replication beyond any single organization.

CONCLUSIONS

Aging theory has matured from evolutionary reasoning through damage catalogues, Hallmark ontology, information models, geroscience, hyperfunction, and network and resilience frameworks. Each has clarified

part of biology; none, by itself, specifies how to act in a measured biological state. We have proposed a control-theoretic framework designed to supply that layer: a latent biological state, a functional viability set, stochastic aging dynamics, intervention vector fields, a minimum-safe-cost definition of biological age, explicit irreversibility coordinates and absorbing boundaries, modality-aware admissible sets, and safety as a first-class constraint. A worked aged-liver example with an analytic Lie bracket and protocol scoring is explicit enough to be reproduced and falsified; an implementation architecture covers measurement, latent-state estimation, perturbation-response learning, safety modeling, constrained model-predictive control, and closed-loop updating; three translational case studies specify state vectors, intervention libraries, constraints, and go/no-go criteria; a validation plan with power analysis and external replication is provided; and twenty pre-specified, falsifiable predictions are enumerated. The framework will succeed or fail empirically. If control-value scoring does not outperform Hallmark, clock, or network baselines, and if order-dependence, controllability boundaries, modality effects, and adaptive-control advantages fail to appear, it will be rejected. The framework supplies aging drug discovery with equations of motion, admissible sets, value functions, reversibility hierarchies, and falsifiable predictions, and its claims are stated so that adequately powered, pre-specified studies can confirm or refute them. The present paper establishes the framework and its predictions; the deeper mathematical and engineering development, including the full geometric-control treatment of noncommutativity and reachability, formal identifiability and estimation of the vector fields, and hybrid-automaton and stochastic-transition models of irreversibility, is reserved for a companion sequel.

ACKNOWLEDGMENTS

The authors thank Charles Cantor for his input and support. The authors thank colleagues at Insilico Medicine and collaborating institutions for valuable discussions.

AUTHOR CONTRIBUTIONS

AZ conceived the idea and proposed the theory and wrote the paper. BM reviewed and contributed to the math, control theory fundamentals and revised the paper.

CONFLICTS OF INTEREST

BM is the advisor and AZ is the founder and CEO of Insilico Medicine, a leading generative AI clinical-stage biotechnology company with a broad clinical and preclinical pipeline of therapeutic assets targeting aging and disease.

FUNDING

No funding was provided for this study.

REFERENCES

1. Gladyshev VN, Anderson B, Barlit H, Barré B, Beck S, Behrouz B, Belsky DW, Chaix A, Chamoli M, Chen BH, Cheng K, Chuprin J, Churchill GA, et al. Disagreement on foundational principles of biological aging. *PNAS Nexus*. 2024; 3:pgae499.
<https://doi.org/10.1093/pnasnexus/pgae499>
PMID:[39660064](https://pubmed.ncbi.nlm.nih.gov/39660064/)
2. Medawar PB. An unsolved problem of biology. London: H. K. Lewis. 1952.
3. Williams GC. Pleiotropy, natural selection, and the evolution of senescence. *Evolution*. 1957; 11:398–411.
<https://doi.org/10.2307/2406060>
4. Hamilton WD. The moulding of senescence by natural selection. *J Theor Biol*. 1966; 12:12–45.
[https://doi.org/10.1016/0022-5193\(66\)90184-6](https://doi.org/10.1016/0022-5193(66)90184-6)
PMID:[6015424](https://pubmed.ncbi.nlm.nih.gov/6015424/)
5. Harman D. Aging: a theory based on free radical and radiation chemistry. *J Gerontol*. 1956; 11:298–300.
<https://doi.org/10.1093/geronj/11.3.298>
PMID:[13332224](https://pubmed.ncbi.nlm.nih.gov/13332224/)
6. Kirkwood TB. Evolution of ageing. *Nature*. 1977; 270:301–4.
<https://doi.org/10.1038/270301a0> PMID:[593350](https://pubmed.ncbi.nlm.nih.gov/593350/)
7. López-Otín C, Blasco MA, Partridge L, Serrano M, Kroemer G. The hallmarks of aging. *Cell*. 2013; 153:1194–217.
<https://doi.org/10.1016/j.cell.2013.05.039>
PMID:[23746838](https://pubmed.ncbi.nlm.nih.gov/23746838/)
8. López-Otín C, Blasco MA, Partridge L, Serrano M, Kroemer G. Hallmarks of aging: An expanding universe. *Cell*. 2023; 186:243–78.
<https://doi.org/10.1016/j.cell.2022.11.001>
PMID:[36599349](https://pubmed.ncbi.nlm.nih.gov/36599349/)
9. de Grey AD, Ames BN, Andersen JK, Bartke A, Campisi J, Heward CB, McCarter RJ, Stock G. Time to talk SENS: critiquing the immutability of human aging. *Ann N Y Acad Sci*. 2002; 959:452–62.
<https://doi.org/10.1111/j.1749-6632.2002.tb02115.x>
PMID:[11976218](https://pubmed.ncbi.nlm.nih.gov/11976218/)
10. Lu Y, Brommer B, Tian X, Krishnan A, Meer M, Wang C, Vera DL, Zeng Q, Yu D, Bonkowski MS, Yang JH, Zhou S, Hoffmann EM, et al. Reprogramming to recover youthful epigenetic information and restore vision. *Nature*. 2020; 588:124–9.
<https://doi.org/10.1038/s41586-020-2975-4>
PMID:[33268865](https://pubmed.ncbi.nlm.nih.gov/33268865/)
11. Yang JH, Hayano M, Griffin PT, Amorim JA, Bonkowski MS, Apostolides JK, Salfati EL, Blanchette M, Munding EM, Bhakta M, Chew YC, Guo W, Yang X, et al. Loss of epigenetic information as a cause of mammalian aging. *Cell*. 2023; 186:305–26.e27.
<https://doi.org/10.1016/j.cell.2022.12.027>
PMID:[36638792](https://pubmed.ncbi.nlm.nih.gov/36638792/)
12. Kennedy BK, Berger SL, Brunet A, Campisi J, Cuervo AM, Epel ES, Franceschi C, Lithgow GJ, Morimoto RI, Pessin JE, Rando TA, Richardson A, Schadt EE, et al. Geroscience: linking aging to chronic disease. *Cell*. 2014; 159:709–13.
<https://doi.org/10.1016/j.cell.2014.10.039>
PMID:[25417146](https://pubmed.ncbi.nlm.nih.gov/25417146/)
13. Blagosklonny MV. Aging and immortality: quasi-programmed senescence and its pharmacologic inhibition. *Cell Cycle*. 2006; 5:2087–102.
<https://doi.org/10.4161/cc.5.18.3288>
PMID:[17012837](https://pubmed.ncbi.nlm.nih.gov/17012837/)
14. Bellman R. Dynamic programming. Princeton: Princeton University Press. 1957.
<https://doi.org/10.1515/9781400835386>
15. Kalman RE. A new approach to linear filtering and prediction problems. *J Basic Eng*. 1960;82:35–45.
<https://doi.org/10.1115/1.3662552>
16. Kirk DE. Optimal control theory: an introduction. Englewood Cliffs (NJ): Prentice-Hall. 1970.
17. Sontag ED. Mathematical control theory: deterministic finite-dimensional systems. New York: Springer. 1998.
<https://doi.org/10.1007/978-1-4612-0577-7>
18. Kalman RE. Mathematical description of linear dynamical systems. *J Soc Ind Appl Math Ser A Control*. 1963; 1:152–92.
<https://doi.org/10.1137/0301010>
19. Nijmeijer H, Schaft A. Nonlinear dynamical control systems. New York: Springer. 1990.
<https://doi.org/10.1007/978-1-4757-2101-0>
20. Isidori A. Nonlinear control systems. 3rd ed. New York: Springer. 1995.
<https://doi.org/10.1007/978-1-84628-615-5>
21. Lamb J, Crawford ED, Peck D, Modell JW, Blat IC, Wrobel MJ, Lerner J, Brunet JP, Subramanian A, Ross KN, Reich M, Hieronymus H, Wei G, et al. The Connectivity Map: using gene-expression signatures to connect small molecules, genes, and disease. *Science*. 2006; 313:1929–35.
<https://doi.org/10.1126/science.1132939>
PMID:[17008526](https://pubmed.ncbi.nlm.nih.gov/17008526/)
22. Subramanian A, Narayan R, Corsello SM, Peck DD, Natoli TE, Lu X, Gould J, Davis JF, Tubelli AA, Asiedu JK, Lahr DL, Hirschman JE, Liu Z, et al. A Next Generation

- Connectivity Map: L1000 Platform and the First 1,000,000 Profiles. *Cell*. 2017; 171:1437–52.e17.
<https://doi.org/10.1016/j.cell.2017.10.049>
PMID:29195078
23. Zhavoronkov A, Ivanenkov YA, Aliper A, Veselov MS, Aladinskiy VA, Aladinskaya AV, Terentiev VA, Polykovskiy DA, Kuznetsov MD, Asadulaev A, Volkov Y, Zholus A, Shayakhmetov RR, et al. Deep learning enables rapid identification of potent DDR1 kinase inhibitors. *Nat Biotechnol*. 2019; 37:1038–40.
<https://doi.org/10.1038/s41587-019-0224-x>
PMID:31477924
 24. Zhavoronkov A, Mamoshina P, Vanhaelen Q, Scheibye-Knudsen M, Moskalev A, Aliper A. Artificial intelligence for aging and longevity research: Recent advances and perspectives. *Ageing Res Rev*. 2019; 49:49–66.
<https://doi.org/10.1016/j.arr.2018.11.003>
PMID:30472217
 25. Zhavoronkov A, Mamoshina P. Deep Aging Clocks: The Emergence of AI-Based Biomarkers of Aging and Longevity. *Trends Pharmacol Sci*. 2019; 40:546–9.
<https://doi.org/10.1016/j.tips.2019.05.004>
PMID:31279569
 26. Zhavoronkov A, Buzdin AA, Garazha AV, Borisov NM, Moskalev AA. Signaling pathway cloud regulation for *in silico* screening and ranking of the potential geroprotective drugs. *Front Genet*. 2014; 5:49.
<https://doi.org/10.3389/fgene.2014.00049>
PMID:24624136
 27. Chen J, Leung GH, Leung H, Ustiugova A, Shneyderman A, Korzinkin M, Gennert D, Pun FW, Aliper A, Ren F, Zhavoronkov A. From clock to clock: Therapeutic target discovery for aging and age-related diseases. *Ageing Res Rev*. 2025; 112:102871.
<https://doi.org/10.1016/j.arr.2025.102871>
PMID:40812388
 28. Kirkwood TB, Holliday R. The evolution of ageing and longevity. *Proc R Soc Lond B Biol Sci*. 1979; 205:531–46.
<https://doi.org/10.1098/rspb.1979.0083> PMID:42059
 29. Harman D. The biologic clock: the mitochondria? *J Am Geriatr Soc*. 1972; 20:145–7.
<https://doi.org/10.1111/j.1532-5415.1972.tb00787.x>
PMID:5016631
 30. Wallace DC. A mitochondrial paradigm of metabolic and degenerative diseases, aging, and cancer: a dawn for evolutionary medicine. *Annu Rev Genet*. 2005; 39:359–407.
<https://doi.org/10.1146/annurev.genet.39.110304.095751> PMID:16285865
 31. de Grey ADNJ, Rae M. Ending aging: the rejuvenation breakthroughs that could reverse human aging in our lifetime. New York: St. Martin's Press. 2007.
 32. Sinclair DA, LaPlante MD. Lifespan: why we age and why we don't have to. New York: Atria Books. 2019.
 33. Ocampo A, Reddy P, Martinez-Redondo P, Platero-Luengo A, Hatanaka F, Hishida T, Li M, Lam D, Kurita M, Beyret E, Araoka T, Vazquez-Ferrer E, Donoso D, et al. *In Vivo* Amelioration of Age-Associated Hallmarks by Partial Reprogramming. *Cell*. 2016; 167:1719–33.e12.
<https://doi.org/10.1016/j.cell.2016.11.052>
PMID:27984723
 34. Gill D, Parry A, Santos F, Okkenhaug H, Todd CD, Hernando-Herraez I, Stubbs TM, Milagre I, Reik W. Multi-omic rejuvenation of human cells by maturation phase transient reprogramming. *Elife*. 2022; 11:e71624.
<https://doi.org/10.7554/eLife.71624> PMID:35390271
 35. Sierra F, Kohanski R. Geroscience and the trans-NIH Geroscience Interest Group, GSIG. *Geroscience*. 2017; 39:1–5.
<https://doi.org/10.1007/s11357-016-9954-6>
PMID:28299635
 36. Barzilai N, Cuervo AM, Austad S. Aging as a Biological Target for Prevention and Therapy. *JAMA*. 2018; 320:1321–2.
<https://doi.org/10.1001/jama.2018.9562>
PMID:30242337
 37. Blagosklonny MV. Aging is not programmed: genetic pseudo-program is a shadow of developmental growth. *Cell Cycle*. 2013; 12:3736–42.
<https://doi.org/10.4161/cc.27188> PMID:24240128
 38. Harrison DE, Strong R, Sharp ZD, Nelson JF, Astle CM, Flurkey K, Nadon NL, Wilkinson JE, Frenkel K, Carter CS, Pahor M, Javors MA, Fernandez E, Miller RA. Rapamycin fed late in life extends lifespan in genetically heterogeneous mice. *Nature*. 2009; 460:392–5.
<https://doi.org/10.1038/nature08221> PMID:19587680
 39. Miller RA, Harrison DE, Astle CM, Baur JA, Boyd AR, de Cabo R, Fernandez E, Flurkey K, Javors MA, Nelson JF, Orihuela CJ, Pletcher S, Sharp ZD, et al. Rapamycin, but not resveratrol or simvastatin, extends life span of genetically heterogeneous mice. *J Gerontol A Biol Sci Med Sci*. 2011; 66:191–201.
<https://doi.org/10.1093/gerona/gdq178>
PMID:20974732
 40. Lamming DW, Ye L, Sabatini DM, Baur JA. Rapalogs and mTOR inhibitors as anti-aging therapeutics. *J Clin Invest*. 2013; 123:980–9.
<https://doi.org/10.1172/JCI64099> PMID:23454761
 41. Kennedy BK, Lamming DW. The Mechanistic Target of Rapamycin: The Grand Conductor of Metabolism and Aging. *Cell Metab*. 2016; 23:990–1003.
<https://doi.org/10.1016/j.cmet.2016.05.009>
PMID:27304501

42. Barnett HAR. The Biology of Life Span: a Quantitative Approach. By L. A. Gavrilov and N. S. Gavrilova (Harwood Academic Publishers). J Inst Actuar. 1992; 119:379–80.
<https://doi.org/10.1017/S0020268100019946>
43. Gompertz B. On the nature of the function expressive of the law of human mortality. Philos Trans R Soc Lond. 1825; 115:513–85.
<http://www.jstor.org/stable/107756>
44. Scheffer M, Bascompte J, Brock WA, Brovkin V, Carpenter SR, Dakos V, Held H, van Nes EH, Rietkerk M, Sugihara G. Early-warning signals for critical transitions. Nature. 2009; 461:53–9.
<https://doi.org/10.1038/nature08227>
PMID:19727193
45. Scheffer M, Carpenter SR, Lenton TM, Bascompte J, Brock W, Dakos V, van de Koppel J, van de Leemput IA, Levin SA, van Nes EH, Pascual M, Vandermeer J. Anticipating critical transitions. Science. 2012; 338:344–8.
<https://doi.org/10.1126/science.1225244>
PMID:23087241
46. Gao J, Barzel B, Barabási AL. Universal resilience patterns in complex networks. Nature. 2016; 530:307–12.
<https://doi.org/10.1038/nature16948>
PMID:26887493
47. Hipp MS, Kasturi P, Hartl FU. The proteostasis network and its decline in ageing. Nat Rev Mol Cell Biol. 2019; 20:421–35.
<https://doi.org/10.1038/s41580-019-0101-y>
PMID:30733602
48. Bahar R, Hartmann CH, Rodriguez KA, Denny AD, Busuttill RA, Dollé ME, Calder RB, Chisholm GB, Pollock BH, Klein CA, Vijg J. Increased cell-to-cell variation in gene expression in ageing mouse heart. Nature. 2006; 441:1011–4.
<https://doi.org/10.1038/nature04844>
PMID:16791200
49. Enge M, Arda HE, Mignardi M, Beausang J, Bottino R, Kim SK, Quake SR. Single-Cell Analysis of Human Pancreas Reveals Transcriptional Signatures of Aging and Somatic Mutation Patterns. Cell. 2017; 171:321–30.e14.
<https://doi.org/10.1016/j.cell.2017.09.004>
PMID:28965763
50. Martinez-Jimenez CP, Eling N, Chen HC, Vallejos CA, Kolodziejczyk AA, Connor F, Stojic L, Rayner TF, Stubbington MJ, Teichmann SA, de la Roche M, Marioni JC, Odom DT. Aging increases cell-to-cell transcriptional variability upon immune stimulation. Science. 2017; 355:1433–6.
<https://doi.org/10.1126/science.aah4115>
PMID:28360329
51. Demaria M, Ohtani N, Youssef SA, Rodier F, Toussaint W, Mitchell JR, Laberge RM, Vijg J, Van Steeg H, Dollé ME, Hoeijmakers JH, de Bruin A, Hara E, Campisi J. An essential role for senescent cells in optimal wound healing through secretion of PDGF-AA. Dev Cell. 2014; 31:722–33.
<https://doi.org/10.1016/j.devcel.2014.11.012>
PMID:25499914
52. Ritschka B, Storer M, Mas A, Heinzmann F, Ortells MC, Morton JP, Sansom OJ, Zender L, Keyes WM. The senescence-associated secretory phenotype induces cellular plasticity and tissue regeneration. Genes Dev. 2017; 31:172–83.
<https://doi.org/10.1101/gad.290635.116>
PMID:28143833
53. López-Otín C, Galluzzi L, Freije JM, Madeo F, Kroemer G. Metabolic Control of Longevity. Cell. 2016; 166:802–21.
<https://doi.org/10.1016/j.cell.2016.07.031>
PMID:27518560
54. Hannum G, Guinney J, Zhao L, Zhang L, Hughes G, Sada S, Klotzle B, Bibikova M, Fan JB, Gao Y, Deconde R, Chen M, Rajapakse I, et al. Genome-wide methylation profiles reveal quantitative views of human aging rates. Mol Cell. 2013; 49:359–67.
<https://doi.org/10.1016/j.molcel.2012.10.016>
PMID:23177740
55. Horvath S. DNA methylation age of human tissues and cell types. Genome Biol. 2013; 14:R115.
<https://doi.org/10.1186/gb-2013-14-10-r115>
PMID:24138928
56. Levine ME, Lu AT, Quach A, Chen BH, Assimes TL, Bandinelli S, Hou L, Baccarelli AA, Stewart JD, Li Y, Whitsel EA, Wilson JG, Reiner AP, et al. An epigenetic biomarker of aging for lifespan and healthspan. Aging (Albany NY). 2018; 10:573–91.
<https://doi.org/10.18632/aging.101414>
PMID:29676998
57. Lu AT, Quach A, Wilson JG, Reiner AP, Aviv A, Raj K, Hou L, Baccarelli AA, Li Y, Stewart JD, Whitsel EA, Assimes TL, Ferrucci L, Horvath S. DNA methylation GrimAge strongly predicts lifespan and healthspan. Aging (Albany NY). 2019; 11:303–27.
<https://doi.org/10.18632/aging.101684>
PMID:30669119
58. Cox DR. Regression models and life-tables. J R Stat Soc Series B Stat Methodol. 1972; 34:187–220.
<https://www.jstor.org/stable/2985181>
59. Aalen O. Nonparametric inference for a family of counting processes. Ann Stat. 1978; 6:701–26.
<https://www.jstor.org/stable/2958850>

60. Kalbfleisch JD, Prentice RL. The statistical analysis of failure time data. 2nd ed. Hoboken (NJ): Wiley. 2002.
<https://doi.org/10.1002/9781118032985>
61. Liu YY, Slotine JJ, Barabási AL. Controllability of complex networks. *Nature*. 2011; 473:167–73.
<https://doi.org/10.1038/nature10011>
PMID:21562557
62. Luria SE, Delbrück M. Mutations of Bacteria from Virus Sensitivity to Virus Resistance. *Genetics*. 1943; 28:491–511.
<https://doi.org/10.1093/genetics/28.6.491>
PMID:17247100
63. Jee J, Rasouly A, Shamovsky I, Akivis Y, Steinman SR, Mishra B, Nudler E. Rates and mechanisms of bacterial mutagenesis from maximum-depth sequencing. *Nature*. 2016; 534:693–6.
<https://doi.org/10.1038/nature18313>
PMID:27338792
64. Lotfollahi M, Klimovskaia Susmelj A, De Donno C, Hetzel L, Ji Y, Ibarra IL, Srivatsan SR, Naghipourfar M, Daza RM, Martin B, Shendure J, McFaline-Figueroa JL, Boyeau P, et al. Predicting cellular responses to complex perturbations in high-throughput screens. *Mol Syst Biol*. 2023; 19:e11517.
<https://doi.org/10.15252/msb.202211517>
PMID:37154091
65. Holford NH, Sheiner LB. Understanding the dose-effect relationship: clinical application of pharmacokinetic-pharmacodynamic models. *Clin Pharmacokinet*. 1981; 6:429–53.
<https://doi.org/10.2165/00003088-198106060-00002>
PMID:7032803
66. Bhalla US, Iyengar R. Emergent properties of networks of biological signaling pathways. *Science*. 1999; 283:381–7.
<https://doi.org/10.1126/science.283.5400.381>
PMID:9888852
67. Sorger PK, Allerheiligen SR, Abernethy DR, Altman RB, Brouwer KL, Califano A, D'argenio DZ, Iyengar R, Jusko WJ, Lalonde RL, Lauffenburger DA, Shoichet BK, Stevens JL, Subramaniam S, van der Graaf PH, Vicini P, Ward RJ. Quantitative and Systems Pharmacology in the Post-genomic Era: New Approaches to Discovering Drugs and Understanding Therapeutic. 2011.
<https://api.semanticscholar.org/CorpusID:16626059>
68. Baker DJ, Wijshake T, Tchkonja T, LeBrasseur NK, Childs BG, van de Sluis B, Kirkland JL, van Deursen JM. Clearance of p16Ink4a-positive senescent cells delays ageing-associated disorders. *Nature*. 2011; 479:232–6.
<https://doi.org/10.1038/nature10600> PMID:22048312
69. Xu M, Pirtskhalava T, Farr JN, Weigand BM, Palmer AK, Weivoda MM, Inman CL, Ogrodnik MB, Hachfeld CM, Fraser DG, Onken JL, Johnson KO, Verzosa GC, et al. Senolytics improve physical function and increase lifespan in old age. *Nat Med*. 2018; 24:1246–56.
<https://doi.org/10.1038/s41591-018-0092-9>
PMID:29988130
70. Banito A, Rashid ST, Acosta JC, Li S, Pereira CF, Geti I, Pinho S, Silva JC, Azuara V, Walsh M, Vallier L, Gil J. Senescence impairs successful reprogramming to pluripotent stem cells. *Genes Dev*. 2009; 23:2134–9.
<https://doi.org/10.1101/gad.1811609> PMID:19696146
71. Mosteiro L, Pantoja C, Alcazar N, Marión RM, Chondronasiou D, Rovira M, Fernandez-Marcos PJ, Muñoz-Martin M, Blanco-Aparicio C, Pastor J, Gómez-López G, De Martino A, Blasco MA, et al. Tissue damage and senescence provide critical signals for cellular reprogramming *in vivo*. *Science*. 2016; 354:aaf4445.
<https://doi.org/10.1126/science.aaf4445>
PMID:27884981
72. De Cecco M, Ito T, Petrashen AP, Elias AE, Skvir NJ, Criscione SW, Caligiana A, Broccoli G, Adney EM, Boeke JD, Le O, Beauséjour C, Ambati J, et al. L1 drives IFN in senescent cells and promotes age-associated inflammation. *Nature*. 2019; 566:73–8.
<https://doi.org/10.1038/s41586-018-0784-9>
PMID:30728521
73. Simon M, Van Meter M, Ablaeva J, Ke Z, Gonzalez RS, Taguchi T, De Cecco M, Leonova KI, Kogan V, Helfand SL, Neretti N, Roichman A, Cohen HY, et al. LINE1 Derepression in Aged Wild-Type and SIRT6-Deficient Mice Drives Inflammation. *Cell Metab*. 2019; 29:871–85.e5.
<https://doi.org/10.1016/j.cmet.2019.02.014>
PMID:30853213
74. McIntyre RL, Molenaars M, Schomakers BV, Gao AW, Kamble R, Jongejan A, van Weeghel M, van Kuilenburg AB, Possemato R, Houtkooper RH, Janssens GE. Anti-retroviral treatment with zidovudine alters pyrimidine metabolism, reduces translation, and extends healthy longevity via ATF-4. *Cell Rep*. 2023; 42:111928.
<https://doi.org/10.1016/j.celrep.2022.111928>
PMID:36640360
75. Kanfi Y, Naiman S, Amir G, Peshti V, Zinman G, Nahum L, Bar-Joseph Z, Cohen HY. The sirtuin SIRT6 regulates lifespan in male mice. *Nature*. 2012; 483:218–21.
<https://doi.org/10.1038/nature10815> PMID:22367546
76. Justice JN, Nambiar AM, Tchkonja T, LeBrasseur NK, Pascual R, Hashmi SK, Prata L, Masternak MM, Kritchevsky SB, Musi N, Kirkland JL. Senolytics in idiopathic pulmonary fibrosis: Results from a first-in-human, open-label, pilot study. *EBioMedicine*. 2019; 40:554–63.
<https://doi.org/10.1016/j.ebiom.2018.12.052>
PMID:30616998

77. Fahy GM, Brooke RT, Watson JP, Good Z, Vasanaawala SS, Maecker H, Leipold MD, Lin DT, Kobor MS, Horvath S. Reversal of epigenetic aging and immunosenescent trends in humans. *Aging Cell*. 2019; 18:e13028. <https://doi.org/10.1111/ace1.13028> PMID:31496122
78. Chen RTQ, Rubanova Y, Bettencourt J, Duvenaud D. Neural ordinary differential equations. In: *Advances in Neural Information Processing Systems*. NeurIPS. 2018:31. https://proceedings.neurips.cc/paper_files/paper/2018/file/69386f6bb1dfed68692a24c8686939b9-Paper.pdf
79. Rajman L, Chwalek K, Sinclair DA. Therapeutic Potential of NAD-Boosting Molecules: The *In Vivo* Evidence. *Cell Metab*. 2018; 27:529–47. <https://doi.org/10.1016/j.cmet.2018.02.011> PMID:29514064
80. Partridge L, Fuentealba M, Kennedy BK. The quest to slow ageing through drug discovery. *Nat Rev Drug Discov*. 2020; 19:513–32. <https://doi.org/10.1038/s41573-020-0067-7> PMID:32467649
81. Pun FW, Leung GH, Leung HW, Liu BH, Long X, Ozerov IV, Wang J, Ren F, Aliper A, Izumchenko E, Moskalev A, de Magalhães JP, Zhavoronkov A. Hallmarks of aging-based dual-purpose disease and age-associated targets predicted using PandaOmics AI-powered discovery engine. *Aging (Albany NY)*. 2022; 14:2475–506. <https://doi.org/10.18632/aging.203960> PMID:35347083
82. Henzinger TA. The theory of hybrid automata. In: *Verification of digital and hybrid systems*. NATO ASI Series. Berlin: Springer. 2000. https://doi.org/10.1007/978-3-642-59615-5_13
83. Olde Loohuis L, Witzel A, Mishra B. Cancer hybrid automata: model, beliefs and therapy. *Inf Comput*. 2014; 236:68–86. <https://doi.org/10.1016/j.ic.2014.01.013>
84. Olde Loohuis L, Witzel A, Mishra B. Cancer hybrid automata. In: *Proceedings of the 4th International Workshop on Hybrid Systems and Biology (HSB)*; 2012. 137–51. <https://doi.org/10.48550/arXiv.1208.3857>
85. Massey SE, Mishra B. Origin of biomolecular games: deception and molecular evolution. *J R Soc Interface*. 2018; 15:20180429. <https://doi.org/10.1098/rsif.2018.0429> PMID:30185543
86. Bryant HE, Schultz N, Thomas HD, Parker KM, Flower D, Lopez E, Kyle S, Meuth M, Curtin NJ, Helleday T. Specific killing of BRCA2-deficient tumours with inhibitors of poly(ADP-ribose) polymerase. *Nature*. 2005; 434:913–7. <https://doi.org/10.1038/nature03443> PMID:15829966
87. Kaelin WG Jr. Synthetic lethality: a framework for the development of wiser cancer therapeutics. *Genome Med*. 2009; 1:99. <https://doi.org/10.1186/gm99> PMID:19863774
88. ICMJE. Recommendations for the conduct, reporting, editing, and publication of scholarly work in medical journals: updated recommendations on disclosures on use of AI-assisted technologies. 2024. <https://www.icmje.org/icmje-recommendations.pdf>

SUPPLEMENTARY MATERIALS

Supplementary Table

Please browse Full Text version to see the data of Supplementary Table 1.

Supplementary Table 1. Full_scored_dataset.

Supplementary Code

Please browse Full Text version to see the data of Supplementary Code.

Supplementary Information

Please browse Full Text version to see the data of Supplementary Information 1–10.