

Dynamical Network Stability Analysis of Multiple Biological Ages Provides a Framework for Understanding the Aging Process

Glen Pridham, MSc^{id} and Andrew D. Rutenberg, PhD*

Department of Physics and Atmospheric Science, Dalhousie University, Halifax, Nova Scotia, Canada.

*Address correspondence to: Andrew D. Rutenberg, PhD. E-mail: adr@dal.ca

Decision Editor: Gustavo Duque, MD, PhD, FRACP, FGSA (Biological Sciences Section)

Abstract

Widespread interest in nondestructive biomarkers of aging has led to a multitude of biological ages that each proffers a “true” health-adjusted individual age. Although each measure provides salient information on the aging process, they are each univariate, in contrast to the “hallmark” and “pillar” theories of aging, which are explicitly multidimensional, multicausal, and multiscale. Fortunately, multiple biological ages can be systematically combined into a multidimensional network representation. The interaction network between these biological ages permits analysis of the multidimensional effects of aging, as well as quantification of causal influences during both natural aging and, potentially, after anti-aging intervention. The behavior of the system as a whole can then be explored using dynamical network stability analysis, which identifies new, efficient biomarkers that quantify long-term resilience scores on the timescale between measurements (years). We demonstrate this approach using a set of 8 biological ages from the longitudinal Swedish Adoption/Twin Study of Aging (SATSA). After extracting an interaction network between these biological ages, we observed that physiological age, a proxy for cardiometabolic health, serves as a central node in the network, implicating it as a key vulnerability for slow, age-related decline. We furthermore show that while the system as a whole is stable, there is a weakly stable direction along which recovery is slow—on the timescale of a human lifespan. This slow direction provides an aging biomarker, which correlates strongly with chronological age and predicts longitudinal decline in health—suggesting that it estimates an important driver of age-related changes.

Keywords: Biological age, Complexity, Eigen analysis, Systems biology

The continued search for a biomarker of aging that quantifies the effects of natural aging and antiaging interventions (1–3) has resulted in a proliferation of biological ages (BAs) (3), including recent epigenetic clocks (4) (such as Refs (5–8)). Each BA estimates an individual’s health-adjusted effective age, which may differ from their chronological age (CA). Each BA uses a model that converts a battery of measurements into a univariate proxy for health, either using regression on CA (7) or a heuristic mapping into a specific measure of health such as risk of death (5,6). Popular BAs have been extensively validated and are often sensitive to mortality risk (9) and the effects of antiaging interventions (2). Although it is tempting to simply pick the best BA for a particular application, or aggregate BAs using a heuristic approach (10,11), this risks missing the effects of multivariate interactions during the aging process.

Aging is putatively an interacting multivariate, multicausal process (12–15), as has borne out explicitly in computational studies (16–18). This puts aging firmly in the purview of complexity science, where network analysis can be used to account for potentially catastrophic confounding effects due to interactions such as feedbacks between biological variables (12). Such confounding effects could explain the conflicting results emerging from antiaging intervention studies (2,19,20).

Learning the underlying network topology can help us to explain these confounding effects, and also to identify vulnerabilities of the biological system. For example, networks with feedback loops may be vulnerable to run-away effects (12), while bottlenecks through high-connectivity nodes may put a network at risk of complete collapse if those vital nodes are damaged (21,22). Networks also provide a useful tool for visualizing and quantifying causal sequences: both during natural aging and, potentially, after interventions.

Fortunately, the complexity of the problem enables an elegant approach for processing and interpreting a network of interacting BAs. Dynamical stability analysis tells us that systems which are mostly stable—such as living organisms—can be understood by the way in which they respond to small perturbations. That is, we look for longitudinal disruptions to homeostasis and subsequent recovery (or lack thereof) (23). The eigen-directions provide a spectrum of fundamental recovery rates (timescales), oriented to completely account for the complex interactions of the network. These rates describe canonical changes so that we can infer long-term behavior by the slowest recovery rates (23,24), which determine system resilience (25). Taking a long time to recover is indicative of weak stability and hence vulnerability to stochastic stressors, whose effects tend to pile up along slow or unstable

Received: August 31 2023; Editorial Decision Date: December 30 2023.

© The Author(s) 2024. Published by Oxford University Press on behalf of The Gerontological Society of America.

This is an Open Access article distributed under the terms of the Creative Commons Attribution License (<https://creativecommons.org/licenses/by/4.0/>), which permits unrestricted reuse, distribution, and reproduction in any medium, provided the original work is properly cited.

eigen-directions. Indeed, prior work on health biomarkers showed that across 4 data sets (2 mice and 2 humans), the dominant risk direction for survival or dementia onset was always the first or second least-stable (slowest) eigen-direction (23). Furthermore, a recent deep learning result from mice has shown the existence of an unstable latent variable (26) with properties similar to the frailty index (FI) (20,27)—notably nonlinear growth with age and sensitivity to antiaging intervention. Since salient aging information naturally condenses into the least stable (slowest recovery) eigen-directions, these least stable eigen-directions are the best way to describe the collective aging behavior of a network of BAs.

By using longitudinal data, we can also potentially infer causality within our network. We have already developed a generic model for homeostasis, which estimates a network of interactions between biomarkers together with steady-state behavior (23). Although our model is trained using observational data (28), we expect that interventions that cause small perturbations will behave similarly to the random stresses, which drive observational data—because these represent the random effects of interventions that individuals experience throughout their lives, such as lifestyle changes, medicine, and disease. Previously, we found that including such directed cause \rightarrow effect links conferred little benefit in reducing the root mean squared error (RMSE) of biomarker trajectories. Here, we revisit this question with a novel quantitative score based on predicting the correct direction of biomarker change. As a result, we obtain a model that predicts both the causal sequence of events occurring during normal aging and a best guess for how interventions will propagate.

We demonstrate that network analysis can leverage the abundance of BAs to answer fundamental questions about what causes aging and about how aging systems are likely to respond to interventions. We apply our approach to longitudinal multivariate BA data, generating a network interactome capable of capturing the coordinated effects of the BAs. By considering BAs from multiple biological scales we can surmise how the effects of aging propagate from DNA to functional decline. Once the network is estimated, we can analyze its eigen-directions to understand how the aggregate effects of multiple BAs affect the organism as time progresses. This yields dominant natural variables, which are the salient features of aging and provide canonical coordinates. We show that the least stable natural variable is an efficient choice for monitoring the aging process. We also show that the interaction network has vulnerabilities, which are consistent with qualitative theories of aging, and is able to describe confounding effects of model interventions.

Method

Model

We model an arbitrary dynamical system near a stable (homeostatic) point as

$$\begin{aligned} \vec{b}_{n+1} &= \vec{b}_n + \mathbf{W}\Delta t_{n+1}(\vec{b}_n - \vec{\mu}_n) + \vec{\varepsilon}_{n+1} \\ \vec{\varepsilon}_{n+1} &\sim N(\mathbf{0}, \mathbf{\Sigma}|\Delta t_{n+1}) \\ \vec{\mu}_n &\equiv \vec{\mu}_0 + \mathbf{\Lambda}\vec{x}_n \end{aligned} \quad (1)$$

where \vec{b}_n represents an individual's set of BAs measured at time timepoint t_n and $\Delta t_{n+1} \equiv t_{n+1} - t_n$ is the measurement interval. Each individual has a different number of total measurements before leaving the study, T_i , and hence eqn

(1) applies to the $T_i - 1$ pairs of sequential measurements; dropout and missing data are discussed below and in [Supplementary Material](#). We refer to this model as the Stochastic Finite-difference (SF) model, reflecting its relationship to the Stochastic Process model (29), which is the generalized continuous version of the SF model (23). The model estimates: an equilibrium position, $\vec{\mu}_n$, a causal, resilience parameter, \mathbf{W} , which captures an interacting recovery network; and a noise term, $\mathbf{\Sigma}$ which implicitly includes additional effects not in the model—such as nonlinear effects and fast dynamical changes. The model is only sensitive to changes which occur slower than the timescale set by Δt_{n+1} , which for this study is approximately 3 years. This means that short-term changes such as due to a flu infection appear in the noise term, $\mathbf{\Sigma}$. \mathbf{W} captures only long-term “resilience,” that is, longitudinal correlations, over the course of years—such as age-related decline. The equilibrium position, $\vec{\mu}_n$, is allowed to vary linearly with respect to a set of covariates for each individual, \vec{x}_n , through $\mathbf{\Lambda}$.

The diagonal elements of \mathbf{W} permit recovery toward $\vec{\mu}_n$, whereas the off-diagonals couple values across BAs,

$$\frac{E(b_{jn+1} - b_{jn})}{E(\Delta t_{n+1})} = W_{jj}E(b_{jn} - \mu_{jn}) + \sum_{k \neq j} W_{jk}E(b_{kn} - \mu_{kn}) \quad (2)$$

where $E(x)$ represents the expectation value of x . (In deriving eqn (2) we assume that the current BA values, b_{jn} , are negligibly correlated with the follow-up time, Δt_{n+1} .) This means that if—through intervention or natural aging—the values of some of the BAs change, b_k , then we will see a change in each different b_j via W_{jk} . The diagonal elements W_{jj} are the marginal recovery rates of b_j ignoring all other BAs, whereas the off-diagonal elements W_{jk} allow interactions between different BAs—hastening or ameliorating their decline. Hence, changes to one BA can propagate into the other BAs, allowing for a central driver of either age-related dysfunction or antiaging treatment. The system only stops when each b_j simultaneously reaches its equilibrium position, μ_{jn} . As we will see, the estimated equilibrium positions can only be reached well beyond the lifespan of normal humans and hence the system drifts indefinitely. The interactions, \mathbf{W} , can be simplified by diagonalization using the eigen-decomposition to yield a set of composite natural aging variables z_k which satisfy

$$\frac{E(z_{kn+1} - z_{kn})}{E(\Delta t_{n+1})} = \lambda_k E(z_{kn} - \tilde{\mu}_{kn}) \quad (3)$$

for $\vec{z} \equiv \mathbf{P}^{-1}\vec{b}$ and $\tilde{\vec{\mu}} \equiv \mathbf{P}^{-1}\vec{\mu}$, where λ_k is the associated eigenvalue. We index the z_k by their sorted eigenvalue strength, such that z_1 has the greatest (closest to $+\infty$) eigenvalue λ_1 . (For simplicity, we drop the tilde notation for the remainder of the paper.) Observe that the natural aging variables, z_k , do not interact: they either increase, decrease or stay the same, depending on the value of $Re(\lambda_k)$ as can be seen by iterating eqn (3). ($Im(\lambda_k) \neq 0$ contributes oscillations, but we focus exclusively on the real parts of eigenvalues in the present study.) The timescale over which these changes occur is set by $|Re(\lambda_k)|^{-1} \equiv |\lambda_k|^{-1}$, that is, the absolute value of the real part of λ_k (λ_k has units of years^{-1}). The inverse timescale, λ_k , determines how quickly the average individual reaches a steady-state. Since the z_k are independent via eqn (3), events or interventions that modulate only $E(z_k)$ will not affect any other $E(z_j)$. The payoff of this approach is twofold: aging

Table 1. Biological Age Summary (9)

Biological Age	Risk Direction*	Scale	Input	Output/Pooling
Frailty index (FI) [†]	Up	Organism	Health deficits [‡]	Mean
Functional aging index (FAI)	Up	Organism	Sensory, grip, pulmonary, and gait [§]	Standardize then average
Cognition	Down	System (brain)	Cognitive testing	PC1
Physiological Age (PhysioAge)	Up	System (cardiometabolic)	Biomarkers and physical exam	PCA then Klemera-Doubal (7)
GrimAge	Up	Epigenetic [¶]	CpGs	Mortality risk [#]
PhenoAge	Up	Epigenetic**	CpGs	Mortality risk ^{††}
Hannum	Up	Epigenetic	CpGs	CA ^{##}
Horvath	Up	Epigenetic	CpGs	CA ^{##}
Telomere	Down	Genetic	$\frac{\text{Telomere length}}{\text{standard deviation}}$	—

*Direction of change with increasing chronological age.

[†]Reserved as an outcome measure of individual health.

[‡]Score from 0 (none) to 1 (full): disability, disease, and self-reported ill-health.

[§]Self-reported hearing/vision, grip strength, lung strength, and gait speed.

^{||}Male: body mass index, waist-to-height ratio, weight, systolic blood pressure, diastolic blood pressure, hemoglobin, serum glucose (log), and apolipoprotein B. Female: hip circumference, waist circumference, systolic blood pressure, serum glucose (log), and triglycerides (log).

[¶]Trained to emulate smoking pack years and plasma proteins: adrenomedullin, beta-2-microglobulin, cystatin C, GDF-15, leptin, PAI-1, and tissue inhibitor metalloproteinases 1 (6).

[#]Linearly transform mortality risk to match mean/standard deviation of chronological age (6).

^{**}Trained to predict time-to-death, which includes a proportional hazard from: albumin, creatinine, serum glucose, C-reactive protein, lymphocytes (%), mean red cell volume, red cell distribution width, alkaline phosphatase, and white blood cell count (5).

^{††}Invert 10-year multivariate mortality risk (Gompertz + proportional hazard with 9 covariates) (5).

^{##}CA = chronological age.

information gets compressed into a few specific variables, and the interconnected system behavior is greatly simplified.

The z_k are able to drive the observed BAs through the mapping $Pz = b$, which will spread out the effects across several BAs since P is often dense (23). In general, the slowest z_k (greatest λ_k) have the slowest recovery ($\lambda_k > 0$ never recover); previously, we observed that the key z_k driving changes in b are always among the slowest (23) (for health biomarkers).

The eigen-decomposition also lets us decompose the network, represented as a matrix of weights W , into a sum of sub (eigen)-networks (matrices),

$$W = \sum_i \lambda_i P_{.i} \otimes P_{.i}^{-1} \tag{4}$$

where $P_{.i}$ is both the i th column of P and the i th eigenvector, and $\bar{x} \otimes \bar{x} \equiv \bar{x}\bar{x}^T$ defines the outer product, \otimes . Each eigenvalue, λ_i , has associated with it a subnetwork, $P_{.i} \otimes P_{.i}^{-1}$. The network is the sum of all subnetworks, weighted by their associated eigenvalues (e.g., Supplementary Figure S2). This permits us to visually analyze the effective network for each z_k using their associated eigenvalue-eigenvector pair, λ_k and $P_{.k} \otimes P_{.k}^{-1}$.

We estimated Λ , W , and Σ using linear regression as described in Supplementary Material. We also iteratively impute the expected model mean for all missed measurements, as described in Missing Data.

Data

We use publicly available longitudinal data from Li et al. (9) Their data are derived from the Swedish Adoption/Twin Study of Aging (SATSA), and include: age, sex, and 9 BAs. The population included $N = 845$ individuals (342 males), average age at entry: 63.6 ± 0.3 years (standard deviation: 8.6, min: 44.9, max: 88.0). Individuals were regularly measured with median $\Delta t = 3$ years (interquartile range: 2.3–3.4 years) and a median number of 4 measurements per person (interquartile range: 3–7 measurements, max: 9). Survival data were not

included in the data set, and patients were instead labeled as dropouts after their last measurement.

The BAs used in the present analysis are summarized in Table 1. We considered 9 BAs from 4 biological scales: genetic, epigenetic, system, and entire organism. BAs were harmonized to the same scale (~years) as follows: Telomere was multiplied by 69.29 and then we added 20.72 to match the standard deviation and mean of CA. Similarly, Cognition was multiplied by 0.9440 then we added 21.34. We model the dynamics of the 8 BAs (“predictors”) and hold out the FI as a longitudinal health outcome, since it is a good predictor of risk of adverse outcome such as morbidity and mortality (18,30) (the FI is also non-Gaussian (27), in contrast to our model assumptions). We fit a more general network including the FI and CA in Supplementary Material.

Statistics and Data Handling

All analysis and statistics were performed using R version 4.1.1 (31). Errors were estimated by bootstrapping using 100 resamples, unless otherwise specified. All statistical tests are z -tests, unless otherwise specified. All error bars are standard errors, unless specified otherwise. Fitting and simulating functions, as well as fitted parameter values, are available on GitHub at https://github.com/GlenPr/stochastic_finite-difference_model.

Preprocessing

Before fitting, we transformed the BAs at each timepoint using principal component analysis (PCA) where the transformation was learned from the first timepoint (except for the diagonal model). The transformation is isomorphic (information-preserving) (23) so we estimated model parameters in PC-space then mapped them into BA-space. The number of PCs to use was selected by minimizing the 632-corrected RMSE, which was 8 (max/information preserving). 632-correction uses a linear mixture of 63.2% out-of-sample test error and 36.8% in-sample training error (23). We used PCA because selecting fewer PCs than BAs can

avoid collinearity—as was done in [Supplementary Material](#) when the FI and CA were included in the network. A priori, Telomere was initially batch adjusted using linear regression (32), we observed that Telomere was normally distributed but included a few extreme outliers (right tail). Since these could be artifacts of the batch adjustment, we excluded all outliers with $p < 10^{-5}$ ($9/6006 \ll 1\%$ of entries).

Model Selection

For initial model selection, we minimize the RMSE and mean absolute error (MAE). We used 632-corrected error values, since these have minimal bias for our model (23). For ties, we maximize the area under the receiver operator characteristic curve (AUC) (33) of the 8 pooled BAs worsening in the next timestep, which is the probability that the prediction will correctly rank individuals who will see an increase in BA as higher than those who will not (34). We select between a fully flexible \mathbf{W} (“FullW”) and 3 simplified versions: the null model (with $\mathbf{W} = 0$), diagonal in BA-space (“DiagW”), or diagonal in principal component-space (“SymW,” which has symmetric \mathbf{W}).

Missing Data

Data were missing due to missed measurements and dropout at an overall rate of 76%. We imputed the missed measurements and considered the effect of imputing dropped patients in [Supplementary Material](#)—the latter made no visible difference to the final results. Excluding dropout, the majority of (“predictor”) BA values were missing (53%), which broke down as the following missingness: 20% (PhysioAge), 23% (Cognition), 27% (FAI), 60% (Telomere), 74% (Horvath), 74% (Hannum), 74% (PhenoAge), and 74% (GrimAge). The FI was missing in 20% of cases.

Missing data were initially imputed by carrying forward the last measurement, then reversed and carried backwards, then we imputed any remaining missingness using the mean of a multivariate Gaussian independently for each timepoint. This initial imputation was replaced at each fit iteration ($\times 5$) by the mean model prediction (expectation-maximization). See [Supplementary Material](#) for full details.

Failure to impute could lead to biased conclusions (35) since most missingness in clinical studies is due in part to poor health (36): here we observed that individuals missing all epigenetic BA measurements were significantly older ($p = 10^{-10}$, Wilcoxon test). Imputed values for these BAs were higher than observed, ostensibly accounting for this effect. The relatively high missingness makes imputation quality important. Imputation quality was visually assessed as good, with realistic dispersion, trajectories, and age-dependence ([Supplementary Figures S3 and S4](#)). The available case analysis had much lower significance levels, but captured most of the coarse-grained features of the imputed analysis ([Supplementary Figures S7 and S8](#)). Multiple imputation may give a better estimate of true effect sizes since it accounts for imputation uncertainty, see [Supplementary Figure S8](#); qualitative results were identical to our primary imputation result. We consider only the singly imputed analysis in the main text. All outcome measures consider only observed values.

Results

We compared several model variants, notably: the full model according to eqn (1) (FullW), the diagonal model eqn (3)

either in the PCA basis (SymW) or with the raw BA variables (DiagW), and the null model with $\mathbf{W} = 0$. The diagonal elements parameterize self-recovery from perturbations, while the off-diagonal elements parameterize interactions between the variables. The symmetrical \mathbf{W} has only undirected links (bidirectional interactions). The RMSE and MAE were worse both for the null model and for the noninteracting model (DiagW)—but did not discriminate the asymmetric \mathbf{W} (FullW) from the symmetric (SymW). This indicates that interactions were present and important for prediction ($\mathbf{W} \neq 0$ and not diagonal). To break the tie between FullW and SymW, we picked the one which best predicted the direction of change in BA at the next timestep (BA went up in 60% of measurements and down in 40%). We found that FullW performed better at predicting this worsening, having both lower MAE at 68% confidence and higher AUC at $p = .1$ using the DeLong test (33) (combined: $.04 \leq p \leq .1$). Our final model (FullW) predicted future values with accuracy $R_{\text{train}}^2 \approx R_{\text{test}}^2 = 0.65 \pm 0.01$, $\text{RMSE}_{632} = 5.75 \pm 0.09$ years, and worsening AUC of 0.764 ± 0.005 .

The interaction network, \mathbf{W} , estimated from the data is presented in [Figure 1](#). \mathbf{W} indicates that PhysioAge is the central node and the primary driver of changes over time (the strongest total outgoing links), with GrimAge as an important secondary, high-connectivity node. Observe that there is a positive feedback loop between the highest-connected nodes, $\text{PhysioAge} \rightarrow \text{GrimAge} \rightarrow \text{PhysioAge}$. Explicit inclusion of CA and the FI into the network does not appreciably change the connectivity of these nodes ([Supplementary Figure S22](#)); nor does the choice of imputation strategy ([Supplementary Figure S7](#)). Note that in [Supplementary Material](#), we confirm that the FI is not connected to Hannum, Horvath, PhenoAge, or GrimAge, as was reported elsewhere using an unrelated statistical model (37). Returning to [Figure 1](#), the high connectivity of PhysioAge would allow it to very quickly propagate dysfunction via eqn (2).

Drift of the BAs with age is the result of the pursuit of equilibrium, μ_i, μ_{0j} ranged from -120 ± 100 to 35 ± 24 years for BAs which decrease with age (Cognition and Telomere) and 127 ± 38 to 190 ± 82 years for the increasing BAs (remaining). (Sex effects were small, ≤ 10 years; [Supplementary Table S1](#).) In all cases, the equilibrium position is far outside of the age distribution of the population, causing them to drift coherently with age in their respective risk directions ([Supplementary Figure S14](#)). In the z -picture this effect is concentrated into z_1 and z_2 which drift the most, and z_3 which saturates around age 90—the remaining z_k quickly equilibrated and stopped changing with age ([Supplementary Figure S13](#)). This is an indication that age-related changes are concentrated into the slowest natural variables, z_1, z_2 , and z_3 , and primarily into z_1 .

The eigenvalues of \mathbf{W} determine system stability, so our focus is on the greatest eigenvalues, which therefore recover slowest (i.e. those closest to zero since they are negative; mean-stability is determined by eqn (3)). The eigenvalues are presented in [Figure 2A](#). Both z_1 and z_2 (green triangles) are notably slower than the slowest diagonal elements W_{11} and W_{22} (orange points). The associated timescales are $|\lambda_1|^{-1} = 127 \pm 53$ years and $|\lambda_2|^{-1} = 44 \pm 8$ years. Observe that both timescales are on the order of a typical human lifespan and are significantly longer than the remaining lifespan of the population, which were all older adults (baseline ages 45–88). The timescales, $|\lambda_k|^{-1}$ determine how quickly the z_k

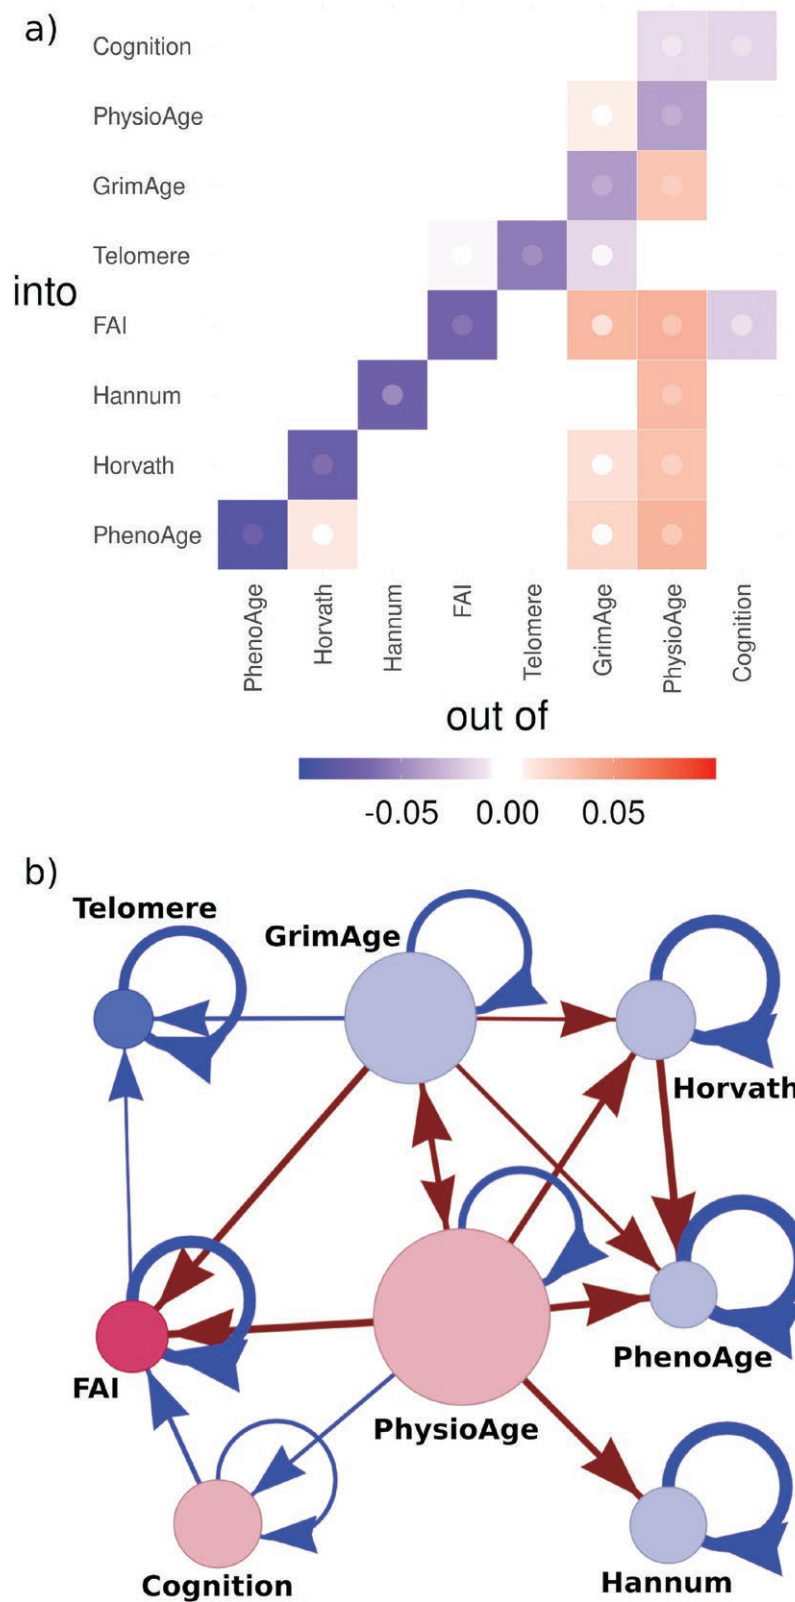


Figure 1. Network interactome. Both representations: matrix (A) and network (B) are equivalent. PhysioAge is the dominant node with the strongest connections, directly driving almost all other BAs (but not Telomere $p = .3$). GrimAge has weaker connections but also has many outgoing connections. All links are significant at $p < .05$. (A) Network weight matrix, W . Our model estimates each interaction parameter in this matrix. Inner point is limit of 95% CI closest to 0: point is most visible for the least significant tiles. Nonsignificant tiles are whited-out ($p > .05$). The elements leaving PhysioAge are typically larger than those leaving GrimAge, and have higher statistical significance (Supplementary Figure S8). Both have weak diagonal recovery W_{ii} . Matrix is rank-ordered by diagonal recovery strength. (B) Network representation. Networks encode conditional dependence structures: variables are conditionally dependent if and only if there is a link directly connecting them. For example, GrimAge and Cognition are conditionally independent, since they only interact via an intermediary (PhysioAge). Node size, $n_k = \sqrt{\sum_{j \neq k} W_{jk}^2}$ (outgoing strength). Node color indicates biological scale (see Table 1).

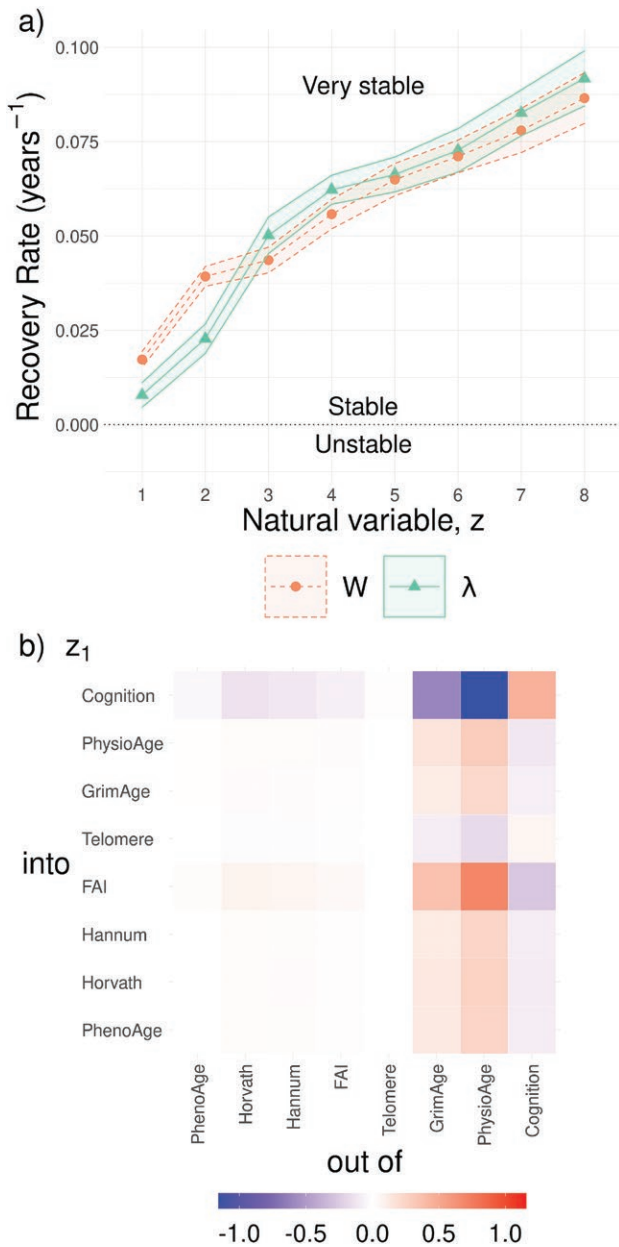


Figure 2. Natural variables of W . Natural variables do not interact, allowing us to analyze their stability. We observed a very weak stability (A), indicative of a slow recovery rate, $\lambda_1 \approx 0$. The associated eigenvector is visualized in (B), comparing to Figure 1A we see that the slowest eigenvector captures the dense outgoing connections from GrimAge and PhysioAge, including feedback loops. (A) Network stability (resilience). Eigenvalues, λ_k , determine recovery rate, $-\lambda_k$. Although the network diagonal (W_{jj}) indicates that some biomarkers recover slowly, the network as a whole recovers even slower along λ_1 and λ_2 . (Eigenvalue rank is used to index the z_k .) (B) Eigen-network of z_1 . Associated with each eigenvalue is an eigenvector. The matrix of eigenvectors, P , is used to generate the natural variables as linear combinations of BAs ($\vec{z} = P^{-1}\vec{b}$). The slowest recovering/least stable direction, z_1 , is predominantly PhysioAge, Cognition and GrimAge, all connected into the remaining BAs. Plotted is $P_1 \otimes P_1^{-1}$, where P_1 is the first eigenvector (eqn (4)). Note the role of well-connected BAs with feedback loops: the z_1 eigen-network has links both above and below the diagonal.

converge to the steady state (eqn (3)). This means that the long-time behavior of the system will depend increasingly on z_1 and z_2 , which will dominate the mapping into the BAs via $\vec{b} = P\vec{z}$.

In Figure 2B we visualize the λ_1 -eigenvector using eqn (4). The λ_1 -eigenvector is centered on a fully-outgoing-connected PhysioAge with feedbacks between GrimAge and Cognition. This means that z_1 represents the collective action of these 3 BAs driving changes in all 8 BAs. Although these 3 BAs already have the slowest marginal recoveries, W_{jj} , the collective action of z_1 is even slower due to interactions between the BAs.

This bottlenecking of aging information into z_1 and, to an extent z_2 , is easily confirmed by looking at the correlation matrix, Figure 3A. z_1 is strongly correlated with almost every BA (weakly with Telomere), and always in the same risk direction. z_2 shares these correlations except for Cognition, suggesting that splitting between z_1 and z_2 is primarily due to differences in cognitive aging rate. Both also had the strongest correlations with CA and the FI of any z_k . Multivariate ANOVA confirmed that the real part of z_1 was the dominant predictor of the FI (59% of the explained variance). Correlations with the FI were concentrated into the lowest z_k , which is clearly demonstrated in the multiply imputed correlations in Supplementary Figure S9 (which accounted for imputation error). In Supplementary Material, we demonstrate that z_1 and z_2 are the furthest from equilibrium, which causes them to drift for the entire human lifespan leading to the observed correlation with CA, for example, Figure 3B. Altogether, it appears that the observed age-related changes, including health, are concentrated into the least stable dimensions, particularly z_1 .

Our model, eqn (1), encodes causal dependence of the current timepoint on the previous timepoint. Hence, we can simulate the dynamics after a hypothetical intervention using eqn (1) and the estimated model parameters. We operationalize interventions as an instantaneous rejuvenation of a targeted BA at a specific CA, in a manner which emulates the switching of mortality risk immediately due to an antiaging intervention (38). We simulated matched case and control populations for various interventions, each population contains 50 000 males and 50 000 females and starts at age 60; initial values were sampled from the fully observed BAs of sex-matched individuals in the age range 55–65. We simulated using eqn (1) with timesteps of 1 year. We include a simple model for the FI as a function of the BAs and CA, $R^2 = 0.30$, to demonstrate how to track the expected change in health as a function of the BAs (details in Supplementary Material).

Having observed the central role of PhysioAge in Figure 1, we simulate the impact of a beneficial intervention administered at age 70, which instantly rejuvenates PhysioAge by 10 years (Figure 4). The intervention causes complex, delayed effects in the other BAs, including an adverse effect: a small, transient telomere shortening (all other BAs improved). This is due to the intervention effect propagating through the network. For example, Telomere worsens (shortens) for about 5 years post-intervention then recovers and ultimately improves after about 10 years post-intervention. Observe that, in contrast to the BAs, the relative FI *continuously improves* with time post-intervention. This is due to the unstable nature of the FI: which grows exponentially with age (27) due to compounding (propagating) secondary damage. Conversely, if we simulate an adverse event, say of disease, which increases PhysioAge by 10 years then we see the same effects with the sign flipped (Supplementary Figure S15). The long-term consequences of the adverse event continue to worsen the relative health (FI) of the case versus control even after the disease.

This is consistent with results from a computational network model of disease, which show the long-term FI-effect is due to secondary, compound (“propagated”) damage (39).

It is much easier to understand the effects of the intervention on PhysioAge using the natural variables, z_k , as shown in Figure 4B. All natural variables, z_k , connected to PhysioAge immediately improve upon the rejuvenation, then the z_k return to normal on the timescale set by $|\lambda_k|^{-1}$, which means that after 20 years all of the fast z_k have mostly recovered. This explains why we should primarily concern ourselves with the greatest eigenvalues, λ_1 and λ_2 , since their respective natural variables are the only ones with lasting, long-term impacts—good or bad.



Figure 3. Natural variable correlates. (A) Spearman correlations. z_1 is strongly correlated with each other BA, CA, and the FI (weakly with Telomere). (B) z_1 is strongly correlated with CA. As such, z_1 is capturing essential aging information, including CA and individual health (as estimated by the FI). (Imputed values are included only for the z_k .)

Intervening directly on the natural variables, z_k , gives a greatly simplified picture (Supplementary Figure S16). Because the z_k don’t interact with each other, the intervention pinpoints one z_k that immediately improves, whereas the other z_j are unaffected. Depending on the stability of the intervened z_k , the effect of the intervention is either gradually lost with age (stable), persists indefinitely (marginal stability/slow dimensions $|\lambda_k| \sim 1 \text{ lifespan}^{-1}$), or improves with increasing age (unstable). In the present study, interventions which improve z_1 are most desirable since they have the strongest relationship with health (Figure 3) and persist for a typical human lifespan. See Supplementary Material for other simulated interventions.

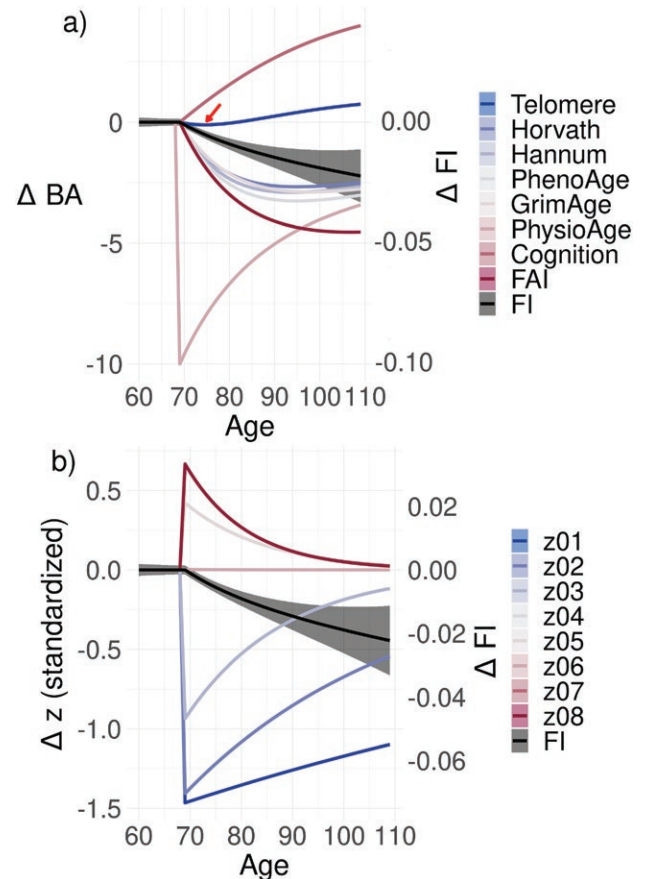


Figure 4. Simulated intervention on PhysioAge. We simulated a hypothetical intervention at age 70, which instantly rejuvenates PhysioAge by 10 years. (A) BA-picture. $\Delta BA \equiv BA_{\text{case}} - BA_{\text{control}}$. We see an immediate rejuvenation of PhysioAge at age 70 due to the intervention, whereas the remaining BAs have complex delayed effects. For example, Telomere temporarily worsens (arrow) then recovers and ultimately rejuvenates. (B) z-picture. $\Delta z \equiv z_{\text{case}} - z_{\text{control}}$. When working with natural variables, the same intervention effect is immediately spread across all z_k which are connected to PhysioAge through $\vec{b} = P\vec{z}$ (P is dense). Each z_k simultaneously responds then the z_k with fast recovery times quickly revert back leaving only the slow recovery times, z_1 and z_2 . For the FI, the relative improvement gets better over time since compound (propagated) damage is avoided by the rejuvenation (the FI is unstable). The FI has its own y scale as indicated. z_k have been standardized for convenient comparison (zero-mean, unit-variance). The sign of z_k is arbitrary due to idiosyncrasies of the eigen-decomposition, but could be aligned using their correlation with age or health. Case and control have been perfectly matched for age, sex, and stochastic effects. See Supplementary Material for other simulated interventions. Band is standard error (often smaller than line width).

Discussion

The ongoing proliferation of new BAs (biological ages) presents 2 opportunities for dynamical network analysis: (i) BAs can be used to generate network interactomes to better understand how age-related changes are naturally orchestrated and for comparison to theory, and (ii) there is an increasing need for a robust method of aggregating multiple BAs. Here, we address both opportunities. Using a dynamical model with minimal assumptions, we are able to estimate an interaction network from a collection of BAs. We can also apply eigen-analysis to the network such that we are able to generate dynamically independent aggregate BAs—natural aging variables. The slow recovery and age-related drift of these variables reflect their underlying importance in quantifying the aging process. We propose that such natural variables are the natural language to communicate the aging process, in the same manner that spectral signal analysis has come to dominate many quantitative disciplines.

The estimated interaction network, \mathbf{W} , encodes conditional dependencies across BAs of the current timestep from the previous timestep. This permits causal predictions. For example, $W_{ij} > 1$ indicates that if b_j is lower than μ_j at timepoint t_n then it will push down the i th variable (b_i) by time t_{n+1} (eqn (2)). We observed that the μ_j are large enough such that the b_j are always pushing each other and are autonomously drifting toward worse health. This drift fills the same role as a “mallostatic” drift with age (23), $\mu(t)$, but without explicit inclusion of time. High outgoing-degree nodes, such as PhysioAge and GrimAge, play a key role since they push the other BAs, such that changes to PhysioAge or GrimAge naturally propagate into the other, downstream BAs (e.g., rejuvenation). These relationships are learned from observational changes, rather than interventional (28). However, we know that individuals experience many interventions throughout their lives due to medical interventions, lifestyle changes, and stressors of living, such as disease. This suggests that our observation, \mathbf{W} , should be consistent with perturbative (small) interventions. In support, we found that we were better able to predict worsening of BAs by including causal relationships.

Using multiple BAs, we were able to estimate an informative network of interactions, which can enhance our knowledge of the aging process. In the present study, we used BAs of varying biological scales ranging from genetic (Telomere) to whole organism (FAI). We observed that PhysioAge, representing primarily cardiometabolic system changes, was the central node with outgoing arrows directly affecting all other BAs (the Telomere link was not significant). This means that changes to PhysioAge will propagate to the other BAs, eqn (2). This permits PhysioAge to drive the other BAs. We observed a weaker, but similarly well-connected effect emanating from GrimAge including feedbacks with PhysioAge. GrimAge had a strong association with mortality for this data set (similar to the FI) (9), and may represent damage. This implies that cardiometabolic, system-level dysfunction is essential to age-related changes, complemented by genetic/epigenetic damage with feedbacks between the 2 scales. Both PhysioAge and GrimAge had the strongest Spearman correlations with CA at 0.90 and 0.79, respectively (Hannum was third with 0.73). This supports the interpretation that age-related changes emerge first in those 2 variables then propagate outwards, causing the correlation with CA to drop as the information gradually attenuates through the network

connections. The key natural variable, z_1 , has a strong association with PhysioAge with important contributions from GrimAge, suggesting that well-connected nodes with feedbacks caused the weak stability, which appears to be primarily related to metabolic functioning.

These observations are consistent with theory. Systems biology informs us that metabolism is a vulnerable point due to its bottleneck (“bow-tie”) through glucose (12,22); its large number of outputs also make it well suited for propagating dysfunction—just as we observed with PhysioAge. Notably, metabolism is considered 1 of 7 “pillar” causes of aging, another is macromolecular damage—which is ostensibly captured by GrimAge (14). The hallmark theory of aging specifically includes epigenetic changes and genomic instability as 2 of 12 “hallmarks”—which GrimAge may be sensitive to—but is considerably less specific toward metabolic changes, grouping cardiometabolic changes into generic changes to intercellular communication (15). This could be an indication that the hallmark theory lacks specificity, although our suite of BAs may be similarly limited. Our suite of available BAs constrains our model to effective dynamics (26), which may differ from the oracle truth—although we can use CA as a catch-all for unaccounted degrees of freedom such as unmeasured BAs (Supplementary Material). Fortunately, BAs are becoming increasingly specific in response to demands for high-dimensional representations of aging (2), such as biological system-specific ages (11). As new BAs emerge, we can continue to use our approach to refine our understanding of the causal relationships underlying aging.

Although the network topology is informative, the dynamical behavior of the system as a whole is obfuscated by its complexity. Eigen-analysis identifies weakly stable, slow-recovery eigenvalues and associated eigenvectors, which can be used to generate aggregate biomarkers, that is, natural variables. The means of natural variables do not interact, giving them simple and intuitive dynamical behavior over time: stable natural variables simply decay toward μ_j on a timescale of eigenvalue⁻¹ according to eqn (3). Previously, we observed that across 4 data sets (2 mouse and 2 human) the dominant natural variable for predicting survival or dementia onset was either the first or second slowest eigenvalue, which were always stable but near 0—specifically they took an entire lifespan to recover (23). Ostensibly, these natural variables capture irreversible changes, such as damage. Across studies, we consistently observe that the greatest eigenvalues (slowest; most positive) are associated with the salient features of aging. In the present study these were z_1 and z_2 which had the strongest correlations with CA and were associated with health. In contrast to ad hoc approaches to finding the best aggregate BA (10,11), the natural variables are essential features of the dynamical system as a whole. As such, they do not interact and they become increasingly dominant with age, as they store information from stochastic events, making them dominate all of the BAs. The implication is that aging becomes increasingly simple and dependent on these essential natural variables with age (effectively leading to reduced dimensionality (18)). This is a consequence of the mapping $\vec{b} = \mathbf{P}\vec{z}$, which leverages redundancies in the BA representation to reduce the dimensionality. This means that aging may be lower dimension than popular theories suggest (14,15). What’s more, the dimensionality of aging may be perturbative in the sense that the first dimension provides the most important information and the subsequent dimensions

provide less and less, consistent with previous computational studies (17,18). In general, we see that the long-time dynamical behavior is dominated by the slowest network eigenvalues, representing the directions of slowest recovery and, by implication, the lowest resilience (25).

There is significant interest in using BAs to quantify the effects of antiaging interventions (1,2), with some going so far as to *define* rejuvenation by its prolonged effect on biological age (2). Our simulated interventions highlight the utility of dynamical stability (eigen) analysis in disambiguating network dynamics and identifying optimal intervention targets. Using our estimated model parameters, we were able to simulate a hypothetical intervention at age 70 that instantly rejuvenates PhysioAge by 10 years. Improvement to a BA postintervention has been observed in a number of experiments with epigenetic BAs (2) and other BAs, such as the FI (20). A consistent problem that emerges from these experiments is that health may improve along one dimension at the expense of another, such as increased tumorigenesis postrejuvenation (2), reduced visual acuity following antiaging treatment via metformin (20), and increased frailty following mTORC2 disruption (19)—which is inhibited by rapamycin treatment. Such pleiotropic effects are mediated through some biological interaction network—albeit typically an unknown one.

We observed a similar pleiotropic effect in our simulated intervention, wherein 7/8 of the BAs showed immediate or delayed rejuvenation but worsening was observed in Telomere (shortening) during a 10-year-long transient effect. The general issue is that the interaction network obscures the effects of interventions on a single BA. Such a seemingly simple intervention perturbs the network, which then adjusts all of the BAs through its connections—leading to delayed and unexpected downstream effects. This problem can be avoided by working in the natural variables, z_k . Because the mapping is linear and invertible, we can easily transform between the BA and z -pictures as needed. When working in the natural variables the intervention is greatly simplified: all stable z_j revert to the control with a timescale $|\lambda_j|^{-1}$, and all unstable z_k improve continuously postintervention with timescale $|\lambda_k|^{-1}$. This simplifies the problem since we immediately know that $|\lambda_j|^{-1}$ is the length of time that a stable z_j will differ from control, so we only need to monitor the key natural variables which exhibit high risk and low recovery: $\lambda_k \gtrsim 0$ —the remaining z_j will quickly forget the perturbation. This requires only that we determine the relevance of each z_k to health, which can be done prior to an interventional study (and may naturally compress into the lowest z_k). Identifying unstable and weakly stable natural variables, and the interventions that modulate them should be a fruitful topic of future research.

An unstable natural variable would be particularly important, although we have now failed to observe an instability in 5 data sets using the SF model, using either BAs or health biomarkers (23). An instability would lead to super-linear growth, such as are observed in the FI (27), dynamical FI (26), and specific plasma proteins (8). Since natural variables drive observed variables via $\vec{b} = \mathbf{P}\vec{z}$, unstable natural variable(s) could be of prime importance since all stable natural variables should eventually equilibrate such that all observed age-related changes are driven by the unstable natural variable(s). Furthermore, amelioration of an unstable natural variable would result in continuous life-long improvements such as we saw with the FI in our simulated interventions. However, the vast majority of health biomarkers change

linearly with age (40), and BAs are typically designed to track CA in units of time, which, by definition, increases linearly with age. The choice of units used to quantify aging may therefore play a role in determining what super-linear growth means, and therefore stability. A second issue is the effect of a population-level picture, which can mask unstable subpopulations, such as those experiencing or transitioning into chronic disease. This presents an opportunity for more powerful statistical models able to capture such individual effects. Note that a slow instability is indistinguishable from linear drift until advanced ages—where data are sparse. Our current perspective is that organisms live most of their lives in a stable regime of approximately linear decline until a tipping point is reached and nonlinear collapse ensues, quickly leading to organism death or chronic disease.

Although linear drift is reasonable population-level behavior for the BAs and z_k , a tipping point leading to super-linear behavior is necessary to make sense of terminal decline. Terminal decline occurs in biomarkers of death wherein immediately prior to death they become much worse, driving up the hazard and dropping the survival probability toward 0 over a short period of time. This effect can be seen empirically, for example, in the FI (41), cognition (42), and gait (42), which show a dramatic change in slope occurring around 3–4 years prior to death. Two dynamical phases are needed to capture this effect, such as those elucidated by the saturating-repair model of aging (43,44). In this model, age-related decline begins as a stable, approximately linear system until repair processes “saturate” and a new super-linear phase begins. This model unifies ideas of critical behavior (26,45,46) with damage and repair (47,48), which have been increasingly used within the aging modeling literature. Our prior results from mice and humans support the important predictions made by the saturating-repair model, including “mallostatic” (23): the correlation between mortality hazard and linear drift rate (43). The linear phase is characterized by a linear increase in the homeostatic steady state and is ostensibly driven by asymmetric transitions such as epigenetic methylation and accumulation of disease (49). We hypothesize that the linear phase serves to push individuals toward tolerance thresholds (tipping points) upon which a super-linear phase ensues, quickly leading to death or disease, for example, due to the saturation of repair processes.

It is important to understand that our model is of slow, linear dynamics in \mathbf{W} at timescales slower than the interval between measurements (years). Faster dynamics are pushed into the noise term, Σ . The effects of the fast dynamics are to cause biomarkers to rapidly change in values from day to day, whereas the slow dynamics estimated by \mathbf{W} are long-term changes over the course of human-equivalent years. The resilience score we estimate using \mathbf{W} represents resilience on the timescale of years, which (we believe) is a good timescale to assess aging. Nevertheless, this is in contrast to typical measurements of biological resilience, which are on the short timescale of weeks and shows a clear age-dependence (45). In contrast, we observed no clear age-dependence for the slow-resilience assessed by our \mathbf{W} -analysis, noting the large error bars (Supplementary Material). A clear age-dependent drop in the eigenvalues of \mathbf{W} could be an indication of saturating repair, but the effect is indirect and bounded at 0 (43,44), which may make it difficult to observe. The correct interpretation of \mathbf{W} -resilience is currently unclear. \mathbf{W} is capturing the long-term decline due

to aging through the longitudinal correlations it causes in the BAs. The smallest eigenvalues have the longest memories and hence they are the natural place for information regarding long term and irreversible changes to build up, making the associated eigenvectors excellent predictors of age-related health.

Our relatively modest predictive performance of worsening (AUC 0.764 ± 0.005) and explained FI variance of approximately 30%, suggest that we have only captured some of the age-related changes. Our model achieved a prediction error on the order of approximately 6 years for BA progression after an average of 3 years of natural aging, representing 65% of the variance. Model error encapsulates the net effect of 4 major sources: missing values, unobserved variables, stochasticity, and model misspecification. The missingness was particularly high in the present study, especially for the epigenetic BAs, which were more likely to be missing for older individuals (thus nonrandom), which could lead to bias—even with good imputation (35). This missingness reduces data quality and quantity, which may explain why we did not achieve statistical significance when comparing only the AUC of FullW versus SymW ($p = .1$). Second, we relied on only 8 BAs to completely predict future health including only sex as a covariate. This limits predictive power and our ability to detect causal relationships because we cannot identify causal connections from unobserved variables (50). Although it is impossible to capture all information, we would hope to find a saturation “elbow” at some larger number of BAs. Third, in addition to intrinsic stochasticity, there is substantial stochasticity owing to nonlab conditions (individual variability) and measurement noise, putting severe limits of predictability. Finally, our model is a local approximation of near-stable homeostasis and does not capture sudden changes, such as may occur due to the emergence of an underlying chronic condition. Although the model could incorporate sudden changes via μ_n , this is only possible when they are specified.

There is a trend toward increasingly specific BAs, which are compatible with nonredundant, multivariate representations (11,18). Our approach complements these representations because it provides both causal structure and generates salient aggregate features. Our approach is not limited to BAs, it works for all continuous-valued, longitudinal biomarkers. For example, emerging ‘omics data, such as dynamical changes to proteomics (8) could be a natural target for dynamical network stability analysis. We have previously applied the approach to physiological biomarkers including blood tests, body weight, blood pressure, and other generic health biomarkers (23). Although correlation analysis is the de facto standard, it has a severe shortcoming in that it estimates unconditional relationships, which therefore cannot represent a true network (they do not satisfy the standard axioms of graph theory (51) and hence their interpretation is ambiguous at best). A network link indicates a conditional relationship given all variables in the network, permitting easy and intuitive interpretation. Our approach is very general, and we suggest analysts consider using it any time they apply correlation analysis to longitudinal data. We think that the greatly enhanced interpretability is worth the modest additional computational burden. For more quantitative researchers, our linear model could easily be replaced with a more complex model, such as a deep neural network (17), and the analysis of independent natural variables could remain the

same. The key natural variables we observe may also be useful measures for secondary analysis, such as looking for early warning signs for the onset of chronic disease (52).

Our central hypothesis is that BAs provide the raw information needed to generate interaction networks, which can then be analyzed as a whole using dynamical stability (eigen) analysis. Our work highlights the utility of approaches borrowed from complexity science and systems biology. Our results are consistent with aging having several of the key features of complex systems: networks, motifs, and feedbacks, which we show play an important role in understanding age-related changes. This is direct evidence of the importance of complex and systems-level thinking (12) for furthering our understanding of aging. In summation, a simple, interpretable model of the dynamics can be leveraged to estimate the essential effects of aging, and infer the effects of perturbative interventions. We demonstrate that analyzing our fitted data gives results consistent with known theory, making it a potential path forward to operationalizing and testing various qualitative theories of aging. Far from being a curse of plenty, the proliferation of established, new, and increasingly specific BAs may be the key to quantifying and understanding the complex multidimensional changes which characterize the aging process.

Supplementary Material

Supplementary data are available at *The Journals of Gerontology, Series A: Biological Sciences and Medical Sciences* online.

Funding

This work was supported by the Natural Sciences and Engineering Research Council of Canada (grant number RGPIN-2019-05888) to A.D.R.

Conflict of Interest

None.

Data Availability

Data are publicly available from Li et al. (9). Our fitting and simulating functions using R are available online at https://github.com/GlenPr/stochastic_finite-difference_model. Exact model parameters are provided in CSV files on GitHub.

Acknowledgments

G.P. built the model and analyzed the data. A.D.R. and G.P. conceived the project. All authors reviewed the manuscript.

References

1. Campisi J, Kapahi P, Lithgow GJ, Melov S, Newman JC, Verdin E. From discoveries in ageing research to therapeutics for healthy ageing. *Nature*. 2019;571(7764):183–192. <https://doi.org/10.1038/s41586-019-1365-2>
2. Zhang B, Trapp A, Kerepesi C, Gladyshev VN. Emerging rejuvenation strategies—reducing the biological age. *Aging Cell*. Published online December 31, 2021;21(1):e13538. <https://doi.org/10.1111/acel.13538>

3. Jylhävä J, Pedersen NL, Hägg S. Biological age predictors. *EBioMedicine*. 2017;21:29–36. <https://doi.org/10.1016/j.ebiom.2017.03.046>
4. Rutledge J, Oh H, Wyss-Coray T. Measuring biological age using omics data. *Nat Rev Genet*. 2022;23(12):715–727. <https://doi.org/10.1038/s41576-022-00511-7>
5. Levine ME, Lu AT, Quach A, et al. An epigenetic biomarker of aging for lifespan and healthspan. *Aging (Milano)*. 2018;10(4):573–591. <https://doi.org/10.18632/aging.101414>
6. Lu AT, Quach A, Wilson JG, et al. DNA methylation GrimAge strongly predicts lifespan and healthspan. *Aging (Milano)*. 2019;11(2):303–327. <https://doi.org/10.18632/aging.101684>
7. Klemra P, Doubal S. A new approach to the concept and computation of biological age. *Mech Ageing Dev*. 2006;127(3):240–248. <https://doi.org/10.1016/j.mad.2005.10.004>
8. Lehallier B, Gate D, Schaum N, et al. Undulating changes in human plasma proteome profiles across the lifespan. *Nat Med*. 2019;25(12):1843–1850. <https://doi.org/10.1038/s41591-019-0673-2>
9. Li X, Ploner A, Wang Y, et al. Longitudinal trajectories, correlations and mortality associations of nine biological ages across 20-years follow-up. *Elife*. 2020;9:e51507. <https://doi.org/10.7554/eLife.51507>
10. Jansen R, Han LK, Verhoeven JE, et al. An integrative study of five biological clocks in somatic and mental health. *Elife*. 2021;10:e59479. <https://doi.org/10.7554/eLife.59479>
11. Sehgal R, Meer M, Shadyab AH, et al. Systems age: a single blood methylation test to quantify aging heterogeneity across 11 physiological systems. *bioRxiv*. Published online July 17, 2023. <https://doi.org/10.1101/2023.07.13.548904>
12. Cohen AA, Ferrucci L, Fülöp T, et al. A complex systems approach to aging biology. *Nat Aging*. 2022;2(7):580–591. <https://doi.org/10.1038/s43587-022-00252-6>
13. Kirkwood TBL. Systems biology of ageing and longevity. *Philos Trans R Soc Lond B Biol Sci*. 2011;366(1561):64–70. <https://doi.org/10.1098/rstb.2010.0275>
14. Kennedy BK, Berger SL, Brunet A, et al. Geroscience: linking aging to chronic disease. *Cell*. 2014;159(4):709–713. <https://doi.org/10.1016/j.cell.2014.10.039>
15. López-Otín C, Blasco MA, Partridge L, Serrano M, Kroemer G. Hallmarks of aging: an expanding universe. *Cell*. 2023;186(2):243–278. <https://doi.org/10.1016/j.cell.2022.11.001>
16. Pierson E, Koh PW, Hashimoto T, et al. Inferring multidimensional rates of aging from cross-sectional data. In: Chaudhuri K, Sugiyama M, eds. *Proceedings of Machine Learning Research*. Vol 89. PMLR; 2019:97–107. <http://proceedings.mlr.press/v89/pierson19a.html>
17. Farrell S, Mitnitski A, Rockwood K, Rutenber AD. Interpretable machine learning for high-dimensional trajectories of aging health. *PLoS Comput Biol*. 2022;18(1):e1009746. <https://doi.org/10.1371/journal.pcbi.1009746>
18. Pridham G, Rockwood K, Rutenber AD. Efficient representations of binarized health deficit data: the frailty index and beyond. *Geroscience*. Published online January 27, 2023;45: 1687–1711. <https://doi.org/10.1007/s11357-022-00723-z>
19. Howlett SE, Rutenber AD, Rockwood K. The degree of frailty as a translational measure of health in aging. *Nat Aging*. 2021;1(8):651–665. <https://doi.org/10.1038/s43587-021-00099-3>
20. Palliyaguru DL, Moats JM, Di Germanio C, Bernier M, de Cabo R. Frailty index as a biomarker of lifespan and healthspan: focus on pharmacological interventions. *Mech Ageing Dev*. 2019;180:42–48. <https://doi.org/10.1016/j.mad.2019.03.005>
21. Albert R, Jeong H, Barabasi AL. Error and attack tolerance of complex networks. *Nature*. 2000;406(6794):378–382. <https://doi.org/10.1038/35019019>
22. Csete M, Doyle J. Bow ties, metabolism and disease. *Trends Biotechnol*. 2004;22(9):446–450. <https://doi.org/10.1016/j.tibtech.2004.07.007>
23. Pridham G, Rutenber AD. Network dynamical stability analysis reveals key “mallostatic” natural variables that erode homeostasis and drive age-related decline of health. *Sci Rep*. 2023;13(1):1–12. <https://doi.org/10.1038/s41598-023-49129-7>
24. Ledder G. *Mathematics for the Life Sciences*. New York: Springer; 2013. <https://doi.org/10.1007/978-1-4614-7276-6>
25. Ives AR. Measuring resilience in stochastic systems. *Ecol Monogr*. 1995;65(2):217–233. <https://doi.org/10.2307/2937138>
26. Avchaciov K, Antoch MP, Andrianova EL, et al. Unsupervised learning of aging principles from longitudinal data. *Nat Commun*. 2022;13(1):6529. <https://doi.org/10.1038/s41467-022-34051-9>
27. Mitnitski A, Rockwood K. Aging as a process of deficit accumulation: its utility and origin. *Interdiscip Top Gerontol*. 2015;40:85–98. <https://doi.org/10.1159/000364933>
28. Dawid AP. Beware of the DAG! In: *Proceedings of Machine Learning Research*. PMLR; 2010:59–86. <http://proceedings.mlr.press/v6/dawid10a.html>
29. Yashin AI, Arbeev KG, Akushevich I, Kulminski A, Akushevich L, Ukraintseva SV. Stochastic model for analysis of longitudinal data on aging and mortality. *Math Biosci*. 2007;208(2):538–551. <https://doi.org/10.1016/j.mbs.2006.11.006>
30. Kojima G, Iliffe S, Walters K. Frailty index as a predictor of mortality: a systematic review and meta-analysis. *Age Ageing*. 2018;47(2):193–200. <https://doi.org/10.1093/ageing/afx162>
31. R Core Team. *R: A Language and Environment for Statistical Computing*. Published online 2021. <https://www.R-project.org/>
32. Berglund K, Reynolds CA, Ploner A, et al. Longitudinal decline of leukocyte telomere length in old age and the association with sex and genetic risk. *Aging (Milano)*. 2016;8(7):1398–1415. <https://doi.org/10.18632/aging.100995>
33. Robin X, Turck N, Hainard A, et al. pROC: an open-source package for R and S+ to analyze and compare ROC curves. *BMC Bioinf*. 2011;12:77. <https://doi.org/10.1186/1471-2105-12-77>
34. Hanley JA, McNeil BJ. The meaning and use of the area under a receiver operating characteristic (ROC) curve. *Radiology*. 1982;143(1):29–36. <https://doi.org/10.1148/radiology.143.1.7063747>
35. Pridham G, Rockwood K, Rutenber AD. Strategies for handling missing data that improve Frailty Index estimation and predictive power: lessons from the NHANES dataset. *GeroScience*. Published online February 1, 2022;44: 897–923. <https://doi.org/10.1007/s11357-021-00489-w>
36. Hardy SE, Allore H, Studenski SA. Missing data: a special challenge in aging research. *J Am Geriatr Soc*. 2009;57(4):722–729. <https://doi.org/10.1111/j.1532-5415.2008.02168.x>
37. Mak JKL, Karlsson IK, Tang B, et al. Temporal dynamics of epigenetic aging and frailty from midlife to old age. *J Gerontol A Biol Sci Med Sci*. Published online October 27, 2023;XX(XX):1–8. <https://doi.org/10.1093/gerona/glad251>
38. Mair W, Goymer P, Pletcher SD, Partridge L. Demography of dietary restriction and death in *Drosophila*. *Science*. 2003;301(5640):1731–1733. <https://doi.org/10.1126/science.1086016>
39. Tobin R, Pridham G, Rutenber AD. Modelling lifespan reduction in an exogenous damage model of genetic disease. *Sci Rep*. 2023;13(1):16304. <https://doi.org/10.1038/s41598-023-43005-0>
40. Sehl ME, Yates FE. Kinetics of human aging: I. Rates of senescence between ages 30 and 70 years in healthy people. *J Gerontol A Biol Sci Med Sci*. 2001;56(5):B198–B208. <https://doi.org/10.1093/gerona/56.5.b198>
41. Stolz E, Mayerl H, Hoogendijk EO, Armstrong JJ, Roller-Wirnsberger R, Freidl W. Acceleration of health deficit accumulation in late-life: evidence of terminal decline in frailty index three years before death in the US Health and Retirement Study. *Ann Epidemiol*. 2021;58:156–161. <https://doi.org/10.1016/j.annepidem.2021.03.008>
42. Stolz E, Mayerl H, Muniz-Terrera G, Gill TM. Terminal decline in physical function in older adults. *J Gerontol A Biol Sci Med Sci*. Published online May 6, 2023;79(1):1–10. <https://doi.org/10.1093/gerona/glad119>
43. Karin O, Agrawal A, Porat Z, Krizhanovsky V, Alon U. Senescent cell turnover slows with age providing an explanation for

- the Gompertz law. *Nat Commun.* 2019;10(1):5495. <https://doi.org/10.1038/s41467-019-13192-4>
44. Yang Y, Karin O, Mayo A, et al. Damage dynamics and the role of chance in the timing of *E coli* cell death. *Nat Commun.* 2023;14(1):2209. <https://doi.org/10.1038/s41467-023-37930-x>
45. Pyrkov TV, Avchaciov K, Tarkhov AE, Menshikov LI, Gudkov AV, Fedichev PO. Longitudinal analysis of blood markers reveals progressive loss of resilience and predicts human lifespan limit. *Nat Commun.* 2021;12(1):2765. <https://doi.org/10.1038/s41467-021-23014-1>
46. Podolskiy D, Molodtsov I, Zenin A, et al. Critical dynamics of gene networks is a mechanism behind ageing and Gompertz law. *arXiv [q-bioMN]*. Published online February 15, 2015. <http://arxiv.org/abs/1502.04307>
47. Farrell S, Kane AE, Bisset E, Howlett SE, Rutenberg AD. Measurements of damage and repair of binary health attributes in aging mice and humans reveal that robustness and resilience decrease with age, operate over broad timescales, and are affected differently by interventions. *Elife.* 2022;11:e77632. <https://doi.org/10.7554/eLife.77632>
48. Farrell SG, Mitnitski AB, Rockwood K, Rutenberg AD. Network model of human aging: frailty limits and information measures. *Phys Rev E.* 2016;94(5-1):052409. <https://doi.org/10.1103/PhysRevE.94.052409>
49. Tarkhov AE, Denisov KA, Fedichev PO. Aging clocks, entropy, and the limits of age-reversal. *bioRxiv*. Published online October 11, 2022:2022.02.06.479300. <https://doi.org/10.1101/2022.02.06.479300>
50. Eberhardt F. Introduction to the foundations of causal discovery. *Int J Data Sci Anal.* 2017;3(2):81–91. <https://doi.org/10.1007/s41060-016-0038-6>
51. Pearl J. *Causality*. 2nd ed. Cambridge: Cambridge University Press; 2009. <https://doi.org/10.1017/CBO9780511803161>
52. Olde Rikkert MGM, Dakos V, Buchman TG, et al. Slowing down of recovery as generic risk marker for acute severity transitions in chronic diseases. *Crit Care Med.* 2016;44(3):601–606. <https://doi.org/10.1097/CCM.0000000000001564>