AGING EPIDEMIOLOGY

The AccelerAge framework: a new statistical approach to predict biological age based on time-to-event data

Marije Sluiskes[1] · Jelle Goeman[1] · Marian Beekman[2] · Eline Slagboom[2,3] · Erik van den Akker[2,4] · Hein Putter[1] · Mar Rodríguez-Girondo[1]

Abstract

Aging is a multifaceted and intricate physiological process characterized by a gradual decline in functional capacity, leading to increased susceptibility to diseases and mortality. While chronological age serves as a strong risk factor for age-related health conditions, considerable heterogeneity exists in the aging trajectories of individuals, suggesting that biological age may provide a more nuanced understanding of the aging process. However, the concept of biological age lacks a clear operationalization, leading to the development of various biological age predictors without a solid statistical foundation. This paper addresses these limitations by proposing a comprehensive operationalization of biological age, introducing the "AccelerAge" framework for predicting biological age, and introducing previously underutilized evaluation measures for assessing the performance of biological age predictors. The AccelerAge framework, based on Accelerated Failure Time (AFT) models, directly models the effect of candidate predictors of aging on an individual's survival time, aligning with the prevalent metaphor of aging as a clock. We compare predictors based on the AccelerAge framework to a predictor based on the GrimAge predictor, which is considered one of the best-performing biological age predictors, using simulated data as well as data from the UK Biobank and the Leiden Longevity Study. Our approach seeks to establish a robust statistical foundation for biological age clocks, enabling a more accurate and interpretable assessment of an individual's aging status.

Keywords Biological age · Accelerated failure time models · Aging · Metabolomics

Introduction

Aging is a multifaceted physiological process marked by functional decline, leading to increased disease risk and mortality [1]. While chronological age is a significant factor in aging-related outcomes [2], there is considerable variation in disease onset and mortality timing among older individuals [3, 4]. This variation suggests that chronological age alone cannot capture the diverse aging rates among individuals. Hence, the concept of 'biological age' has been proposed, aiming to be a holistic measure reflecting one's overall position in the lifespan [5].

Individual aging, which is what biological age aims to capture, is often described using the metaphor of a clock:

✉ Marije Sluiskes
  m.h.sluiskes@lumc.nl

✉ Mar Rodríguez-Girondo
  m.rodriguez_girondo@lumc.nl

  Jelle Goeman
  j.j.goeman@lumc.nl

  Marian Beekman
  m.beekman@lumc.nl

  Eline Slagboom
  p.slagboom@lumc.nl

  Erik van den Akker
  e.b.van_den_akker@lumc.nl

  Hein Putter
  h.putter@lumc.nl

[1] Medical Statistics, Biomedical Data Sciences, Leiden University Medical Center, Leiden, The Netherlands

[2] Molecular Epidemiology, Biomedical Data Sciences, Leiden University Medical Center, Leiden, The Netherlands

[3] Max Planck Institute for the Biology of Ageing, Cologne, Germany

[4] Pattern Recognition and Bioinformatics, Delft University of Technology, Delft, The Netherlands

as if each of us possesses some latent clock that is ticking slowly but surely towards death, with a rate that varies between individuals. This clock paradigm is not new; Alex Comfort, considered a founding father of the biological age prediction field [6], used the term 'age clock' in 1969 [7]. Apparently, clocks align with our conceptual framework for aging, as evident in recent publications where the terms '(biological) age predictors' and 'aging clocks' are used interchangeably [8].

The quest for reliable and valid biological age predictors spans decades [7, 9, 10]. Where early attempts utilized low-dimensional clinical variables as 'candidate markers' (of aging), recent developments focus on high-dimensional markers, particularly epigenetic age predictors using DNA methylation data as the predictor variables and chronological age as the outcome [11, 12]. The difference between predicted age and true chronological age was found to be associated with various age-related outcomes and hence interpreted as a measure of biological aging [13]. Since then, various other chronological age-trained biological age predictors have been developed using different sources of omics [14–18].

As chronological age-trained approaches rely on an untestable assumption [19], a second generation of biological age clocks, focusing on time-to-mortality as the outcome of interest, has emerged. These mortality-trained predictors outperform first-generation chronological age-trained ones across various health conditions, physical and cognitive performance, age-related clinical phenotypes, and frailty measures [20–23]. Notable examples include PhenoAge [24] and GrimAge [25, 26]. PhenoAge, a DNAm-based predictor, predicts 'phenotypic age,' a composite of nine clinical measures associated with time-to-mortality and chronological age. GrimAge is constructed by regressing time-to-mortality on age, sex, and a set of DNAm-based surrogate markers of plasma proteins and smoking pack-years (eight in the initial version and twelve in version 2 [26]).

Despite their advancements, existing biological age predictors lack a clear operationalization of biological age, impacting their statistical rigour and evaluation. When deciding on a prediction approach, defining the estimand, the measure or quantity of interest that a predictor should predict, is crucial. The concept of an estimand guides the choice of statistical methods and techniques to use. Clearly defining what measurable variable should be predicted is especially vital in biological age prediction, where biological age is a latent, i.e. not directly observable concept. Without a precise definition, different interpretations of 'biological age' may arise. Therefore, the starting point of any biological age prediction approach should be to operationalize this latent concept into measurable variables or indicators: the estimands.

Existing biological age predictors, whether age-trained or mortality-trained, may benefit from a more thorough operationalization of the concept of biological age: currently, the estimand remains ambiguous for most of them. The absence of a formal operationalization leads to a perceived ad hoc nature in constructing predictors. One consequence of this ad hoc nature is that these predictors are constructed using statistical models that are not in line with the conceptual framework of aging-as-a-clock that is so ubiquitous within the aging field. It also poses challenges in evaluating the predictors' predictive claims, as it is unclear what measurable quantity should be captured by the predictors. Presently, biological age predictors are generally only evaluated and compared by investigating to what extent the chronological age-independent part of the prediction (usually denoted by the symbol Δ and defined as the residuals after regressing predicted biological age on chronological age) is associated with time-to-mortality and other aging-related outcomes such as common health conditions, physical and cognitive performance, age-related clinical phenotypes and frailty [20–23]. While necessary, this is insufficient; biological age predictions should not only be meaningful on a relative scale (as is the case when only evaluating Δ), but the predictions should also be directly interpretable on an (absolute) age-scale. To assess the performance of biological age predictors on this characteristic, additional evaluation measures need to be considered.

Given the gaps and limitations of current biological age predictors, namely the lack of a proper operationalization, misalignment with the aging-as-a-clock paradigm, and challenges in evaluation, the aim of this paper is threefold. Firstly, we propose a clear operationalization of biological age. Secondly, we present the *AccelerAge* framework: an new approach to predict biological age based on lifespan, that follows directly from the proposed operationalization of biological age. The AccelerAge framework is based on Accelerated Failure Time (AFT) models [27]. Unlike current second-generation prediction approaches, which rely on Proportional Hazard (PH) models [28], AFT models model the effect of candidate markers of aging on one's survival time directly, aligning with the prevalent clock metaphor in aging research. The appeal—but underuse—of AFT models in the context of aging research has been noted before [29], but not in the context of biological age prediction. Thirdly, we present two new evaluation measures for biological age predictions, which directly evaluate the biological age predictions themselves. To illustrate our framework and our novel evaluation measures, we build predictors based on the AccelerAge framework and compare their performance to a predictor similar to GrimAge, currently considered the best-performing. Evaluation is conducted using simulated data, as well as data from the UK Biobank and the Leiden Longevity Study. With this we hope to contribute to a more solid statistical foundation for biological age clocks.

Biological age operationalization

Biological age should be a measure of one's position in their total lifespan or healthspan [5]. One of the earliest papers discussing biological age, from Benjamin [9], already defined biological age as such: "Biologic age may be defined once more through an example: A man (white, American, of upper cultural level) who has lived seventy years has a life expectancy of about nine years. He has, therefore, a chance to celebrate his seventy-ninth birthday. These are statistical figures based on national mortality reports of many thousands of cases. The individual is not considered except as part of the population. The examination of this man shows that he has a favorable heredity, sound organs, and that he functions like a much younger man. The age estimation (to be described later on) makes him 15 years younger, that is to say, his condition corresponds to that of a man of 55, which would be his approximate biologic age."

This idea to define biological age through life expectancy and a comparison with some reference population is intuitively appealing. However, a formal operationalization is required, such that a proper prediction approach can be decided upon, resulting in a predictor that predicts this operationalized concept. Our proposal for a formal mathematical definition of biological age is given below.

Denote by B biological age, by C chronological age, by X a (set of) true marker(s) of biological age (defined as markers that, conditional on C, are associated with B) and by T age-at-death. Under our proposed operationalization, B is defined through some function of residual life. We consider mean residual life (mrl): define mean residual life at age t as $\mathrm{mrl}(t) = E(T - t | T > t)$. Now we can define biological age as follows: individual i, with chronological age $C = c_i$ and marker value $X = x_i$, has $B = b_i$ if

$$E(T - c_i | T > c_i, X = x_i) = E(T - b_i | T > b_i). \tag{1}$$

This can also be written as: $B = b_i$ if

$$\mathrm{mrl}(c_i | x_i) = \mathrm{mrl}(b_i). \tag{2}$$

Hence, we assume that biological age is closely related to expected residual life, and is defined with respect to some reference population: given a prediction of $\mathrm{mrl}(c_i | x_i)$ and a population life table of some reference population, biological age is determined by checking which chronological age within the population corresponds to a mean residual life value of $\mathrm{mrl}(c_i | x_i)$. For example, a heavy smoker aged 50 might have a life expectancy of 20 years given his marker values X_i, while in the general male population a life expectancy of 20 years corresponds with an age of 57. Then the heavy smoker's biological age is defined as 57. This is visualized in Fig. 1.

This operationalization hence suggests a two-step approach for prediction of B: the first step is to obtain $\widehat{\mathrm{mrl}}(c_i | x_i)$ by regression-based estimation of $\mathrm{mrl}(t | X)$; the second step is to obtain $\hat{b}_i$ using $\mathrm{mrl}(t)$, which is generally taken from some external source.

Under this operationalization, someone's biological age is per definition determined with respect to some reference population. This allows for a meaningful and direct interpretation of predictions: a biological age of 50 means that someone has "the same life expectancy as the average person with chronological age 50 *within the reference population*". If the reference

Fig. 1 Illustration of step 2 of our biological age operationalization: going from a mean residual life prediction to a biological age prediction. The black line denotes mean residual life at chronological age t within some reference population. Someone with an estimated mean residual life of 20 years has a corresponding biological age of 57

population changes, someone's mean residual life prediction will not change, but the resulting biological age prediction will. A possible disadvantage of this is the flip side of the same coin: choosing a reference population might not always be straightforward and/or the appropriate life table might not always be available. However, this can be circumvented by using the training data set itself as a reference population, if it is large enough and covers a wide enough chronological age range.

For convenience' sake we have considered mean residual life (mrl) here, but one could also use some other function of residual life, e.g. median residual life (medrl), without loss of generality. When considering time-to-mortality as the outcome of interest, we have found the differences between mrl and medrl to be negligible. The choice for a particular function of residual life can be based on practical arguments: e.g. availability of a corresponding life table (most life tables provide mrl) and computational speed (medrl is faster to compute).

There exist several statistical models that can be used to obtain conditional residual life estimates (e.g. $\widehat{\text{mrl}}(t|X)$ or $\widehat{\text{medrl}}(t|X)$) using time-to-event data. In this paper we will do so using Accelerated Failure Time models, using chronological age as the timescale t, which we believe to be a natural fit to the 'aging as a clock' paradigm. In contrast, for the Grim-Age predictor the more frequently used semi-parametric Cox Proportional Hazards model is used, using time-on-study as the timescale t. We describe both models in more detail in the next section.

Proportional hazards and accelerated failure time

When working with survival data, two commonly encountered regression models are Proportional Hazards (PH) and Accelerated Failure Time (AFT) models. They differ in the assumption on how predictor variables act on one's survival. In this section we describe both models in detail.

Proportional Hazard (PH) models, due to Cox [28], are the most commonly used models for survival analysis. In this model, survival is modeled through the hazard function $h(t)$, also known as the instantaneous failure rate. PH-based models assume that the exponent of the linear predictors has a multiplicative effect on the hazard:

$$h(t|X) = h_0(t) \times \exp(\beta^T X), \tag{3}$$

where $h_0(t)$ denotes the baseline hazard. This implies the effect of the linear predictors on the survival curve $S(t)$ is given by:

$$S(t|X) = S_0(t)^{[\exp(\beta^T X)]}, \tag{4}$$

where $S_0(t)$ denotes the baseline survival. The semi-parametric Cox PH model is the most frequently encountered PH model; it is semi-parametric because it does not make assumptions about the specific form of the baseline hazard function, only about the effect of covariates on this hazard.

AFT models provide an alternative to PH models for the modelling of survival data [27]. In contrast to PH models, the AFT approach models survival times directly. AFT models assume that the effect of covariates on the baseline survival curve is to shrink or stretch this curve, i.e. accelerate or decelerate it by a factor $\exp(\beta^T X)$:

$$S(t|X) = S_0(t \times \exp(\beta^T X)). \tag{5}$$

In contrast to PH models, which are often semiparametric (i.e. no parametric distribution is assumed for the baseline hazard $h_0(t)$), it is common to fit parametric AFT models, which assume that the survival times T follow a known parametric distribution. In principle parametric models are more restrictive than semi-parametric models. However, for models where the event of interest is known to follow a certain distribution, it can be an advantage. For adult mortality this is the case: it has long been known that adult mortality (of many different species, amongst which humankind) is accurately described by the Gompertz distribution [30], with baseline hazard function $h_0(t) = a \exp(bt)$. A lack of fit has been reported for the extreme old, known as 'late-life mortality deceleration'—this phenomenon was in fact already reported by Gompertz himself [31]—but others have found this deceleration to be negligible until over 100 years of age [32]. A more recent study concluded that this mortality deceleration is notoriously difficult to prove [33]. Although it is sometimes stated that the Gompertz distribution can only be parameterized as a PH model, this is not correct: it can also be parameterized as an AFT model [34]. This is illustrated in section 1 of the Supplementary Information.

When modelling survival data, a choice has to be made about the appropriate timescale t. If subjects are followed from some well-defined event, e.g. randomization in a clinical trial, the relevant timescale is time-on-study and all subjects enter at time $t_{start} = 0$. In the context of cohort studies, however, it has long been argued that chronological age is the preferred timescale [35, 36]. Nevertheless, the PH model used in the construction of the GrimAge predictor uses time-on-study at the timescale t (we refer section 2 of the Supplementary Information for more details).

The appeal of AFT models in the context of biological age prediction is that the factor $\exp(\beta^T X)$ in Eq. (5) can be directly interpreted as an individual aging rate. If $\exp(\beta^T X)$ is greater than 1, a subject experiences accelerated aging: time t is multiplied by $\exp(\beta^T X)$. If $\exp(\beta^T X)$ is smaller than 1, a subject experiences decelerated aging. AFT models therefore tie in nicely with the clock paradigm, as they provide an

intuitive measure of accelerated aging (a faster ticking clock) or decelerated aging (a slower ticking clock).

In Fig. 2 we visualize the effect of covariates on the hazard rate and the survival curve for both AFT and PH models. In the top row, the black line represents the baseline survival curve. In the bottom row, the black shade represents the corresponding baseline hazard. The artificial plateau in midlife (where no one dies, i.e. the hazard is zero and hence the survival curve remains constant) was added to more clearly illustrate the difference between AFT and PH models. The topleft panel shows the effect on the survival curve for someone with 'beneficial' covariates under the PH model: i.e. $\exp(\beta^T X)$ is negative. At every point in time, this person experiences a lower hazard, because the baseline hazard is multiplied by a factor $\exp(\beta^T X)$, in line with Eq. (3). Hence, at every given age baseline survival is shifted upwards, because it is raised to the power $-\exp(\beta^T X)$: it is as if this person is protected by some shield. The location of the plateau in midlife, however, does not change. The topright panel shows the effect on survival for someone with beneficial covariates under the AFT model. Now the survival curve is stretched out in the horizontal direction, in line with Eq. (5): the biological age clock of this person is ticking a factor $\exp(\beta^T X)$ slower than the clock of the baseline population. As a result this person also experiences the hazard-free period in midlife at a later chronological age. The AFT model is hence a more natural fit to the aging-as-a-clock concept than the PH model.

The AccelerAge framework

In this section we introduce our new statistical framework for biological age prediction, using an Accelerate Failure Time model with chronological age as the timescale t. This framework is in line with our suggested operationalization for biological age as given in section Biological age operationalization: i.e. biological age is based on (mean) residual life and determined relative to some reference population. This suggests a two-step approach for prediction of B:

1. Get $\widehat{\text{mrl}}(c_i|x_i)$ by regression-based estimation of $\text{mrl}(t|X)$;
2. Obtain $\hat{b}_i$ using $\text{mrl}(t)$ (generally available from some external source) so that $\text{mrl}(\hat{b}_i) = \widehat{\text{mrl}}(c_i|x_i)$.

Fig. 2 Illustration of the effect of markers on baseline hazard and survival under the assumption of Proportional Hazards (left panels) and Accelerated Failure Time (right panels)
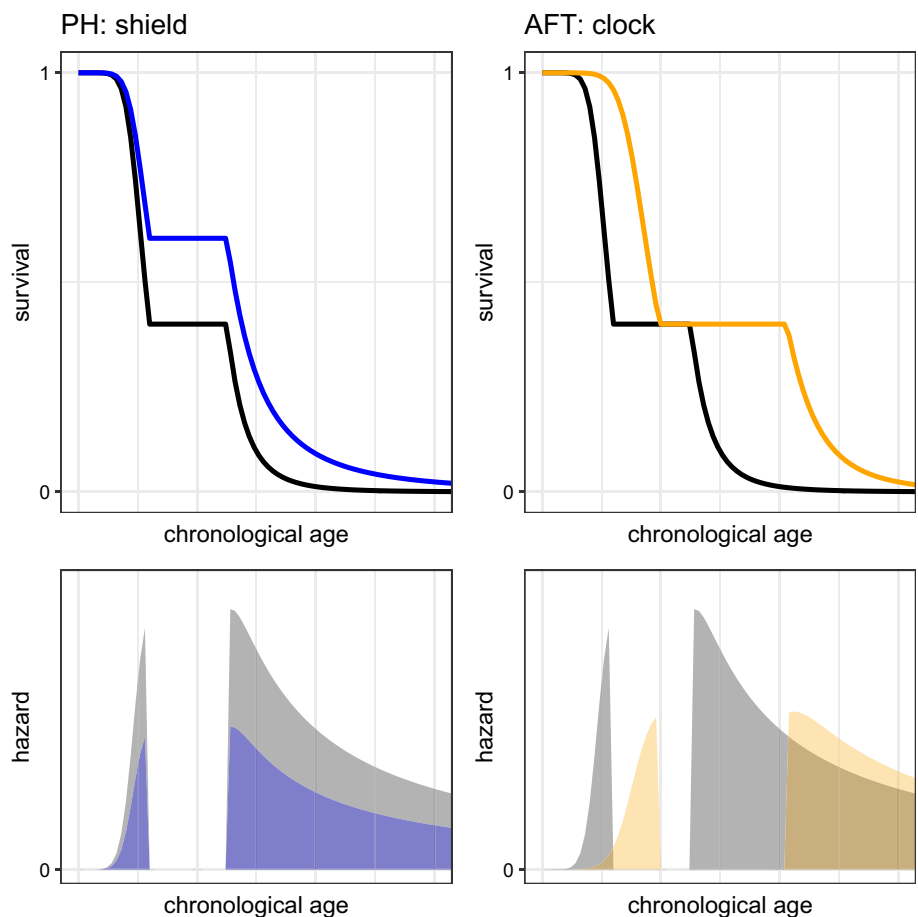
The first step in arriving at a prediction for biological age is to define a model for mean residual life mrl($t|X$), given that a subject has survived until time t. Here, X denotes a (set of) true marker(s) of aging. We choose an approach via the survival function. In terms of the survival function $S(t)$, mrl($t|X$) can be expressed as:

$$\text{mrl}(t|X) = \int_t^\infty \frac{S(u|X)}{S(t|X)} du. \tag{6}$$

How exactly X affects $S(t|X)$ depends on the underlying model. Since we assume the AFT assumption to hold, this relationship is given by Eq. (5).

For step 2, one must have access to the population life table of some reference population, in order to translate the mean residual life prediction to a biological age prediction. We suggest to use the life tables as provided by the national statistics office of the country where the data was gathered. Alternatively, in some cases the data set used to fit the model in step 1 could also be used to construct a life table.

We call this approach the AccelerAge framework, to emphasize its close link with Accelerated Failure Time models. We chose the term 'framework'—in the sense of a conceptual structure—to emphasize the fact that 'AccelerAge' refers both to the operationalization of biological age as well as to the modelling approach that is followed to construct a predictor. Hence, we distinguish between a *framework* (the structure used to define biological age), a prediction *approach* (all modelling steps required to arrive at a prediction) and a *predictor* (a fitted version of a particular framework, fitted on a particular dataset, following all steps of the modelling approach).

In this paper we deliberately present a framework instead of a predictor: our AccelerAge framework can be applied to any type of time-to-event data using any kind of predictor variables. The AccelerAge framework is based on a proper operationalization of biological age, is in line with the 'aging as a clock' paradigm and its predictions are directly interpretable on an age-scale. The names 'Grim-Age' and 'PhenoAge' are generally considered to refer to predictors: these names refer to a specific fitted model, which can be applied to a specific set of predictor variables (DNA methylation data) to arrive at a prediction. Nevertheless, in principle the modelling approaches that were followed to construct the original GrimAge and PhenoAge can also be used in other settings (e.g. fitted on different candidate predictor variables) to produce 'GrimAge-type' and 'PhenoAge-type' predictors. An example that stayed close to the original is GrimAge2, which was constructed using the exact same modelling approach as was used for the original GrimAge predictor, but starting from a slightly larger set of candidate predictor variables [26].

Evaluation measures

The standard approach in the evaluation of biological age clocks is to check whether the chronological age-independent part of a biological age prediction $\hat{\Delta}$ (also known as the 'age acceleration' and obtained by regressing $\hat{B}$ on C) is associated with mortality and other aging-related outcomes, such as frailty or cardiovascular disease. However, the biological age prediction $\hat{B}$ should not only be meaningful in relation to some other time, which is what is done when evaluating $\hat{\Delta}$: $\hat{B}$ itself should also be meaningful. In this section we therefore introduce two new evaluation measures of biological age predictions which evaluate this: discrimination and calibration. Both are routinely evaluated aspects of (clinical) prediction models based on survival data [37].

To check to what extent individuals with a higher predicted biological age indeed die sooner—in other words, to what extent the predictor is able to discriminate—we propose to consider the concordance (also called C-index) of a predictor. The concordance can be interpreted as the fraction of pairs in the data where the observation with the higher observed residual life also has the lower biological age. We use Uno's C-index, which does not depend on the study-specific censoring distribution [38].

To check to which extent the biological age predictions are on the proper scale—in other words, to what extent the predictor is well-calibrated—we propose to use calibration plots. For any biological age prediction, it is possible to obtain the corresponding X-year mortality probability for this age using the population life table. Individuals can then be grouped based on their predicted X-year mortality probability in N equally sized groups. Per group, the mean predicted mortality probability can be compared with the true observed mortality rate within this group. If the predictor is well-calibrated, these correspond closely.

Simulation study set-up

We conducted a simulation study to check the predictive performance and behavior of predictors fitted using the AccelerAge framework and of a predictor fitted using the same approach as was used to construct GrimAge. GrimAge is currently considered the best-performing biological age predictor, as the age-independent part of GrimAge predictions has been found to be associated with more aging-related outcomes than the age-independent part of Pheno-Age-predictons [20, 22]. We compared the performance of these predictors under various realistic circumstances. In simulation studies, predetermined parameters govern the

data-generating process, which ensures a known underlying "truth" (i.e. true biological age) and facilitates easier evaluation and comparison.

We generated data under a similar study design as that of real data sets used to train longitudinal biological age clocks, namely prospective cohort studies: people enter the study at a random chronological age C and are then followed-up over time. For each individual in our simulated data set we generate their marker values X, their age-of-death T (the distribution of which depends on the chosen baseline) and the chronological age C at which we would include them in our study. By excluding individuals for which $T < C$, we mimic the selection process that also takes place in real prospective cohort studies: people who have already died cannot be observed.

Data was generated assuming two true markers X_1 and X_2, constant over each person's lifetime and following a standard normal distribution at birth. Although the assumption that markers are constant over one's lifetime is not a realistic one (most omics-based markers change over the course of a lifetime), it is in line with the original approach followed to construct the GrimAge predictor, in which the the composite markers are age-adjusted before inclusion in the PH model.

We considered three data-generating mechanisms: one in which the baseline hazard $h_0(t)$ follows a Gompertz distribution and the PH assumption holds (referred to as Gompertz-PH), one in which $h_0(t)$ follows a Gompertz distribution and the AFT assumption holds (referred to as Gompertz-AFT) and one in which $h_0(t)$ follows a Weibull distribution (referred to as Weibull), which per definition can be parameterized as both an AFT and a PH model at the same time. We included the Gompertz scenarios because the Gompertz distribution is known to accurately describe (human) lifespan. We included the Weibull scenario to illustrate the effect of a misspecified baseline when fitting a parametric AFT model.

We independently generated observations of markers X_1 and X_2 and chronological age C as follows:

- $X_1 \sim N(0, \sigma_1^2)$ — biomarker, constant over one's lifetime;
- $X_2 \sim N(0, \sigma_2^2)$ — biomarker, constant over one's lifetime;
- $C \sim U(c_{min}, c_{max})$ — chronological age at which individual would enter the study.

We used the following parameter values: $\sigma_1 = \sigma_2 = 1$, $c_{min} = 20$ and $c_{min} = 80$.

Next we generated age-at-death T under each of the three data-generating mechanisms. For a given individual i, under the PH-assumption, age-at-death t_i given X_i (where X_i here denotes a vector of x_{i1} and x_{i2}) can be drawn as follows:

$$t_i = H_0^{-1}\left[\frac{-\log(U)}{\exp(\beta^T X_i)}\right], \qquad (7)$$

where U follows a uniform distribution on the interval from 0 to 1. Under the AFT-assumption, t_i given X_i can be generated as follows:

$$t_i = \frac{H_0^{-1}[-\log(U)]}{\exp(\beta^T X_i)}. \qquad (8)$$

We provide the derivation of Eqs. (7) and (8) in section 3 of the Supplementary Information. The expression for $H_0^{-1}(t)$ depends on the chosen baseline distribution: Gompertz for the first two scenarios, Weibull for the third. We chose the parameters of these distributions such that the resulting event times approximately resembled human lifespan: for Gompertz, $a = \exp(-9)$ and $b = 0.085$ (where the baseline hazard is given by $h_0(t) = a\exp(bt)$), for Weibull, $\lambda = 34^{-10}$ and $\nu = 8$ (using the operationalization as given in Bender et al. [39], where $h_0(t) = \lambda\nu t^{\nu-1}$). For Gompertz-PH, $\beta = (0.3, 0.3)$, for Gompertz-AFT, $\beta = (0.05, 0.05)$ and for Weibull, $\beta = (0.35, 0.35)$. These coefficients cannot be directly compared between the three scenarios, since on a PH-scale the interpretation of an effect size is different than on an AFT-scale, but they were chosen such that the resulting age-of-death distribution was comparable between the three scenarios, as can be seen in Fig. 3

If for a given individual age-at-death t_i was smaller than chronological age c_i, this individual is not observed, since he or she has already died at the age that we otherwise would have observed them. Those cases were removed. For the remaining individuals, we determined their survival curve $S_i(t)$ via Eqs. (4) or (5).

We consider median residual life (medrl) instead of mean residual life (mrl) because medrl is considerably faster to compute than mrl, since there is no integration involved. A pilot simulation was conducted considering both mrl and medrl: results were very similar under the settings considered in this simulation. For each individual in the data set we determined $\widehat{medrl}(t = c_i|X_i)$ as follows. First, determine the time $t = t_{i,med}$ at which survival is half the current value $S_i(t = c_i)$. Next, subtract c_i from $t_{i,med}$ to obtain expected median residual life medrl$(c_i|X_i)$.

The final step in simulating biological age B involves a population life table for median residual life, medrl(t). The life tables were constructed using the true parameter values of the three different data-generating mechanisms. Finally, for each subject i, we determined their biological age by checking in the population life table for which t the median residual life prediction of a particular individual i, $\widehat{medrl}(c_i|X_i)$, is equal to the population's medrl(t).

We assumed individuals were followed-up for a period of 20 years. So if $(t_i - c_i) > 20$, the age-of-death of this

Fig. 3 Histogram of the ages-of-death (generated at birth) for different quantiles of the linear predictor for each of the three data-generating mechanisms considered

individual is censored and he/she is observed until age $c_i + 20$. If $(t_i - c_i) < 20$, this individual is observed until their age-of-death t_i. Figure 4 contains plots of chronological age against biological age for a realization of each of the three data-generating mechanisms described.

We generated training data sets of sizes $n_{obs} = 500$, 2500, 5000, 7500 and 10,000 under each of the three data-generating mechanisms (Gompertz-AFT, Gompert-PH and Weibull) and generated test sets of size $n_{test} = 5000$.

Three different prediction approaches were used to fit three different biological age predictors on the simulated training datasets: (1) a predictor based on the Acceler-Age framework with a Gompertz baseline (AccelerAge-Gompertz), (2) a predictor that, like the AccelerAge framework, works via residual life and life tables but assumes proportional hazards (PH-semipar) and (3) a predictor constructed following the approach that was used to construct the original GrimAge predictor (GrimAge-type). An elaborate description of the approach we took to fit this predictor can be found in section 2 of the Supplementary Information. Note that of these three prediction approaches, GrimAge-type is the only one that uses

time-on-study as timescale t and relies on an ad hoc transformation to an age scale.

AccelerAge-Gompertz is correctly specified for the Gompertz-AFT data-generation mechanism. PH-semipar is correctly specified for the Gompertz-PH data-generating mechanism. Since PH-semipar is semiparametric, it is also correctly specified for the Weibull data-generation mechanism. GrimAge-type is not based on an underlying definition or operationalization of biological age, so it is not clear under which data-generation mechanism it would be correctly specified. But since it uses a Cox PH regression, it is to be expected that it will do better under the Gompertz-PH and Weibull data-generating mechanisms than under the Gompertz-AFT data-generating mechanism.

We use root-mean-square error (RMSE) as the performance measure: $\text{RMSE} = \sqrt{\frac{1}{n_{obs}} \Sigma_{i=1}^{n_{obs}} \left(\hat{b}_i - b_i \right)^2}$, where for individual i $\hat{b}_i$ denotes predicted biological age and b_i denotes true biological age. The simulation study sample size is $n_{sim} = 200$ (i.e. 200 simulated data sets of size 500, 2500, 5000, 7500 and 10,000 for the three different data-generation mechanisms is $200 \times 4 \times 3 = 2400$ simulated

Fig. 4 Plots of chronological age against (true) biological age for a simulated dataset of size $n_{obs} = 5000$ for each of the three data-generating mechanisms considered
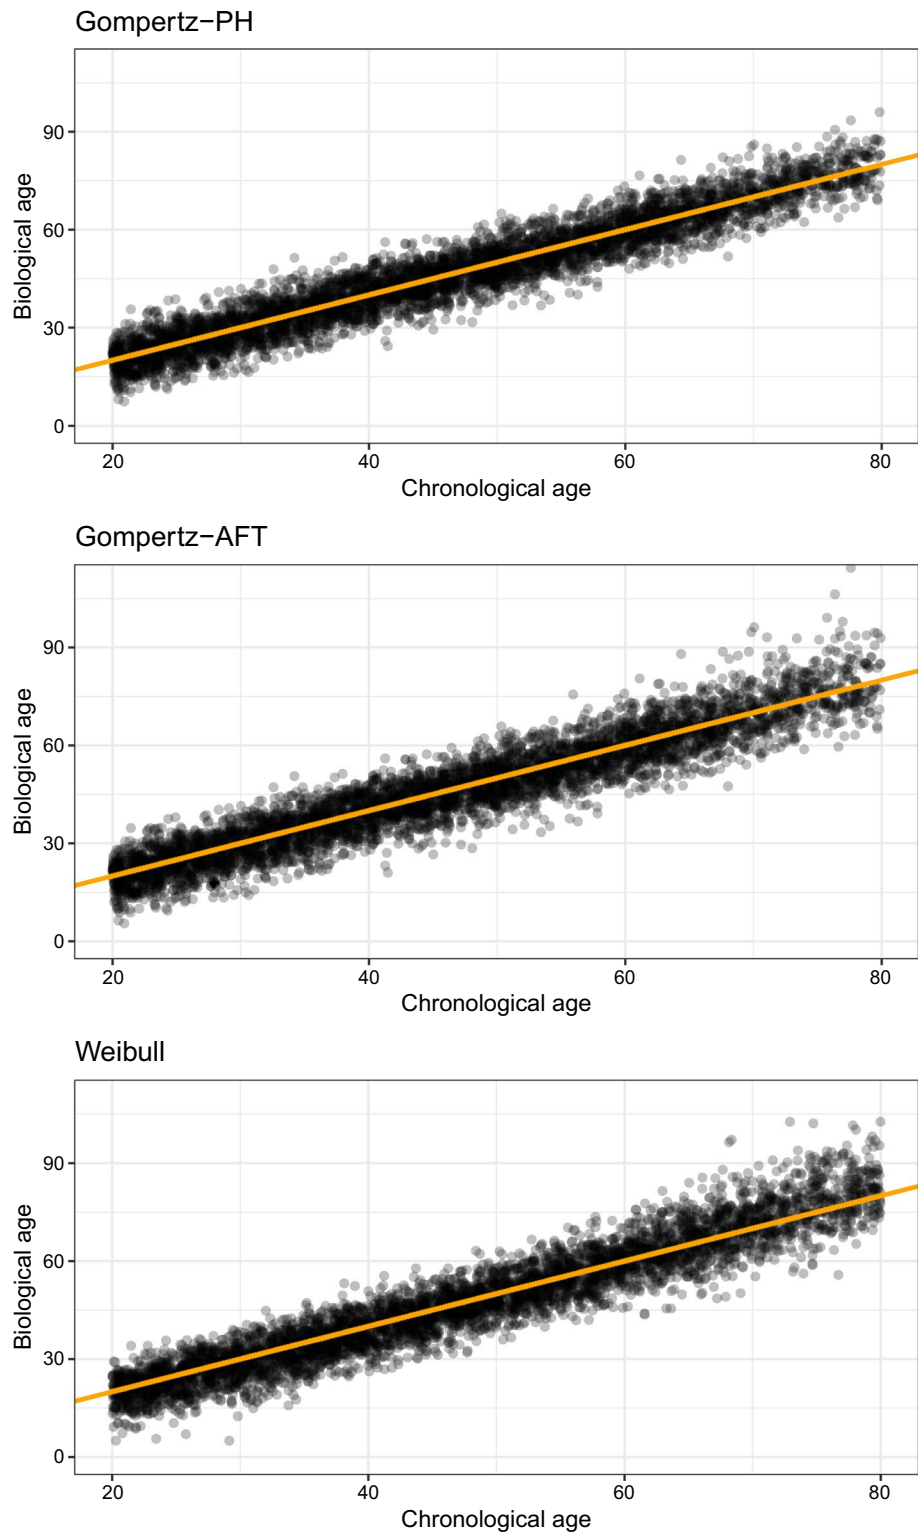
training data sets in total). In a few cases the Gompertz AFT model could not be fitted on the simulated data set, as it is a numerically delicate model to fit: an error was thrown that the Hessian was singular or that the model did not converge. This was the case for 9 of the 2400 simulated training data sets. We left those data sets out. For each data-generation mechanism and for each n_{obs}-size we report the average RMSE over the n_{sim} repetitions.

All analyses were performed using R version 4.1.0. The parametric AFT models were fitted using the eha package version 2.10.3 [40].

Simulation study results

The results for the Gompertz-PH, Gompertz-AFT and Weibull data-generating mechanisms are presented in Figs. 5, 6 and 7, respectively.

Fig. 5 Performance of the three different biological age predictors in terms of the root-mean-square error under the Gompertz-PH data-generating mechanism. Results are reported for data sets of varying sizes (n_{obs} = 500, 2500, 5000, 7500 and 10,000) as the average root-mean-square error over a simulation sample size of $n_{sim} = 200$
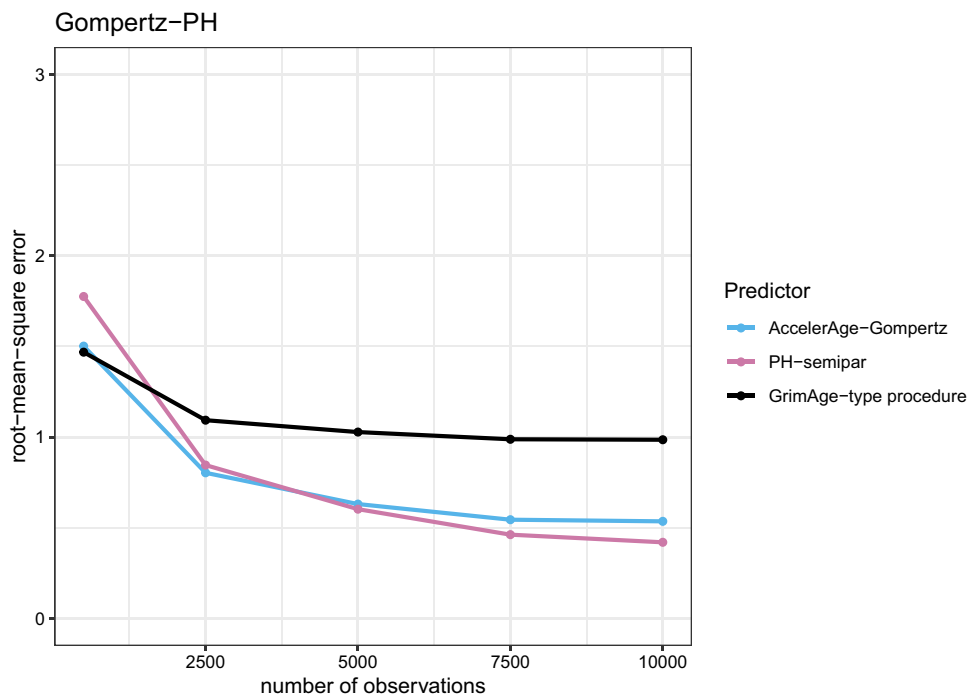
Fig. 6 Performance of the three different biological age predictors in terms of the root-mean-square error under the Gompertz-AFT data-generating mechanism. Results are reported for data sets of varying sizes (n_{obs} = 500, 2500, 5000, 7500 and 10,000) as the average root-mean-square error over a simulation sample size of $n_{sim} = 200$

Fig. 7 Performance of the three different biological age predictors in terms of the root-mean-square error under the Weibull data-generating mechanism. Results are reported for data sets of varying sizes ($n_{obs} = 500$, 2500, 5000, 7500 and 10,000) as the average root-mean-square error over a simulation sample size of $n_{sim} = 200$

Figure 5 shows that under the Gompertz-PH data-generating mechanism, GrimAge-type does best if the training data sample size is small ($n_{obs} = 500$): the RMSE is lowest. This is likely due to the fact that GrimAge-type uses a Cox PH model to estimate the effect of candidate markers on mortality, which matches the used data-generating mechanism here. GrimAge-type does not need to estimate the population baseline survival $S_0(t)$ to arrive at biological age predictions due to its ad hoc transformation of the linear predictors. The other predictors do: 500 observations in a training data set (of which a substantial number censored) apparently do not suffice to properly estimate baseline survival. This is advantageous to GrimAge-type. However, whereas the RMSE of the other predictors keep decreasing with increasing training dataset size, GrimAge-type's performance stops improving. AccelerAge-Gompertz and PH-semipar perform similarly and outperform GrimAge-type when the size of the training data is larger than approximately 1,500 samples. Although PH-semipar is correctly specified under this data-generating mechanism, this predictor also needs enough events to properly estimate $S_0(t)$. And indeed, as n_{obs} increases, PH-semipar eventually outperforms AccelerAge-Gompertz. Even though AccelerAge-Gompertz is misspecified under this data-generating mechanism, its performance is still quite good. This can likely be attributed to the fact that it assumes the correct underlying baseline distribution and the effect sizes β are relatively small.

In Fig. 6 it can be seen that under the Gompertz-AFT data-generating mechanism, the corresponding correctly specified predictor (AccelerAge-Gompertz) performs best

from the start. GrimAge-type again does okay for smaller training data sets but here its performance also quickly stops increasing.

Figure 7 shows that under the Weibull data-generating mechanism, GrimAge-type performs worst of all predictors with a large margin. This is likely due to the fact that the Cox PH regression in GrimAge's second step uses time-on-study as the timescale instead of chronological age. This affects its performance. This mismodeling was less of an issue for the Gompertz-PH and Gompertz-AFT scenarios in Figs. 5 and 6, because the Gompertz distribution belongs to the so-called 'exponential family'. For this family mismodeling of time in a Cox PH model does not matter (for an elaborate discussion, see Thiébaut and Bénichou[36]). The performance of the AccelerAge predictor and PH-semipar is in line with expectations: for larger data sets, the correctly specified PH-semipar performs best, closely followed by AccelerAge-Gompertz.

Real data illustration

In this section we evaluate and compare the performance of a predictor fitted using our newly proposed AccelerAge framework and a predictor fitted using the GrimAge framework on real data. We use data from the UK Biobank (UKB). The UKB is a large population-based prospective cohort study. Between 2006 and 2010, more than 500,000 participants aged 37–73 were recruited from different sites across the United Kingdom. Participants' health is being

followed long-term. The study contains extensive phenotypic and genotypic detail about its participants, including biological measurements and biomarker values, and longitudinal follow-up for a wide range of health-related outcomes, provided through a linkage with medical and health records [41]. Participants are free to withdraw from the UK Biobank at any time and request that their data is no longer used.

We use blood-based metabolic biomarkers as predictor variables. The metabolic biomarkers were quantified using the high-throughput nuclear magnetic resonance (NMR) targeted metabolomics platform of Nightingale Health Ltd. (Helsinki, Finland), known for its high repeatability over time and absence of batch effects [42, 43] Per EDTA plasma sample, 249 metabolic measures are provided (of which 168 in absolute concentration units and 81 ratios). The majority of these biomarkers relate to lipoprotein metabolism. In the UKB the first tranche of NMR-metabolomics data became available in March 2021, for more than 120,000 samples (from approximately 118,000 participants at baseline and 4,000 at repeat assessment, of which 1,500 at both) [44]. We only included subjects with measurements at baseline. Metabolic variables for which measurements were missing for more than 1 percent of all samples were excluded (excluding 7 of 249). Subjects with missing measurements in the remaining 242 metabolic variables were also excluded (excluding 1,647 of 104,296). This left a sample of size N = 102,649, of which 47,293 men and 55,356 women. Mean age at recruitment was 56.3 years (sd = 8.1, IQR = 50–63; see Figure 1 of the Supplementary Information for a histogram). Table 1 of the Supplementary Information contains information on the prevalence of various (chronic) diseases among the included participants.

The absolute concentrations of the metabolic variables measured at baseline are known to be 5–10% diluted in the UKB data [45]. To still allow for validation of our fitted biological age predictors in an external dataset, we decided to scale and center each metabolic variable prior to analysis.

Prospective mortality (i.e. time-to-mortality) is the outcome of interest. Follow-up data until November 2021 was available. Median follow-up time was 13.3 years (IQR = 12.5–14.0). During follow-up 7,629 participants died. No participant was followed for more than 16.9 years. Since at inclusion no participant was older than 71, the population survival curve of the participants is only well-defined from age 40 to age 85. That the population survival curve is not well-defined at its right tail has significant negative consequences for the semiparametric PH-semipar predictor considered in the simulation study. As there is no information in the data on how the right tail of the survival function looks like in reality, these predictors cannot properly estimate baseline survival $S_0(t)$. We therefore decided to exclude PH-semipar from this real data illustration. This leaves two predictors: one based on our our AccelerAge Gompertz framework fitted on the metabolic biomarkers (referred to as metabo-Accelerage-Gompertz) and one based on the GrimAge approach, as described in the previous section, but now fitted on the metabolic biomarkers (referred to as metabo-GrimAge).

To construct population life tables, necessary for the residual-life based biological age prediction approaches, we

Fig. 8 Comparison of survival curves of the included UKB participants, stratified by sex, and of the UK general population (the period life table of 2018–2020), stratified by sex, as provided by the Office for National Statistics. The population survival curves were scaled such that they only start decreasing from age 40 onward, to avoid an unfair comparison due to the immortal time bias present in the UKB data

used data from the United Kingdom's Office for National Statistics [46]. When comparing survival in the UKB population with that of the general population, it becomes apparent that the UKB participants on average live longer. This is shown in Fig. 8, using the most recent period life tables (2018–2020), stratified by sex, from the Office for National Statistics as a comparison. We scaled these curves such that they only start decreasing from age 40 onward, to avoid an unfair comparison due to the immortal time bias present in the UKB data, as no participant was included before age 40. This figure also illustrates the limited age range of the UKB participants' survival curves. We also compared survival of our subset of UKB participants, which consisted of those with metabolic biomarker measurements, with that of the entire UKB sample, but no difference in survival could be noted (see Figure 3 of the Supplementary Information). The fact that UKB participants are not representative of the general population, but seem to belong to a healthier subset, has been noted before; see e.g. Fry et al. [47].

The fact that we use sex-stratified population survival curves means that a man and a woman who have the same predicted mean residual life, will have a different predicted biological age: women live longer, so the woman's biological age will be higher.

Given the large sample size available, we randomly selected 60,000 individuals to serve as the training data set and the remaining 42,649 individuals to serve as the test data set. As the metabolic variables are often strongly correlated, we fit our predictors on the training data using as predictor variables the first 22 principal components of the metabolic variables (together explaining at least 95 percent of the variance) and sex. Time-to-all-cause-mortality is the outcome.

We evaluated the biological age clocks using both standard and new evaluation measures, as introduced in section Evaluation measures. To calculate the concordance, we translated predicted biological age back to predicted mean residual life for metabo-AccelerAge-Gompertz, via the sex-specific baseline survival curves. For metabo-GrimAge it is unclear whether predictions can be directly translated to a predicted mean/median residual life value. We did try this, but it resulted in a lower concordance than when using the predicted GrimAges directly. We hence went for the option most beneficial to metabo-GrimAge. To make the calibration plots, we considered 5-year mortality and placed participants

in five equally sized groups based on their predicted 5-year mortality.

The results for the 'traditional' evaluation of biological age clocks, i.e. to what extent Δ is associated with all-cause mortality in a Cox PH model that also includes chronological age and sex, can be found in Table 1. It can be seen that the Δ of both metabo-AccelerAge-Gompertz and metabo-GrimAge is strongly associated with time to all-cause mortality.

Table 2 contains the results for one of our new proposed evaluation measures, the concordance. Both biological age predictors have a higher concordance than chronological age.

Calibration plots for metabo-AccelerAge-Gompertz and metabo-GrimAge can be found in Fig. 9. AccelerAge-Gompertz is better calibrated: predicted all-cause mortality risk is closer to observed all-cause mortality risk. This result can be understood when plotting chronological age against predicted biological age, as done in Figs. 10 and 11. The metabo-GrimAge predictions are—by design—centered around the line $C = B$. For this particular data set that does not make sense, since we know that the UKB population lives longer than the average UK population (Fig. 8), so their biological ages should on average be somewhat lower than the line $C = B$. The predicted biological ages of our metabo-AccelerAge-Gompertz predictor are on average also lower than the chronological ages.

We validated our findings that metabo-AccelerAge-Gompertz and metabo-GrimAge perform equally well on discrimination but that metabo-AccelerAge-Gompertz is better calibrated in an external data set, the Leiden Longevity Study (LLS). The LLS tracks long-lived Dutch siblings of Caucasian origin, their offspring and the partners of the offspring. Participants were recruited between March 2002 and May 2006. Registry-based follow-up data until November 2021 was available. We used data on the offspring and the partners (N = 2312). Participants who dropped out (N = 10) or had missing values on the 242 included metabolic variables (N = 46) were excluded, resulting in 1007 men and 1249 women with a mean age of 59.2 years (sd = 6.7, IQR = 54.7–63.9; see Figure 2 of the Supplementary Information for a histogram) at inclusion. Median follow-up time was 16.3 years (IQR = 15.3–17.1) and 313 deaths were observed. Table 1 of the Supplementary Information

Table 1 Hazard ratios for all-cause mortality associated with a standard unit increase in Δ in a Cox PH model adjusted for chronological age and sex, evaluated on the test set of the UKB data

	HR (95% CI)	p-value
Δ: metabo-GrimAge	1.61 (1.56–1.65)	$< 2e{-}16$
Δ: metabo-AccelerAge-Gompertz	1.56 (1.52–1.60)	$< 2e{-}16$

Table 2 Concordance (Uno's C) of the biological age predictions evaluated on the test set of the UKB data

|	Uno'c C (s.e.)
Chronological age	0.703 (0.007)
Biological age: metabo-GrimAge	0.739 (0.009)
Biological age: metabo-AccelerAge-Gompertz	0.739 (0.009)

Fig. 9 Calibration metabo-GrimAge and metabo-AccelerAge-Gompertz based on 5-year survival in the test set of the UKB data

Fig. 10 Chronological age plotted against predicted biological age in the test set of the UKB data for metabo-AccelerAge-Gompertz. The orange line has slope 1: it denotes where chronological age is equal to predicted biological age

contains information on the prevalence of various (chronic) diseases among the included participants.

Table 3 shows to what extent age acceleration Δ is associated with all-cause mortality in a Cox PH model that also includes chronological age and sex on the LLS data. The Δ-values of both biological age predictors are still significantly associated with mortality beyond chronological age. Nevertheless, the hazard ratios are much lower.

Table 4 contains the concordance of the biological age predictors on the LLS data. The two predictors still both discriminate slightly better than chronological age, but the difference is smaller.

Figure 12 contains the calibration plots for the two predictors on the LLS data. The metabo-AccelerAge-Gompertz predictor is still better calibrated than metabo-GrimAge, although the difference is smaller. In this case the

Fig. 11 Chronological age plotted against predicted biological age in the test set of the UKB data for metabo-GrimAge. The orange line has slope 1: it denotes where chronological age is equal to predicted biological age

Table 3 Hazard ratios for all-cause mortality associated with a standard unit increase in Δ in a Cox PH model adjusted for chronological age and sex, evaluated on the LLS data

	HR (95% CI)	p-value
Δ: metabo-GrimAge	1.39 (1.23–1.57)	1.58e−7
Δ: metabo-AccelerAge-Gompertz	1.35 (1.21–1.52)	1.82e−7

Table 4 Concordance (Uno's C) of the biological age predictions evaluated on the LLS data

	Uno'c C (s.e.)
Chronological age	0.723 (0.016)
Biological age: metabo-GrimAge	0.739 (0.014)
Biological age: metabo-AccelerAge-Gompertz	0.740 (0.014)

biological age predictions of metabo-AccelerAge-Gompertz are slightly too low, especially for the group with the highest predicted mortality probability: the true event rates are slightly higher than the predicted probabilities.

Discussion

In this paper we presented a new statistical framework for biological age prediction, AccelerAge, based on Accelerated Failure Time models. We proposed this new framework because most existing biological age predictors are not based on a formal operationalization of the concept of biological age, are not in line with the aging-as-a-clock idea and are

challenging to evaluate, because the estimand is ambiguous. In an attempt to overcome these limitations, we started by formally defining the concept of biological age via residual life, and subsequently based the AccelerAge framework on this operationalization. The discussion on what (biological) aging exactly entails is vivid and ongoing, but our proposed definition can serve as a starting point for further discussion. We explained why biological age predictors based on AFT models are in line with the ubiquitous clock metaphor in the aging field, while predictors based on Proportional Hazards models are not. Besides the more natural interpretation of biological age predictors based on AFT models, another advantage of AFT models is that they are robust to covariate omission: neglecting true covariates might lead to a distribution outside the parametric family considered, but it does not affect the regression-part of the model. This is not true for the PH model [34]. Another appealing aspect of the AccelerAge framework is that the AFT-model is fitted using chronological age as the timescale t instead of time-on-study. This is the preferred timescale in the context of cohort studies [35, 36]. Finally, in this paper we introduced two new evaluation measures into the context of biological age clocks, namely concordance and discrimination. This allows for a broader evaluation and validation of biological age predictors. AccelerAge predictions are directly interpretable on an age-scale, which means that the biological age prediction itself is meaningful, not just the age-independent part Δ.

Our UK Biobank application illustrates that metabo-AccelerAge-Gompertz is a worthy competitor to our Grim-Age-implementation, metabo-GrimAge, in terms of prediction performance. Using the often-used evaluation measure

Fig. 12 Calibration metabo-GrimAge and metabo-Acceler-Age-Gompertz based on 5-year survival in the LLS data

of checking whether the age-independent part of a biological age prediction (Δ) is significantly associated with prospective mortality, the hazard ratio for metabo-AccelerAge-Gompertz was slightly lower than that for our GrimAge-implementation, but metabo-AccelerAge-Gompertz was better calibrated in both the UKB data and the LLS data, used as the external validation set. On concordance the two predictors scored similarly. This implies that opting for the approach with a more robust methodological foundation and aligned with hypothesized biological reality does not require any sacrifices in performance.

The real data application highlights the need for having more than one evaluation measure in biological age prediction. A true biological age predictor should be directly interpretable on an age-scale. Calibration plots can be used to assess the ability of a predictor to do so; checking whether Δ is significantly associated with time-to-mortality cannot. In addition, only relying on the association of Δ with age-related outcomes might paint a too optimistic picture of the (clinical) usefulness of a biological age predictor. Although we found that both metabolite-based biological age clocks performed better than chronological age on discrimination, the differences with chronological age were small. One potential reason for this is that we only considered the metabolic variables as measured by the Nightingale NMR platform as predictor variables. Using more of the human metabolome, or using other sources of omics, might capture more of the aging process and hence further improve the discriminative ability of a predictor.

Our AccelerAge framework can be applied to any kind of time-to-event outcome and any type of predictor variables.

Naturally, the conclusions resulting from the comparison with a GrimAge-based predictor might change if different predictor variables or outcomes are considered. We plan to extend the framework to also allow for regularization. Incorporating multiple predictor types, e.g. multiple omics, would be possible with a group-wise penalty term. In addition, it would be interesting to see how AccelerAge develops within individuals over time: this would e.g. also make it possible to investigate if certain (early-life) interventions contribute to achieving a lower predicted biological age, either immediately or later in life.

There are several limitations to our work. The AccelerAge framework is only based on lifespan, not healthspan, because healthspan tables are generally not available. This can be considered a (too) narrow view of what it means to age. In addition, our operationalization requires that one has access to a life table of the reference population of interest. This might not always be the case, in particular not for outcomes other than mortality. In certain cases it might be feasible to simply construct a table based on the training data set itself, but then the ability to detect whether the reference population differs from the sample is lost. The reliance of the AccelerAge approach on life tables can in itself be seen as a limitation, as it forces one to decide on a reference population. If the reference population changes, the resulting biological age predictions also change. However, we would like to point out that *any* prediction approach that makes predictions on an age-scale uses a reference population, albeit implicit: if not, the resulting age cannot be interpreted. For instance, GrimAge uses the mean and standard deviation of chronological age in the training data to obtain predictions in

units of years. The AccelerAge is different from most other biological age approaches in that it first predicts remaining lifespan, instead of directly prediction (biological) age, and that it makes the use of a reference population explicit.

We illustrated the AccelerAge framework using a fully parametric AFT model, based on the Gompertz distribution. As it has long been known that the Gompertz distribution describes adult mortality well, this parametric approach sufficed for our real data application, in which time-to-mortality was our only outcome of interest. AFT models can also be fit using a semiparametric [48–51] or flexible parametric [52] approach. We initially included these two approaches in our simulation study as well: results can be found in section 5 of the Supplementary Information. Our semiparametric approach, based on the weighted least squares method, is described in section 6 of the Supplementary Information. We found that especially the fitting of flexible parametric AFT models can be inconsistent and suffer from convergence issues. However, as developing robust flexible parametric AFT models is an area of active research [53], we believe the flexible parametric AccelerAge approach will soon become an appealing alternative to our fully parametric approach. It should however be noted that, in order to fit any AccelerAge model that is not fully parametric, the data must cover the whole of human lifespan to properly estimate the baseline survival curve or hazard function (which must be fully specified, because it needs to be integrated over). In the UK Biobank data, the oldest included participant was 85. Semiparametric prediction approaches therefore would have no data to estimate baseline survival or hazard at ages after 85.

In conclusion, our work represents a substantial advancement in the field of biological age research. By introducing AccelerAge, a new AFT-based statistical framework to predict biological age based on a solid operationalization of biological age, and incorporating previously underutilized evaluation measures, namely discrimination and concordance, we have laid a robust statistical foundation for the development and validation of biological age clocks.

Declarations

References

1. López-Otín C, Blasco MA, Partridge L, Serrano M, Kroemer G. Hallmarks of aging: an expanding universe. Cell. 2023.
2. Kennedy BK, Berger SL, Brunet A, Campisi J, Cuervo AM, Epel ES, Franceschi C, Lithgow GJ, Morimoto RI, Pessin JE. Geroscience: linking aging to chronic disease. Cell. 2014;159(4):709–13.
3. Lowsky DJ, Olshansky SJ, Bhattacharya J, Goldman DP. Heterogeneity in healthy aging. J Gerontol Ser A Biomed Sci Med Sci. 2014;69(6):640–9.
4. Partridge L, Deelen J, Slagboom PE. Facing up to the global challenges of ageing. Nature. 2018;561(7721):45–56.
5. Jylhävä J, Pedersen NL, Hägg S. Biological age predictors. EBioMedicine. 2017;21:29–36.
6. Moreira T. Anticipatory measure: Alex Comfort, experimental gerontology and the measurement of senescence. Stud Hist Philos Sci Part C Stud Hist Philos Biol Biomed Sci. 2019;77: 101179.
7. Comfort A. Test-battery to measure ageing-rate in man. The Lancet. 1969;294(7635):1411–5.
8. Rutledge J, Oh H, Wyss-Coray T. Measuring biological age using omics data. Nat Rev Genet. 2022; 1–13.
9. Benjamin H. Biologic versus chronologic age. J Gerontol. 1947;2(3):217–27.
10. Hollingsworth JW, Hashizume A, Jablon S. Correlations between tests of aging in Hiroshima subjects-an attempt to define "physiologic age''. Yale J Biol Med. 1965;38(1):11.

11. Horvath S. DNA methylation age of human tissues and cell types. Genome Biol. 2013;14(10):1–20.

12. Hannum G, Guinney J, Zhao L, Zhang L, Hughes G, Sadda S, Klotzle B, Bibikova M, Fan J-B, Gao Y. Genome-wide methylation profiles reveal quantitative views of human aging rates. Mol Cell. 2013;49(2):359–67.

13. Marioni RE, Shah S, McRae AF, Chen BH, Colicino E, Harris SE, Gibson J, Henders AK, Redmond P, Cox SR. DNA methylation age of blood predicts all-cause mortality in later life. Genome Biol. 2015;16(1):1–12.

14. Peters MJ, Joehanes R, Pilling LC, Schurmann C, Conneely KN, Powell J, Reinmaa E, Sutphin GL, Zhernakova A, Schramm K. The transcriptional landscape of age in human peripheral blood. Nat Commun. 2015;6(1):1–14.

15. Van Den Akker EB, Trompet S, Barkey Wolf JJ, Beekman M, Suchiman HED, Deelen J, Asselbergs FW, Boersma E, Cats D, Elders PM. Metabolic age based on the BBMRI-NL 1H-NMR metabolomics repository as biomarker of age-related disease. Circul Genom Precis Med. 2020;13(5):541–7.

16. Kristic J, Vuckovic F, Menni C, Klaric L, Keser T, Beceheli I, Pucic-Bakovic M, Novokmet M, Mangino M, Thaqi K. Glycans are a novel biomarker of chronological and biological ages. J Gerontol Ser A. 2014;69(7):779–89.

17. Tanaka T, Biancotto A, Moaddel R, Moore AZ, Gonzalez-Freire M, Aon MA, Candia J, Zhang P, Cheung F, Fantoni G. Plasma proteomic signature of age in healthy humans. Aging Cell. 2018;17(5):12799.

18. Galkin F, Mamoshina P, Aliper A, Putin E, Moskalev V, Gladyshev VN, Zhavoronkov A. Human gut microbiome aging clock based on taxonomic profiling and deep learning. iScience. 2020;3(6): 101199.

19. Sluiskes MH, Goeman JJ, Beekman M, Slagboom PE, Putter H, Rodriguez-Girondo M. Clarifying the biological and statistical assumptions of cross-sectional biological age predictors. bioRxiv, 2023;2023-01.

20. Hillary RF, Stevenson AJ, McCartney DL, Campbell A, Walker RM, Howard DM, Ritchie CW, Horvath S, Hayward C, McIntosh AM. Epigenetic measures of ageing predict the prevalence and incidence of leading causes of death and disease burden. Clin Epigenetics. 2020;12(1):1–12.

21. Maddock J, Castillo-Fernandez J, Wong A, Cooper R, Richards M, Ong KK, Ploubidis GB, Goodman A, Kuh D, Bell JT. DNA methylation age and physical and cognitive aging. J Gerontol Ser A. 2020;75(3):504–11.

22. McCrory C, Fiorito G, Hernandez B, Polidoro S, O'Halloran AM, Hever A, Ni Cheallaigh C, Lu AT, Horvath S, Vineis P. GrimAge outperforms other epigenetic clocks in the prediction of age-related clinical phenotypes and all-cause mortality. J Gerontol Ser A. 2021;76(5):741–9.

23. Kuiper LM, Polinder-Bos HA, Bizzarri D, Vojinovic D, Vallerga CL, Beekman M, Dollé ME, Ghanbari M, Voortman T, Reinders MJ, et al. Epigenetic and metabolomic biomarkers for biological age: a comparative analysis of mortality and frailty risk. J Gerontol Ser A 2023;137.

24. Levine ME, Lu AT, Quach A, Chen BH, Assimes TL, Bandinelli S, Hou L, Baccarelli AA, Stewart JD, Li Y. An epigenetic biomarker of aging for lifespan and healthspan. Aging (Albany NY). 2018;10(4):573.

25. Lu AT, Quach A, Wilson JG, Reiner AP, Aviv A, Raj K, Hou L, Baccarelli AA, Li Y, Stewart JD. DNA methylation GrimAge strongly predicts lifespan and healthspan. Aging (Albany NY). 2019;11(2):303.

26. Lu AT, Binder AM, Zhang J, Yan Q, Reiner AP, Cox SR, Corley J, Harris SE, Kuo P-L, Moore AZ. Dna methylation grimage version 2. Aging (Albany NY). 2022;14(23):9484.

27. Wei L-J. The accelerated failure time model: a useful alternative to the Cox regression model in survival analysis. Stat Med. 1992;11(14–15):1871–9.

28. Cox DR. Regression models and life-tables. J Roy Stat Soc Ser B (Methodol). 1972;34(2):187–202.

29. Swindell WR. Accelerated failure time models provide a useful statistical framework for aging research. Exp Gerontol. 2009;44(3):190–200.

30. Olshansky SJ, Carnes BA. Ever since Gompertz. Demography. 1997;34(1):1–15.

31. Gompertz B. On the nature of the function expressive of the law of human mortality, and on a new mode of determining the value of life contingencies. Philos Trans R Soc Lond. 1825;115:513–85.

32. Gavrilov LA, Gavrilova NS. Mortality measurement at advanced ages: a study of the Social Security Administration Death Master File. N Am Actuar J. 2011;15(3):432–47.

33. Böhnstedt M, Gampe J, Putter H. Information measures and design issues in the study of mortality deceleration: findings for the gamma-Gompertz model. Lifetime Data Anal. 2021;27(3):333–56.

34. Hougaard P. Fundamentals of survival data. Biometrics. 1999;55(1):13–22.

35. Korn EL, Graubard BI, Midthune D. Time-to-event analysis of longitudinal follow-up of a survey: choice of the time-scale. Am J Epidemiol. 1997;145(1):72–80.

36. Thiébaut AC, Bénichou J. Choice of time-scale in Cox's model analysis of epidemiologic cohort data: a simulation study. Stat Med. 2004;23(24):3803–20.

37. Steyerberg EW, Vickers AJ, Cook NR, Gerds T, Gonen M, Obuchowski N, Pencina MJ, Kattan MW. Assessing the performance of prediction models: a framework for some traditional and novel measures. Epidemiology (Cambridge, Mass). 2010;21(1):128.

38. Uno H, Cai T, Pencina MJ, D'Agostino RB, Wei L-J. On the C-statistics for evaluating overall adequacy of risk prediction procedures with censored survival data. Stat Med. 2011;30(10):1105–17.

39. Bender R, Augustin T, Blettner M. Generating survival times to simulate Cox proportional hazards models. Stat Med. 2005;24(11):1713–23.

40. Broström G. Event history analysis with R. 2nd ed. Boca Raton: Chapman and Hall/CRC; 2022. https://doi.org/10.1201/97804 29503764.

41. Sudlow C, Gallacher J, Allen N, Beral V, Burton P, Danesh J, Downey P, Elliott P, Green J, Landray M. UK Biobank: an open access resource for identifying the causes of a wide range of complex diseases of middle and old age. PLoS Med. 2015;12(3):1001779.

42. Soininen P, Kangas AJ, Würtz P, Suna T, Ala-Korpela M. Quantitative serum nuclear magnetic resonance metabolomics in cardiovascular epidemiology and genetics. Circul Cardiovasc Genet. 2015;8(1):192–206.

43. Würtz P, Kangas AJ, Soininen P, Lawlor DA, Davey Smith G, Ala-Korpela M. Quantitative serum nuclear magnetic resonance metabolomics in large-scale epidemiology: a primer on-omic technologies. Am J Epidemiol. 2017;186(9):1084–96.

44. Julkunen H, Cichońska A, Tiainen M, Koskela H, Nybo K, Mäkelä V, Nokso-Koivisto J, Kristiansson K, Perola M, Salomaa V. Atlas of plasma nmr biomarkers for health and disease in 118,461 individuals from the UK biobank. Nat Commun. 2023;14(1):604.

45. Allen NE, Arnold M, Parish S, Hill M, Sheard S, Callen H, Fry D, Moffat S, Gordon M, Welsh S. Approaches to minimising the epidemiological impact of sources of systematic and random variation that may affect biochemistry assay data in uk biobank. Wellcome Open Res. 2021;5:222.

46. ONS: National life tables: UK. https://www.ons.gov.uk/peopl epopulationandcommunity/birthsdeathsandmarriages/lifeexpect ancies/datasets/nationallifetablesunitedkingdomreferencetables. [Online; accessed 04-January-2023] (2021)

47. Fry A, Littlejohns TJ, Sudlow C, Doherty N, Adamska L, Sprosen T, Collins R, Allen NE. Comparison of sociodemographic and health-related characteristics of UK Biobank participants with those of the general population. Am J Epidemiol. 2017;186(9):1026–34.

48. Buckley J, James I. Linear regression with censored data. Biometrika. 1979;66(3):429–36.

49. Tsiatis AA. Estimating regression parameters using linear rank tests for censored data. Ann Stat. 1990;354–372.

50. Ying Z. A large sample study of rank estimation for censored regression data. Ann Stat. 1993;76–99.

51. Stute W. Consistent estimation under random censorship when covariables are present. J Multivar Anal. 1993;45(1):89–103.

52. Pang M, Platt RW, Schuster T, Abrahamowicz M. Spline-based accelerated failure time model. Stat Med. 2021;40(2):481–97.

53. Crowther MJ, Royston P, Clements M. A flexible parametric accelerated failure time model and the extension to time-dependent acceleration factors. Biostatistics. 2023;24(3):811–31.