

A recipe for accurate estimation of lifespan brain trajectories, distinguishing longitudinal and cohort effects

Øystein Sørensen^{a,*}, Kristine B. Walhovd^{a,b}, Anders M. Fjell^{a,b}

^a Center for Lifespan Changes in Brain and Cognition, Department of Psychology, University of Oslo, Norway

^b Department of Radiology and Nuclear Medicine, Oslo University Hospital, Norway

ARTICLE INFO

Keywords:

Aging
Cohort effects
Generalized additive mixed models
Lifespan brain research
Longitudinal analysis
MRI
R

ABSTRACT

We address the problem of estimating how different parts of the brain develop and change throughout the lifespan, and how these trajectories are affected by genetic and environmental factors. Estimation of these lifespan trajectories is statistically challenging, since their shapes are typically highly nonlinear, and although true change can only be quantified by longitudinal examinations, as follow-up intervals in neuroimaging studies typically cover less than 10% of the lifespan, use of cross-sectional information is necessary. Linear mixed models (LMMs) and structural equation models (SEMs) commonly used in longitudinal analysis rely on assumptions which are typically not met with lifespan data, in particular when the data consist of observations combined from multiple studies. While LMMs require a priori specification of a polynomial functional form, SEMs do not easily handle data with unstructured time intervals between measurements. Generalized additive mixed models (GAMMs) offer an attractive alternative, and in this paper we propose various ways of formulating GAMMs for estimation of lifespan trajectories of 12 brain regions, using a large longitudinal dataset and realistic simulation experiments. We show that GAMMs are able to more accurately fit lifespan trajectories, distinguish longitudinal and cross-sectional effects, and estimate effects of genetic and environmental exposures. Finally, we discuss and contrast questions related to lifespan research which strictly require repeated measures data and questions which can be answered with a single measurement per participant, and in the latter case, which simplifying assumptions that need to be made. The examples are accompanied with R code, providing a tutorial for researchers interested in using GAMMs.

1. Introduction

Large datasets with structural magnetic resonance images (MRIs) of participants whose ages span from early childhood to late adulthood provide ample opportunities to study lifespan brain trajectories. Important questions such data can contribute to answering include how brain structure is related to aging, how the aging effect is modified by genetics and environmental exposures, and at which age critical events like maximum volume or maximum rate of change occur. Lifespan brain trajectories are nonlinear and differ between regions, as illustrated in Fig. 1 for volumes of cerebral white matter, cortex, and hippocampus for 4352 observations of 2017 healthy participants from the Center for Lifespan Changes in Brain and Cognition longitudinal studies (Fjell et al., 2017; Walhovd et al., 2016), henceforth referred to as the LCBC data. Modeling the type of nonlinear effects shown in Fig. 1 using linear mixed models (LMMs) (Laird and Ware, 1982) with polynomials typically leads to poor fits at least over parts of the lifespan, and is highly dependent on manual selection of terms (Fjell et al., 2010; Sørensen et al., 2021).

Structural equation models (SEMs) may be better able to estimate nonlinear trajectories, e.g., with a latent basis model (McArdle and Epstein, 1987; Meredith and Tisak, 1990), but SEMs require that the time intervals between measurements for all participants take on a small set of unique values (Newsom, 2015; Oud and Jansen, 2000), an assumption which may be hard to satisfy with lifespan data (see Fig. 3). Generalized additive mixed models (GAMMs)¹ (Lin and Zhang, 1999) offer an attractive alternative, typically yielding good fit over the full lifespan in an automated and data-driven manner. This is illustrated in Fig. 2, comparing a GAMM to LMMs with quadratic and cubic polynomials for the effect of age on cerebellum cortex volume. See Box 1 for a definition of these and other key terms used in this paper. Similar to often researched structures like hippocampus and cerebral white matter, cerebellum cortex is characterized by a nonlinear age trajectory. In contrast to the GAMM, neither of the LMMs capture the steep increase seen in early childhood, and the cubic LMM predicts an increase in cerebellum

¹ We will use the common abbreviation "GAMM", although strictly speaking only additive mixed models (AMMs) are used in this paper. If necessary, all models described can be straightforwardly generalized, e.g. to logistic or Poisson regression.

* Corresponding author.

E-mail address: oystein.sorensen@psykologi.uio.no (Ø. Sørensen).

Cohort effect: The effect of birth year (cohort) on the relationship between a set of explanatory variables and an outcome of interest.

Cross-sectional effect: The effect of age on an outcome of interest at a particular point in time, across participants with different birth dates.

Longitudinal effect: The effect of increasing age for participants belonging to a given birth cohort.

Linear mixed models (LMMs): Linear regression models used for data with hierarchical structure, e.g. longitudinal data with multiple measurements per individual.

Fixed effects: Regression parameters in mixed models which are common for all participants.

Random effects: Regression parameters in mixed models which are unique to each participant, used to model correlation between repeated measurements.

Polynomial model: Linear regression model which includes the first n powers of an explanatory variable x as distinct variables.

Quadratic model: A polynomial model containing the first two powers of an explanatory variable x , on the form $\beta_0 + \beta_1x + \beta_2x^2$.

Cubic model: A polynomial model containing the first three powers of an explanatory variable x , on the form $\beta_0 + \beta_1x + \beta_2x^2 + \beta_3x^3$.

Generalized additive model (GAM): A linear regression model in which the outcome is modeled as an unknown smooth function of the explanatory variables.

Generalized additive mixed model (GAMM): An extension of GAMs to data with hierarchical structure, containing random effects.

Smoothing parameter: Parameter controlling the degree of nonlinearity (wiggleness) of a function estimated by a GAM/GAMM.

Cubic regression splines: A set of cubic polynomials, each of which is defined over a small part of the x -axis and is zero elsewhere.

Thin-plate regression splines: A set of functions, each of which represents a given nonlinear shape over the full x -axis.

Smooth function: A nonlinear function estimated by GAMs/GAMMs represented as a weighted sum of (e.g., cubic or thin-plate) regression splines.

Box 1. Key terms used in this paper, defined in the context of longitudinal data analysis.

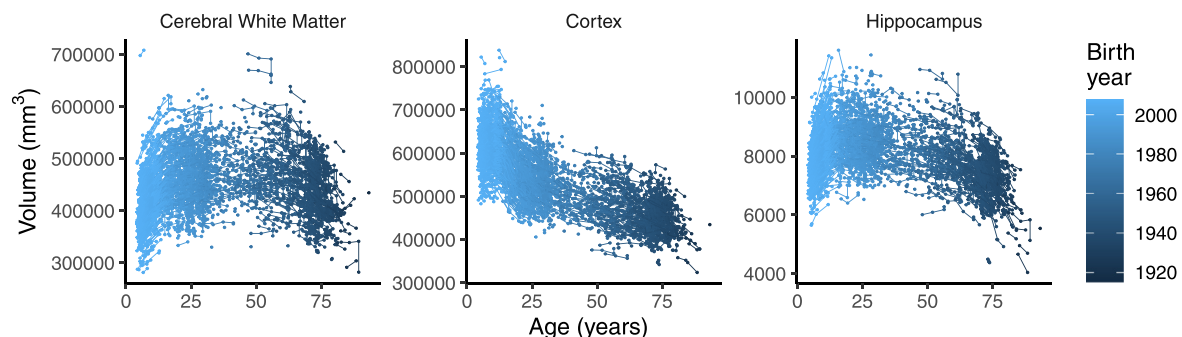


Fig. 1. Lifespan brain development is highly nonlinear. Cerebral white matter, cortex and hippocampal volumes from 4352 MRI scans of 1017 participants in the LCBC data. The color scale indicates the birth cohort to which the participant belongs. Dots represent observations and lines connecting the dots indicate repeated observations of the same individual.

cortex volume in old age, whereas the GAMM adequately captures the decline seen in the data. In addition, both the quadratic and cubic model estimate cerebellum cortex volume to reach its maximum at the age of around 25, while the GAMM instead estimates the maximum to occur around 14 years of age, and the latter seems to be in better agreement with the data. Figs. 1 and 2 and all subsequent plots were created in R (R Core Team, 2019) with ggplot2 (Wickham, 2016).

The goal of this paper is to provide clear recommendations for optimal estimation of lifespan trajectories of brain development and aging. To this end, several aspects need consideration. First, as has been emphasized by a large number of authors, when analyzing data with repeated observations over time, care must be taken to distinguish within-individual and between-individual effects, which for the purpose of this paper are longitudinal and cross-sectional effects (Curran and Bauer, 2011; Hoffman, 2007; Hoffman and Stawski, 2009; Morrell et al., 2009; Sliwinski et al., 2010; Thompson et al., 2011). Individual change can only be assessed with repeated measurements, but how important are longitudinal data when the task is to estimate trajectories spanning many times the maximum follow-up interval realistically attainable in a neuroimaging study? Large datasets combined




Fig. 2. Comparison of LMMs and GAMMs for lifespan data. Comparison of LMMs with quadratic and cubic terms and a GAMM, fitted to lifespan cerebellum cortex volume. Black dots represent observations and black lines connecting the dots indicate repeated observations of the same individual.

from different studies, either conducted by the same group as for the LCBC data or by multiple groups participating in a data-sharing consor-

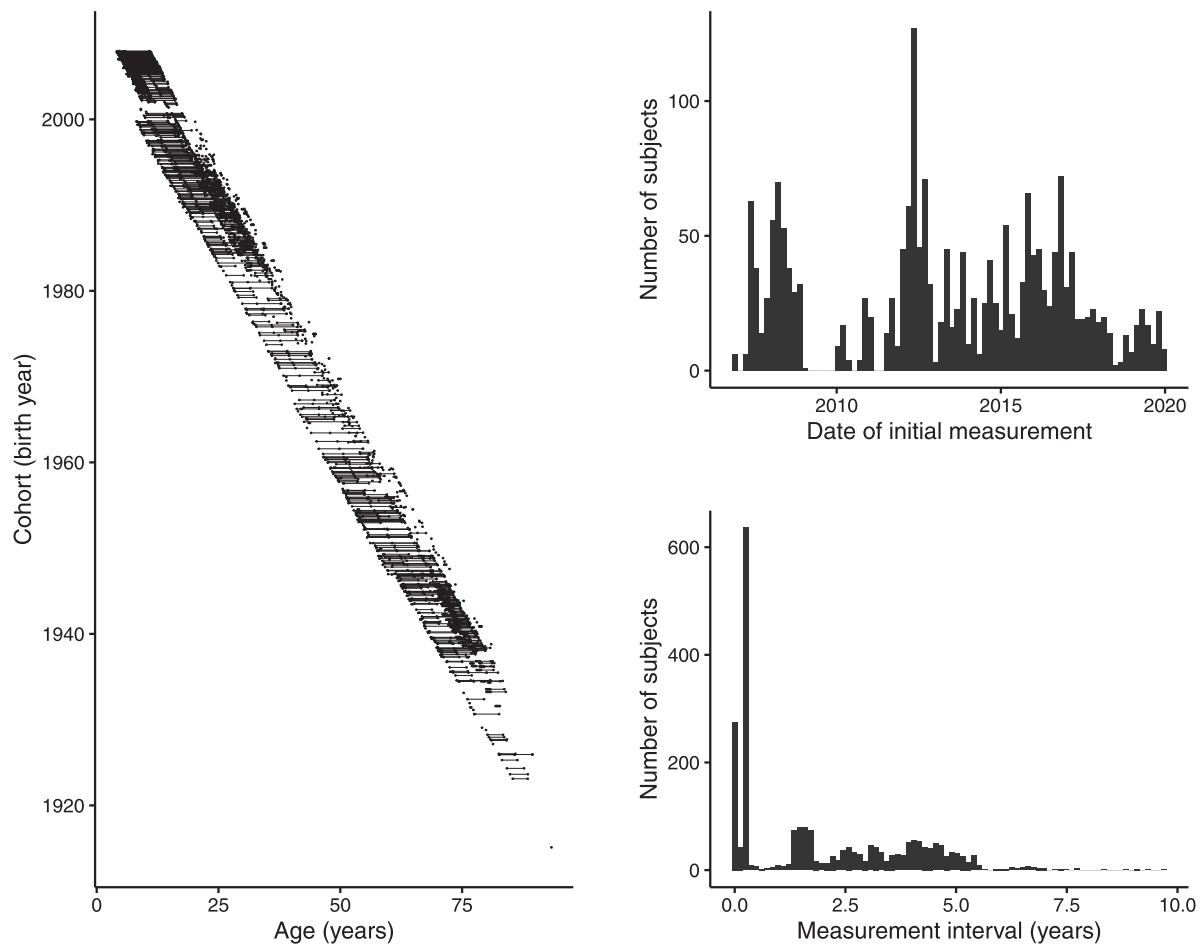


Fig. 3. Characteristics of lifespan data. The plots show data from 4352 MRI scans of 2017 participants in the LCBC data. Left: Scatter plot of age and cohort. Connected dots show repeated measurements of the same participant. Top right: Histogram of date of initial measurement for the same participants. Bottom right: Histogram of time (in years) between measurements. The peak at zero corresponds to participants scanned twice on the same day, with different scanners, and the highest peak corresponds to participants with 10–11 weeks between measurements.

tium like Lifebrain (Walhovd et al., 2018) or a meta-analysis network like ENIGMA (Bearden and Thompson, 2017; Thompson et al., 2017), present further challenges for longitudinal modeling as the number of measurements per participant and the time intervals between measurements are typically highly varying. All of these issues are illustrated for the LCBC data in Fig. 3. While GAMMs flexibly handle data with varying follow-up intervals, the statistical literature on use of GAMMs for longitudinal analysis has almost exclusively focused on cases in which each participant has been followed over the full time range under consideration, from a common baseline (Brumback and Rice, 1998; Durbán et al., 2005; Edwards et al., 2005; Gu and Ma, 2005; Ke and Wang, 2001; Lambert et al., 2001; Sullivan et al., 2015). There is hence a need for an understanding of how GAMMs should be optimally used in lifespan brain research, with short follow-up intervals and varying dates of initial measurement as shown in Fig. 3.

The outline of this paper is as follows. In Section 2 we introduce GAMMs formally and define three different candidate models for estimating lifespan brain trajectories. We also describe simulation experiments conducted in order to compare these GAMMs in a realistic setting for estimating 12 different brain regions. In Section 3 the simulation results are presented, and next we show two example applications demonstrating how GAMMs can be used for estimating lifespan brain trajectories and the effect of genetic variations on the trajectories. Accompanying R code provides a tutorial for researchers interested in using GAMMs. In Section 4 we discuss the results taking into regard currently available longitudinal studies. We contrast questions that strictly require

longitudinal data to questions that under some simplifying assumptions may be answered with a single measurement per participant. Finally, in Section 5 we conclude by presenting recommendations for how to use GAMMs in lifespan brain research.

2. Materials and methods

2.1. Longitudinal and cross-sectional effects

The effect of age on an outcome in a population can be completely explained by longitudinal and cohort effects, with the former representing the effect of aging for participants in a given birth cohort and the latter determining how the longitudinal effects differ between participants belonging to different birth cohorts (Diggle et al., 2002, Ch. 1.1). Cross-sectional effects are the effects of age across cohorts when considered at a particular point in time, as illustrated in Fig. 4. In the absence of cohort effects the cross-sectional and longitudinal effects are identical. Age-independent cohort effects result in different slopes for the longitudinal and cross-sectional effects, while age-cohort interactions lead to longitudinal effects whose slopes depend on age. Selective survival, by which life expectancy is correlated with the dependent variable, leads to population changes over time and hence are part of the longitudinal effects (Baltes, 1968). In contrast, with sampling bias, by which the probability of recruitment or the probability of dropout before the end of the study depends on the outcome variable, the sample is not repre-

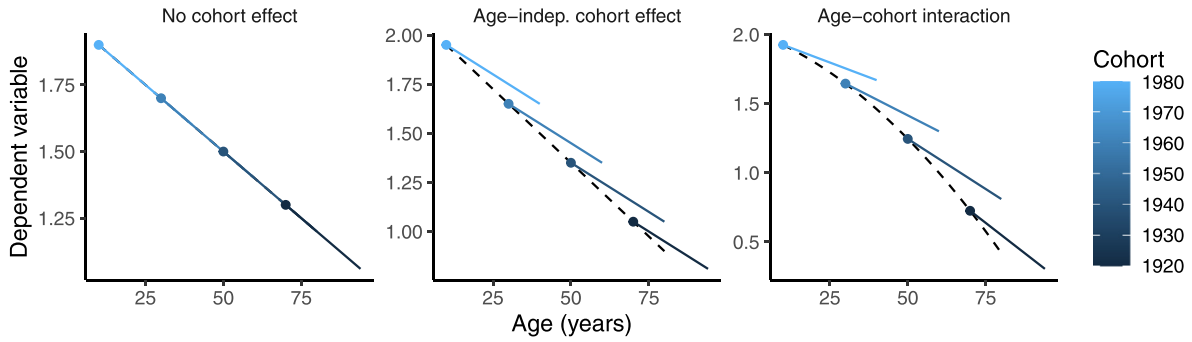


Fig. 4. Cohort effects. Illustration of the impact of cohort effects in a hypothetical dataset. Dashed lines show the cross-sectional age effect in 1990, colored dots show four cohorts of participants whose age in 1990 was 10, 30, 50, and 70 years, respectively, and the blue lines show longitudinal age effects for each cohort. In the left plot, there are no cohort effects, and hence longitudinal and cross-sectional effects coincide. In the center plot, the cohort effects are independent of age, and the longitudinal effects differ by an offset but the effect of aging is identical across cohorts, as seen by the parallel blue lines. In the right plot, the cohort effects depend on age, and in this case also the slope of the longitudinal effect varies between cohorts. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

sentative of the population under study and biased estimates may result (Molenberghs and Fitzmaurice, 2009).

2.2. Generalized additive models

Generalized additive models (GAMs) (Hastie and Tibshirani, 1986) model the effect of a variable x on an outcome y with smooth functions $f(x)$, constructed as weighted sums of K basis functions $b_1(x), b_2(x), \dots, b_K(x)$ with weights $\beta_1, \beta_2, \dots, \beta_K$, i.e., $f(x) = \sum_{k=1}^K \beta_k b_k(x)$. Commonly used basis functions are cubic regression splines and thin-plate regression splines (Wood, 2003), and the number of basis functions is typically chosen large enough to allow a wide range of nonlinear patterns to be estimated, while small enough to allow computational efficiency. For a GAM with a single smooth term, $y = f(x) + \epsilon$, the estimate given n observations is computed by finding the values of β_1, \dots, β_K minimizing the criterion

$$\sum_{i=1}^n \left[y_i - \sum_{k=1}^K \beta_k b_k(x_i) \right]^2 + \lambda \int_a^b \left[\sum_{k=1}^K \beta_k b_k''(x) \right]^2 dx.$$

The first term is the least squares criterion using the basis functions as explanatory variables, and the second term represents the wiggleness of $f(x)$ as measured by its squared second derivative over some range $[a, b]$, typically the minimum and maximum values of x in the sample. The smoothing parameter λ controls the extent to which wiggleness is penalized, striking a balance between overfitting (too low λ , too wiggly $f(x)$) and underfitting (too high λ , too smooth $f(x)$). For data with repeated measurements, GAMs can be extended to GAMMs by the inclusion of random effects. A key insight allowing use of LMM software for efficient fitting of GAMMs is that the penalized smooth terms may be decomposed into a fixed effect part representing unpenalized linear functional forms with zero second derivative, and a random effect part representing penalized nonlinear functional forms with non-zero second derivative (Lin and Zhang, 1999; Wood, 2004; 2010). The variance of the random effects is proportional to $1/\lambda$, and this allows the smoothing parameter to be estimated as a mixed model variance component.

2.3. Generalized additive mixed models for longitudinal data

In this section we present three different models for estimating lifespan brain trajectories.

Consider a dataset of n participants indexed $i = 1, \dots, n$, assume an outcome y_{ij} has been measured m_i times in participant i , with timepoints indexed by $j = 1, \dots, m_i$, and let a_{ij} denote the age of participant i at her/his j th timepoint. The question of interest is how the outcome varies as a function of age, and this can be modeled with the GAMM

$$y_{ij} = \beta_0 + f(a_{ij}) + b_{0i} + \epsilon_{ij}, \quad (1)$$

where $f(a_{ij})$ is the effect of age, β_0 is the intercept, b_{0i} is the random intercept for participant i , and ϵ_{ij} is a random noise term. Both b_{0i} and ϵ_{ij} are assumed to be normally distributed, $b_{0i} \sim N(0, \sigma_b^2)$ and $\epsilon_{ij} \sim N(0, \sigma^2)$, with σ_b representing the between-participant standard deviation and σ the within-participant residual standard deviation. We do not consider random slopes, due to the low number of repeated measurements in the typical applications considered in this paper, although this could be included with an additional term $b_{1i}a_{ij}$ in (1). With sufficient data, use of random slopes is recommended, as it relaxes the assumptions on the covariance structure of repeated measurements (Fitzmaurice et al., 2011, Ch. 19).

In the presence of cohort effects, the term $f(a_{ij})$ represents some weighted combination of cross-sectional and longitudinal effects, and hence cannot be interpreted as either. The typical method of correcting for this in LMMs is by splitting the age term into a_{i1} representing age at first measurement, and t_{ij} representing time since baseline (Fitzmaurice et al., 2011; Zeger and Liang, 1992) (see Mehta and West (2000) for an equivalent method in SEMs). Extending this to a GAMM yields

$$y_{ij} = \beta_0 + f(a_{i1}, t_{ij}) + b_{0i} + \epsilon_{ij}, \quad (2)$$

where $f(a_{i1}, t_{ij})$ is a smooth bivariate function of baseline age and time. Considering the plots in Fig. 1, using a bivariate function seems necessary for estimating lifespan trajectories, as the direction of change clearly depend on baseline age. In model (2) the longitudinal effect of aging t from a baseline a_{i1} is given by $f(a_{i1}, t)$ keeping a_{i1} constant, while the cross-sectional effect of varying baseline age a is given by $f(a, 0)$.

Model (2) has some important limitations, however. First, an assumption in the LMMs motivating its definition is that all participants have identical baseline dates. Second, when participants are followed over a short period compared to the full lifespan, the values of t_{ij} vary between zero and some maximum which is much lower than the maximum age, which might make estimation of nonlinear longitudinal effects challenging. We hence introduce an alternative GAMM modeling cohort effects by including birth date c_i ,

$$y_{ij} = \beta_0 + f(a_{ij}) + \beta(a_{ij})c_i + b_{0i} + \epsilon_{ij}, \quad (3)$$

in which $f(a_{ij})$ is defined as for (1), while $\beta(a_{ij})c_i$ is a varying-coefficient term (Hastie and Tibshirani, 1993) representing the main effect of cohort (birth date) as a function of age. The longitudinal effect of aging a for a participant belonging to cohort c_i is given by $f(a) + \beta(a)c_i$ keeping c_i constant. The cross-sectional effect of age a at date d is given by $f(a) + \beta(a)c$, with $c = d - a$ representing the birth date of participants of age a at date d , and hence both a and c are varying in this case. Model (3) is not identified if all participants have identical measurement dates, since then birth date and age are perfectly collinear, i.e., $c_i = d_j - a_{ij}$ where d_j is the common date of the j th timepoint. However, as illus-

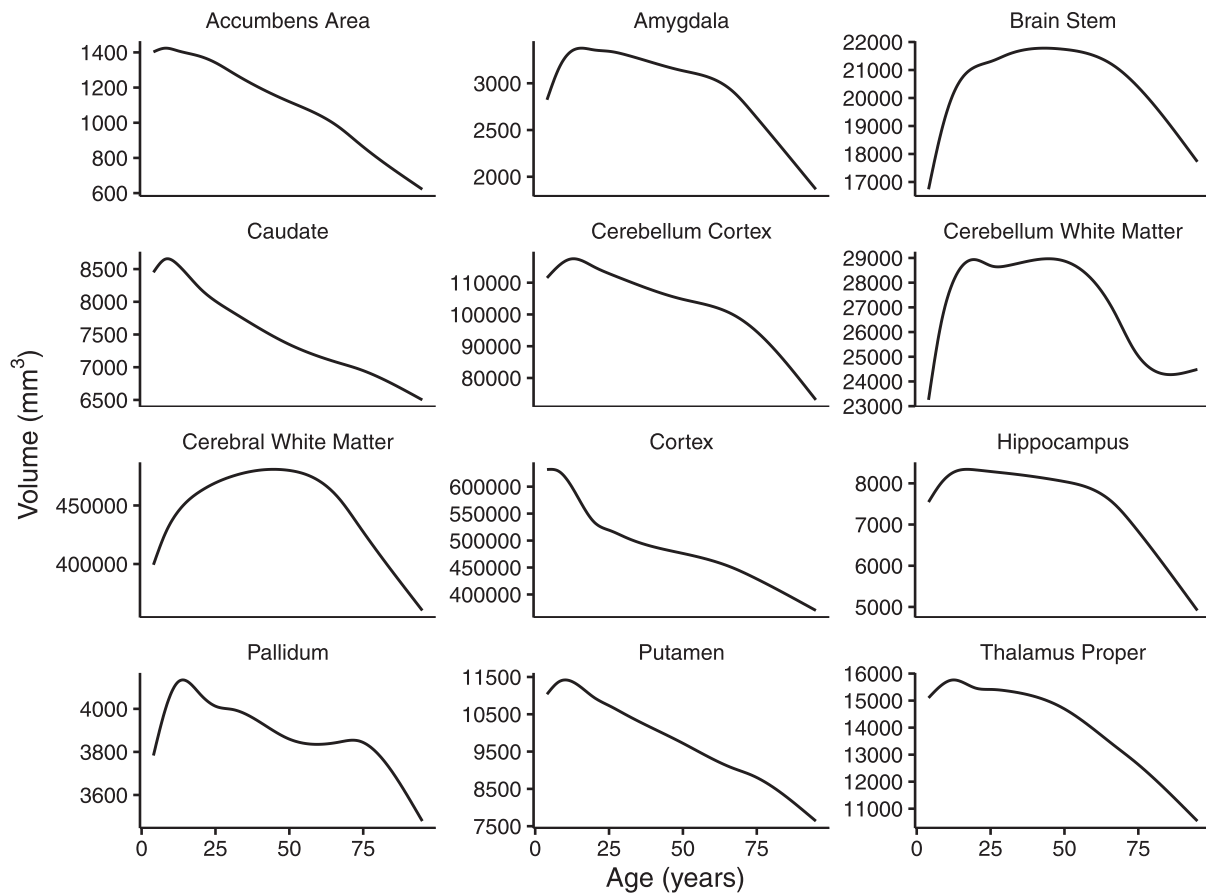


Fig. 5. Lifespan curves. Characteristic curves of 12 brain regions, estimated from the LCBC data and used in simulation experiments.

trated in Fig. 3 (right), both the dates of initial measurements and the times between measurements may be highly varying in lifespan data, and this variability helps identifying the estimates of model (3).

The effect of a time-invariant variable x_i on the age trajectory can be estimated by adding an interaction term to models (1)-(3). If x_i is a continuous variable, the interaction may be a varying-coefficient term of the form introduced in model (3). For (1) and (3) it would be $\beta_x(a_{ij})x_i$, and for (2) it would be $\beta_x(a_i, t_{ij})x_i$, where $\beta_x(\cdot)$ in both cases is a smooth function representing the effect of x_i as a function of either age or baseline age and time. On the other hand, if x_i is a categorical variable with L unique values, it can be encoded as $L - 1$ dummy variables x_{2i}, \dots, x_{Li} , with a varying-coefficient term $\beta_l(a_{ij})x_i$ associated with the l th level for models (1) and (3), and similarly $\beta_l(a_i, t_{ij})x_i$ for model (2). Each varying-coefficient term now represents how the trajectory for the l th level differs from the trajectory for the baseline level. An example is given in Section 3.2.2.

2.4. Simulation experiments

In order to compare the GAMMs (1)-(3), characteristic lifespan curves were estimated for 12 brain regions with the LCBC data, using GAMMs on the form (1), with additional covariates sex, scanner, and total intracranial volume (ICV). Volumes were estimated with FreeSurfer 6.0 (Dale et al., 1999; Fischl et al., 2002; Reuter et al., 2012), and detailed sample characteristics are presented in the Supplementary Section S1. The curves, shown in Fig. 5, were used as ground truths from which measurements were sampled. For each region, three cases were considered: no cohort effects, age-independent cohort effects, and age-cohort interactions. In the latter two cases, cohort effects were added to the characteristic curves as illustrated in Fig. 6 for cerebral white mat-

ter, cortex, and hippocampus, and in Supplementary Section S2.2 for the remaining regions. Data were generated with $n = 1000$ participants, and the number of timepoints m_i for each was multinomially distributed with equal probability of 1, 2, or 3 timepoints. The time between two measurements of a given participant was uniformly distributed between 1 and 6 years, which combined with the maximum number of 3 timepoints set the maximum possible follow-up interval to $(3 - 1) \times 6 = 12$ years. Baseline age was uniformly distributed between 4 and 90 years, and the date of initial measurement was uniformly distributed over 10 years, from 1st January 2000 to 1st January 2010. The simulations were repeated with identical dates of initial measurement. Random intercepts b_{0i} and residuals ϵ_{ij} were sampled from normal distributions with mean zero and standard deviations equal to 50% and 20% of the sample standard deviation of the region's volume, respectively, similar to what was observed in the LCBC data. Datasets for each of the $12 \times 3 = 36$ combinations of regions and cohort effects were randomly sampled 1000 times.

Six models were fitted to each dataset, as summarized by Table 1. The two formulations of model (2) differ in that for (2b), the term $f(a_{i1}, t_{ij})$ is a smooth bivariate function of a_{i1} and t_{ij} defined through the tensor product construction of Wood et al. (2012) (see the left plot in Fig. 11 for an illustration), whereas (2a) uses the stricter formulation $f(a_{i1}, t_{ij}) = f(a_{i1}) + \beta(a_{i1})t_{ij}$, where $f(a_{i1})$ is the main effect of age and $\beta(a_{i1})$ is a varying-coefficient term (Hastie and Tibshirani, 1993) representing the effect of time as a function of age. Model (2a) thus assumes that the effect of time is linear, with a slope that depends on baseline age. Models (3a) and (3b) differ in that (3b) allows the cohort effect to depend on age, with a term $\beta(a_{ij})c_i$ as shown in (3), while (3a) assumes that the cohort effect is age-independent, i.e., $\beta(a_{ij}) = \beta$ for all a_{ij} . See Supplementary Section S2.1 for precise mathematical definitions of the models.




Fig. 6. Simulated cohort effects. Cohort effects used in simulation studies, illustrated for cerebral white matter, cortex, and hippocampus. (Cerebral WM = Cerebral White Matter).

Table 1

Models used in simulation experiments with GAMMs. The 'Identifier' column describes the name used to identify the model in the simulation results presented in Section 3.1 and in the Supplementary Material.

Identifier	Description
(1a)	Model (1) without random effects, using only the first timepoint.
(1b)	Model (1) fitted to the complete data.
(2a)	Model (2) with a varying-coefficient term for the interaction between baseline age and time.
(2b)	Model (2) with a two-dimensional smooth function for jointly modeling the effect of baseline age and time.
(3a)	Model (3) with linear age-independent cohort effects.
(3b)	Model (3) with a varying-coefficient term allowing cohort-age interactions.

3. Results

3.1. Simulation experiments

Fig. 7 shows root-mean-square error (RMSE) of longitudinal estimates for each of the first 12 years after baseline ages 10, 35, and 60 years in the case of no cohort effects. Overall, model (1) with longitudinal data had the most accurate fits, but the two variants of model (3) were close. The two variants of model (2), on the other hand, had poorer fits than the other models, even for times very close after baseline, for which the data contained a large number of observations. Fig. 8 shows the results in the presence of age-cohort interactions. Now model (1) with or without longitudinal data had higher RMSE than model (2) for baseline ages 35 and 60 years, but lower RMSE for baseline age 10 years. Model (3) had the lowest RMSE for all baseline ages in this case. Model (3b), which allows cohort-age interactions, had better overall performance than model (3a), which only contains the cohort effect as a single offset term. Results for age-independent cohort effects and for other regions were similar, and are shown in Supplementary Section S2.3.1.

Table 2 shows the RMSE of the longitudinal estimates 12 years ahead averaged over each year and over baseline ages of 10, 35, and 60 years, along with its bias-variance decomposition $RMSE^2 = Bias^2 + Variance$ (e.g., Hastie et al. (2009)). The bias here quantifies the systematic error made by the model at any given baseline age and time, while the variance represents how much the model fit differs from one dataset to

another². In the absence of cohort effects, model (1b) utilizing longitudinal data had the lowest RMSE and variance across all regions. The two formulations of model (3) had bias close to (1b), but higher variance. This can be attributed to the fact that the cohort terms in (3a) and (3b) are unnecessary in this case, and increase the variance because they increase the number of parameters to be estimated. Models (2a) and (2b) had the highest RMSE and bias for all regions in the absence of cohort effects. With age-independent cohort effects, model (3a) which includes the cohort effect as a single offset term had lower RMSE and bias than the other models across all regions. In this case, model (1b) had the lowest variance, at the cost of a much higher bias than models (3a) and (3b). The two formulations of model (1) and the two formulations of model (2) had considerably higher RMSE than either formulation of model (3), for all regions. Finally, with cohort-age interactions, model (3b) had lower bias and RMSE than the other models for all regions. Model (2b), which contains the sufficient terms to capture such interaction effects, had higher RMSE than both versions of model (3) across all regions. In this case the two formulations of model (1) also performed poorly.

Fig. 9 shows estimated trajectories of hippocampal volume from a baseline age 10 years, from a random subset of 100 out of the total 1000 models fitted to the simulated datasets. While the estimates of

² Table 2 shows the square root of the average squared bias and the square root of the average variance, but for ease of presentation we will refer to these quantities as "bias" and "variance", respectively.

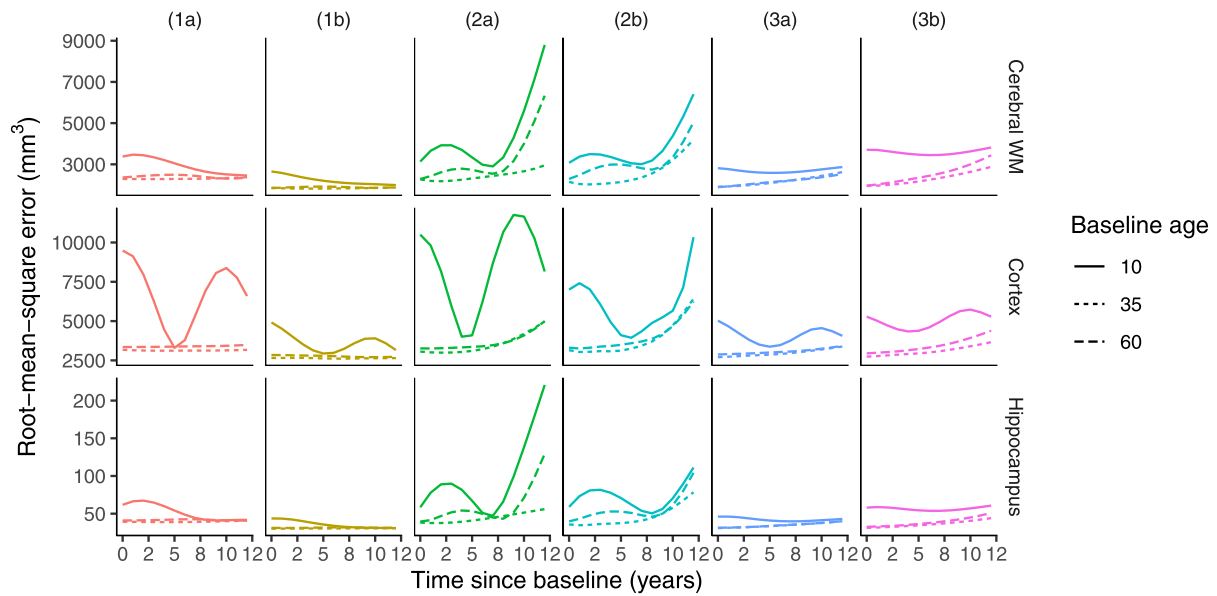


Fig. 7. Longitudinal estimates with no cohort effects. Simulation results in the case of no cohort effects, showing the RMSE of the predicted value after baseline ages 10, 35, and 60 years. For any given time t along the x -axis, the curves represent the RMSE of the predicted longitudinal effect of t years of increased age since baseline. Column headers specify the model fitted to the data, as defined in Table 1.

Table 2

RMSE, bias, and variance of longitudinal estimates averaged over the next 12 years following baseline ages 10, 35, and 60 years. Mean-square error, squared bias, and variance of the prediction were averaged over all Monte Carlo samples for each baseline age and time, and then averaged over baseline ages and times. The square roots of these averages are shown for each region, model, and cohort effect. That is, each cell in column 3 (no cohort effect, RMSE) represents the average of a subplot in Fig. 7, and similarly each cell in column 9 (cohort-age interaction, RMSE) represents the average of a subplot in Fig. 8.

Region	Model	No cohort effect			Age-indep. cohort effect			Cohort-age interaction		
		RMSE	Bias	$\sqrt{\text{Var.}}$	RMSE	Bias	$\sqrt{\text{Var.}}$	RMSE	Bias	$\sqrt{\text{Var.}}$
Cerebral White Matter	(1a)	2575.	933.2	2401.	4454.	3758.	2392.	5071.	4435.	2460.
Cerebral White Matter	(1b)	1995.	419.5	1952.	3801.	3259.	1957.	4487.	4028.	1979.
Cerebral White Matter	(2a)	3639.	2493.	2652.	3789.	2710.	2650.	3842.	2776.	2658.
Cerebral White Matter	(2b)	3281.	1536.	2901.	3646.	2251.	2870.	3726.	2302.	2931.
Cerebral White Matter	(3a)	2376.	407.0	2342.	2339.	418.3	2303.	3129.	2091.	2329.
Cerebral White Matter	(3b)	2885.	641.9	2814.	2858.	662.2	2781.	2860.	666.7	2782.
Cortex	(1a)	4843.	3446.	3405.	6144.	5109.	3414.	7408.	6598.	3369.
Cortex	(1b)	3069.	1128.	2856.	5044.	4164.	2848.	6205.	5509.	2858.
Cortex	(2a)	5907.	4580.	3734.	6060.	4761.	3750.	6263.	4988.	3789.
Cortex	(2b)	4877.	2330.	4287.	5627.	3593.	4332.	5658.	3580.	4384.
Cortex	(3a)	3438.	1160.	3238.	3410.	1093.	3232.	4697.	3351.	3293.
Cortex	(3b)	3966.	1215.	3777.	4045.	1120.	3889.	4052.	1129.	3893.
Hippocampus	(1a)	45.01	19.20	40.73	73.93	61.70	40.74	83.46	72.65	41.11
Hippocampus	(1b)	32.73	7.817	31.80	60.90	51.83	31.99	71.86	64.41	31.87
Hippocampus	(2a)	77.75	64.06	44.09	80.04	66.48	44.60	80.91	67.50	44.64
Hippocampus	(2b)	60.92	35.30	49.68	69.19	46.32	51.41	68.12	44.75	51.39
Hippocampus	(3a)	37.55	7.463	36.82	38.00	6.732	37.41	49.72	33.12	37.10
Hippocampus	(3b)	45.08	9.353	44.12	45.52	8.306	44.78	45.99	8.622	45.20

model (3) follow the true effect over the full 12-year period, a large proportion of the estimates of model (2) are close to straight lines. Due to the simulated dropout after the first or second timepoint, combined with the unstructured time intervals between measurements, the average follow-up time in the data is only 3.5 years, and thus much lower than the maximum follow-up of 12 years. A consequence of this data structure, which resembles the LCBC data shown in Fig. 3, is that for the two formulations of model (2), there is not enough data to estimate the effect of time beyond the first few years after baseline, even though the maximum follow-up interval is 12 years. With a limited amount of data for timepoints further than 3–4 years from baseline, the second derivative penalization used by GAMMs pulls the estimates towards straight lines, which have zero second derivatives and hence are not penalized, an effect which is clearly seen in Fig. 9.

Simulation results with identical baseline dates shown in Supplementary Section S2.4 are practically identical to those described in this section, suggesting that the issue of varying baseline dates is not critical for GAMMs when the maximum variation is ten years. Instead, as Fig. 9 shows, the main challenge with estimating longitudinal effects using the GAMM (2) is caused by the fact that time t_{ij} for most participants spans a short period compared to the full follow-up interval, making estimation of nonlinear effects increasingly challenging as time since baseline increases. For estimation of cross-sectional effects, the differences between models were smaller. However, the two versions of model (2) still showed the poorest performance, while model (1b) and the two versions of model (3) showed the best performance. Detailed results for estimation of cross-sectional effects are shown in Supplementary Section S2.3.2.

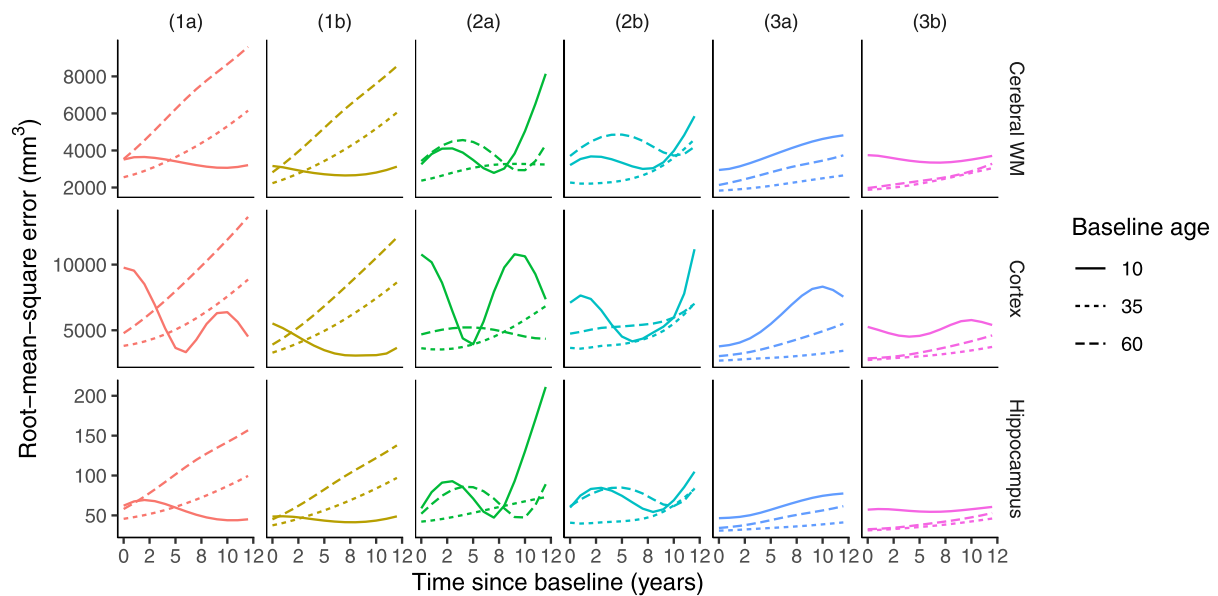


Fig. 8. Longitudinal estimates with cohort-age interactions. Simulation results in the case of age-cohort interactions, showing the RMSE of the predicted value after baseline ages 10, 35, and 60 years. For any given time t along the x-axis, the curves represent the RMSE of the predicted longitudinal effect of t years of increased age since baseline. Column headers specify the model fitted to the data, as defined in Table 1.

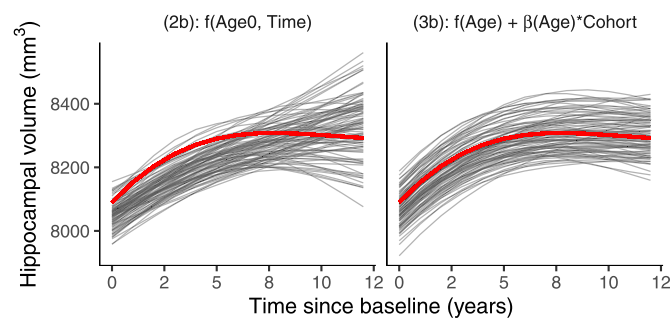


Fig. 9. Sample fits. A random sample of 100 fits in the case of cohort-age interaction for hippocampal volume. Thin lines show estimated longitudinal effects from baseline age 10, and the thick red lines show the true values. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

3.2. Example applications

In this section we show how GAMMs can be applied to the study of lifespan brain development, with example R code using the packages `mgcv` (Wood, 2017) and `gamm4` (Wood and Scheipl, 2017). Some parts of the code are omitted for ease of presentation, and can be found in the Supplementary Section S3.

3.2.1. Modeling lifespan volume trajectories

We first consider the hippocampal volumes shown in Fig. 1 (right). The data are organized in long format in the dataframe `dat`, with each row representing one timepoint of a single participant, containing the following variables:

- `ID`: Unique participant ID.
- `Age`: Age in years at MRI session.
- `Hippocampus`: Estimated hippocampal volume in mm^3 .
- `ICV_z`: Estimated total intracranial volume, standardized to have zero mean and unit variance.
- `Sex`: Participant sex, coded as "Female" and "Male".
- `Scanner`: Factor variable indicating which scanner was used for MRI.
- `Age_bl`: Age in years at initial MRI session.
- `Time`: Time in years since initial MRI session.
- `Birth_Date_z`: Decimal number of years between birth date and 1st January 1970.

The transformed variables with suffix `_z` were created because the algorithms used in the models to be fitted are most stable when the explanatory variables are of similar magnitude.

Models not separating longitudinal and cross-sectional effects We start by fitting a GAM with only the first timepoint of each participant, using the `gam()` function from `mgcv`. By default, `gam()` uses generalized cross-validation (Golub et al., 1979) for smoothing, but for comparison with mixed models we specify that restricted maximum likelihood (REML) should be used, with the argument `method = "REML"`. The smooth function

corresponding to the term $f(a_{ij})$ in model (1) is specified with `s(Age, k = 20, bs = "cr")`, where we use $k = 20$ cubic regression (`bs = "cr"`) splines. The default is thin-plate splines (`bs = "tp"`), but in our experience cubic regression splines typically require half the computing time, without yielding poorer fit. ICV, sex, and scanner are added as additional covariates. Throughout this section, the names of the fitted models correspond to the model identifiers in Table 1.

```
library(mgcv)
# Keep only first timepoint (Time = 0)
baseline_data <- subset(dat, Time == 0)
# Fit GAM to data
mod1a <- gam(Hippocampus ~ s(Age, k = 20, bs = "cr") + ICV_z + Sex + Scanner,
             data = baseline_data, method = "REML")
```

The estimated smooth function can be immediately visualized with the `plot()` function. The output is not shown, but see the curve labeled `mod1a` in Fig. 12 (left).

```
plot(mod1a)
```

We can check that the number of splines is sufficiently high with `k.check()`, which implements a permutation algorithm described in Wood (2017, Ch. 5.9). A significant p -value and estimated degrees of freedom (edf) close to the maximum degrees of freedom k' , indicates that more splines are required. As shown in the output below, k seems sufficiently high. The maximum number of degrees freedom $k' = 19$ is one less than the number of cubic regression splines, because one degree of freedom is used to center the smooth function such that it has zero mean over the range of age values in the data (Wood, 2017, Ch. 5.4.1).

```
k.check(mod1a)
##          k'      edf  k-index p-value
## s(Age) 19 7.862725 0.9946078    0.38
```

Next, we fit model (1) using the complete data. This can be achieved both with the `gam()` and `gamm()` functions from `mgcv`, and the `gamm4()` function from `gamm4`. For the type of longitudinal data considered here, with a low number of repeated measurements of a large number of participants, `gam()` is very slow compared to `gamm()` and `gamm4()`. We opt for the latter, as it is typically faster and more numerically stable. The main difference from the model with only cross-sectional data, is that we now specify a random intercept with the argument `random = ~(1|ID)`.

```
library(gamm4)
mod1b <- gamm4(Hippocampus ~ s(Age, k = 20, bs = "cr") + ICV_z + Sex + Scanner,
              data = dat, random = ~(1|ID))
```

On a MacBook Pro, fitting this GAMM took 1.7 seconds, while fitting the GAM above took less than 0.08 seconds. The `gamm4()` function returns a list with two elements, named `mer` and `gam`. The element `mer` contains information from the `lmer()` function in the `lme4` package (Bates et al., 2015) used in the numerical computations, and is useful for studying the random effect distributions. The element `gam` contains information about the smooth functions, and is useful for studying smooth terms and parametric fixed effects.

```
# Print information about model object:
str(mod1b, max.level = 1)

## List of 2
## $ mer:Formal class 'lmerMod' [package "lme4"] with 13 slots
## $ gam:List of 32
## ..- attr(*, "class")= chr "gam"
```

The `lme4` package is automatically loaded when `gamm4` is loaded, and its accessor functions can be used to study the random effect distributions. The summary below shows that the between-participant variation, $\hat{\sigma}_b = 601 \text{ mm}^3$, is much larger than the within-participant variation, $\hat{\sigma} = 133 \text{ mm}^3$. Note that for `mod1a` which does not have random intercepts, this information is not available. The second line in the output ($\hat{\sigma}_\lambda = 21 \text{ mm}^3$) is related to the formulation of smooth functions as random effects, and the estimated smoothing parameter is given by $\hat{\lambda} = \hat{\sigma}^2 / \hat{\sigma}_\lambda^2 = 38.7$.

```
VarCorr(mod1b$mer)
## Groups   Name                Std.Dev.
## ID      (Intercept)         601.378
## Xr      s(Age)              21.457
## Residual                            133.410
```

Modeling cohort effects Next, we take cohort effects into account by fitting two versions of model (3), `mod3a` which contains a linear cohort effect term, and `mod3b` which contains a varying-coefficient term $\beta(a_{ij})$ consisting of five cubic regression splines, allowing the cohort effect to depend on age.

```
mod3a <- gamm4(Hippocampus ~ s(Age, k = 20, bs = "cr") +
              Birth_Date_z + ICV_z + Sex + Scanner, data = dat, random = ~(1|ID))
mod3b <- gamm4(Hippocampus ~ s(Age, k = 20, bs = "cr") +
              s(Age, by = Birth_Date_z, bs = "cr", k = 5) + ICV_z + Sex + Scanner,
              data = dat, random = ~(1|ID))
```




Fig. 10. Estimated cohort effect. Estimated cohort effect on hippocampal volume as a function of age. Dashed lines show 95% across-the-function CIs, which have the property that the true function is expected to lie within the CI over 95% of the *x*-axis. Dotted lines show 95% simultaneous CIs, which have the property that the true function is expected to be completely confined within the CI 95% of the time under repeated sampling from the population.

We extract the cohort effect estimated by `mod3a` from the matrix `p.table` in the object returned by `mgcv`'s `summary()` function. A 95% confidence interval (CI) is computed by adding the standard error multiplied by the 2.5% and 97.5% quantiles of the standard normal distribution to the estimate. As the output shows, the estimated cohort effect from this model is a negative offset of 1.25 mm³ per birth year, with 95% CI [-4.29, 1.78] mm³.

```
# Extract birth date coefficient from table of parametric estimates
coef_info <- summary(mod3a$gam)$p.table["Birth_Date_z", ]

# Compute 95 % confidence intervals and add them to coef_info
confints <- sapply(c(conf.low = .025, conf.high = .975), function(p)
  coef_info[["Estimate"]] + qnorm(p) * coef_info[["Std. Error"]])

(coef_info <- c(coef_info, confints))

## Estimate Std. Error t value Pr(>|t|) conf.low conf.high
## -1.2544 1.5468 -0.8109 0.4175 -4.2862 1.7774
```

For example, since the variable `Birth_Date_z` represents the number of years between the participant's birth date and 1st January 1970, we can estimate the offset effect of being born in 1970 compared to being born in 1920 by multiplying the estimates in `coef_info` by 50 years³. This is illustrated in the line below, which shows that the estimate is -62.7 mm³ with 95% CI [-214, 89.0] mm³. The upper and lower limits of the CI are small, but not negligible compared to the sample average of 8065 mm³.

```
coef_info[c("Estimate", "conf.low", "conf.high")] * 50

## Estimate conf.low conf.high
## -62.72 -214.31 88.87
```

Next, the varying-coefficient term in `mod3b` is extracted from the matrix `s.table` in the object returned by `mgcv`'s `summary()` function. Its estimated degrees of freedom is 2, implying that the cohort effect is estimated as a straight line defined by an intercept and a slope. Its *p*-value of 0.0506 also suggests that there is some evidence of an age-dependent cohort effect.

```
summary(mod3b$gam)$s.table["s(Age):Birth_Date_z", ]

## edf Ref.df F p-value
## 2.00000 2.00000 2.98508 0.05064
```

The estimated cohort effect can be plotted with the code shown below, using the `plot()` function for `gam` objects. This term is numbered 2 because it was entered as the second smooth term in the formula specifying the model. The argument `scale = 0` ensures that the y-axis limits are adjusted to the term to be plotted, rather than also covering the full range of the term representing the main effect of age.

```
plot(mod3b$gam, select = 2, scale = 0)
```

A slightly modified version of the resulting plot is shown by the solid and dashed lines in Fig. 10. The fact that the estimated cohort effect averaged over all ages is slightly negative is in agreement with `mod3a`, which estimated a negative but non-significant cohort effect. The CIs shown by the `plot()` function have the property that under repeated sampling from the population, the true function will on average be confined within the upper and lower limits over 95% of the *x*-axis (Marra and Wood, 2012; Nychka, 1988). These across-the-function CIs will contain the true function less than 95% of the time under repeated sampling from the population, which explains why the upper limit in Fig. 10 is well below zero despite the *p*-value being larger than 0.05. Simultaneous CIs, on the other hand, would fully contain the complete function 95% of the time under repeated sampling, and can be constructed using a simulation-based approach (Ruppert et al., 2003; Simpson, 2016) shown in the Supplementary Section S3.1. These simultaneous CIs are shown as the dotted lines in Fig. 10, and are wider than the across-the-function CIs. The fact that the upper limit is very close to zero for high ages and the lower limit never is above zero, is in agreement with the *p*-value being approximately 0.05.

³ Since the cohort effect estimated by `mod3a` does not interact with age, the result applies to any set of birth dates separated by 50 years.

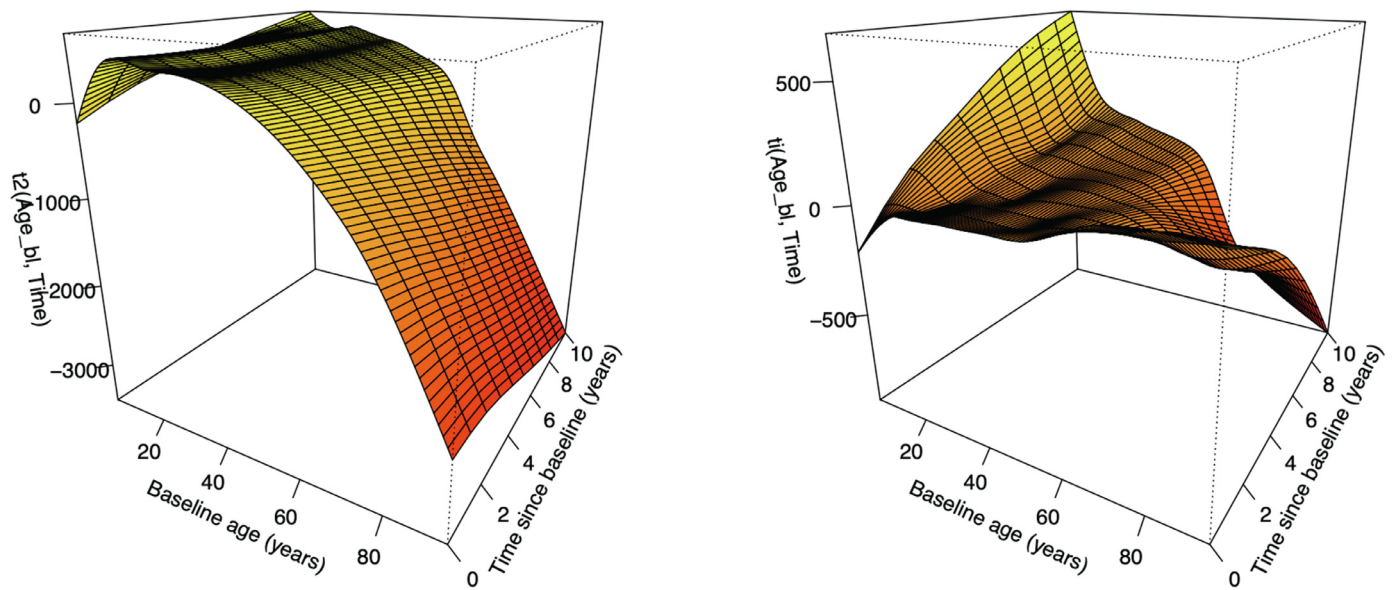


Fig. 11. Two-dimensional smooth functions. Left: tensor product term $t_2(\text{Age_bl}, \text{Time})$ in mod2b , representing the total effect of baseline age and time. Right: tensor interaction term $t_i(\text{Age_bl}, \text{Time})$ in mod2b_ti , representing the interaction effect between baseline age and time.

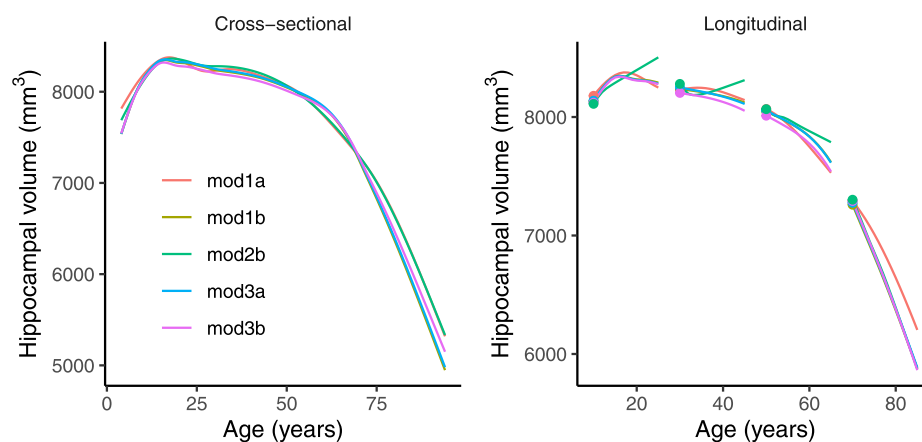


Fig. 12. Model estimates. Estimated cross-sectional and longitudinal effects from each of the five models considered in Section 3.2.1. The cross-sectional estimates are computed for 1st January 2010. The longitudinal estimates are computed 15 years ahead from baseline ages of 10, 30, 50, and 70 years.

The age-dependent cohort effects estimated by mod3b imply that a participant born in 1920 is expected to have a 131 mm^3 lower hippocampal volume at age 20 than a participant born in 1970 at age 20 (CI: $[-407, 145] \text{ mm}^3$). Conversely, a participant born in 1920 is expected to have a 340 mm^3 higher hippocampal volume than a participant born in 1970, at age 70 (CI: $[-5.8, 685] \text{ mm}^3$). R code for computing these estimates is shown in Supplementary Section S3.1.1. As for mod3a , these results suggest that a cohort effect cannot be ruled out, despite the term not being significant, since cohort effects of relatively large magnitude are contained within the 95% CIs.

Modeling baseline age and time since baseline Finally, we fit model (2), using the $t_2()$ function to create a two-dimensional smooth term (Wood et al., 2012). The argument $k = c(20, 5)$ specifies that 20 cubic regression splines are used for the effect of baseline age, while only 5 are used for the effect of time, as this term does not span more than 11 years. The construction of the two-dimensional function involves forming products of all combinations of splines for baseline age and time, implying that the total number of degrees of freedom used by the term equals $20 \times 5 - 1 = 99$, where the one degree of freedom subtracted has been used for imposing a sum-to-zero constraint. Fitting mod2b below took ≈ 90 seconds on a MacBook Pro.

```
mod2b <- gamm4(Hippocampus ~ t2(Age_bl, Time, k = c(20, 5), bs = "cr") +
               ICV_z + Sex + Scanner, data = dat, random = ~(1|ID))
```

Model (2) in Section 2.3 could alternatively be formulated with three smooth terms: the main effects of age and time, and their interaction. This allows significance testing of each term separately. Functionality for fitting such a model is provided by mgcv 's $t_i()$ function, representing a two-dimensional tensor interaction term in which the main effects have been removed (Wood, 2006). As it is not available in gamm4 , the $\text{gamm}()$ function needs to be used. The syntax is very similar, except that the random intercept is specified with $\text{random} = \text{list}(\text{ID} = \sim 1)$ and the use of REML rather than the default marginal maximum likelihood is specified with $\text{method} = \text{"REML"}$ for comparability with the models fitted with $\text{gamm4}()$.

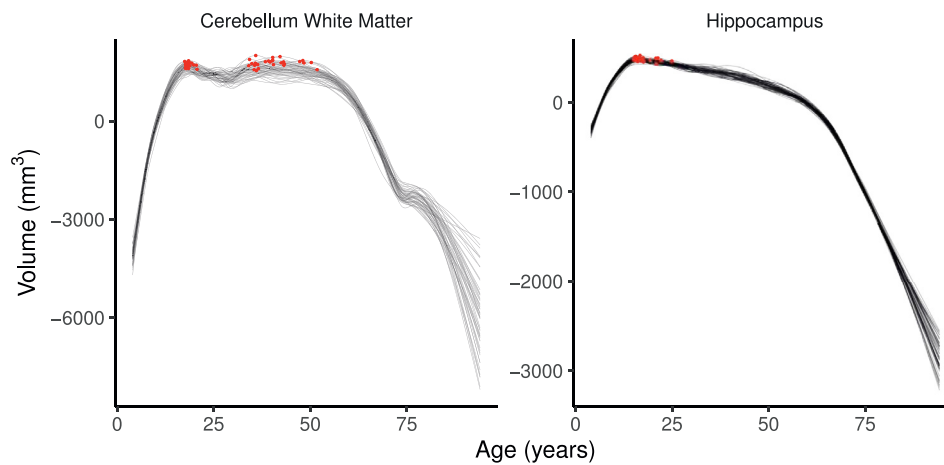


Fig. 13. Posterior samples. The plots show 50 samples from the posterior distributions of curves for lifespan volumes of cerebellum white matter and hippocampus. Red dots indicate the maximum of each curve.

```
# Alternative formulation with tensor interaction terms
mod2b_ti <- gamm(Hippocampus ~ s(Age_bl, k = 20, bs = "cr") +
  s(Time, k = 5, bs = "cr") +
  ti(Age_bl, Time, k = c(20, 5), bs = "cr") +
  ICV_z + Sex + Scanner,
  data = dat, random = list(ID = ~ 1), method = "REML")
```

Information about the smooth terms can again be obtained from `s.table` returned by `summary()`, and shows that all terms are significant. Interestingly, the main effect of time is estimated to be linear, as can be seen by its single degree of freedom, while the two-dimensional interaction term is highly nonlinear.

```
summary(mod2b_ti$gam)$s.table
```

##	edf	Ref.df	F	p-value
## s(Age_bl)	7.648	7.648	138.26	0
## s(Time)	1.000	1.000	45.44	0
## ti(Age_bl,Time)	23.074	23.074	67.46	0

Two-dimensional smooth terms can also be visualized with the `plot()` function, for which a perspective plot is produced by setting `scheme = 1`.

```
plot(mod2b$gam, scheme = 1) # Plot full tensor product
plot(mod2b_ti$gam, select = 3, scheme = 1) # Plot tensor interaction, term #3
```

The resulting plots are shown in Fig. 11. Considering the left part of the plot, the cross-sectional effect is visualized along the baseline age axis, with a trajectory similar to the lifespan hippocampal volume shown in Fig. 5. The longitudinal effect, plotted along the time axis, is positive for low baseline ages and negative for higher baseline ages. The tensor interaction term plotted in the right part of Fig. 11 shows that the effect of time is positive in the youngest participants, quite flat in adults, and negative in the oldest participants. Note that the left and right plots in Fig. 11 are not comparable. Since the right plot is a pure interaction term, the direction of the estimated total longitudinal effect cannot be evaluated based on Fig. 11 (right) alone, but also needs to take the main effect of time into account.

Comparison of model fits

Fig. 12 shows estimated cross-sectional and longitudinal effects from the five models estimated in this section. The cross-sectional effects are estimated for 1st January 2010, and are all quite similar. Model (1a), estimated with only baseline measurements, indicates a less steep growth during childhood, and also exhibits some wiggleness between the age of 20 and the age of 50. The longitudinal effects are estimated 15 years ahead from baseline ages of 10, 30, 50, and 70 years. Models (1a) and (1b) do not distinguish longitudinal and cross-sectional effects, and hence have identical estimates in both plots. The estimates from model (2b) are quite different from those of the other models, except for a baseline age of 70. Given the simulation results of Section 3.1, we suspect that the estimates from (2b) are not accurate. The longitudinal estimates of models (3a) and (3b) are close to the estimates of model (1b), again suggesting that the cohort effects in these data are moderate.

Estimating age at maximum volume A question of interest when estimating lifespan curves, is the age at which critical points occur, e.g. the age of maximum volume, maximum growth, or maximum decline. Point estimates of such critical ages can be read directly from the fits, if necessary after computing derivatives, but an assessment of their statistical uncertainty is not directly available. A Bayesian view of the smoothing introduced by Kimeldorf and Wahba (1970) lets us achieve this. Letting $\hat{\beta}$ denote the estimated regression parameters, including spline weights, and $\hat{\Sigma}_{\beta}$ their covariance matrix, the posterior distribution of the true coefficients β is now a normal distribution with mean $\hat{\beta}$ and covariance $\hat{\Sigma}_{\beta}$, $\beta|y \sim N(\hat{\beta}, \hat{\Sigma}_{\beta})$ (Wood, 2017, Ch. 6.10). By sampling from this posterior distribution we can make confidence statements about any quantity derived from the smooth functions. As an example, Fig. 13 shows 50 samples from the posterior distribution of volume curves for cerebellum white matter and hippocampus, with the maximum of each marked with a red dot. Even from these small samples it is evident that there is high uncertainty about the age at which cerebellum white matter volume is maximal, while there is less uncertainty about hippocampal volume.

By sampling 20000 curves and locating the age at maximum for each, we obtained posterior distributions of the age at maximum for each region. Fig. 14 shows 95% highest posterior density intervals computed from the posterior distributions for all 12 regions, using the HDInterval package

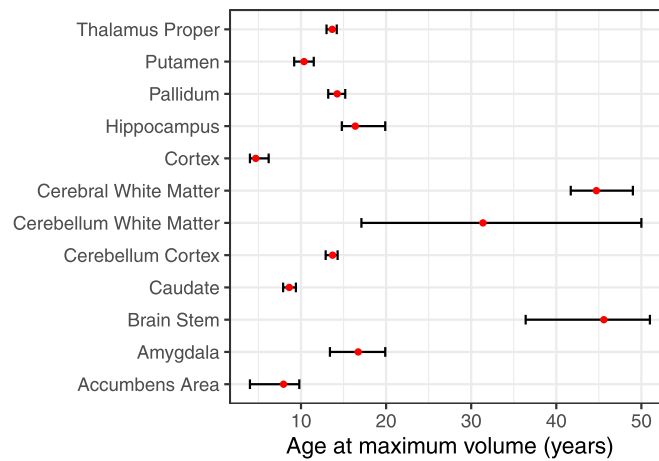


Fig. 14. Age at maximum volume. 95% highest posterior density intervals for the age at maximum volume of 12 brain regions. Red dots show posterior means.

(Meredith and Kruschke, 2019). The plot shows that the uncertainty about the location of the maximum is highly variable between regions, with very narrow intervals for, e.g. caudate, cerebellum cortex, and thalamus proper, and wide intervals and high uncertainty, e.g. for the brain stem and cerebellum white matter.

3.2.2. Interaction effects on lifespan volume trajectories

We now demonstrate how factor-smooth interactions can be used to study how lifespan brain volumes are affected by a categorical variable. As an example application we consider the apolipoprotein E (APOE) $\epsilon 4$ allele, which is a known risk factor for Alzheimer’s disease (Corder et al., 1993; Genin et al., 2011), and study how lifespan trajectories of cerebellum cortex volume differ between carriers of zero, one, or two APOE $\epsilon 4$ alleles. Similar models were used by Walhovd et al., 2020, who studied the impact of the APOE $\epsilon 4$ allele on lifespan hippocampal volume.

The data are still contained in a dataframe named `dat`, with identical structure to the data used in Section 3.2.1, except that the variable representing hippocampal volume now is replaced by the variable `Cerebellum` representing cerebellum cortex volume in mm^3 . In addition, a new variable `Gene_APOEnE4` represents the total number of APOE $\epsilon 4$ alleles. After excluding participants without information about APOE status, `dat` contained 2707 observations of 1139 participants. Of these, 764 (1,838 observations) had zero alleles, 341 (789 observations) had 1 allele, and 34 (80 observations) had two alleles.

Factor smooths In order to estimate the interaction effects, the variable `Gene_APOEnE4` needs to be coded as an ordered factor. This is done with the following code.

```
dat$Gene_APOEnE4 <- ordered(dat$Gene_APOEnE4)
# Print the levels of the ordered factor
levels(dat$Gene_APOEnE4)

## [1] "0" "1" "2"
```

A factor-smooth interaction is defined by `s(Age, by = Gene_APOEnE4, k = 10, bs = "cr")`. For an ordered factor variable with L levels, this term creates $L - 1$ smooth functions, each representing the difference between the trajectory associated with the l th level ($l = 2, \dots, L$) and the trajectory associated with the baseline level $l = 1$. The difference between trajectories does not include pure offset effects, and hence the main effect of the ordered factor must be added. In this case, two smooth factor interaction terms are created, associated with 1 or 2 APOE $\epsilon 4$ alleles. In contrast, if `Gene_APOEnE4` was a numeric variable, `gamm4()` would estimate a varying-coefficient term as used in Section 3.2.1 treating the number of alleles as a continuous variable, and if `Gene_APOEnE4` was a factor variable a single smooth term would be independently estimated for each of the three factor levels and the main effect term `s(Age, k = 10, bs = "cr")` would have to be omitted for the model to be identified.

```
mod <- gamm4(Cerebellum ~ s(Age, k = 10, bs = "cr") +
  s(Age, by = Gene_APOEnE4, k = 10, bs = "cr") +
  Gene_APOEnE4 + ICV_z + Sex + Scanner, data = dat, random = ~(1|ID))
```

Again, the matrix `s.table` returned by `summary()` can be used to study the smooth terms. The first term represents the trajectory for non-carriers, and the two terms starting with `s(Age):Gene_APOEnE4` represent the difference between trajectories of carriers of one or two alleles to carriers of zero alleles, respectively. From the p -values for the interaction terms, it is clear that there is no evidence that the shape of the lifespan cerebellum white matter volume depends on APOE $\epsilon 4$ status.

```
summary(mod$gam)$s.table

##           edf Ref.df      F p-value
## s(Age)      8.436  8.436 103.456  0.000
## s(Age):Gene_APOEnE41 1.000  1.000  0.424  0.515
## s(Age):Gene_APOEnE42 2.446  2.446  2.196  0.187
```

The main effects of APOE $\epsilon 4$ status can be extracted from `p.table` returned by `summary()`. We use R’s `grep()` function to retain only terms containing the pattern `"Gene_APOE"`. The estimates are less than one standard error from zero and not significant, indicating that there is no




Fig. 15. Factor smooth interactions. Estimated lifespan trajectories of cerebellum cortex volume for males and females with 0, 1, or 2 APOE ϵ 4 alleles.

evidence for an offset effect of APOE ϵ 4 status on cerebellum cortex volume. The suffixes .L ('linear') and .Q ('quadratic') are a consequence of how R treats ordered factors, and represent the offset effect of having one or two alleles, respectively, relative to having zero alleles.

```
ptab <- summary(mod$gam)$p.table
# Extract terms containing "Gene_APOE"
ptab[grep("Gene_APOE", rownames(ptab)), ]

##           Estimate Std. Error t value Pr(>|t|)
## Gene_APOEnE4.L   -766.3     1103.5 -0.6944  0.4875
## Gene_APOEnE4.Q   -509.8       754.3 -0.6758  0.4992
```

Prediction from GAMMs Creating predictions from GAMMs aids interpretation of the estimated effects, and we illustrate it here by comparing the estimated lifespan cerebellum cortex volumes for participants with zero, one, or two APOE ϵ 4 alleles. First, a grid over which to compute the predictions is created. Using `expand.grid()`, all combinations of ages between 4 and 94 years with a spacing of 0.1 years, number of APOE ϵ 4 alleles, and sexes are generated. The `predict()` function requires all variables in the model to be defined, and we hence set `ICV_z` equal to the sample mean and `Scanner` arbitrarily to "ousAvanto", which is one of the scanners used in the LCBC data. Other values of `ICV_z` and `Scanner` would shift the resulting curves vertically, but the interpretation would not change.

```
# Create grid with all combinations of ages and APOE ε4 alleles
grid <- expand.grid(
  Age = seq(from = 4, to = 94, by = .1), Gene_APOEnE4 = ordered(0:2),
  Sex = factor(c("Female", "Male")), ICV_z = 0, Scanner = "ousAvanto")
```

Next, predictions are computed at all values of the grid and `ggplot2` (Wickham, 2016) is used to plot the predicted values.

```
grid$fit <- predict(mod$gam, newdata = grid) # Compute fit on grid

library(ggplot2) # Plot with grouping by sex and number of alleles
ggplot(grid, aes(x = Age, y = fit, group = interaction(Gene_APOEnE4, Sex),
  color = Gene_APOEnE4, linetype = Sex)) + geom_line()
```

A slightly modified version of the resulting plot is shown in Fig. 15. Note that since `Sex` is a parametric term, it merely shifts the curves vertically, without changing their shapes. In this example, none of the interaction effects were significant, but the plots still show how smooth interaction terms create different functional shapes depending on the number of APOE ϵ 4 alleles.

4. Discussion

This paper has highlighted that GAMMs are well-suited for estimating lifespan brain trajectories. However, the issue of potential cohort effects requires careful consideration, and direct translation of LMM formulations used for separating longitudinal and cross-sectional effects has potential pitfalls. In Section 2.3 we defined the "age model" (1) which ignores cohort effects, the "age-time model" (2) which is a direct extension of LMMs commonly used to separate longitudinal and cross-sectional effects, and the alternative "age-cohort model" (3) which includes participant birth date as a model term. These models' abilities to accurately estimate longitudinal and cross-sectional effects were compared in realistic simulation experiments reported in Section 3.1. Not surprisingly, in the absence of cohort effects the age model was most accurate, with the version using longitudinal data performing better than the equivalent model with only cross-sectional data. With cohort effects, on the other hand, the age cohort model was most accurate. More importantly, the age-time model – which may be seen as a "classic" model used to separate longitudinal and cohort effects – consistently performed worse than the age-cohort model, and as shown in Table 2 it had both higher bias and variance across the simulated samples. As also suggested by Figs. 7 and 8, the age-time model is not able to estimate nonlinear longitudinal effects beyond short follow-up intervals. On the other hand, in the special case of linear longitudinal effects, from which the age-time model originates, longitudinal effects for any time after baseline can in principle be accurately estimated with follow-up intervals of arbitrary length. Interpretation of the terms in the age-time model as longitudinal and cross-sectional effects also requires that all participants have equal dates of initial measurement. However, simulations reported in Supplementary Section S2.4 suggest that varying baseline dates have a very small effect on the accuracy of the age-time model in the settings considered here.

While limitations will vary with regard to study specific characteristics, we find it important to emphasize, in light of the present findings, that age-time models will never be able to make accurate estimates of lifespan trajectories, or even trajectories for any substantial part of the lifespan, unless a majority of participants have been followed over the whole interval of interest. Currently available human cohorts with longitudinal imaging data do not span the desired intervals. Furthermore, reaching acceptable power is impossible if dates of initial measurement are to be contained within a small fraction of time. A look at some of the most powerful and impressive combined cross-sectional and longitudinal studies of brain changes with age, suggests that fractional follow-up interval and range of variation in initial measurement dates realistically need to be accommodated in all statistical models. For instance, even the ABCD study (Casey et al., 2018) which utilizes numerous scan sites to track development in thousands of children at very similar age, need to allow for some variation in initial date of measurement, and follow-up intervals are so far limited to a couple of years. While there luckily are major and most impressive studies that contain information on participant samples for many decades, such as the Whitehall study (Filippini et al., 2014), the Baltimore Longitudinal Study of Aging (Tian et al., 2015), the Betula (Gorbach et al., 2017), or the Lothian Birth Cohort study (Cox et al., 2018) these still await longitudinal imaging data (Filippini et al., 2014), or typically have MRI data only for a small fraction of the time, less than a decade (Cox et al., 2018; Gorbach et al., 2017; Tian et al., 2015), sometimes with scan waves being completed across several years (Gorbach et al., 2017). We note this not as a critique of any study, but as a reminder that statistical models at the very least need to accommodate the realistic situation for the best possible data.

4.1. When is a single measurement per participant sufficient?

The dangers of using cross-sectional data have repeatedly been pointed out in the quantitative psychology literature. For example, mediation analysis using purely cross-sectional data is likely to lead to biased and misleading estimates under realistic conditions (Cole and Maxwell, 2003; Lindenberger et al., 2011; Lindenberger and Pötter, 1998). In mediation analysis, the goal is to understand the causal paths through which one or more variables x influence an outcome y , directly or through one or more mediating variables m . Since a cause precedes its effects, carefully designed longitudinal data collection as well as models capable of utilizing this information are then necessary (Collins et al., 1998). Longitudinal data is also required to understand how within-individual change differs from between-individual change and how within-individual change is correlated across multiple processes (Lindenberger et al., 2011; Molenaar, 2004). Traditionally, such studies have been conducted by following a group of participants of similar age over a number of waves, e.g. Cox et al. (2020); Raz et al. (2010, 2005).

While the above mentioned cautions about use of cross-sectional data are completely justified, they do not necessarily extrapolate to estimation of lifespan trajectories. If the goal is to estimate the population effect of aging on the volume of one or more brain regions, potentially including interaction effects of static trait variables like genetic variations or education level (after completed education), a single measurement per participant may be sufficient. One example is when the strong assumption of no cohort effects is made. If it holds, cross-sectional and longitudinal effects are equal, and both can be accurately estimated by the age model using purely cross-sectional data. However, with sufficient variation in baseline dates, the age-cohort model is in principle able to estimate longitudinal and cross-sectional effects using a single measurement per participant. This approach has been used in studies of cognitive aging, which allows estimation of aging and cohort effects without the risk of confounding by retest effects (Horn and Donaldson, 1976; Salthouse, 2013; 2014; 2019; Schaie et al., 1973). In practice, however, we have experienced that cohort-age models become more

stable and accurate with longitudinal data. In particular, both the additional variation provided by repeated measurements with heterogeneous follow-up intervals and the correlation between repeated measurements of the same participant likely contribute to better separation of age effects and cohort effects. Furthermore, as shown in Section 3.2.1, a GAMM using longitudinal data also estimates the between-individual variation and the within-individual variation, quantifying the extent to which differences between participants are due to systematic variation and noise, respectively.

4.2. Limitations and future directions

The GAMMs studied in this paper also have some limitations. The age-cohort model is not identified if all participants have been measured at the exact same dates, since age and cohort then are perfectly collinear. This is also true with longitudinal data, and emphasizes the fact that these models are developed for heterogeneous data, typically combined from multiple studies.

There is also a need for methodological development related to estimation of correlated change between regions. In principle, this could be done by fitting GAMMs separately for each region, and using the correlation of random slopes across regions as an estimate of correlated change. A more principled approach is offered by joint modeling frameworks for LMMs (Fieuws and Verbeke, 2004; 2006), which in this case would amount to fitting a single hierarchical GAMM (Pedersen et al., 2019) for the lifespan trajectories of multiple regions, with interaction terms distinguishing trajectories for each region and random effect structures modeling the within-individual level and change correlation between region trajectories. However, the fact that the extent of correlated change between any pair of brain regions is likely to vary across the lifespan would also need to be taken into account, e.g. by modeling correlations as functions of age. Combined with the need for three or more timepoints to accurately estimate random slopes, it may currently be challenging to obtain a sufficient amount of longitudinal data for fitting such models.

A limitation of GAMMs is that the parameters of the estimated smooth terms typically are not interpretable, and quantities such as rates of change, initial levels, and final asymptotic levels have to be inferred from the estimated functions, e.g. as demonstrated for age at maximum volume in Figs. 13 and 14 in Section 3.2.1. When prior knowledge about the functional form of the phenomenon under study is available, nonlinear mixed models are an attractive alternative, offering directly interpretable parameters of substantive interest (Davidian, 2009; Lindstrom and Bates, 1990; Ram and Grimm, 2016). Applications include modeling of learning curves (Cudeck, 1996; Ghisletta et al., 2010), modeling the effect of preschool instructions on academic achievement using sigmoidal functions (Grimm and Ram, 2009), and modeling of rate of change and acceleration in a lexical retrieval task (Grimm et al., 2013).

While we have considered GAMMs for estimating how time-invariant variables interact with lifespan trajectories in Section 3.2.2, interaction with time-dependent variables may also be of interest. Although time-dependent interaction variables can be used within the framework considered here, the interpretation of the estimated effects becomes more challenging. If only the value of the time-dependent variable at the given timepoint affects the outcome, the effect can be interpreted in exactly the same way as for a time-invariant variable (Fitzmaurice et al., 2011, Ch. 13.5). In many applications, however, it may be more plausible to assume that also the variable's change since the previous timepoint contains relevant information, and in this case the models used in Section 3.2.2 will have biased estimates. Continuous-time SEMs (Driver and Voelkle, 2018; Oud and Jansen, 2000) may be useful for this purpose, as they allow regressing a time-dependent process (the outcome of interest) on the value of another time-dependent process (the interaction variable) at earlier times, without the restrictive assumption of equally spaced time intervals imposed by ordinary SEMs. However, estimation of nonlinear smooth functions within continuous-time SEMs has

not been reported in the literature, and is likely to be both computationally challenging and require large longitudinal datasets.

5. Conclusion

GAMMs are attractive tools for estimating lifespan brain trajectories, which flexibly handle the nonlinear effects and variable follow-up intervals and measurement dates characteristic of lifespan data. If cohort effects are negligible, age models on the form (1) yield the most accurate estimates, and in this case a single measurement per participant may even be sufficient. More realistically, cohort effects are likely to be present, and in this case the age-cohort model (3) which directly models the effect of birth cohort is able to accurately estimate longitudinal and cross-sectional effects. On the other hand, the age-time model (2) which separates the effect of age into a baseline term and a time term as is common with LMMs, yields poor estimates of longitudinal effects. With sufficient variation of measurement dates and follow-up intervals, we thus recommend the age-cohort model for estimating lifespan brain trajectories. On the other hand, for time-structured data containing little variation in measurement dates, the age-cohort model is not identified and the age-time model seems to be the best option, with the caveat that estimated longitudinal effects may not be reliable for times larger than the average follow-up interval, as will also be apparent from the confidence intervals.

The R packages `mgcv` and `gamm4` provide efficient software for fitting GAMMs, and are complemented by additional packages enabling easy visualization and interpretation of summary statistics.

Data and Code Availability Statement

Code for running simulation experiments is available in Supplementary Section S4 and at the Open Science Framework, <https://osf.io/xnuhz/>. The LCBC data are available from the corresponding author on reasonable request, given appropriate ethical and data protection approvals.

Ethics Statement

Collection of the LCBC was approved by a Norwegian Regional Committee for Medical and Health Research Ethics South. Written informed consent was obtained from all participants of age at majority, and the parents or guardian for children below this age.

Declaration of Competing Interest

The authors report no competing interests.

Credit authorship contribution statement

Øystein Sørensen: Conceptualization, Methodology, Software, Formal analysis, Writing - original draft, Writing - review & editing, Visualization. **Kristine B. Walhovd:** Writing - original draft, Writing - review & editing. **Anders M. Fjell:** Conceptualization, Writing - original draft, Writing - review & editing, Visualization, Supervision.

Acknowledgement

The authors thank Andreas Brandmaier, Paolo Ghisletta, and Magne Thoresen for helpful discussions. The data collection was supported by the [European Research Council](#) under grant agreements [283634](#), [725025](#) (to A.M.F.) and [313440](#) (to K.B.W.), the [Norwegian Research Council](#) (to A.M.F., K.B.W.), The National Association for Public Health's dementia research program, Norway (to A.M.F) and center support from the University of Oslo.

Supplementary material

Supplementary material associated with this article can be found, in the online version, at [10.1016/j.neuroimage.2020.117596](https://doi.org/10.1016/j.neuroimage.2020.117596).

References

- Baltes, P.B., 1968. Longitudinal and cross-sectional sequences in the study of age and generation effects. *Hum. Dev.* 11 (3), 145–171. doi:[10.1159/000270604](https://doi.org/10.1159/000270604).
- Bates, D., Mächler, M., Bolker, B., Walker, S., 2015. Fitting linear mixed-effects models using lme4. *J. Stat. Softw.* 67 (1), 1–48. doi:[10.18637/jss.v067.i01](https://doi.org/10.18637/jss.v067.i01).
- Bearden, C.E., Thompson, P.M., 2017. Emerging global initiatives in neurogenetics: The enhancing neuroimaging genetics through meta-analysis (ENIGMA) consortium. *Neuron* 94 (2), 232–236. doi:[10.1016/j.neuron.2017.03.033](https://doi.org/10.1016/j.neuron.2017.03.033).
- Brumback, B.A., Rice, J.A., 1998. Smoothing spline models for the analysis of nested and crossed samples of curves. *J. Am. Stat. Assoc.* 93 (443), 961–976. doi:[10.1080/01621459.1998.10473755](https://doi.org/10.1080/01621459.1998.10473755).
- Casey, B.J., Cannonier, T., Conley, M.I., Cohen, A.O., Barch, D.M., Heitzeg, M.M., Soules, M.E., Teslovich, T., Dellarco, D.V., Garavan, H., Orr, C.A., Wager, T.D., Banich, M.T., Speer, N.K., Sutherland, M.T., Riedel, M.C., Dick, A.S., Bjork, J.M., Thomas, K.M., Charani, B., Mejia, M.H., Hagler, D.J., Cornejo, M.D., Sicut, C.S., Harms, M.P., Dosenbach, N.U.F., Rosenberg, M., Earl, E., Bartsch, H., Watts, R., Polimeni, J.R., Kuperman, J.M., Fair, D.A., Dale, A.M., 2018. The adolescent brain cognitive development (ABCD) study: Imaging acquisition across 21 sites. *Developmental Cognitive Neuroscience, The Adolescent Brain Cognitive Development (ABCD) Consortium: Rationale, Aims, and Assessment Strategy* 32, 43–54. doi:[10.1016/j.dcn.2018.03.001](https://doi.org/10.1016/j.dcn.2018.03.001).
- Cole, D.A., Maxwell, S.E., 2003. Testing mediational models with longitudinal data: questions and tips in the use of structural equation modeling. *J. Abnorm. Psychol.* 112 (4), 558–577. doi:[10.1037/0021-843x.112.4.558](https://doi.org/10.1037/0021-843x.112.4.558).
- Collins, L.M., Graham, J.J., Flaherty, B.P., 1998. An alternative framework for defining mediation. *Multivariate Behav. Res.* 33 (2), 295–312. doi:[10.1207/s15327906mbr3302_5](https://doi.org/10.1207/s15327906mbr3302_5).
- Corder, E., Saunders, A., Strittmatter, W., Schmechel, D., Gaskell, P., Small, G., Roses, A., Haines, J., Pericak-Vance, M., 1993. Gene dose of apolipoprotein e type 4 allele and the risk of Alzheimer's disease in late onset families. *Science* 261 (5123), 921–923. doi:[10.1126/science.8346443](https://doi.org/10.1126/science.8346443).
- Cox, S.R., Allerhand, M., Ritchie, S.J., Muñoz Maniega, S., Valdés, M., Harris, S.E., Dickie, D.A., Anlagan, D., Aribisala, B.S., Morris, Z., Sherwood, R., Abbott, N.J., Starr, J.M., Bastin, M.E., Wardlaw, J.M., Deary, I.J., 2018. Longitudinal serum s100 β and brain aging in the Lothian Birth Cohort 1936. *Neurobiology of Aging* 69, 274–282. doi:[10.1016/j.neurobiolaging.2018.05.029](https://doi.org/10.1016/j.neurobiolaging.2018.05.029).
- Cox, S.R., Harris, M.A., Ritchie, S.J., Buchanan, C.R., Hernández, M.V., Corley, J., Taylor, A.M., Madole, J.W., Harris, S.E., Whalley, H.C., McIntosh, A.M., Russ, T.C., Bastin, M.E., Wardlaw, J.M., Deary, I.J., Tucker-Drob, E.M., 2020. Three major dimensions of human brain cortical ageing in relation to cognitive decline across the 8th decade of life. *bioRxiv* doi:[10.1101/2020.01.19.911420](https://doi.org/10.1101/2020.01.19.911420).
- Cudeck, R., 1996. Mixed-effects models in the study of individual differences with repeated measures data. *Multivariate Behavioral Research* 31 (3), 371–403. doi:[10.1207/s15327906mbr3103_6](https://doi.org/10.1207/s15327906mbr3103_6).
- Curran, P.J., Bauer, D.J., 2011. The disaggregation of within-person and between-person effects in longitudinal models of change. *Annu. Rev. Psychol.* 62 (1), 583–619. doi:[10.1146/annurev.psych.093008.100356](https://doi.org/10.1146/annurev.psych.093008.100356).
- Dale, A.M., Fischl, B., Sereno, M.I., 1999. Cortical surface-based analysis. *Neuroimage* 9 (2), 179–194. doi:[10.1006/nimg.1998.0395](https://doi.org/10.1006/nimg.1998.0395).
- Davidian, M., 2009. Non-linear Mixed Effects Models. In: Fitzmaurice, G., Davidian, M., Verbeke, G., Molenberghs, G. (Eds.), *Longitudinal Data Analysis*. Chapman and Hall, London, U.K., pp. 291–316.
- Diggle, P., Heagerty, P., Liang, K.-Y., Zeger, S., 2002. *Analysis of longitudinal data 2nd edition*. Oxford University Press, Oxford, U.K..
- Driver, C.C., Voelkle, M.C., 2018. Hierarchical Bayesian continuous time dynamic modeling. *Psychol. Methods* 23 (4), 774–799. doi:[10.1037/met0000168](https://doi.org/10.1037/met0000168).
- Durbán, M., Harezlak, J., Wand, M.P., Carroll, R.J., 2005. Simple fitting of subject-specific curves for longitudinal data. *Stat. Med.* 24 (8), 1153–1167. doi:[10.1002/sim.1991](https://doi.org/10.1002/sim.1991).
- Edwards, L.J., Stewart, P.W., MacDougall, J.E., Helms, R.W., 2005. A method for fitting regression splines with varying polynomial order in the linear mixed model. *Stat. Med.* 25 (3), 513–527. doi:[10.1002/sim.2232](https://doi.org/10.1002/sim.2232).
- Fieuws, S., Verbeke, G., 2004. Joint modelling of multivariate longitudinal profiles: pitfalls of the random-effects approach. *Stat. Med.* 23 (20), 3093–3104. doi:[10.1002/sim.1885](https://doi.org/10.1002/sim.1885).
- Fieuws, S., Verbeke, G., 2006. Pairwise fitting of mixed models for the joint modeling of multivariate longitudinal profiles. *Biometrics* 62 (2), 424–431. doi:[10.1111/j.1541-0420.2006.00507.x](https://doi.org/10.1111/j.1541-0420.2006.00507.x).
- Filippini, N., Zsoldos, E., Haapakoski, R., Sexton, C.E., Mahmood, A., Allan, C.L., Topiwala, A., Valkanova, V., Brunner, E.J., Shipley, M.J., Auerbach, E., Moeller, S., Uğurbil, K., Xu, J., Yacoub, E., Andersson, J., Bijstervosch, J., Clare, S., Griffanti, L., Hess, A.T., Jenkinson, M., Miller, K.L., Salimi-Khorshidi, G., Sotiropoulos, S.N., Voets, N.L., Smith, S.M., Geddes, J.R., Singh-Manoux, A., Mackay, C.E., Kivimäki, M., Ebmeier, K.P., 2014. Study protocol: the whitehall II imaging sub-study. *BMC Psychiatry* 14 (1). doi:[10.1186/1471-244x-14-159](https://doi.org/10.1186/1471-244x-14-159).
- Fischl, B., Salat, D.H., Busa, E., Albert, M., Dieterich, M., Haselgrove, C., van der Kouwe, A., Killiany, R., Kennedy, D., Klaveness, S., Montillo, A., Makris, N., Rosen, B., Dale, A.M., 2002. Whole brain segmentation. *Neuron* 33 (3), 341–355. doi:[10.1016/s0896-6273\(02\)00569-x](https://doi.org/10.1016/s0896-6273(02)00569-x).

- Fitzmaurice, G.M., Laird, N.M., Ware, J.H., 2011. *Applied longitudinal analysis*, 2nd edition John Wiley & Sons, Inc., Hoboken, New Jersey, USA.
- Fjell, A.M., Idland, A.-V., Sala-Llonch, R., Watne, L.O., Borza, T., Brækhus, A., Lona, T., Zetterberg, H., Blennow, K., Wyller, T.B., Walhovd, K.B., 2017. Neuroinflammation and tau interact with amyloid in predicting sleep problems in aging independently of atrophy. *Cerebral Cortex* 28 (8), 2775–2785. doi:10.1093/cercor/bhx157.
- Fjell, A.M., Walhovd, K.B., Westlye, L.T., Østby, Y., Tamnes, C.K., Jernigan, T.L., Gamst, A., Dale, A.M., 2010. When does brain aging accelerate? Dangers of quadratic fits in cross-sectional studies. *NeuroImage* 50 (4), 1376–1383. doi:10.1016/j.neuroimage.2010.01.061.
- Genin, E., Hannequin, D., Wallon, D., Sleegers, K., Hiltunen, M., Combarros, O., Bulldo, M.J., Engelborghs, S., Deyn, P.D., Berr, C., Pasquier, F., Dubois, B., Tognoni, G., Fiévet, N., Brouwers, N., Bettens, K., Arosio, B., Coto, E., Zompo, M.D., Mateo, I., Epelbaum, J., Frank-Garcia, A., Helisalmi, S., Porcellini, E., Pilotto, A., Forti, P., Ferri, R., Scarpini, E., Siciliano, G., Solfrizzi, V., Sorbi, S., Spalletta, G., Valdivieso, F., Vepäläinen, S., Alvarez, V., Bosco, P., Mancuso, M., Panza, F., Nacmias, B., Bossù, P., Hanon, O., Piccardi, P., Annoni, G., Seripa, D., Galimberti, D., Licastro, F., Soininen, H., Dartigues, J.-F., Kambh, M.I., Broeckhoven, C.V., Lambert, J.C., Amouyel, P., Campion, D., 2011. APOE and Alzheimer disease: a major gene with semi-dominant inheritance. *Mol. Psychiatry* 16 (9), 903–907. doi:10.1038/mp.2011.52.
- Ghisletta, P., Kennedy, K.M., Rodrigue, K.M., Lindenberger, U., Raz, N., 2010. Adult age differences and the role of cognitive resources in perceptual-motor skill acquisition: Application of a multilevel negative exponential model. *The Journals of Gerontology: Series B* 65B (2), 163–173. doi:10.1093/geronb/gbp126.
- Golub, G.H., Heath, M., Wahba, G., 1979. Generalized cross-validation as a method for choosing a good ridge parameter. *Technometrics* 21 (2), 215–223. doi:10.1080/00401706.1979.10489751.
- Gorbach, T., Pudas, S., Lundquist, A., Orädd, G., Josefsson, M., Salami, A., de Luna, X., Nyberg, L., 2017. Longitudinal association between hippocampus atrophy and episodic-memory decline. *Neurobiology of Aging* 51, 167–176. doi:10.1016/j.neurobiolaging.2016.12.002.
- Grimm, K., Zhang, Z., Hamagami, F., Mazocco, M., 2013. Modeling nonlinear change via latent change and latent acceleration frameworks: Examining velocity and acceleration of growth trajectories. *Multivariate Behavioral Research* 48 (1), 117–143. doi:10.1080/00273171.2012.755111.
- Grimm, K.J., Ram, N., 2009. Nonlinear growth models in mplus and SAS. *Structural Equation Modeling: A Multidisciplinary Journal* 16 (4), 676–701. doi:10.1080/10705510903206055.
- Gu, C., Ma, P., 2005. Generalized nonparametric mixed-effect models: computation and smoothing parameter selection. *Journal of Computational and Graphical Statistics* 14 (2), 485–504. doi:10.1198/106186005x47651.
- Hastie, T., Tibshirani, R., 1986. Generalized additive models. *Statistical Science* 1 (3), 297–310. doi:10.1214/ss/1177013604.
- Hastie, T., Tibshirani, R., 1993. Varying-coefficient models. *Journal of the Royal Statistical Society: Series B (Methodological)* 55 (4), 757–779. doi:10.1111/j.2517-6161.1993.tb01939.x.
- Hastie, T., Tibshirani, R., Friedman, J., 2009. Model assessment and selection. In: Hastie, T., Tibshirani, R., Friedman, J. (Eds.), *The Elements of Statistical Learning*. Springer, New York, pp. 219–259. doi:10.1007/978-0-387-84858-7.
- Hoffman, L., 2007. Multilevel models for examining individual differences in within-person variation and covariation over time. *Multivariate Behav Res* 42 (4), 609–629. doi:10.1080/00273170701710072.
- Hoffman, L., Stawski, R.S., 2009. Persons as contexts: evaluating between-person and within-person effects in longitudinal analysis. *Res. Hum. Dev.* 6 (2–3), 97–120. doi:10.1080/15427600902911189.
- Horn, J.L., Donaldson, G., 1976. On the myth of intellectual decline in adulthood. *American Psychologist* 31 (10), 701–719. doi:10.1037/0003-066x.31.10.701.
- Ke, C., Wang, Y., 2001. Semiparametric nonlinear mixed-effects models and their applications. *J. Am. Stat. Assoc.* 96 (456), 1272–1298. doi:10.1198/016214501753381913.
- Kimeldorf, G.S., Wahba, G., 1970. A correspondence between Bayesian estimation on stochastic processes and smoothing by splines. *The Annals of Mathematical Statistics* 41 (2), 495–502.
- Laird, N.M., Ware, J.H., 1982. Random-effects models for longitudinal data. *Biometrics* 38 (4), 963–974. doi:10.2307/2529876.
- Lambert, P.C., Abrams, K.R., Jones, D.R., Halligan, A.W.F., Shennan, A., 2001. Analysis of ambulatory blood pressure monitor data using a hierarchical model incorporating restricted cubic splines and heterogeneous within-subject variances. *Stat. Med.* 20 (24), 3789–3805. doi:10.1002/sim.1172.
- Lin, X., Zhang, D., 1999. Inference in generalized additive mixed models by using smoothing splines. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* 61 (2), 381–400. doi:10.1111/1467-9868.00183.
- Lindenberger, U., von Oertzen, T., Ghisletta, P., Hertzog, C., 2011. Cross-sectional age variance extraction: What's change got to do with it? *Psychol. Aging* 26 (1), 34–47. doi:10.1037/a0020525.
- Lindenberger, U., Pötter, U., 1998. The complex nature of unique and shared effects in hierarchical linear regression: implications for developmental psychology. *Psychol. Methods* 3 (2), 218–230. doi:10.1037/1082-989x.3.2.218.
- Lindstrom, M.J., Bates, D.M., 1990. Nonlinear mixed effects models for repeated measures data. *Biometrics* 46 (3), 673. doi:10.2307/2532087.
- Marra, G., Wood, S.N., 2012. Coverage properties of confidence intervals for generalized additive model components. *Scand. J. Stat.* 39 (1), 53–74. doi:10.1111/j.1467-9469.2011.00760.x.
- McArdle, J.J., Epstein, D., 1987. Latent growth curves within developmental structural equation models. *Child Development* 58 (1), 110. doi:10.2307/1130295.
- Mehta, P.D., West, S.G., 2000. Putting the individual back into individual growth curves. *Psychol. Methods* 5 (1), 23–43. doi:10.1037/1082-989x.5.1.23.
- Meredith, M., Kruschke, J., 2019. HDInterval: Highest (posterior) density intervals. <https://CRAN.R-project.org/package=HDInterval>. R package version 0.2.2.
- Meredith, W., Tisak, J., 1990. Latent curve analysis. *Psychometrika* 55 (1), 107–122. doi:10.1007/bf02294746.
- Molenaar, P.C.M., 2004. A manifesto on psychology as idiographic science: bringing the person back into scientific psychology, this time forever. *Measurement: Interdisciplinary Research & Perspective* 2 (4), 201–218. doi:10.1207/s15366359mea0204_1.
- Molenberghs, G., Fitzmaurice, G., 2009. Incomplete data: Introduction and overview. In: Fitzmaurice, G., Davidian, M., Verbeke, G., Molenberghs, G. (Eds.), *Longitudinal Data Analysis*. Chapman and Hall, London, U.K., pp. 395–408.
- Morrell, C.H., Brant, L.J., Ferrucci, L., 2009. Model choice can obscure results in longitudinal studies. *The Journals of Gerontology Series A: Biological Sciences and Medical Sciences* 64A (2), 215–222. doi:10.1093/geronb/gln024.
- Newsom, J.T., 2015. *Longitudinal structural equation modeling*. Routledge.
- Nychka, D., 1988. Bayesian confidence intervals for smoothing splines. *J. Am. Stat. Assoc.* 83 (404), 1134–1143. doi:10.1080/01621459.1988.10478711.
- Oud, J.H.L., Jansen, R.A.R.G., 2000. Continuous time state space modeling of panel data by means of SEM. *Psychometrika* 65 (2), 199–215. doi:10.1007/bf02294374.
- Pedersen, E.J., Miller, D.L., Simpson, G.L., Ross, N., 2019. Hierarchical generalized additive models in ecology: an introduction with mgcv. *PeerJ* 7, e6876. doi:10.7717/peerj.6876.
- R Core Team, 2019. R: A language and environment for statistical computing. Vienna, Austria. <https://www.R-project.org/>
- Ram, N., Grimm, K., 2016. Using simple and complex growth models to articulate developmental change: Matching theory to method. *International Journal of Behavioral Development* 31 (4), 303–316. doi:10.1177/0165025407077751.
- Raz, N., Ghisletta, P., Rodrigue, K.M., Kennedy, K.M., Lindenberger, U., 2010. Trajectories of brain aging in middle-aged and older adults: regional and individual differences. *Neuroimage* 51 (2), 501–511. doi:10.1016/j.neuroimage.2010.03.020.
- Raz, N., Lindenberger, U., Rodrigue, K.M., Kennedy, K.M., Head, D., Williamson, A., Dahle, C., Gerstorf, D., Acker, J.D., 2005. Regional brain changes in aging healthy adults: general trends, individual differences and modifiers. *Cerebral Cortex* 15 (11), 1676–1689. doi:10.1093/cercor/bhi044.
- Reuter, M., Schmansky, N.J., Rosas, H.D., Fischl, B., 2012. Within-subject template estimation for unbiased longitudinal image analysis. *Neuroimage* 61 (4), 1402–1418. doi:10.1016/j.neuroimage.2012.02.084.
- Ruppert, D., Wand, M.P., Carroll, R.J., 2003. *Semiparametric regression*. Cambridge University Press, Cambridge, U.K.
- Salthouse, T.A., 2013. Within-cohort age-related differences in cognitive functioning. *Psychological Science* 24 (2), 123–130. doi:10.1177/0956797612450893.
- Salthouse, T.A., 2014. Why are there different age relations in cross-sectional and longitudinal comparisons of cognitive functioning? *Current Directions in Psychological Science* 23 (4), 252–256. doi:10.1177/0963721414535212.
- Salthouse, T.A., 2019. Trajectories of normal cognitive aging. *Psychology and Aging* 34 (1), 17–24. doi:10.1037/pag0000288.
- Schaie, K.W., Labouvie, G.V., Buech, B.U., 1973. Generational and cohort-specific differences in adult cognitive functioning: a fourteen-year study of independent samples. *Dev. Psychol.* 9 (2), 151–166. doi:10.1037/h0035093.
- Simpson, G., 2016. Simultaneous intervals for smooths revisited. <https://fromthebottomoftheheap.net/2016/12/15/simultaneous-interval-revisited/> (accessed July 10, 2020).
- Slivinski, M., Hoffman, L., Hofer, S.M., 2010. Evaluating convergence of within-person change and between-person age differences in age-heterogeneous longitudinal studies. *Res. Hum. Dev.* 7 (1), 45–60. doi:10.1080/15427600903578169.
- Sørensen, Ø., Brandmaier, A.M., Macià, D., Ebmeier, K., Ghisletta, P., Kievit, R.A., Mowinckel, A.M., Walhovd, K.B., Westerhausen, R., Fjell, A., 2021. Meta-analysis of generalized additive models in neuroimaging studies. *NeuroImage* 224, 117416. doi:10.1016/j.neuroimage.2020.117416.
- Sullivan, K.J., Shadish, W.R., Steiner, P.M., 2015. An introduction to modeling longitudinal data with generalized additive models: applications to single-case designs. *Psychol. Methods* 20 (1), 26–42. doi:10.1037/met000020.
- Thompson, P.M., Andreassen, O.A., Arias-Vasquez, A., Bearden, C.E., Boedhoe, P.S., Brouwer, R.M., Buckner, R.L., Buitelaar, J.K., Bulayeva, K.B., Cannon, D.M., Cohen, R.A., Conrod, P.J., Dale, A.M., Deary, I.J., Dennis, E.L., de Reus, M.A., Desrivieres, S., Dima, D., Donohoe, G., Fisher, S.E., Fouché, J.-P., Francks, C., Frangou, S., Franke, B., Ganjgahi, H., Garavan, H., Glahn, D.C., Grabe, H.J., Guadalupe, T., Gutman, B.A., Hashimoto, R., Hibar, D.P., Holland, D., Hoogman, M., Pol, H.E.H., Hosten, N., Jahanshad, N., Kelly, S., Kochunov, P., Kremen, W.S., Lee, P.H., Mackey, S., Martin, N.G., Mazoyer, B., McDonald, C., Medland, S.E., Morey, R.A., Nichols, T.E., Paus, T., Pausova, Z., Schmaal, L., Schumann, G., Shen, L., Sisodiya, S.M., Smit, D.J.A., Smoller, J.W., Stein, D.J., Stein, J.L., Toro, R., Turner, J.A., van den, H., Martijn, P., van den, H., Odile, L., van Erp, T.G.M., van Rooij, D., Veltman, D.J., Walter, H., Wang, Y., Wardlaw, J.M., Whelan, C.D., Wright, M.J., Ye, J., 2017. ENIGMA and the individual: predicting factors that affect the brain in 35 countries worldwide. *Neuroimage* 145, 389–408. doi:10.1016/j.neuroimage.2015.11.057.
- Thompson, W.K., Hallmayer, J., O'Hara, R., 2011. Design considerations for characterizing psychiatric trajectories across the lifespan: Application to effects of APOE-ε4 on cerebral cortical thickness in Alzheimer's disease. *AJP* 168 (9), 894–903. doi:10.1176/appi.ajp.2011.10111690.
- Tian, Q., Studenski, S.A., Resnick, S.M., Davatzikos, C., Ferrucci, L., 2015. Midlife and late-life cardiorespiratory fitness and brain volume changes in late adulthood: Results

- from the baltimore longitudinal study of aging. *J Gerontol A Biol Sci Med Sci* 71 (1), 124–130. doi:[10.1093/gerona/glv041](https://doi.org/10.1093/gerona/glv041).
- Walhovd, K.B., Fjell, A.M., Sørensen, Ø., Mowinckel, A.M., Reinbold, C.S., Idland, A.-V., Watne, L.O., Franke, A., Dobricic, V., Kilpert, F., Bertram, L., Wang, Y., 2020. Genetic risk for alzheimer's disease predicts hippocampal volume through the lifespan. *Neurology Genetics* 6 (5). doi:[10.1212/NXG.0000000000000506](https://doi.org/10.1212/NXG.0000000000000506).
- Walhovd, K.B., Fjell, A.M., Westerhausen, R., Nyberg, L., Ebmeier, K.P., Lindenberger, U., Bartres-Faz, D., Baare, W.F.C., Siebner, H.R., Henson, R., 2018. Healthy minds 0–100 years: optimising the use of European brain imaging cohorts (“lifebrain”). *European Psychiatry* 47, 76–77. doi:[10.1016/j.eurpsy.2017.10.005](https://doi.org/10.1016/j.eurpsy.2017.10.005).
- Walhovd, K.B., Krogsrud, S.K., Amlie, I.K., Bartsch, H., Bjørnerud, A., Due-Tønnessen, P., Grydeland, H., Hagler, D.J., Håberg, A.K., Kremen, W.S., Ferschmann, L., Nyberg, L., Panizzon, M.S., Rohani, D.A., Skranes, J., Storsve, A.B., Søsnes, A.E., Tamnes, C.K., Thompson, W.K., Reuter, C., Dale, A.M., Fjell, A.M., 2016. Neurodevelopmental origins of lifespan changes in brain and cognition. *Proceedings of the National Academy of Sciences* 113 (33), 9357–9362. doi:[10.1073/pnas.1524259113](https://doi.org/10.1073/pnas.1524259113).
- Wickham, H., 2016. *ggplot2: Elegant Graphics for Data Analysis*. Springer-Verlag, New York, New York, USA.
- Wood, S., Scheipl, F., 2017. *gamm4: Generalized additive mixed models using 'mgcv' and 'lme4'*. <https://CRAN.R-project.org/package=gamm4>. R package version 0.2–5.
- Wood, S.N., 2003. Thin plate regression splines. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* 65 (1), 95–114. doi:[10.1111/1467-9868.00374](https://doi.org/10.1111/1467-9868.00374).
- Wood, S.N., 2004. Stable and efficient multiple smoothing parameter estimation for generalized additive models. *J. Am. Stat. Assoc.* 99 (467), 673–686. doi:[10.1198/016214504000000980](https://doi.org/10.1198/016214504000000980).
- Wood, S.N., 2006. Low-rank scale-invariant tensor product smooths for generalized additive mixed models. *Biometrics* 62 (4), 1025–1036. doi:[10.1111/j.1541-0420.2006.00574.x](https://doi.org/10.1111/j.1541-0420.2006.00574.x).
- Wood, S.N., 2010. Fast stable restricted maximum likelihood and marginal likelihood estimation of semiparametric generalized linear models. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* 73 (1), 3–36. doi:[10.1111/j.1467-9868.2010.00749.x](https://doi.org/10.1111/j.1467-9868.2010.00749.x).
- Wood, S.N., 2017. *Generalized additive models: An introduction with r, 2nd edition* Chapman and Hall/CRC, Boca Raton, Florida, USA.
- Wood, S.N., Scheipl, F., Faraway, J.J., 2012. Straightforward intermediate rank tensor product smoothing in mixed models. *Stat. Comput.* 23 (3), 341–360. doi:[10.1007/s11222-012-9314-z](https://doi.org/10.1007/s11222-012-9314-z).
- Zeger, S.L., Liang, K.-Y., 1992. An overview of methods for the analysis of longitudinal data. *Stat Med* 11 (14–15), 1825–1839. doi:[10.1002/sim.4780111406](https://doi.org/10.1002/sim.4780111406).