

Multilevel growth curve models that incorporate a random coefficient model for the level I variance function

Harvey Goldstein,¹ George Leckie,¹ Christopher Charlton,¹ Kate Tilling² and William J Browne¹

Statistical Methods in Medical Research
0(0) 1–14

© The Author(s) 2017

Reprints and permissions:

sagepub.co.uk/journalsPermissions.nav

DOI: 10.1177/0962280217706728

journals.sagepub.com/home/smm



Abstract

Aim: To present a flexible model for repeated measures longitudinal growth data within individuals that allows trends over time to incorporate individual-specific random effects. These may reflect the timing of growth events and characterise within-individual variability which can be modelled as a function of age.

Subjects and methods: A Bayesian model is developed that includes random effects for the mean growth function, an individual age-alignment random effect and random effects for the within-individual variance function. This model is applied to data on boys' heights from the Edinburgh longitudinal growth study and to repeated weight measurements of a sample of pregnant women in the Avon Longitudinal Study of Parents and Children cohort.

Results: The mean age at which the growth curves for individual boys are aligned is 11.4 years, corresponding to the mean 'take off' age for pubertal growth. The within-individual variance (standard deviation) is found to decrease from 0.24 cm² (0.50 cm) at 9 years for the 'average' boy to 0.07 cm² (0.25 cm) at 16 years. Change in weight during pregnancy can be characterised by regression splines with random effects that include a large woman-specific random effect for the within-individual variation, which is also correlated with overall weight and weight gain.

Conclusions: The proposed model provides a useful extension to existing approaches, allowing considerable flexibility in describing within- and between-individual differences in growth patterns.

Keywords

Heteroscedasticity, multilevel model, repeated measures, variance model, Avon Longitudinal Study of Parents and Children

1 Introduction

In modelling growth and developmental change over time, the data structure typically consists of repeated measures within individuals and a suitable relationship, for example based upon polynomials or smoothing splines, relates the serial measurements to age. Random effect or multilevel statistical models are commonly used to fit such data, where the coefficients, in for example a polynomial age function, vary across individuals.¹ A particular concern in the development of such models has been to make them more general and hence more flexible and able to describe change over time more realistically. One recent addition is the use of smooth spline functions to describe growth while allowing each individual to have their own growth 'origin'.² This has been extended³ to the multivariate case where several measurements are modelled simultaneously. Another extension is to allow the residual, within individual, variation to be modelled

¹Centre for Multilevel Modelling, University of Bristol, Bristol, UK

²Department of Social and Community Medicine, University of Bristol, Bristol, UK

Corresponding author:

Harvey Goldstein, Graduate School of Education, University of Bristol, 35 Berkeley Square, Bristol BS8 1JA, UK.

Email: h.goldstein@bristol.ac.uk

as a function of time and other covariates, where the parameters of this model component may be individual specific and allowed to correlate with other individual random effects.^{4,5} The present paper brings together these developments within a single framework and applies the model to data both on adolescent height growth and changes in weight during pregnancy.

In the standard, especially linear, modelling framework, a constant (homoscedastic) residual variance is generally assumed. When such models are extended to multilevel models, including random effect growth curve models, to account for clustering of the repeated measures (level 1) within individuals (level 2), this assumption becomes an assumption of a constant level 1 variance. In many scenarios, this assumption is violated and thus models that account for heteroscedasticity are required and these are common in the literature. In some cases, log or similar transformations of the response variable are used,³ but in many cases, there is substantive interest in the variance itself so that it will generally be more satisfactory to model the variance explicitly. For example, the within-individual occasion-to-occasion variation in variables such as weight or blood pressure may be related to behavioural or clinical factors.

More explicitly, an extension in the multilevel setting is to allow the variability at level 1 itself to vary across level 2 units through the addition of random effects to an explicit sub-model for the level 1 variance. To this end, Hedeker et al.⁴ consider variance component models where the logarithm of the level 1 variance is a linear function of explanatory variables with an additional simple random effect that can vary across level 2 units and is allowed to correlate with the conventional level 2 random effects describing between-individual variation in growth patterns. They show how to obtain maximum likelihood estimates and implement this in the stand-alone MIXREGLS software⁶ which is also accessible from within Stata.⁷ Lee and Nelder⁸ also consider variance component models that allow a single random effect in the variance function, but do not allow this to correlate with the level 2 random effects. Rast et al.⁹ present a Bayesian version of this model extended to allow random coefficients in both the mean and variance function and where the two sets of random effects are allowed to covary. They develop their model in the context of a repeated measures data set. Leckie⁷ and Brunton-Smith et al.¹⁰ consider similar models in the context of cross-sectional two-level and cross-classified data again fitting their models using Markov Chain Monte Carlo (MCMC) methods.

In the present paper, we likewise consider a general model for the level 1 variance that allows the inclusion of several random coefficients varying at level 2 (that of the individual), in two-level repeated measures growth curve models. These random coefficients are also allowed to covary with conventional level 2 random coefficients of age and potentially other explanatory variables in the model predicting growth. The model is developed within the general framework of two-level repeated measures data, but is not limited to such models. The model incorporates the case where there may be further level 2 random effects within the definition of a predictor variable.

One application of this model is to allow each individual to have their own ‘origin’ for age. The model then corresponds to the SITAR (Superimposition by Translation and Rotation) model² for the case of adolescent growth in height where the timing of adolescent growth varies across individuals so that chronological age can be considered as a time-shifted proxy for an underlying ‘maturity’. Another example might be in studies where previously measured characteristics, such as attitudes, are individual level predictors and where the institutions, such as schools or hospitals, within which such characteristics, were assessed have mean differences that are assumed directly to influence the individual characteristics. Including a random effect varying at the institution level would then be expected to provide a better specified model and have an interpretation of substantive interest.

Another application of our model is where individuals exhibit extremely small or extremely large occasion to occasion, within-individual, variability around an overall growth trajectory. Thus, events such as repeated episodes of illness may be associated with patterns of normal growth followed by periods where growth shows less variability, and where the sensitivity of growth to such events may vary across individuals. In such cases, allowing the within-individual variation to vary across individuals as a function of covariates such as illness, will allow the detection of such effects that themselves may be associated with later life events. In the following sections, we set out the formal model followed by applications.

2 A two-level model with complex level 1 variation

We consider first the extension of a standard two-level random coefficients model to contain both fixed and random effects within the level 1 variance function. Let y_{ij} denote the response at occasion i ($i = 1, \dots, n_j$) for

individual j ($j = 1, \dots, J$). The model can be written as follows

$$\begin{aligned}
 y_{ij} &= \mathbf{x}_{1ij}\boldsymbol{\beta} + \mathbf{z}_{1ij}\mathbf{u}_{1j} + e_{ij} \\
 e_{ij} &\sim N(0, \sigma_{e_{ij}}^2) \\
 \log(\sigma_{e_{ij}}^2) &= \mathbf{x}_{2ij}\boldsymbol{\gamma} + \mathbf{z}_{2ij}\mathbf{u}_{2j} \\
 \begin{pmatrix} \mathbf{u}_{1j} \\ \mathbf{u}_{2j} \end{pmatrix} &\sim N\left\{\begin{pmatrix} 0 \\ 0 \end{pmatrix}, \begin{pmatrix} \boldsymbol{\Omega}_{\mathbf{u}_1} & \\ & \boldsymbol{\Omega}_{\mathbf{u}_2} \end{pmatrix}\right\}
 \end{aligned} \tag{1}$$

The log-linear form for the level 1 variance function guarantees that it remains positive. Here, we distinguish between the traditional vector of level 2 random effects associated with the mean of the response (\mathbf{u}_{1j}) and the vector of random effects associated with the level 1 variance function (\mathbf{u}_{2j}) and allow all of these to covary. The model specifies that the log of the level 1 variance ($\sigma_{e_{ij}}^2$) is a linear function of explanatory variables, such as age, socioeconomic status or ethnic group. Here, $\boldsymbol{\Omega}_{\mathbf{u}_1}$ is the covariance matrix for the traditional random effects associated with the mean function, and $\boldsymbol{\Omega}_{\mathbf{u}_2}$ is the covariance matrix of the random effects in the level 1 variance function, with $\boldsymbol{\Omega}_{\mathbf{u}_1\mathbf{u}_2}$ representing their cross-covariances. We note that \mathbf{x}_{1ij} and \mathbf{x}_{2ij} , respectively the explanatory variable vectors for the growth function and the level 1 variance function, may have variables in common, for example time, but there may also be variables in \mathbf{x}_{2ij} that do not influence the predicted response and hence are modelled solely in the level 1 variance function. The \mathbf{z}_{1ij} , and \mathbf{z}_{2ij} are the explanatory variables for the level 2 random effects and are typically a subset of the variables included in \mathbf{x}_{1ij} and \mathbf{x}_{2ij} .

2.1 Estimation

To estimate the parameters of this general model, we use a Bayesian MCMC algorithm, where we assume independent diffuse (uninformative) priors. The details of the MCMC algorithm are given in Appendix 1. The modular nature of MCMC algorithms means that the steps in the algorithm are readily generalised to models with three or more levels and to cross-classifications and other non-hierarchical data structures. All models fitted in this paper are written in C++ code that is implemented in templates in the statistical software package, StatJR.¹¹ The relevant templates are available from the third author (c.charlton@bristol.ac.uk).

3 Age-shifted growth curve models

In this section, we consider extending equation (1) to incorporate the SITAR age-shifted model.² This model was developed in the context of modelling adolescent height and can be written, with a slight change of notation, as

$$\begin{aligned}
 y_{ij} &= \beta_{0j} + f(x_{ij}^*) + e_{ij} \\
 x_{ij}^* &= e^{\gamma_j}(x_{ij} - v_j) \\
 \beta_{0j} &= \beta_0 + u_{0j}
 \end{aligned} \tag{2}$$

where $f(\cdot)$ is a suitable growth function of transformed age x_{ij}^* . The original authors use a natural cubic spline, but any suitable function can be used, such as the ‘Preece-Baines’ curve¹² or a regression spline (see below). The default distributional assumptions are

$$\begin{pmatrix} \beta_{0j} \\ \gamma_j \\ v_j \end{pmatrix} \sim N\left\{\begin{pmatrix} \beta_0 \\ \gamma_0 \\ v_0 \end{pmatrix}, \begin{pmatrix} \sigma_\beta^2 & & \\ \sigma_{\beta\gamma} & \sigma_\gamma^2 & \\ \sigma_{\beta v} & \sigma_{\gamma v} & \sigma_v^2 \end{pmatrix}\right\}, \quad e_{ij} \sim N(0, \sigma_e^2)$$

The term β_{0j} is a subject-specific ‘size’ random effect. The term γ_j is a random effect determining the relative scale along which a ‘standard’ growth pattern occurs for individual j , and v_j is an age-shifted ‘tempo’ random effect representing an individual-specific ‘origin’, representing differences in the timing of the adolescent growth spurt. Cole et al.² use the nlme package¹³ in R¹⁴ to obtain maximum likelihood estimates for the parameters of this model. They also suggest that a log transformation of height and/or age can provide a somewhat better fit, but for ease of interpretation we shall use height itself as the response variable in our example.

This model has just three random effects and contrasts with more traditional random effect growth curve models, for example using polynomial functions, where the number of random effects is generally three or four for non-linear growth periods such as adolescence,¹ and typically the underlying mean curve is a higher order polynomial or other smooth function, similar to the spline used by Cole et al.² To illustrate the essential difference between the traditional model and the SITAR model, we can write a relatively simple traditional model with an underlying third-degree polynomial and random coefficients up to the second-order term as

$$\begin{aligned} y_{ij} &= \beta_{0j} + \beta_{1j}x_{ij} + \beta_{2j}x_{ij}^2 + \beta_{3j}x_{ij}^3 + e_{ij} \\ \beta_{0j} &= \beta_0 + u_{0j}, \quad \beta_{1j} = \beta_1 + u_{1j}, \quad \beta_{2j} = \beta_2 + u_{2j} \\ e_{ij} &\sim N(0, \sigma_e^2) \end{aligned} \quad (3)$$

In our example below, we will see that a higher order polynomial is needed, but for exposition, we restrict ourselves to the simpler version given by equation (3). For the SITAR model, for comparison purposes with the same number of level 2 random effects, we also choose an underlying third-degree polynomial for x_{ij}^* rather than a smoothing spline, so that we have

$$y_{ij} = \beta_{0j} + \beta_1 e^{\gamma_j} (x_{ij} - v_j) + \beta_2 e^{2\gamma_j} (x_{ij} - v_j)^2 + \beta_3 e^{3\gamma_j} (x_{ij} - v_j)^3 + e_{ij} \quad (4)$$

Both the traditional and the SITAR model in this case have three random effects at level 2 but with different interpretations, although when the order of the polynomial is reduced to one the interpretations will be similar except for the origin random effect (v_j) which has no counterpart in the traditional model. The SITAR model has the property, at least in the case of data such as adolescent growth, that the random effects have intuitively appealing interpretations, although this does not imply anything inherently more ‘biological’ in the formulation.

In fact, despite the apparent simplicity of equation (4), it may not necessarily be a good fit to any given data set. We therefore generalise the model by adding further random effects, still incorporating an age shift but allowing for a more flexible modelling of the ‘scale’. Thus equation (4) becomes

$$\begin{aligned} y_{ij} &= \beta_{0j} + \beta_{1j}(x_{ij} - v_j) + \beta_{2j}(x_{ij} - v_j)^2 + \beta_{3j}(x_{ij} - v_j)^3 + e_{ij} \\ \beta_{0j} &= \beta_0 + u_{0j}, \quad \beta_{1j} = \beta_1 + u_{1j}, \quad \beta_{2j} = \beta_2 + u_{2j}, \quad \beta_{3j} = \beta_3 + u_{3j} \end{aligned} \quad (5)$$

We now have five rather than three random effects at the individual level (level 2) where the multiplicative single random effect e^{γ_j} and its powers are replaced by the three correlated additive random effects u_{1j} , u_{2j} and u_{3j} .

In our first example, we compare the fit of these models to a real data set. Rather than using a natural cubic spline as in Cole et al.,² we shall use a regression spline in our analyses as this provides more flexibility for our purposes and is straightforward to implement while providing sufficient flexibility to graduate growth. At this point, it is again worth pointing out that the choice of growth curve function is essentially empirical in the sense that it should graduate growth smoothly with enough complexity to allow the modelling of key events such as peak height velocity (PHV). In our view, none of the curves suggested in the literature for growth curve functions has any special ‘biological’ as opposed to ‘mathematical’ interpretation that suggests that they should be treated preferentially.

We fit all our models using MCMC with default priors as specified in the previous section. The random effects u_{0j}, \dots, u_{3j} , γ_j and v_j are sampled using Metropolis steps similar to that for the variance function terms.

4 An example using height growth data

The data set we use is the same one used by Cole et al.¹⁵ It consists of 1725 height measurements between 8 and 18 years of age for a sample of 101 boys with measurements made approximately six months apart. Age is centred at the mean age of 12.66 years. Only 1% of the measurements are less than five months apart and as Browne and Goldstein¹⁶ showed, it is not until measurements are about two months apart that serial correlation effects are to be found in height growth data, and hence we do not attempt to fit an autocorrelation function for the residuals. The details of the appropriate MCMC steps for doing so, however, are given by Browne and Goldstein¹⁶ and involve a further Metropolis step in the algorithm.

We first look at a polynomial model extension to the original SITAR model (which corresponds to model (4)). This new model has a fifth order polynomial spline, with eight parameters describing growth, chosen after

exploring different order polynomials, and can be written as

$$\begin{aligned}
 y_{ij} &= \beta_{0j} + \beta_1 e^{\gamma_j} z_{ij} + \beta_2 e^{2\gamma_j} z_{ij}^2 + \beta_3 e^{3\gamma_j} z_{ij}^3 + \beta_4 e^{4\gamma_j} z_{ij}^4 + \beta_5 e^{5\gamma_j} z_{ij}^5 + \alpha_0 s_{0ij}^2 + \alpha_1 s_{1ij}^2 + e_{ij} \\
 z_{ij} &= x_{ij} - v_j, \\
 s_{0ij} &= \begin{cases} 0 & \text{if } z_{ij} < 0 \\ z_{ij} & \text{if } z_{ij} \geq 0 \end{cases} \\
 s_{1ij} &= \begin{cases} 0 & \text{if } z_{ij} \geq 0 \\ z_{ij} & \text{if } z_{ij} < 0 \end{cases} \\
 \left(\begin{array}{c} \beta_{0j} \\ \gamma_j \\ v_j \end{array} \right) &\sim N \left\{ \left(\begin{array}{c} \beta_0 \\ \gamma_0 \\ v_0 \end{array} \right), \left(\begin{array}{ccc} \sigma_\beta^2 & & \\ \sigma_{\beta\gamma} & \sigma_\gamma^2 & \\ \sigma_{\beta v} & \sigma_{\gamma v} & \sigma_v^2 \end{array} \right) \right\}, \quad e_{ij} \sim N(0, \sigma_e^2)
 \end{aligned} \tag{6}$$

The terms s_{0ij} and s_{1ij} are the regression splines that join smoothly at the shifted age origins. For each model, for the MCMC algorithm, three chains were run, each with different sets of starting values. Convergence was assumed when the chains all converged to the same posterior distribution as judged by inspecting the parameter estimates. In practice, it was sufficient to specify good starting values for the fixed coefficients (using for example estimates from a simple model with no random effects) for the chains to show good mixing properties. The default priors described above were used here including a uniform prior for the covariance matrix at level 2. Different models were tried with different join points and orders for the spline terms and the choice of a single ‘forward’ spline s_0 and a single ‘backward’ spline s_1 joining at shifted age = 0 was found to be adequate, as judged using the Deviance Information Criterion (DIC) statistics, and the results for these choices are reported here.

Table 1 shows the results of fitting the model to this data set.

The estimate for γ_0 implies additional multipliers for the coefficients of the (shifted) age polynomial. For example, the additional multiplier for the coefficient of the linear term is approximately $\exp(0.13) = 1.12$

Table 1. Edinburgh growth data.

Parameter	Estimate	Standard error	Cole et al. ² estimate
α_0	-0.10	0.41	
α_1	-3.03	0.64	
σ_{u0}^2			48.7
σ_{uv}^2			0.97
$\sigma_{u\gamma}^2$			0.013
Ω_u	$\begin{pmatrix} 51.0 & 3.14 & -0.13 \\ 3.14 & 1.98 & 0.17 \\ -0.13 & 0.17 & 0.03 \end{pmatrix}$		
$R_u = \text{Cor}(\Omega_u)$	$\begin{pmatrix} 1.0 & 0.31 & -0.11 \\ 0.31 & 1.0 & 0.70 \\ -0.11 & 0.70 & 1.0 \end{pmatrix}$		
δ_0	0.13	0.05	0.13
β_0	147.1	0.7	133.4
β_1	3.11	0.17	
β_2	1.97	0.55	
β_3	-0.26	0.04	
β_4	-0.01	0.002	
β_5	0.002	0.0004	
v_0	-0.92	0.14	-1.48
σ_e^2	0.68	0.03	0.51
DIC (PD)	4541.9 (298.5)		

Note: MCMC with adaptation length 25,000, burnin = 25,000, iterations = 25,000; Extended SITAR model using regression splines.

and so the total fixed effect becomes $\beta_1 \exp(0.13) = 3.11 \times 1.12 = 3.5$. The mean age at which the growth curves are aligned from this model is $12.66 + 0.92 = 13.58$ which is close to the reported mean age for UK boys at PHV.¹⁷ If we consider a boy with $v_j = 0$, that is at a shifted age of 0.92 years later than the mean (-0.92) and ‘average’ growth function random effects at the mean (i.e. all equal to zero), then his velocity at centred age $x_{ij} = 0$ (i.e. the mean age of 12.66), is estimated to be 3.5 cm/year. If we use the fixed coefficients to estimate the age when the growth acceleration (that is the second differential of the growth curve) is zero (found by solving the corresponding cubic equation), this is calculated as 2.18, so that for the average value $v_0 = -0.92$, the centred age of maximum velocity for the adolescent growth period is $2.18 - 0.92 = 1.26$, that is an actual age of 13.9 years with a corresponding velocity estimate of 8.8 cm/year. The level 1 standard deviation is $\sqrt{0.68} = 0.83$ cm and it also appears that we require only one spline function since the estimate for α_0 is non-significant. The results obtained by Cole et al.¹⁵ using maximum likelihood estimation, for the variance and covariance parameters and the mean age shift and scaling parameters are given in the final column and are not too dissimilar. We note that the level 1 variance is somewhat smaller for the Cole et al. model which may reflect the fact that Cole et al. used maximum likelihood estimation.

Figure 1 displays the predicted median growth velocity, using actual age around the age of maximum growth velocity. Figure 2 displays the predicted median growth velocity, using the shifted age around the age of maximum growth velocity.

Note that the ‘sharpness’ of the peak in Figure 2 results from the two spline terms joining at $z_{ij} = 0$, while the velocity remains a continuous function of adjusted age. The age-shifted model essentially displays the individual-specific median velocity centered on the shifted age, which is higher than that estimated by averaging the velocity at the median age, but is consistent with other estimates based upon individual growth curves¹⁷

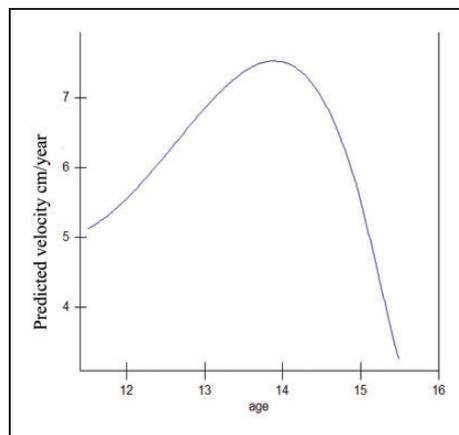


Figure 1. Estimated median growth velocity curve by age. Model (7) with $z_{ij} = x_{ij}$.

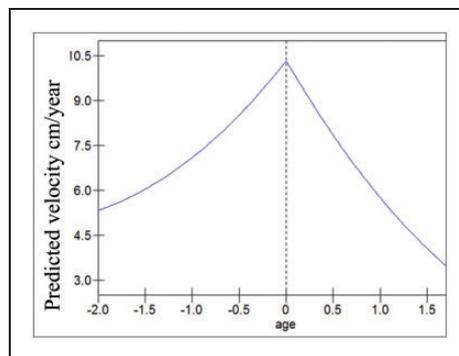


Figure 2. Estimated median growth velocity curves by adjusted age. Model (7) with $z_{ij} = x_{ij} - v_j$.

We now fit the more flexible model, in terms of random coefficients, namely

$$\begin{aligned}
 y_{ij} &= \beta_{0j} + \beta_{1j}z_{ij} + \beta_{2j}z_{ij}^2 + \beta_{3j}z_{ij}^3 + \beta_{4j}z_{ij}^4 + \beta_{5j}z_{ij}^5 + \alpha_0s_{0ij}^2 + \alpha_1s_{1ij}^2 + e_{ij} \\
 z_{ij} &= x_{ij} - v_j \\
 s_{0ij} &= \begin{cases} 0 & \text{if } z_{ij} < 0 \\ z_{ij} & \text{if } z_{ij} \geq 0 \end{cases} \\
 s_{1ij} &= \begin{cases} 0 & \text{if } z_{ij} \geq 0 \\ z_{ij} & \text{if } z_{ij} < 0 \end{cases} \\
 \left(\begin{array}{c} \beta_{0j} \\ \beta_{1j} \\ \beta_{2j} \\ \beta_{3j} \\ v_j \end{array} \right) &\sim N \left\{ \left(\begin{array}{c} \beta_0 \\ \beta_1 \\ \beta_2 \\ \beta_3 \\ v_0 \end{array} \right), \left(\begin{array}{ccccc} \sigma_{u0}^2 & & & & \\ \sigma_{u01} & \sigma_{u1}^2 & & & \\ \sigma_{u02} & \sigma_{u12} & \sigma_{u2}^2 & & \\ \sigma_{u03} & \sigma_{u13} & \sigma_{u23} & \sigma_{u3}^2 & \\ \sigma_{u0v} & \sigma_{u1v} & \sigma_{u2v} & \sigma_{u3v} & \sigma_v^2 \end{array} \right) \right\}, \quad e_{ij} \sim N(0, \sigma_e^2)
 \end{aligned} \tag{7}$$

The results are given in Table 2.

Adding random coefficients for the fourth and fifth degree terms did not substantially change the inferences and has therefore not been used here. Again, we only appear to require one spline function, although this time it is the forward spline function. We also note that the level 1 standard deviation is greatly reduced (from $\sqrt{0.68} = 0.83$ cm to $\sqrt{0.28} = 0.53$ cm) and the DIC statistic is considerably smaller than before, with a greater number of effective parameters (PD). These extra parameters arise from the presence of five as opposed to three random effects, with the linear and cubic polynomial coefficients highly correlated ($\rho_{u13} = -0.96$). Recalling that our age scale is centred on 12.66, we also note that the mean growth curve alignment age is now $12.66 - 1.27 = 11.4$ years.

Table 2. Edinburgh growth data.

Parameter	Estimate	Standard error
α_0	-4.67	0.16
α_1	0.01	0.18
Ω_u	$\begin{pmatrix} 53.5 & & & & \\ 0.41 & 0.55 & & & \\ -0.29 & -0.01 & 0.08 & & \\ -0.01 & -0.01 & 0.001 & 0.0002 & \\ 1.99 & -0.35 & 0.001 & 0.003 & 3.17 \end{pmatrix}$	
$R_u = \text{Cor}(\Omega_u)$	$\begin{pmatrix} 1.0 & & & & \\ 0.08 & 1.0 & & & \\ -0.14 & -0.10 & 1.0 & & \\ -0.10 & -0.96 & 0.25 & 1.0 & \\ 0.15 & -0.27 & 0.002 & 0.12 & 1.0 \end{pmatrix}$	
β_0	162.6	0.70
β_1	10.31	0.10
β_2	2.01	0.16
β_3	0.26	0.008
β_4	0.0011	0.0006
β_5	-0.0012	0.00011
v_0	1.27	0.18
σ_e^2	0.28	0.03
DIC (PD)	3095.5 (432.2)	

Note: MCMC with adaptation length 25,000, burnin = 25,000, iterations = 25,000.
Extended SITAR model with random coefficients.

Table 3. Edinburgh growth data.

Parameter	Model A		Model B		
	Estimate	Standard error	Estimate	Standard error	
α_0	-5.46	0.08	-4.47	0.44	
α_1	-0.72	0.09	0.24	0.42	
Ω_u	$\begin{pmatrix} 53.2 & & & & & & \\ 0.10 & 0.52 & & & & & \\ -0.16 & 0.0001 & 0.005 & & & & \\ 0.005 & -0.007 & 0.0003 & 0.0002 & & & \\ 2.05 & -0.41 & 0.02 & 0.005 & 3.50 & & \\ 0.18 & -0.05 & -0.008 & 0.00001 & 0.01 & 0.18 & \end{pmatrix}$				
$R_u = \text{Cor}(\Omega_u)$	$\begin{pmatrix} 1.0 & & & & & & \\ 0.02 & 1.0 & & & & & \\ -0.32 & 0.002 & 1.0 & & & & \\ 0.06 & -0.76 & 0.32 & 1.0 & & & \\ 0.15 & -0.30 & 0.13 & 0.18 & 1.0 & & \\ 0.06 & -0.17 & -0.26 & 0.01 & 0.01 & 1.0 & \end{pmatrix}$				
β_0	163.1	0.64	163.4	0.51	
β_1	10.3	0.09	10.3	0.08	
β_2	2.71	0.08	1.72	0.43	
β_3	0.26	0.007	0.26	0.006	
β_4	0.0041	0.0005	0.0044	0.00005	
β_5	-0.00081	0.00007	-0.00077	0.00007	
ν_0	1.37	0.18	1.38	0.18	
δ_0	-1.29	0.08	-1.21	0.08	
δ_1	-0.051	0.025	0.015	0.017	
δ_2	-0.023	0.006	-0.032	0.007	
DIC (PD)	2849.7	(478.3)	2836.1	(478.5)	

Note: MCMC with adaptation length 25,000, burnin = 25,000, iterations = 25,000. Time taken for model A with 3 parallel chains, using a 2.8 Ghz PC running under Windows 8 = 5.5 min. Model with random coefficients and complex variance function with shifted age in the variance function (model A) and actual (centred) age in the variance function (Model B).

5 An example using repeated measures of weight in pregnancy

The data are from the Avon Longitudinal Study of Parents and Children (ALSPAC, Boyd et al.¹⁸). A total of 14,541 pregnant women were recruited, who were living in a defined area of Avon including the city of Bristol during their pregnancy and had an expected delivery date between 1 April 1991 and 31 December 1992. The study website contains details of all the data that are available through a fully searchable data dictionary (<http://www.bris.ac.uk/alspac/researchers/data-access/data-dictionary>). Ethical approval for the study was obtained from the ALSPAC Ethics and Law Committee and the Local Research Ethics Committees.

We restrict analysis to singleton live term births (≥ 37 weeks' gestation). All weight measurements, which were taken routinely as part of antenatal care by midwives or obstetricians, were abstracted from obstetric records by six trained research midwives. The number of measurement occasions within individuals varies from 2 to 15 and we have, for present purposes, used only those with at least 12 measurements: this reduces the available number of individual pregnancies from 9447 to 475, provides more stable estimates and reduces any possible biases arising from selective dropout: those with few measurements may be individuals experiencing periods of illness and in any case will contribute less information to the analysis. One modelling possibility that we have not pursued is to introduce the number of measurements as a covariate, together possibly including interactions with some of the polynomial terms. Gestational age is measured in weeks centered at 30 weeks and the response is weight in kilograms.

Table 4. ALSPAC data.

Parameter	Estimate	Standard error
α_0	0.00079	0.00019
Ω_u	$\begin{pmatrix} 151.8 & & \\ 0.468 & 0.027 & \\ 2.48 & 0.019 & 0.388 \end{pmatrix}$	
$R_u = \text{Cor}(\Omega_u)$	$\begin{pmatrix} 1 & & \\ 0.23 & 1 & \\ 0.32 & 0.19 & 1 \end{pmatrix}$	
β_0	73.70	0.40
β_1	0.503	0.009
β_2	-0.0069	0.0011
β_3	-0.00043	0.000060
δ_0	-0.111	0.035
δ_1	-0.011	0.0032
DIC (PD)	17166.9 (1215.7)	

Note: Pregnancy weight regression spline model with random coefficients, against gestational age centered at 30 weeks. MCMC estimation with diffuse priors. Burnin = 1000, iterations = 5000. Adaptation length = 5000.

Our first model fits a two-level regression spline. Several models have been explored up to two join points and up to a quartic polynomial, and the following cubic with coefficients random up to the linear, and a single join point at 30 weeks provides a good fit. Let y_{ij} denote the weight of mother j ($j = 1, \dots, J$) at occasion i ($i = 1, \dots, n_j$) and $x_{ij} \equiv z_{ij}$ denotes centred gestational age. The model can be written as

$$\begin{aligned}
 y_{ij} &= \beta_{0j} + \beta_{1j}x_{ij} + \beta_{2j}x_{ij}^2 + \beta_{3j}x_{ij}^3 + \alpha_0s_{0ij}^2 + e_{ij} \\
 \beta_{0j} &= \beta_0 + u_{0j}, \beta_{1j} = \beta_1 + u_{1j} \\
 \log(\sigma_{eij}^2) &= \delta_0 + \delta_1x_{ij} + w_j \\
 \begin{pmatrix} u_{0j} \\ u_{1j} \\ w_j \end{pmatrix} &\sim N \left\{ \begin{pmatrix} 0 \\ 0 \\ 0 \end{pmatrix}, \begin{pmatrix} \sigma_{u0}^2 & & \\ \sigma_{u01} & \sigma_{u1}^2 & \\ \sigma_{u0w} & \sigma_{u1w} & \sigma_w^2 \end{pmatrix} \right\}, \quad e_{ij} \sim N(0, \sigma_e^2) \\
 s_{0ij} &= \begin{cases} 0 & \text{if } x_{ij} < 0 \\ x_{ij} & \text{if } x_{ij} \geq 0 \end{cases}
 \end{aligned} \tag{9}$$

Table 4 shows the results of fitting this model.

We see, for example, that the predicted weight at age 30 weeks is 73.7 kilograms with a between-individual standard deviation of $\sqrt{151.8} = 12.3$ cm.

We have also fitted a model allowing a different age origin for each pregnancy, but the estimated variance for these is close to zero and non-significant. Likewise, there is no significant random coefficient for the term γ_1 in the variance function.

We see significant variation in the slope as well as the intercept across pregnancies. There is a decreasing within-pregnancy variance with gestational age and the within-pregnancy variance also varies considerably between women with a standard deviation on the log scale of $\sqrt{0.388} = 0.62$, compared with a mean -0.11 . There is a moderate correlation of 0.32 between the random effect in the within-pregnancy variance function and the intercept in the growth component of the model – so women who weigh more also have more fluctuation of weight during pregnancy. Modelling the response as $\log(\text{weight})$ does not remove the dependency of the within-pregnancy variance on age.

To illustrate the heterogeneity of the data, we have plotted in Figure 3 the predicted 5th, 50th and 95th percentiles based upon the parameters in Table 4 including the level 2 variance and the level 1 variance estimated at $w_j = 0$.

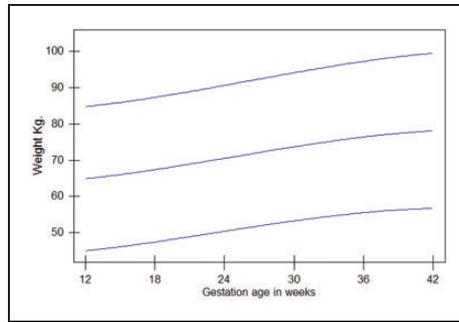


Figure 3. Predicted 5, 50, 95 percentiles of weight.

6 Discussion

Newly developed modelling procedures such as SITAR provide a framework that allows for the individual timing of growth to be incorporated within a single modelling framework to allow efficient parameter estimation. The use of MCMC methods as proposed by Willemssen et al.³ and in the present paper, provides a flexible and straightforward way of incorporating further components such as complex variance functions with random coefficients. Previous attempts to allow for individual timing of growth events, such as the age at PHV have relied on two stage procedures that involve preliminary estimation of growth events that are then used to adjust individual age scales, and these are both inefficient and biased unless procedures are used that recognise the uncertainty associated with the first stage estimates.

This paper extends existing models by allowing several random effects in the level 1 variance function associated with the random coefficients in the mean function, and in particular introduces the use of complex variance functions into the modelling of growth curves. This allows us to study the relationship between within-individual variability and other growth or individual level parameters.

The MCMC algorithm is flexible, but at least in the adolescent height growth example, it was important to provide reasonable starting values for the fixed growth curve coefficients and a long adaptation period (25,000 iterations) as well as a long burn in (25,000) in order for the MCMC algorithm to find the mode of the posterior distribution and the Metropolis algorithm to then tune correctly. It is also important to run several chains with different random starting points to check convergence. A Bayesian formulation also allows informative prior distributions to be used, and this is a topic for further research.

We have considered models with normally distributed responses, but these can be extended to ordered or unordered responses using ‘latent normal’ modelling as described by Goldstein et al.¹⁹ An extra step in the MCMC algorithm samples from an underlying multivariate normal distribution and the remaining steps are as we have described. The case of a binary response is discussed in detail in Browne and Goldstein.¹⁶ We can also extend our model to several higher levels, cross-classified and multiple membership structures as well as the multivariate case with several responses.¹ All of these extensions involve extra, but standard, steps in the MCMC algorithm. Further work could look at the computational issues that might arise in such models and we plan to study this further.

The ability to model random effects in the level 1 variance function allows us to explore how this variance relates to variation among units at higher levels. We would also stress that there is often substantive interest in modelling the level 1 variance, and this is not merely a device to avoid misspecification arising from omitted variables which might affect the model for the mean response.⁷ It is also worth noting that the variance function may in some applications be quite complex so that, for example, it might be modelled by a spline function. Where the data structure has more levels or cross classifications, we can likewise explore associations between random effects for units at these levels or classifications.¹⁰ In addition, we can include covariates in the variance as well as the mean function that are of substantive interest, such as for example social class or ethnic background.

We have used data with reasonably large (12 or more) level 1 measurements per subject but, since the dimension of Ω_{it} is at most 6, at least in principle we could obtain estimates with fewer numbers of level 1 measurements, so long as there were adequate numbers of individuals with more than six measurements. Extending the work of Hedeker et al.,⁴ a topic for further research is how different distributions of measurements per individual, their spacing and the number of individuals influence power and efficiency.

In our growth example, where there is a height asymptote for each individual, we have shown how modelling the level 1 variance on the shifted age scale is more realistic than a constant variance. It is also more realistic than modelling on the actual age scale since once individuals reach adulthood where the level 1 variation should simply reflect measurement error, we still in that case obtain estimates that result from between individual variation. For height, the level 1 variation is relatively small so that modelling this creates only small changes to the remaining parameters. In other cases, however, such as modelling weight, an understanding of how the level 1, within-individual, variance changes on an individually age-shifted scale, especially in the transition from adolescence to adulthood, will be more important.

In our example of weight change during pregnancy, the within-pregnancy standard deviation changes relatively little over the course of a pregnancy, and conclusions based on the fixed effects are likely to be robust to not including this change over pregnancy in the random effects. However, there was considerable between-woman variation in the within-pregnancy standard deviation, and this could be particularly relevant if using these models to monitor individuals.

A useful further extension of our model is to the case where we wish to predict future measurements, for example of morbidity. In such cases, we can set up a joint model where the future measurement is at level 2 and has a random effect that is allowed to covary with the other level 2 random effects. For details, see Goldstein and Kounali²⁰ and Sayers et al.²¹

A further area for exploration is in the choice of informative priors for the parameters, especially where there may be indications from existing research about likely values.

Although we have not encountered parameter identifiability problems where the models for the mean and the variance have (shifted) age in common, these problems might occur for other variables that appear in both these parts of the overall model, and this is another area for further work. Using the STATJR¹¹ software, all the models used were fitted within six minutes at most.

Acknowledgements

We are very grateful to Tim Cole for permission to use the growth data set and for helpful comments. We are extremely grateful to all the families who took part in the pregnancy study, the midwives for their help in recruiting them, and the whole ALSPAC team, which includes interviewers, computer and laboratory technicians, clerical workers, research scientists, volunteers, managers, receptionists and nurses.

Declaration of conflicting interests

The author(s) declared no potential conflicts of interest with respect to the research, authorship, and/or publication of this article.

Funding

The author(s) disclosed receipt of the following financial support for the research, authorship, and/or publication of this article: This research was funded partly under the LEMMA3 project, a node of the UK Economic and Social Research Council's (ESRC) National Centre for Research Methods (grant number RES-576-25-0035), partly by a further ESRC grant (grant number ES/K007246/1), and partly by the UK Medical research Council (MRC) (grant number MR/N027485/1). The MRC and the Wellcome Foundation (Grant ref: 102215/2/13/2) and the University of Bristol provide core support for ALSPAC.

References

1. Goldstein H. *Multilevel statistical models*, 4th ed. Chichester: Wiley, 2011.
2. Cole TJ, Donaldson MD and Ben-Shlomo Y. SITAR – a useful instrument for growth curve analysis. *Int J Epidemiol* 2010; **39**: 1558–1566.
3. Willemsen SP, Eilers PHC, Steegers-Theunissen RPM, et al. A multivariate Bayesian model for embryonic growth. *Stat Med* 2015; **34**: 1351–1365.
4. Hedeker D, Mermelstein R and Demirtas H. An application of a mixed-effects location scale model for analysis of momentary assessment (EMA) data. *Biometrics* 2008; **64**: 627–634.
5. Leckie G, French R, Charlton C, et al. Modelling heterogeneous variance-covariance components in two level multilevel models. *J Educ Behav Stat* 2014; **39**: 307–332.

6. Hedeker D and Nordgren R. “MIXREGLS: a program for mixed-effects location scale analysis. *J Stat Softw* 2013; **52**: 1–38.
7. Leckie G. runmixregls – a program to run the MIXREGLS mixed-effects location scale software from within Stata. *J Stat Softw* 2014; **59** (Code Snippet 2): 1–41.
8. Lee Y and Nelder JA. Double hierarchical generalised linear models (with discussion). *Appl Stat* 2006; **55**: 139–185.
9. Rast P, Hofer SM and Sparks C. Modeling individual differences in within-person variation of negative and positive affect in a mixed effects location scale model using BUGS/JAGS. *Multivariate Behav Res* 2012; **47**: 177–200.
10. Brunton-Smith I, Sturgis P and Leckie G. Detecting, and understanding interviewer effects on survey data using a cross-classified mixed-effects location scale model. *J R Stat Soc* 2017; **180**: 551–568.
11. Charlton CMJ, Michaelides DT, Parker RMA, et al. Stat-JR software. Centre for Multilevel Modelling, University of Bristol and Electronics and Computer Science, University of Southampton, www.bristol.ac.uk/cmm/software/statjr/ (2013, accessed 11 April 2017).
12. Preece MA and Baines MJ. A new family of mathematical models describing the human growth curve. *Ann Hum Biol* 1978; **5**: 1–24.
13. Pinheiro J, Bates D, DebRoy S, and the R development core team. Nmlle: linear and nonlinear mixed effects models. R package version 3, 2013, pp.1–113.
14. R Core Team. *A language and environment for statistical computing. R foundation for statistical computing*. Austria: Vienna, www.R-project.org (2013, accessed 11 April 2017).
15. Cole TJ, Pan H and Butler GE. A mixed effects model to estimate timing and intensity of pubertal growth from height and secondary sexual characteristics. *Ann Hum Biol* 2014; **41**: 76–83.
16. Browne WJ and Goldstein H. MCMC sampling for a multilevel model with non-independent residuals within and between cluster units. *J Educ Behav Stat* 2010; **35**: 453–473.
17. Tanner J and Whitehouse R. Clinical longitudinal standards for height, weight, height velocity and weight velocity and stages of puberty. *Arch Dis Child* 1976; **51**: 170–179.
18. Boyd A, Golding J, Macleod J, et al. Cohort profile: the ‘Children of the 90s’ – the index offspring of the Avon Longitudinal Study of Parents and Children. *Int J Epidemiol* 2012 2013; **42**: 111–127.
19. Goldstein H, Carpenter J, Kenward M, et al. Multilevel models with multivariate mixed response types. *Stat Model* 2009; **9**: 173–197.
20. Goldstein H and Kounali D. Multivariate multilevel modelling of childhood growth, numbers of growth measurements and adult characteristics. *J R Stat Soc A* 2009; **172**: 599–613.
21. Sayers A, Heron J, Smith AD, et al. Joint modelling compared with two stage methods for analysing longitudinal data and prospective outcomes: a simulation study of childhood growth and BP. *Stat Meth Med Res* 2014. DOI: 10.1177/0962280214548822.
22. Browne WJ and Draper D. Implementation and performance issues in the Bayesian and likelihood fitting of multilevel models. *Comput Stat* 2000; **15**: 391–420.

Appendix I. MCMC algorithm steps

We detail below the MCMC algorithm for our general two-level model. The modular nature of MCMC algorithms means that the steps in the algorithm are readily generalised to models with three or more levels and to cross-classifications and to different response types. Where there is overlap between the current algorithm and that proposed in Browne and Goldstein,¹⁶ we shall for brevity refer the reader to that paper. In what follows, we focus on the additional steps required to implement novel extensions proposed in the current paper, namely the inclusion of random effects within the level 1 variance function.

Step 1: For the level 2 random effects (residuals), u_{1j} the conditional posterior distribution is given by

$$p(u_{1j} | y_{ij}, \Omega_{u1}^*, \sigma_{eij}^2, u_{2j}) \propto \exp(u_{1j}^T (\Omega_{u1}^*)^{-1} u_{1j}) \prod_{i=1}^{n_j} \left(\frac{1}{\sigma_{eij}^2} \right) \exp \left[-\frac{1}{2\sigma_{eij}^2} (y_{ij} - X_{1ij}\beta - Z_{1ij}u_{1j})^2 \right]$$

where

$$\Omega_{u1}^* = \Omega_{u1} - \Omega_{u12} \Omega_{u2}^{-1} \Omega_{u12}^T$$

which is the covariance matrix for u_1 conditioned on u_2 and is derived from equation (1) This is analogous to the step given in Browne and Goldstein¹⁶ with $\Omega_{u_1}^*$ replacing Ω_u and results in a Gibbs sampling step from a multivariate normal distribution.

Step 2: For the fixed effects, we assume a diffuse prior distribution $p(\beta) \propto 1$ so that

$$p(\beta | y_{ij}, \sigma_{eij}^2, u_{1j}) \propto \prod_{i,j} \sigma_{eij}^{-2} \exp \left[-0.5 \sigma_{eij}^{-2} (y_{ij} - Z_{1ij} u_{1j} - X_{1ij} \beta)^2 \right]$$

And we sample from

$$\begin{aligned} \beta &\sim MVN(\hat{\beta}, \hat{D}_\beta) \\ \hat{D}_\beta &= \left[\sum_{i,j} X_{1ij}^T \sigma_{eij}^{-2} X_{1ij} \right]^{-1} \\ \hat{\beta} &= \left[\sum_{i,j} X_{1ij}^T \sigma_{eij}^{-2} X_{1ij} \right]^{-1} \left[\sum_{i,j} \sigma_{eij}^{-2} X_{1ij}^T (y_{ij} - Z_{1ij} u_{1j}) \right] \end{aligned}$$

The level 1 residuals are obtained by subtraction using the first line of equation (1).

Step 3: For the fixed coefficients within the level 1 variance function, we use a random walk Metropolis step which is the same as that given by Browne and Goldstein¹⁶ noting here that the level 1 variance function will additionally contain level 2 random effects.

Step 4: Sampling the level 2 covariance matrix, Ω_u , is a standard conjugate inverse Wishart Gibbs sampling step as given by Browne and Goldstein¹⁶ and we omit details.

Step 5: The random effects within the level 1 variance function are the u_{2j} . Here, we adopt a single site updating strategy using another random walk Metropolis step for each element of u_{2j} . For example, in the simple case where there is just one random effect in the mean function and one random effect in the variance function, we adopt the following step:

To sample each u_{2j} , we condition on the corresponding u_{1j} . At iteration t , we propose a new value u_{2j}^* with proposal distribution $N(u_{2j}^{(t-1)}, \sigma_p^2)$ where σ_p^2 may be obtained using an adaptation period where the value is tuned.²² We accept the proposed value according to the Metropolis ratio

$$\frac{\prod_{i,j} \sigma_{eij}^{-2} (*) \exp \left[-0.5 \sigma_{eij}^{-2} (*) (y_{ij} - Z_{1ij} u_{1j} - X_{1ij} \beta)^2 \right]}{\prod_{i,j} \sigma_{eij}^{-2} (t-1) \exp \left[-0.5 \sigma_{eij}^{-2} (t-1) (y_{ij} - Z_{1ij} u_{1j} - X_{1ij} \beta)^2 \right]} \times \frac{\varnothing(u_{2j}^*)}{\varnothing(u_{2j}^{(t-1)})}$$

where $\sigma_{eij}^2 (*) = \exp(X_{2ij} \delta + u_{2j}^*)$, $\sigma_{eij}^2 (t-1) = \exp(X_{2ij} \delta + u_{2j}^{(t-1)})$ and \varnothing is the conditioned normal density function of u_{2j} given u_{1j} and $u_{2j}^{(t-1)}$ is the current value.