

Estimating reducible stochastic differential equations by conversion to a least-squares problem

Oscar García

Dasometrics, Concón, Chile

Abstract

Stochastic differential equations (SDEs) are increasingly used in longitudinal data analysis, compartmental models, growth modelling, and other applications in a number of disciplines. Parameter estimation, however, currently requires specialized software packages that can be difficult to use and understand. This work develops and demonstrates an approach for estimating reducible SDEs using standard nonlinear least squares or mixed-effects software. Reducible SDEs are obtained through a change of variables in linear SDEs, and are sufficiently flexible for modelling many situations. The approach is based on extending a known technique that converts maximum likelihood estimation for a Gaussian model with a nonlinear transformation of the dependent variable into an equivalent least-squares problem. A similar idea can be used for Bayesian maximum a posteriori estimation. It is shown how to obtain parameter estimates for reducible SDEs containing both process and observation noise, including hierarchical models with either fixed or random group parameters. Code and examples in *R* are given. Univariate SDEs are discussed in detail, with extensions to the multivariate case outlined more briefly. The use of well tested and familiar standard software should make SDE modelling more transparent and accessible.

Keywords: stochastic processes, longitudinal data, growth curves, compartmental models, mixed-effects, *R*

Contents

1	Introduction	2
2	Transformed Gaussian model	3
2.1	Example 1	3
2.2	Generalized formulation	4
2.3	Maximum likelihood	4
2.4	Solution of Example 1	5
2.5	Correlation	6
2.6	Maximum a posteriori estimation	6

Email address: garcia@dasometrics.net (Oscar García)

¹ORCID: 0000-0002-8995-1341

3	Stochastic differential equations	7
3.1	Estimation	9
3.2	Computational details	10
3.2.1	Example 2	12
3.3	Hierarchical models	13
3.3.1	Fixed local parameters	15
	Example 3.	16
3.3.2	Random local parameters	16
	Example 4.	17
3.4	Multivariate SDEs	17
4	Conclusions	19
	Acknowledgements	20
Appendix A	Example 1, Section 2.4	20
Appendix B	Functions for ML parameter estimation in SDE models, Section 3.2	21
Appendix C	Example 2, Section 3.2.1	23
Appendix C.1	Additive process noise	23
Appendix C.2	Multiplicative process noise	24
Appendix D	ML estimation in hierarchical SDE models, Section 3.3	25
Appendix E	Example 3, Section 3.3.1	26
Appendix F	Example 4, Section 3.3.2	29
	References	32

1. Introduction

Stochastic differential equation (SDE) models incorporate random elements into ordinary differential equations, representing the effects of a noisy environment on rates of change. In addition to their traditional applications in physics and engineering (Gardiner 1985), SDEs are becoming routine in finance (Kunita 2010, Kloeden and Platen 1992, Sec. 7.3) and pharmacokinetics (Kristensen et al 2005, Klim et al 2009, Donnet and Samson 2013), and are increasingly used in fields like econometrics (Paige and Allen 2010, Bu et al 2010, 2016), animal growth (Lv and Pitchford 2007, Strathe et al 2009, Filipe et al 2010), oncology (Sen 1989, Favetto and Samson 2010), and forestry (García 1983, Broad and Lynch 2006, Batho and García 2014). Other applications include Artzrouni and Reneke (1990), Cleur (2000), Driver et al (2017). In particular, deterministic growth curves are commonly fitted to repeated measurements assuming that observation error is the only source of randomness (Davidian and Giltinan 2003), but an SDE growth model can be more realistic (Hotelling 1927, Donnet et al 2010). In addition to environmental or process noise, some models include also observation errors (García 1983, Donnet et al 2010, Donnet and Samson

2013). Another extension of the basic SDE is to hierarchical models with repeated measurements, where some parameters may be specific to each individual or group. Such *local* parameters can be fixed numbers (García 1983) or, more often, are viewed as random in mixed-effects models (Donnet and Samson 2008, Klim et al 2009, Picchini et al 2010, Driver et al 2017).

Parameter estimation for general SDEs is a difficult problem, with most algorithms employing approximations to maximum likelihood (e. g., Nielsen et al 2000, Bishwal 2008, Picchini et al 2010, Donnet and Samson 2013). Bayesian methods have also been proposed (Donnet et al 2010, Brouste et al 2014, King et al 2016, Whitaker et al 2017). *R* packages implementing estimation for various types of SDE models include *ctsem*, *dynr*, *mixedsde*, *msde*, *pomp*, *sde*, *Sim.DiffProc*, and *yuima* (CRAN 2017); *pomp* and *yuima* include Bayesian procedures. An alternative method is presented here for SDEs that can be reduced to linear by a change of variables. The approach is based on an extension of ideas of Box and Cox (1964) and Furnival (1961), which convert maximum-likelihood (ML) estimation for a transformed Gaussian model into a least-squares problem. The same principle can be used to obtain Bayesian maximum a posteriori (MAP) estimates. The sum of squares can be minimized with standard statistical software like the `nls` function in *R*, or `nlme` for mixed-effects formulations (Pinheiro and Bates 2000).

The general methodology is described in the next section, followed by its application to SDE models with observation errors in Section 3. Examples in *R* are given. The SDE models may include hierarchical structures with group fixed or random effects (Sec. 3.3). To simplify, only univariate SDEs are discussed in full detail, with extensions to vector-valued multiresponse models treated more briefly in Section 3.4. A Conclusions section closes the article.

2. Transformed Gaussian model

2.1. Example 1

Venables and Ripley (2002, Sec. 8.7) describe a dataset containing measurements of the concentration of a chemical GAG in the urine of 314 children of various ages. Assume that one wants an equation summarizing the relationship between GAG concentration (variable y) and age (variable x). Plotting shows that the relationship is nonlinear, so that a simple linear regression $y = \beta_0 + \beta_1 x$ is unsatisfactory. A common strategy is to try regressions on transformations of x and/or y , such as logarithms or various positive or negative powers. We shall find “good” exponents for powers in both x and y .

Instead of a simple power transformation, it is slightly more convenient to use

$$y^{(\lambda)} = \begin{cases} \frac{y^\lambda - 1}{\lambda} & \text{if } \lambda \neq 0 \\ \ln y & \text{if } \lambda = 0 \end{cases} \quad (1)$$

(Box and Cox 1964). This is an increasing function of y for any λ , it is continuous in λ at $\lambda = 0$, and includes the logarithmic transformation as a special case. Physicists call eq. (1) a *generalized logarithm* (Martinez et al 2008).

The model is then

$$y_i^{(\lambda_y)} = \beta_0 + \beta_1 x_i^{(\lambda_x)} + \epsilon_i. \quad (2)$$

Assuming that the ϵ_i are independent and normally distributed, it is possible to write down the likelihood function and use a general optimization algorithm to find the maximum likelihood (ML) estimates for the parameters $\beta_0, \beta_1, \lambda_x, \lambda_y$, and $\sigma^2 = \text{Var}[\epsilon_i]$. However, the optimization can be

time-consuming and ill-conditioned, requiring good starting points to achieve convergence. Box and Cox (1964) showed a way of converting the likelihood maximization into a least-squares solution, which is more efficient and stable. Furnival (1961) had used the same idea to devise an index for comparing dependent variable transformations. Specifically, Box and Cox (1964) considered various one- and two-parameter transformations on the left-hand side and a linear right-hand side, solving by ordinary linear least squares. Estimation in eq. (2) is similar, except that with a free λ_x it requires a non-linear least-squares procedure.

The next two sections develop an extension of the Box–Cox–Furnival approach, needed for the application to SDEs. The ML estimation for Example 1, both by the method of Box and Cox (1964) and by direct maximization of the likelihood, is presented in Section 2.4.

2.2. Generalized formulation

Consider a vector of n observations $\mathbf{y} = (y_1, \dots, y_n)$, and a transformation

$$\mathbf{z} = \boldsymbol{\varphi}(\mathbf{y}, \boldsymbol{\theta}_1)$$

that may depend on a vector of unknown parameters $\boldsymbol{\theta}_1$. Assume that $\boldsymbol{\varphi}$ is one-to-one and differentiable over the admissible domain, and that the transformed variables follow a Gaussian nonlinear model

$$z_i = f_i(\boldsymbol{\theta}_2) + \epsilon_i ,$$

where the ϵ_i are independent normally distributed with mean 0 and a common variance σ^2 . In Box and Cox (1964) and in Section 2.1 the transformations are of the simpler form $z_i = \varphi(y_i, \boldsymbol{\theta}_1)$, f_i takes the form $f(\mathbf{x}_i, \boldsymbol{\theta}_2)$ where \mathbf{x}_i is a vector of predictors, and the parameters in $\boldsymbol{\theta}_1$ and in $\boldsymbol{\theta}_2$ are different. Those restrictions are not imposed here. In particular, we denote by $\boldsymbol{\theta}$ the vector containing the union of the elements from $\boldsymbol{\theta}_1$ and $\boldsymbol{\theta}_2$, some of which may be the same, and write

$$\mathbf{z} = \boldsymbol{\varphi}(\mathbf{y}, \boldsymbol{\theta}) \tag{3}$$

and

$$z_i = f_i(\boldsymbol{\theta}) + \epsilon_i , \tag{4}$$

There is no loss in generality from the assumption of homoscedastic uncorrelated errors in eq. (4), because any other correlation structure can be accommodated through a linear transformation (Sec. 2.5).

2.3. Maximum likelihood

The following extends the results of Box and Cox (1964) and Furnival (1961) to the more general formulation of eqns. (3)–(4):

Proposition 1. *Let*

$$J = J(\mathbf{y}, \boldsymbol{\theta}) = \text{abs} \left(\left| \frac{\partial \boldsymbol{\varphi}}{\partial \mathbf{y}} \right| \right) = \text{abs} \left(\left| \left\{ \frac{\partial \varphi_i}{\partial y_j} \right\} \right| \right) \tag{5}$$

be the absolute value of the Jacobian determinant of the transformation (3), and write

$$u_i = \frac{\epsilon_i}{J^{1/n}} = \frac{\varphi_i(\mathbf{y}, \boldsymbol{\theta}) - f_i(\boldsymbol{\theta})}{J(\mathbf{y}, \boldsymbol{\theta})^{1/n}} . \tag{6}$$

Then, the ML estimate of $\boldsymbol{\theta}$ is the value $\hat{\boldsymbol{\theta}}$ that minimizes the sum of squares $S(\boldsymbol{\theta}) = \sum u_i^2$. The ML estimate of σ^2 is

$$\hat{\sigma}^2 = J(\mathbf{y}, \hat{\boldsymbol{\theta}})^{2/n} \frac{S(\hat{\boldsymbol{\theta}})}{n} . \quad (7)$$

Proof. The probability density of \mathbf{y} , and hence the likelihood in terms of the original observations, is the product of the density of \mathbf{z} and the absolute value of the Jacobian of the transformation (e. g., Wilks 1962, Sec. 2.8),

$$L = \frac{1}{(2\pi)^{n/2} \sigma^n} \exp \left(-\frac{\sum \epsilon_i^2}{2\sigma^2} \right) J . \quad (8)$$

Substituting $\epsilon_i = J^{1/n} u_i$ (from eq. (6)) and $\sigma = J^{1/n} \sigma_u$,

$$L = \frac{1}{(2\pi)^{n/2} \sigma_u^n} \exp \left(-\frac{\sum u_i^2}{2\sigma_u^2} \right) . \quad (9)$$

This has the form of a normal density, and therefore the standard normal theory applies (e. g., Seber and Wild 2003, Sec. 2.2.1). Briefly, equating to 0 the derivative of L with respect to σ_u^2 it is found that, for given $\boldsymbol{\theta}$, the likelihood is maximized by

$$\sigma_u^2 = \frac{\sum u_i^2}{n} . \quad (10)$$

Substituting into eq. (9),

$$\max_{\sigma_u^2} L = \left(\frac{n}{2\pi \sum u_i^2} \right)^{n/2} \exp(-n/2) .$$

Clearly, the ML estimate $\hat{\boldsymbol{\theta}}$ can be calculated minimizing over $\boldsymbol{\theta}$ the sum of squares $\sum u_i(\mathbf{y}, \boldsymbol{\theta})^2$. Equation (10) and $\sigma^2 = J^{2/n} \sigma_u^2$ give eq. (7). \square

An alternative variance estimator with $n - p$ instead of n in the denominator of eq. (7) is often used, which generally has less bias but a larger mean square error. It may be worth noting that in general the u_i are not Gaussian, even though the likelihood takes that form.

2.4. Solution of Example 1

Eq. (2) corresponds to (3)–(4) with $z_i = y_i^{(\lambda_y)}$ and $\boldsymbol{\theta} = (\beta_0, \beta_1, \lambda_x, \lambda_y)$.

In this instance the Jacobian matrix is diagonal, and the determinant is the product of the diagonal elements $\partial z_i / \partial y_i$. Then, as shown by Furnival (1961) and by Box and Cox (1964),

$$J^{1/n} = \dot{y}^{\lambda-1} , \quad (11)$$

where \dot{y} denotes the geometric mean of the y_i .

R commands for finding the solution are given in Appendix A. First, the exponents are fixed as $\lambda_x = 1$ and $\lambda_y = 0$, and a simple linear regression $\ln y = \beta_0 + \beta_1 x$ is used to produce suitable starting points for the optimization. Then, the sum of squares of the u_i from eq. (6) is minimized with the nonlinear regression procedure `nls`, making use of eq. (11). The ML estimates are found to be

$$\beta_0 = 3.314 , \quad \beta_1 = -0.3502 , \quad \lambda_x = 0.4249 , \quad \lambda_y = 0.1032 .$$

From eq. (7), the ML variance estimate is $\sigma^2 = 0.3839$.

A more convenient and parsimonious relationship can be obtained by rounding $\lambda_x = 0.5, \lambda_y = 0$, and using the linear regression

$$\ln y = 3.366 - 0.5069\sqrt{x} .$$

It is also shown in Appendix A that the same estimates are obtained by direct minimization of the negative log-likelihood from eq. (8), confirming the validity of the Box-Cox approach.

2.5. Correlation

Suppose that the ϵ_i in eq. (4) are not independent. Let the variance-covariance matrix of the vector $\boldsymbol{\epsilon} = (\epsilon_1, \dots, \epsilon_n)$ be

$$\text{Var}(\boldsymbol{\epsilon}) = \sigma^2 \mathbf{C} , \quad (12)$$

where the matrix \mathbf{C} may depend on unknown parameters. The Cholesky factorization gives \mathbf{C} as the product of a lower-triangular matrix \mathbf{L} and its transpose:

$$\mathbf{C} = \mathbf{L}\mathbf{L}' . \quad (13)$$

It follows that the elements of $\mathbf{L}^{-1}\boldsymbol{\epsilon}$ are uncorrelated, and so are those of $\mathbf{L}^{-1}\mathbf{z}$ (e. g., Seber and Wild 2003, Sec. 2.1.4, where the upper-triangular Cholesky factor $\mathbf{U} = \mathbf{L}'$ is used). Therefore, with the substitution $\boldsymbol{\varphi}(\mathbf{y}, \boldsymbol{\theta}) \rightarrow \mathbf{L}^{-1}\boldsymbol{\varphi}(\mathbf{y}, \boldsymbol{\theta})$ in eq. (3) we are in the same situation as before. Eq. (5) becomes

$$J = \text{abs} \left(\left| \frac{\partial \mathbf{L}^{-1}\boldsymbol{\varphi}}{\partial \mathbf{y}} \right| \right) = \text{abs} \left(\left| \left\{ \frac{\partial \varphi_i}{\partial y_j} \right\} \right| \right) / |\mathbf{L}| , \quad (14)$$

with

$$|\mathbf{L}| = \prod_{i=1}^n l_{ii}$$

since the determinant of a triangular matrix is the product of the elements on the diagonal, and those elements l_{ii} in \mathbf{L} are positive. In practice it is better to use logarithms to prevent numerical overflow or underflow:

$$|\mathbf{L}|^{1/n} = \exp \left(\frac{1}{n} \sum_{i=1}^n \ln l_{ii} \right) . \quad (15)$$

With this J , the vector of u_i for the sum of squares is now

$$\mathbf{u} = \frac{\mathbf{L}^{-1}\boldsymbol{\epsilon}}{J^{1/n}} \quad (16)$$

(c. f. eq. (6)).

2.6. Maximum a posteriori estimation

In Bayesian MAP estimation the likelihood is weighted by an *a priori* parameter distribution density:

$$\max_{\boldsymbol{\theta}, \sigma^2} L(\boldsymbol{\theta}, \sigma^2) p(\boldsymbol{\theta}) .$$

If L is Gaussian, this has the same form as eq. (8) with $p(\boldsymbol{\theta})$ in place of J . Therefore, analogously to Sec. 2.3, the MAP estimate of $\boldsymbol{\theta}$ is the value $\hat{\boldsymbol{\theta}}$ that minimizes the sum of squares $S(\boldsymbol{\theta}) = \sum (\epsilon_i/p^{1/n})^2$, and the estimate of σ^2 is $\hat{\sigma}^2 = p(\hat{\boldsymbol{\theta}})^{2/n} S(\hat{\boldsymbol{\theta}})/n$. Of course, this can be combined with variable transformations as above.

3. Stochastic differential equations

The behavior of systems that evolve in time is often modelled by a relationship between a magnitude x and its rate of change, given by a differential equation $dx/dt = f(x)$. In many situations the rate is subject to perturbations, due to varying environmental conditions, that cause substantial deviations from the predicted trajectory. It is then more realistic to include noise, represented by a random function $\epsilon(t)$:

$$\frac{dx}{dt} = f(x) + g(x)\epsilon(t) . \quad (17)$$

In particular, ϵ may be assumed to be *white noise*, where the values $\epsilon(t)$ have the same distribution for all t , with mean 0, and values for different t 's are uncorrelated. Such an SDE was proposed by Langevin in 1908 as a simpler formulation of Einstein's 1905 theory of Brownian motion. Hotelling (1927) used a similar model for logistic population growth. By the middle of the 20th Century it had been found that a rigorous mathematical treatment of the theory becomes surprisingly complex. Textbooks go over extensive advanced mathematical background material before dealing with SDEs (Allen 2007, Arnold 1974, Gardiner 1985, Henderson and Plaschko 2006, Kloeden and Platen 1992, Øksendal 2003). For modelling, an intuitive understanding may suffice.

One problem with eq. (17) is that a completely uncorrelated $\epsilon(t)$ would be discontinuous everywhere and have infinite variance. The way around that is to work instead with its integral $W(t)$, known as a Brownian motion or Wiener process. Increments ΔW represent the accumulated white noise over a time interval Δt , and are uncorrelated for non-overlapping intervals. Because of this, the expected value of ΔW is 0 and the variance is proportional to $|\Delta t|$; the proportionality constant is standardized as 1. As a sum of a potentially large number of uncorrelated increments, $W(t)$ is Gaussian. The white noise in eq. (17) would correspond to the limit of $\Delta W/\Delta t$ as Δt tends to 0 but, although $W(t)$ is continuous, its fine structure is exceedingly rough (it is not differentiable), and such limit does not exist. SDEs are therefore conventionally written in terms of differentials:

$$dX(t) = f[X(t), t] dt + g[X(t), t] dW(t) , \quad (18)$$

also called a continuous-time diffusion model. The capital letters emphasize the fact that $X(t)$ and $W(t)$ are random variables. In general, f and g can depend on time, in a *time-variant* SDE, although it is more common to assume that behavior does not change over time (*time-invariant* or *autonomous*). The function f is sometimes called the drift or trend, and g the diffusion function or volatility.

Another difficulty lies with the definition of the integral of the second term on the right-hand side of eq. (18). There X and W are stochastic processes and, unlike what happens with ordinary functions, when representing the integral as the limit of a sum the location within a time interval where $X(t)$ is evaluated makes a difference. Two definitions are commonly used, that of Ito, with $X(t)$ evaluated at the start of each interval, and that of Stratanovich, based on the interval mid-point. We will use only volatilities g independent of X , in which case Ito and Stratanovich coincide, and the rules of ordinary calculus hold. A useful fact is that, for any sufficiently well-behaved function g ,

$$\delta = \int g(t) dW(t) \text{ is Gaussian with } E[\delta] = 0 \text{ and } \text{Var}[\delta] = \int g(t)^2 dt . \quad (19)$$

More generally, the state of the system may be specified by several variables, and then $X(t)$ is a vector. To simplify the presentation only one-dimensional time-invariant SDEs will be discussed

in detail, mentioning briefly the main changes needed for time-variant models. Most results can be extended to multiple state variables by substituting appropriate vectors and matrices (Sec. 3.4).

In general, explicit forms for the probability distributions associated to eq. (18) are not available, and parameter estimation is difficult (Nielsen et al 2000, Bishwal 2008, Donnet and Samson 2013). Linear SDEs, however, are more tractable and can be explicitly solved. The same is true for any process that can be reduced to a linear SDE through a sufficiently smooth nonlinear transformation (Kloeden and Platen 1992, Sec. 4.3). The model used here is a linear SDE

$$dY = (\beta_0 + \beta_1 Y) dt + \sigma_p dW , \quad (20)$$

where Y is a one-to-one differentiable transformation $Y = \varphi(X, \boldsymbol{\lambda})$ of the original variable X , parameterized by a one- or higher-dimensional vector $\boldsymbol{\lambda}$. The constants $\beta_0, \beta_1, \sigma_p$, and the elements in $\boldsymbol{\lambda}$, are parameters to be estimated. There are no additional difficulties if a parameter appears both in the transformation and in the rest of the model, so that, more generally,

$$Y = \varphi(X, \boldsymbol{\theta}) , \quad (21)$$

where $\boldsymbol{\theta}$ is a vector including all the model parameters. In the literature, Y and dW are often written as $Y(t)$ and $dW(t)$, or as Y_t and dW_t , respectively.

The data are observations x_i of $X(t_i)$ at n occasions $t_1 < t_2 < \dots < t_n$. Optionally, the observations may be subject to measurement or sampling errors such that

$$y_i = \varphi(x_i, \boldsymbol{\theta}) = Y(t_i) + \epsilon_i ; \quad i = 1, \dots, n , \quad (22)$$

with independent normal errors $\epsilon_i \sim N(0, \sigma_m^2)$. The initial condition $Y(t_0)$ at $t_0 < t_1$ may be a known or unknown constant, or a Gaussian random variable. In general, $Y(t_0) = y_0 + \epsilon_0$, with $\epsilon_0 \sim N(0, \sigma_0^2)$ and independent of the other ϵ_i (σ_0 may be 0).

In a time-variant model, β_0, β_1 and σ_p can be functions of t , possibly containing unknown parameters. The transformation φ could also depend on t (Bu et al 2016).

The linear SDE in eq. (20) is known as the Ornstein-Uhlenbeck process in physics, or the Vasicek model in financial economics. Parameter estimation for linear SDEs, including multivariate and mixed-effects models with observation errors, has been implemented in the *R* packages *PSM* by Klim et al (2009) and *ctsem* by Driver et al (2017). Reducible SDEs have been used by García (1983), Lv and Pitchford (2007), Paige and Allen (2010), Filipe et al (2010), and Bu et al (2010, 2016).

A power transformation $Y = X^\lambda$ in eq. (21) produces a stochastic version of the Bertalanffy-Richards growth model (von Bertalanffy 1949, Richards 1959), which is very flexible and includes other commonly used models as special cases: von Bertalanffy's animal growth equation with $\lambda = 1/3$, the Mistcherlich or monomolecular with $\lambda = 1$, and the logistic with $\lambda = -1$. The Box-Cox transformation $Y = X^{(\lambda)}$ from eq. (1) includes in addition the Gompertz model, $\lambda = 0$ (García 1983, Seber and Wild 2003, sections 7.3 and 7.5.3). García (2008) discusses a "double Box-Cox" transformation, with two shape parameters, that covers nearly all the sigmoid growth curves from the literature.

In some applications the emphasis lies on modelling the volatility, as the trend may not be important if the system operates near the steady state (Bu et al 2010, 2016). A volatility function $g(X, t)$ can be converted to one that does not depend on X by using the Lamperti transform

$$Y = \int \frac{dX}{g(X, t)}$$

(see Moller and Madsen 2010, for details). In particular, an SDE with multiplicative error $\sigma X dW$ can be converted to one with additive error σdW by a logarithmic transformation.

3.1. Estimation

From linearity, $Y(t)$ and the y_i are Gaussian. The means, variances and covariances needed for the likelihood can be obtained by integrating eq. (20). Assume for now that $\beta_1 \neq 0$, the changes needed if $\beta_1 = 0$ are indicated later. As in linear ordinary differential equations, let us multiply throughout by the integrating factor $e^{-\beta_1 t}$. Using the rule for differentiation of a product,

$$\begin{aligned} d \left[e^{-\beta_1 t} (Y + \beta_0/\beta_1) \right] &= e^{-\beta_1 t} d(Y + \beta_0/\beta_1) + d(e^{-\beta_1 t})(Y + \beta_0/\beta_1) \\ &= e^{-\beta_1 t} [dY - (\beta_0 + \beta_1 Y) dt] , \end{aligned}$$

so that

$$d \left[e^{-\beta_1 t} (Y + \beta_0/\beta_1) \right] = \sigma_p e^{-\beta_1 t} dW . \quad (23)$$

Integrating both sides between t_0 and t_i ,

$$e^{-\beta_1 t_i} [Y(t_i) + \beta_0/\beta_1] - e^{-\beta_1 t_0} [Y(t_0) + \beta_0/\beta_1] = \sigma_p \int_{t_0}^{t_i} e^{-\beta_1 t} dW(t) .$$

Finally, substituting $Y(t_i) = y_i - \epsilon_i$ from eq. (22), and solving for y_i ,

$$y_i = -\beta_0/\beta_1 + e^{\beta_1(t_i-t_0)} (y_0 + \beta_0/\beta_1 - \epsilon_0) + \epsilon_i + \delta_{0i} ; \quad i = 1, \dots, n , \quad (24)$$

where, from (19),

$$\delta_{0i} = \sigma_p \int_{t_0}^{t_i} e^{\beta_1(t_i-t)} dW(t)$$

is a Gaussian random variable. For a time-variant model, similar results are obtained replacing $e^{-\beta_1 t}$ with the integrating factor $\exp[-\int \beta(t) dt]$, and moving $\sigma_p(t)$ inside the integral (Kloeden and Platen 1992, Sec. 4.4). This can be generalized to the multidimensional case (Sec. 3.4).

The means, variances and covariances determining the Gaussian probability density of a sample (y_1, \dots, y_n) can be obtained from eq. (24) (García 1983). The methods of Sec. 2 can then be used, noting that x_i and y_i here correspond to y_i and z_i in Sec. 2. A sum of squares $\sum u_i^2$, calculated from the observations x_1, \dots, x_n , can be minimized to compute ML or MAP estimates for the parameters $\boldsymbol{\lambda}, \beta_0, \beta_1, \sigma_p, \sigma_m$, and possibly y_0 and σ_0 .

The covariance matrix of the y_i is dense, all the y_i are correlated because the integrals in δ_{0i} involve overlapping time intervals. It is more convenient to introduce an additional transformation from \mathbf{y} to \mathbf{z} that produces diagonal or tri-diagonal covariance matrices, simplifying computation. The idea is to consider changes between consecutive observations, integrating eq. (23) from t_{i-1} to t_i instead of from t_0 to t_i . The stochastic integrals are then independent, because the time intervals do not overlap. If there is an observation error, it affects the observed changes over the adjacent intervals, so that these are correlated, but intervals further apart remain independent. The relationship between consecutive observations is obtained substituting $i-1$ for 0 in eq. (24):

$$y_i = -\beta_0/\beta_1 + e^{\beta_1 \Delta_i} (y_{i-1} + \beta_0/\beta_1 - \epsilon_{i-1}) + \epsilon_i + \delta_i ; \quad i = 1, \dots, n , \quad (25)$$

where $\Delta_i \equiv t_i - t_{i-1}$ and

$$\delta_i = \sigma_p \int_{t_{i-1}}^{t_i} e^{\beta_1(t_i-t)} dW(t) .$$

From (19), the δ_i are independent Gaussian random variables with $E[\delta_i] = 0$ and

$$\text{Var}[\delta_i] = \sigma_p^2 \int_{t_{i-1}}^{t_i} e^{2\beta_1(t_i-t)} dt = \frac{\sigma_p^2}{2\beta_1} \left(e^{2\beta_1\Delta_i} - 1 \right) .$$

Define

$$z_i = y_i + \beta_0/\beta_1 - e^{\beta_1\Delta_i} (y_{i-1} + \beta_0/\beta_1) ; \quad i = 1, \dots, n . \quad (26)$$

The Jacobian of the transformation $\mathbf{y} \rightarrow \mathbf{z}$ is 1, so that the z_i can be directly substituted for the y_i in the likelihood. From eq. (25),

$$z_i = -e^{\beta_1\Delta_i} \epsilon_{i-1} + \epsilon_i + \delta_i .$$

Therefore,

$$E[z_i] = 0 ; \quad i = 1, \dots, n \quad (27a)$$

$$\text{Var}[z_1] = e^{2\beta_1\Delta_1} \sigma_0^2 + \sigma_m^2 + \frac{\sigma_p^2}{2\beta_1} \left(e^{2\beta_1\Delta_1} - 1 \right) \quad (27b)$$

$$\text{Var}[z_i] = e^{2\beta_1\Delta_i} \sigma_m^2 + \sigma_m^2 + \frac{\sigma_p^2}{2\beta_1} \left(e^{2\beta_1\Delta_i} - 1 \right) ; \quad i = 2, \dots, n \quad (27c)$$

$$\text{Cov}[z_i, z_{i-1}] = \text{Cov}[z_{i-1}, z_i] = -e^{\beta_1\Delta_i} \sigma_m^2 ; \quad i = 2, \dots, n \quad (27d)$$

$$\text{Cov}[z_i, z_j] = 0 \quad \text{if} \quad |i - j| > 1 \quad (27e)$$

(García 1983, Seber and Wild 2003, Sec. 7.5.3). It can be seen that the z_i are the conditional residuals $y_i - E[y_i|y_{i-1}]$ and, as expected, they are uncorrelated if $\sigma_m = 0$, or are negatively correlated with the adjacent z_{i-1} and z_{i+1} otherwise.

If $\beta_1 = 0$, the relevant equations are:

$$\begin{aligned} dY &= \beta_0 dt + \sigma_p dW \\ d(Y - \beta_0 t) &= \sigma_p dW \\ y_i - \epsilon_i - \beta_0 t_i - y_{i-1} + \epsilon_{i-1} + \beta_0 t_{i-1} &= \delta_i \\ z_i = y_i - y_{i-1} - \beta_0 \Delta_i &= \epsilon_i - \epsilon_{i-1} + \delta_i \\ \text{Var}[z_1] &= \sigma_m^2 + \sigma_0^2 + \sigma_p^2 \Delta_1 \\ \text{Var}[z_i] &= 2\sigma_m^2 + \sigma_p^2 \Delta_i ; \quad i = 2, \dots, n \\ \text{Cov}[z_i, z_{i-1}] &= -\sigma_m^2 ; \quad i = 2, \dots, n \end{aligned}$$

3.2. Computational details

To apply the methods of Section 2, it is necessary to express the covariance matrix as the product of an unknown parameter σ^2 and a matrix \mathbf{C} , see eq. (12). One could take $\sigma^2 = \sigma_m^2$ or $\sigma^2 = \sigma_p^2$, but this would fail if $\sigma^2 = 0$. A better option is

$$\sigma^2 = \sigma_m^2 + \sigma_p^2 . \quad (28)$$

Then, with the re-parametrization

$$\sigma_m^2 = \sigma^2 \eta, \quad \sigma_p^2 = \sigma^2(1 - \eta), \quad \sigma_0^2 = \sigma^2 \eta_0, \quad (29)$$

if $\beta_1 \neq 0$ the non-zero elements of \mathbf{C} are

$$c_{11} = e^{2\beta_1 \Delta_1} \eta_0 + \eta + (1 - \eta) \frac{e^{2\beta_1 \Delta_1} - 1}{2\beta_1} \quad (30a)$$

$$c_{ii} = e^{2\beta_1 \Delta_i} \eta + \eta + (1 - \eta) \frac{e^{2\beta_1 \Delta_i} - 1}{2\beta_1}; \quad i = 2, \dots, n \quad (30b)$$

$$c_{i,i-1} = c_{i-1,i} = -e^{\beta_1 \Delta_i} \eta; \quad i = 2, \dots, n. \quad (30c)$$

If $\beta_1 = 0$,

$$c_{11} = \eta + \eta_0 + (1 - \eta) \Delta_1 \quad (31a)$$

$$c_{ii} = 2\eta + (1 - \eta) \Delta_i, \quad i = 2, \dots, n \quad (31b)$$

$$c_{i,i-1} = c_{i-1,i} = -\eta, \quad i = 2, \dots, n. \quad (31c)$$

The ML method is invariant under re-parametrization (Zacks 1971, Sec. 5.1). In the optimization one needs to ensure that $0 \leq \eta \leq 1$.

It remains to find \mathbf{L} such that $\mathbf{C} = \mathbf{L}\mathbf{L}'$, and then $\ln J = \sum \ln |\partial \varphi_i / \partial x_i| - \sum \ln l_{ii}$, and $\mathbf{u} = \mathbf{L}^{-1} \mathbf{z} / J^{1/n}$ (see equations (13)–(16)). In *R* this could be done using the *base* functions `chol` and `baksolve`, or the corresponding sparse-matrix functions from package *Matrix*. We show instead a generic method that makes efficient use of the structure of \mathbf{C} .

Because \mathbf{C} is tri-diagonal, \mathbf{L} has non-zero elements only on the diagonal and sub-diagonal, and $\mathbf{C} = \mathbf{L}\mathbf{L}'$ gives

$$\begin{aligned} c_{11} &= l_{11}^2 \\ c_{ii} &= l_{i,i-1}^2 + l_{ii}^2; \quad i = 2, \dots, n \\ c_{i,i-1} &= l_{i,i-1} l_{i-1,i-1}; \quad i = 2, \dots, n. \end{aligned}$$

Therefore, \mathbf{L} can be calculated sequentially from

$$\begin{aligned} l_{11} &= \sqrt{c_{11}} \\ \text{for } i &= 2, \dots, n \\ l_{i,i-1} &= c_{i,i-1} / l_{i-1,i-1} \\ l_{ii} &= \sqrt{c_{ii} - l_{i,i-1}^2}. \end{aligned}$$

Similarly, from $\mathbf{L}^{-1} \mathbf{z} = \mathbf{v}$ and $\mathbf{z} = \mathbf{L}\mathbf{v}$,

$$\begin{aligned} z_1 &= l_{11} v_1 \\ z_i &= l_{i,i-1} v_{i-1} + l_{ii} v_i; \quad i = 2, \dots, n, \end{aligned}$$

which can be solved as

$$\begin{aligned} v_1 &= z_1 / l_{11} \\ \text{for } i &= 2, \dots, n \\ v_i &= (z_i - l_{i,i-1} v_{i-1}) / l_{ii}. \end{aligned}$$

The two loops can be combined into one. The R function `logdet.and.v` in Appendix B uses this to compute $\ln|\mathbf{L}|$ and $\mathbf{v} = \mathbf{L}^{-1}\mathbf{z}$.

Now the vector \mathbf{u} can be computed with function `uvector` from Appendix B, given user-supplied functions for the transformation $x \rightarrow y$ and its derivative. The ML parameter estimates are found by minimizing the sum of squares $\sum u_i^2$. The same function `uvector` may be used to calculate estimates for σ_p, σ_m and σ_0 and the maximized log-likelihood value.

If it is known that there are no observation errors ($\sigma_m = 0$), the z_i are uncorrelated and the calculations can be greatly simplified.

MAP estimates can be produced by adjoining a prior distribution as shown in Section 2.6. A prior on η such as a Beta distribution might be useful to force it away from the extremes $\eta = 0$ and $\eta = 1$.

3.2.1. Example 2

Pinheiro and Bates (2000, Ch. 6, Example 1) present data consisting of 6 height-age observations in each of 14 loblolly pine trees. This `Loblolly` dataset is included in the standard R data collection (R Development Core Team 2009). In this example we use only the first tree, tree #301. The height $h_i = H(t_i)$ is in feet, and the age or time t is in years.

A suitable model is

$$\frac{dH^c}{dt} = b(a^c - H^c) .$$

Integration gives the Bertalanffy-Richards growth curve

$$H = a \{1 - [1 - (h_0/a)^c] \exp[-b(t - t_0)]\}^{1/c}$$

that, as explained above, for small c approximates the Gompertz curve used by Pinheiro and Bates (2000), and for $c = -1$ coincides with the logistic from the original study. The parameter a is the curve upper asymptote, b is a time scale factor, and c determines the location of the inflection point. Assume $H(0) = 0$.

Now introduce additive process noise,

$$dH^c = b(a^c - H^c) dt + \sigma_p dW , \quad (32)$$

and measurement error,

$$h_i^c = H(t_i)^c + \epsilon_i , \quad i = 1, \dots, 6 ,$$

with the ϵ_i independent normally distributed with mean 0 and variance σ_m^2 .

The calculations are shown in detail in Appendix C. To facilitate convergence, initial estimates were first obtained fixing $\eta = 0.5$ (η is the relative measurement variance from eqns. (28)–(29)), minimizing the sum of squares of \mathbf{u} using `uvector` and `nls`. Then, the sum of squares was minimized with η constrained to be between 0 and 1. The result was $a = 72.55$, $b = 0.0967$, $c = 0.5024$, $\eta = 1.000$, $\sigma_p = 0$, $\sigma_m = 0.04865$, with a maximized log-likelihood value of -3.988.

The result is on the boundary $\eta = 1$, suggesting that most of the variability arises from measurement errors, with an indication of ill-conditioning and little effect of η on the other parameter estimates. It has been found that with small datasets like this one it is often difficult to discriminate between environmental and observational sources of error (García 1983).

As an alternative to the additive noise, consider multiplicative process noise:

$$dH^c = b(a^c - H^c)(dt + \sigma_p dW) = b(a^c - H^c) dt + b\sigma_p(a^c - H^c) dW .$$

The Lamperti transform suggests

$$Y = \varphi(H, \boldsymbol{\theta}) = \ln|a^c - H^c|.$$

Under the Stratanovich interpretation of the SDE the ordinary calculus rules hold, and

$$dY = -b dt + b\sigma_p dW$$

(ignoring an inconsequential change of sign in W). The measurement errors are assumed to be independent, of the form

$$y_i = Y(t_i) + \epsilon_i, \quad \epsilon_i \sim N(0, \sigma_m^2).$$

There seem to be good reasons to prefer Stratanovich to Ito in this instance (Kloeden and Platen 1992, Section 6.1). For the more common Ito interpretation, it is found from Ito's formula (Gardiner 1985, eq. (4.3.14)) that the equations differ only by a change of parameters: $dY = -(b + \frac{1}{2}b^2\sigma_p^2)dt + b\sigma_p dW$. We keep the previous formulation.

The ML estimation for the multiplicative noise version gave $a = 77.11$, $b = 0.08405$, $c = 0.54946$, $\eta = 1.000$, $\sigma_p = 0$, $\sigma_m = 0.01577$, with a maximized log-likelihood value of -3.568 (details in Appendix C). The log-likelihood is not significantly different from the one for the additive model, and on a graph the two curves are quite close (Figure 1).

Multiplicative models for most classical growth curves can be derived using Box-Cox transformations, as special cases of either $Y = \ln[-(H/a)^{(c_1)}]^{(c_2)}$ or $Y = \ln[-(1 - H/a)^{(c_1)}]^{(c_2)}$ (García 2008). The Bertalanffy-Richards family corresponds to $c_2 = 0$

3.3. Hierarchical models

Often the data consists of several measurements on each of a number of individuals or sampling or experimental units (*units*, for short). For instance, the height measurements at various ages on each of 14 trees in the `Loblolly` dataset of Example 2. This is known as panel, repeated measures, or longitudinal data, and gives rise to hierarchical or multilevel models; the case of two hierarchical levels is discussed here. Some parameters may vary among units (*local*), while others are common to all units (*global*), possibly after a re-parametrization of the original model. Local parameters may be treated simply as fixed unknown values, called nuisance parameters if they are not the main object of interest (García 1983). More commonly, in mixed-effects models, the locals are viewed as “random”, possibly reflecting sampling from some hypothetical super-population (Pinheiro and Bates 2000, Snijders 2003). Mixed-effects SDE examples include Donnet and Samson (2008), Klim et al (2009), Picchini et al (2010), Driver et al (2017), and Whitaker et al (2017).

In the derivations of Section 2, let us substitute a double subscript ij for i , indicating observation $i = 1, \dots, n_j$ in unit $j = 1, \dots, m$. Then, the observations become $\mathbf{y} = (y_{11}, \dots, y_{n_m, m})$, or $\mathbf{y} = (\mathbf{y}_1, \dots, \mathbf{y}_m)$, where \mathbf{y}_j is the vector of n_j observations y_{ij} in unit j . With these notational changes, it is seen that the arguments and results of Section 2 are still valid, provided that σ^2 is common to all units (global). However, some computations can be simplified:

It is commonly assumed that the units are statistically independent. Assume also that the transformations involve only observations in the same unit, that is, $z_{ij} = \varphi_{ij}(\mathbf{y}_j, \boldsymbol{\theta}_j)$. The parameters $\boldsymbol{\theta}_j$ can vary from unit to unit. Then, assuming that σ^2 is common to all units, eq. (16) becomes

$$\mathbf{u} = \frac{(\mathbf{L}_1^{-1}\boldsymbol{\epsilon}_1(\mathbf{y}_1, \boldsymbol{\theta}_1), \dots, \mathbf{L}_m^{-1}\boldsymbol{\epsilon}_m(\mathbf{y}_m, \boldsymbol{\theta}_m))}{J^{1/n}} \quad (33)$$




Figure 1: Data and fitted models for Example 2. Continuous curve: additive process noise. Dashed curve: multiplicative process noise.

where J in eq. (14) can be obtained from

$$J = \text{abs} \left(\left| \frac{\partial \mathbf{L}^{-1} \boldsymbol{\varphi}}{\partial \mathbf{y}} \right| \right) = \text{abs} \left(\prod_{j=1}^m \left| \frac{\partial \mathbf{L}_j^{-1} \boldsymbol{\varphi}_j}{\partial \mathbf{y}_j} \right| \right) .$$

Applying this to the SDEs, one can compute $\mathbf{v}_j = \mathbf{L}_j^{-1} \mathbf{z}_j$ and $\ln J_j$ for each unit as before (Section 3.2), and obtain $\mathbf{u} = \mathbf{v}/J^{1/n}$ with $\mathbf{v} = (\mathbf{v}_1, \dots, \mathbf{v}_m)$ and $\ln J = \sum \ln J_j$.

To ensure a global σ^2 under realistic circumstances, assume that σ_m can be written as the product of a possibly local multiplier $\mu_m(\boldsymbol{\theta}_j) = \mu_{mj}$ and a global parameter σ_M , and similarly for σ_0 and σ_p :

$$\sigma_m = \mu_{mj} \sigma_M , \quad \sigma_0 = \mu_{0j} \sigma_Z , \quad \sigma_p = \mu_{pj} \sigma_P . \quad (34)$$

Then, define

$$\sigma^2 = \sigma_M^2 + \sigma_P^2 , \quad \eta = \sigma_M^2 / \sigma^2 , \quad \eta_0 = \sigma_Z^2 / \sigma^2 , \quad (35)$$

so that in eqs. (27) we substitute

$$\sigma_m^2 = \sigma^2 \mu_{mj}^2 \eta , \quad \sigma_0^2 = \sigma^2 \mu_{0j}^2 \eta_0 , \quad \sigma_p^2 = \sigma^2 \mu_{pj}^2 (1 - \eta)$$

to obtain the analogous of eq. (30) for unit j

$$c_{11j} = e^{2\beta_{1j}\Delta_{1j}} \mu_{0j}^2 \eta_0 + \mu_{mj}^2 \eta + (1 - \eta) \mu_{pj}^2 \frac{e^{2\beta_{1j}\Delta_{1j}} - 1}{2\beta_{1j}} \quad (36a)$$

$$c_{ii j} = e^{2\beta_{1j}\Delta_{ij}} \mu_{mj}^2 \eta + \mu_{mj}^2 \eta + (1 - \eta) \mu_{pj}^2 \frac{e^{2\beta_{1j}\Delta_{ij}} - 1}{2\beta_{1j}} , \quad i = 2, \dots, n \quad (36b)$$

$$c_{i,i-1,j} = c_{i-1,i,j} = -e^{\beta_{1j}\Delta_{ij}} \mu_{mj}^2 \eta , \quad i = 2, \dots, n \quad (36c)$$

with $\Delta_{ij} = t_{ij} - t_{i-1,j}$.

In the case $\beta_{1j} = 0$, eqs. (31) become

$$c_{11j} = \mu_{mj}^2 \eta + \mu_{0j}^2 \eta_0 + (1 - \eta) \mu_{pj}^2 \Delta_{1j} \quad (37a)$$

$$c_{ii j} = 2\mu_{mj}^2 \eta + (1 - \eta) \mu_{pj}^2 \Delta_{ij} , \quad i = 2, \dots, n \quad (37b)$$

$$c_{i,i-1,j} = c_{i-1,i,j} = -\mu_{mj}^2 \eta , \quad i = 2, \dots, n . \quad (37c)$$

An extension of the function `uvector` to handle hierarchical models is listed in Appendix D.

3.3.1. Fixed local parameters

Viewing both global and local parameters as unknown constants, estimation is relatively simple in R , where function `nls` allows vector-valued parameters. This feature seems to be documented only in the last example of the `nls` help page, which appears also on page 219 of Venables and Ripley (2002).

The method performs reasonably well for at least several hundred parameters. Presumably, `nls` exploits special structure through matrix partitioning strategies, as in Griewank and Toint (1982), García (1983), or Soo and Bates (1992). It was verified that the approach gives exactly the same results as the method of García (1983), which is based on direct optimization of the log-likelihood (software and documentation available at <http://forestgrowth.unbc.ca/sde>). Computation times are much longer than those of García's compiled Fortran code, but that is not an impediment with modern computing hardware.

Example 3. Consider fitting Richards SDE models simultaneously to all the 14 trees of the Loblolly dataset of Example 2. The parametrization can be important here, so re-write eq. (32) using the Box-Cox transformation of eq. (1), ensuring that a and b are proper scale parameters:

$$Y = (H/a)^{(c)}$$

$$dY = -Y d(bt) + \sigma_P dW(bt) = -bY dt + \sqrt{|b|}\sigma_P dW(t) .$$

Assume that the measurement error is negligible compared to the process noise.

See Appendix E for details of the calculations.

First, take a as local, i. e., the asymptotes a_j vary from tree to tree. The ML estimates are found to be:

$$a_1 = 68.37, \dots, a_{14} = 78.84, b = 0.09472, c = 0.4918, \sigma_P = 0.03359 ,$$

with a log-likelihood of -88.4. For comparing models with different numbers of parameters one can use Akaike's Information Criterion $AIC = 210.8$, or Schwarz's Bayesian Information Criterion $BIC = 252.1$. These indices penalize twice the negative log-likelihood by subtracting a quantity that increases with the number of estimated parameters.

Trying with a global and b local gives

$$a = 73.08, b_1 = 0.08912, \dots, b_{14} = 0.1031, c = 0.4916, \sigma_P = 0.03231 ,$$

with log-likelihood = -85.2, $AIC = 204.3$, $BIC = 245.6$. The log-likelihood (or equivalently, the AIC or BIC) indicates that this model fits the data slightly better than the one with a local.

Taking both a and b as locals gave worse AIC and BIC values than those for the one-local versions. Other structures could be defined by re-parametrization, substituting functions of other global and local parameters for a , b and c (García 1983).

3.3.2. Random local parameters

A more popular alternative is to view local parameters as, in some sense, random. This is known as a *nonlinear mixed-effects model* (Pinheiro and Bates 2000). The easiest situation to understand and justify is where the units are thought to be a random sample from a population where the parameters follow certain frequency function. We adopt the usual assumption of Gaussian local parameter distributions. In mixed-effects modelling terminology, the units are sometimes called *groups* (of observations), the global parameters and the means of the locals are *fixed effects*, and the deviations of the local parameters from their means are *random effects*.

With this approach the number of unknown parameters is much reduced, instead of one local parameter value for each unit it is only necessary to estimate a mean and a variance (and possibly covariances among locals). On the other hand, there are additional model assumptions, and the estimation algorithms are more complicated. The model may be questionable when, as is often the case, rather than being a simple random sample the units are chosen to have a more efficient coverage of a range of conditions.

A mixed-effects SDE model can be estimated by maximum likelihood using standard packages and the u_{ij} computations above. In *R* one can use function `nlme` from the package of the same name, or `nlmer` from package *lme4*. These functions do not support bounds on parameter values, so that in SDEs with observation errors it would be necessary to do multiple runs with given values of η , or alternatively, to embed the mixed-effects estimation within a one-dimensional optimization over η (e. g., using `optimize`).

Example 4.. The b -local version from Example 3 was fitted as a mixed-effects model using the `nlme` *R* package (Appendix F). It was necessary to relax the default tolerance in order to achieve convergence. The ML estimates for the global parameters were $a = 73.43$ and $c = 0.4938$. In this formulation b is a random variable, so that it does not make sense to speak of “estimates”, but the software provides “predictions” $b_1 = 0.08973, \dots, b_{14} = 0.09964$, and a mean 0.09381.

The log-likelihood was -101.78 , which is not comparable to that in Example 3 because there are less free parameters in the mixed-effects approach. The AIC and BIC should be more informative. The AIC = 213.6 compared to the AIC = 204.3 in Example 3 suggests that the fixed-parameters model is better than the mixed-effects version. However, the opposite is true for the BIC that penalizes the number of parameters more heavily, 225.7 *vs.* 245.6.

3.4. Multivariate SDEs

Some applications use a system of SDEs or, equivalently, a vector-valued SDE (e. g., García 1984, Klim et al 2009, Picchini and Ditlevsen 2011, Donnet and Samson 2013, Driver et al 2017). Analogously to the single-variable case, we assume that there is a random p -dimensional vector \mathbf{Y} that follows a linear SDE

$$d\mathbf{Y} = (\mathbf{A}\mathbf{Y} + \mathbf{b}) dt + \mathbf{S} d\mathbf{W} . \quad (38)$$

Here \mathbf{A} and \mathbf{S} are $p \times p$ matrices, and \mathbf{b} is a p -vector, any or all of them dependent on parameters in $\boldsymbol{\theta}$, and \mathbf{W} is a vector of p independent Wiener processes.

The observations \mathbf{x}_i at time t_i are related to \mathbf{Y} by a one-to-one transformation

$$\mathbf{y}_i = \boldsymbol{\varphi}(\mathbf{x}_i, \boldsymbol{\theta}) , \quad (39)$$

and may be subject to observation error according to

$$\mathbf{y}_i = \mathbf{Y}(t_i) + \boldsymbol{\epsilon}_i , \quad \boldsymbol{\epsilon}_i \sim N(\mathbf{0}, V) . \quad (40)$$

The errors $\boldsymbol{\epsilon}_i$ are uncorrelated across observation times, and independent of the process noise \mathbf{W} . The initial conditions are $\mathbf{Y}(t_0) = \mathbf{y}_0 + \boldsymbol{\epsilon}_0$, with $\boldsymbol{\epsilon}_0 \sim N(\mathbf{0}, \mathbf{V}_0)$. All this can apply to units with different parameter values $\boldsymbol{\theta}_j$ and expected initial conditions \mathbf{y}_{0j} ; non-essential indices will be omitted.

A useful example of transformation $\mathbf{y} = \boldsymbol{\varphi}(\mathbf{x}, \boldsymbol{\theta})$ is

$$y_k = x_1^{c_{k1}} x_2^{c_{k2}} \dots x_p^{c_{kp}} , \quad k = 1, \dots, p ,$$

which can be conveniently denoted as $\mathbf{y} = \mathbf{x}^{\mathbf{C}}$, defining $\mathbf{x}^{\mathbf{C}} \equiv \exp(\mathbf{C} \ln \mathbf{x})$. This gives a flexible multivariate analog of the Bertalanffy-Richards SDE (García 1984).

In the time-invariant case the matrix exponential $e^{-\mathbf{A}t}$ is an integrating factor (see Moler and Van Loan 2003, or the *expm* package in *R* for definitions and computation of the matrix exponential). For non-singular \mathbf{A} ,

$$d[e^{-\mathbf{A}t} (\mathbf{Y} + \mathbf{A}^{-1}\mathbf{b})] = e^{-\mathbf{A}t} \mathbf{S} d\mathbf{W}$$

(c. f. eq. (23)). Then,

$$e^{-\mathbf{A}t_i} [\mathbf{Y}(t_i) + \mathbf{A}^{-1}\mathbf{b}] - e^{-\mathbf{A}t_{i-1}} [\mathbf{Y}(t_{i-1}) + \mathbf{A}^{-1}\mathbf{b}] = \int_{t_{i-1}}^{t_i} e^{-\mathbf{A}t} \mathbf{S} d\mathbf{W}(t) .$$

Multiplying by $e^{\mathbf{A}t_i}$ and substituting $\mathbf{Y}(t_i) = \mathbf{y}_i - \boldsymbol{\epsilon}_i$,

$$\mathbf{z}_i \equiv \mathbf{y}_i + \mathbf{A}^{-1}\mathbf{b} - e^{\mathbf{A}\Delta_i}(\mathbf{y}_{i-1} + \mathbf{A}^{-1}\mathbf{b}) = \boldsymbol{\epsilon}_i - e^{\mathbf{A}\Delta_i}\boldsymbol{\epsilon}_{i-1} + \boldsymbol{\delta}_i \quad (41)$$

with $t_i - t_{i-1} \equiv \Delta_i$ and

$$\begin{aligned} \boldsymbol{\delta}_i &= \int_{t_{i-1}}^{t_i} e^{\mathbf{A}(t_i-t)} \mathbf{S} d\mathbf{W}(t) , \\ \mathbb{E}[\boldsymbol{\delta}_i] &= 0 , \\ \text{Var}[\boldsymbol{\delta}_i] &= \int_{t_{i-1}}^{t_i} e^{\mathbf{A}(t_i-t)} \mathbf{S} \mathbf{S}' e^{\mathbf{A}(t_i-t)'} dt = \int_0^{\Delta_i} e^{\mathbf{A}t} \mathbf{S} \mathbf{S}' e^{\mathbf{A}t'} dt \end{aligned} \quad (42)$$

similarly to eq. (19) (Gardiner 1985, Sec. 4.4.6, and Sec. 4.4.9 for the time-variant case).

Therefore, eq. (41) defines a transformation $\{\mathbf{y}_i\} \rightarrow \{\mathbf{z}_i\}$ with unit Jacobian, such that the \mathbf{z}_i (or \mathbf{z}_{ij}) are Gaussian with

$$\mathbb{E}[\mathbf{z}_i] = \mathbf{0} ; \quad i = 1, \dots, n \quad (43a)$$

$$\text{Var}[\mathbf{z}_1] = \mathbf{V} + e^{\mathbf{A}\Delta_0} \mathbf{V}_0 e^{\mathbf{A}\Delta_0'} + \text{Var}[\boldsymbol{\delta}_1] \quad (43b)$$

$$\text{Var}[\mathbf{z}_i] = \mathbf{V} + e^{\mathbf{A}\Delta_i} \mathbf{V} e^{\mathbf{A}\Delta_i'} + \text{Var}[\boldsymbol{\delta}_i] ; \quad i = 2, \dots, n \quad (43c)$$

$$\text{Cov}[\mathbf{z}_i, \mathbf{z}_{i-1}] = -e^{\mathbf{A}\Delta_i} \mathbf{V} ; \quad i = 2, \dots, n . \quad (43d)$$

García (1984) obtained explicit likelihoods for systems where oscillations are not admissible, for which $\mathbf{A} = \mathbf{P}^{-1}\boldsymbol{\Lambda}\mathbf{P}$, $e^{\mathbf{A}\Delta_i} = \mathbf{P}^{-1}e^{\boldsymbol{\Lambda}\Delta_i}\mathbf{P}$, and $\boldsymbol{\Lambda}$ is a diagonal matrix of real eigenvalues. That includes dynamical systems with monotonic state variables, for instance, forest plantations where diameter and height increases and number of trees decreases over time. There it is convenient to take the elements of $\boldsymbol{\Lambda}$ and \mathbf{P} instead of those of \mathbf{A} as parameters to be estimated.

For general \mathbf{A} , it is found through integration by parts of eq. (42) that:

Proposition 2. $\text{Var}[\boldsymbol{\delta}_i]$ satisfies

$$\mathbf{A} \text{Var}[\boldsymbol{\delta}_i] + \text{Var}[\boldsymbol{\delta}_i] \mathbf{A}' = e^{\mathbf{A}\Delta_i} \mathbf{S} \mathbf{S}' e^{\mathbf{A}\Delta_i'} - \mathbf{S} \mathbf{S}' . \quad (44)$$

Proof. Making use of the fact that \mathbf{A} and $e^{\mathbf{A}t}$ commute,

$$\begin{aligned} d e^{\mathbf{A}t} \mathbf{S} \mathbf{S}' e^{\mathbf{A}t'} &= d (e^{\mathbf{A}t} \mathbf{S}) (e^{\mathbf{A}t'} \mathbf{S}')' \\ &= (\mathbf{A} e^{\mathbf{A}t} \mathbf{S} dt) (e^{\mathbf{A}t'} \mathbf{S}')' + (e^{\mathbf{A}t} \mathbf{S}) (\mathbf{A} e^{\mathbf{A}t'} \mathbf{S}' dt)' \\ &= \mathbf{A} (e^{\mathbf{A}t} \mathbf{S} \mathbf{S}' e^{\mathbf{A}t'} dt) + (e^{\mathbf{A}t} \mathbf{S} \mathbf{S}' e^{\mathbf{A}t'} dt) \mathbf{A}' . \end{aligned}$$

Integrating between 0 and Δ_i and using eq. (42),

$$e^{\mathbf{A}\Delta_i} \mathbf{S} \mathbf{S}' e^{\mathbf{A}\Delta_i'} - \mathbf{S} \mathbf{S}' = \mathbf{A} \text{Var}[\boldsymbol{\delta}_i] + \text{Var}[\boldsymbol{\delta}_i] \mathbf{A}' .$$

□

Eq. (44) is known as a *continuous-time Lyapunov equation*. A recent review of solution methods is contained in Simoncini (2016, Sec. 5).

The likelihood for the linear SDE can also be derived less directly by writing it as a product of the conditional densities of \mathbf{y}_i given the previous $\mathbf{y}_0, \dots, \mathbf{y}_{i-1}$. The conditional expectations and variance-covariance matrices can be computed recursively using the Kalman filter (e. g., Donnet and Samson 2013, Sec. 4.2.1).

Regardless of how the likelihood is produced, the Furnival-Box-Cox device following the methods of Sec. 3 can then be used to reduce the likelihood maximization to minimizing a sum of squares. Multidimensionality does not introduce additional conceptual issues, although the details and notation can be complicated. MAP estimation as in Section 2.6 is also feasible.

Taking advantage of the ML re-parametrization invariance, it is preferable to estimate Cholesky factors instead of optimizing directly over variance-covariance matrices like \mathbf{V} or \mathbf{V}_0 (e. g., García 1984). This ensures that the matrix will be positive-definite, as it should. The case of \mathbf{S} in eq. (38) is more delicate, because the likelihood function depends only on the product $\mathbf{S}\mathbf{S}'$, or on the Cholesky factor of $\mathbf{S}\mathbf{S}'$. This means that a full square matrix \mathbf{S} is unidentifiable, that is, different \mathbf{S} 's produce the same observation distribution. Therefore, a triangular matrix should replace \mathbf{S} in the optimization. Driver et al (2017) directly assume that \mathbf{S} in eq. (38) is triangular; it is not always clear how the issue is handled in other software.

4. Conclusions

Furnival (1961) and Box and Cox (1964) showed how maximum likelihood estimation involving nonlinear transformations of Gaussian dependent variables can be achieved by minimizing a sum of squares. The technique has many potential applications, the focus here is on stochastic differential equation modelling. A similar idea can be applied to Bayesian maximum a posteriori estimates.

SDEs are more realistic than traditional approaches in longitudinal data analysis and other applications, but wider acceptance has been hampered by mathematical and computational complexity. Linear SDEs allow explicit solutions, and are suitable for compartmental and other models (Cleir 2000, Kristensen et al 2005, Klim et al 2009, Favetto and Samson 2010, Cuenod et al 2011, Driver et al 2017, Seber and Wild 2003, Sec. 8.8). However, they may not be appropriate in more general situations. Much more flexibility is achieved by allowing nonlinear transformations of variables in linear SDEs. The resulting so-called reducible SDEs are mathematically tractable, and extensions of the Furnival-Box-Cox method make parameter estimation possible using standard nonlinear least-squares or mixed-effects software.

Mixed-effects modelling has become a *de facto* requirement for publication in many scientific journals, but the underlying assumptions are not always justifiable. Often, a model with group fixed-effects may be preferable. A somewhat obscure feature of the `nls` function in *R* makes the implementation of such fixed parameters models particularly simple.

The proposed methods have been discussed and demonstrated in detail for univariate SDEs. Results match those computed by direct maximization of the likelihood functions. Extensions to the multivariate case are outlined more briefly. There are no conceptual difficulties in these, although the best choice of implementation specifics is not obvious, and likely to be problem-dependent.

Being able to use well tested and familiar standard software should make SDE modelling more transparent and accessible. Within the class of reducible SDEs, the methodology can also provide more flexibility in model specification compared to special-purpose software packages. Ease of use and understanding may contribute to a wider adoption of these models.

Acknowledgements

I am grateful to an anonymous reviewer and to Dr. Lance Broad for thoughtful and detailed comments that contributed to improve the text.

Appendices — *R* computer code

Appendix A. Example 1, Section 2.4

```
> bc <- function(y, lambda){ # Box-Cox transformation
+   if (abs(lambda) < 1e-9) log(y)
+   else (y^lambda - 1) / lambda
+ }
> x <- MASS::GAGurine$Age # data
> y <- MASS::GAGurine$GAG
> lm(log(y) ~ I(x - 1)) # find suitable starting points
> # (with lambda.y = 0, lambda.x = 1)

Call:
lm(formula = log(y) ~ I(x - 1))

Coefficients:
(Intercept)      I(x - 1)
      2.8522      -0.1139

> # Minimize sum of squares:
> gm <- exp(mean(log(y))) # geometric mean
> nls(rep(0, length(y)) ~ (bc(y, lambda.y) - beta0 - beta1 *
+ bc(x, lambda.x)) / gm^(lambda.y - 1), start = c(beta0=2.9,
+ beta1=-0.11, lambda.x=1, lambda.y=0))
Nonlinear regression model
  model: rep(0, length(y)) ~ (bc(y, lambda.y) - beta0 - beta1 * bc(x,
    lambda.x))/gm^(lambda.y - 1)
  data: parent.frame()
    beta0    beta1 lambda.x lambda.y
 3.3142  -0.3502   0.4249   0.1032
residual sum-of-squares: 3214

Number of iterations to convergence: 10
Achieved convergence tolerance: 1.808e-06
> # Estimate for sigma:
> n <- length(y)
> gm^(0.1032 - 1) * sqrt(3214 / n)
[1] 0.383853
>
> # Round the lambdas for a more parsimonious model:
> lm(log(y) ~ sqrt(x))

Call:
lm(formula = log(y) ~ sqrt(x))
```

Coefficients:

(Intercept)	sqrt(x)
3.3665	-0.5069

The results are verified by direct minimization of the negative log-likelihood from eq. (8),

$$-\ln L = \frac{n}{2} \ln(2\pi\sigma^2) + \frac{\sum \epsilon_i^2}{2\sigma^2} - \ln J :$$

```
> neglogL <- function(theta){# theta = (beta0,beta1,lambda.x,lambda.y,sigma)
+ (n/2) * log(2 * pi * theta[5]^2) + sum((bc(y, theta[4]) - theta[1] -
+ theta[2] * bc(x, theta[3]))^2) / (2*theta[5]^2) - n * (theta[4] - 1) *
+ log(gm)}
> optim(par = c(3, -0.4, 0.4, 0.1, 0.4), fn = neglogL)
$par
[1] 3.3140471 -0.3502047 0.4248816 0.1031103 0.3837414

$value
[1] 810.6863

$counts
function gradient
      422      NA

$convergence
[1] 0

$message
NULL
```

Appendix B. Functions for ML parameter estimation in SDE models, Section 3.2

The following function computes $\ln|\mathbf{L}|$ and $\mathbf{v} = \mathbf{L}^{-1}\mathbf{z}$.

```
logdet.and.v <- function(
# Calculate log(det(L)) and v = L^-1 z
  cdiag, # vector with the diagonal elements of C, c[i, i]
  csub, # vector with sub-diagonal c[i, i-1] for i > 1
  z)
{
  v <- z
  ldiag <- sqrt(cdiag[1])
  logdet <- log(ldiag)
  v[1] <- z[1] / ldiag
  for (i in 2:length(z)) {
    lsub <- csub[i] / ldiag
    ldiag <- sqrt(cdiag[i] - lsub^2)
    logdet <- logdet + log(ldiag)
    v[i] <- (z[i] - lsub * v[i - 1]) / ldiag
  }
  list(logdet = logdet, v = v)
}
```

The function `uvector` that follows computes the vector \mathbf{u} , given `logdet.and.v` and user-supplied functions `phi(x, theta)` for the transformation $x \rightarrow y$ and `phiprime(x, theta)` for its derivative. Once $\sum u_i$ is minimized, `uvector` can be used to calculate the ML estimates for σ_p, σ_m and σ_0 , and the maximized log-likelihood value, if `final = TRUE` and the optimal parameter values are passed in the arguments.

```
uvector <- function(
  # If final = FALSE (default), returns vector whose sum of squares
  #   should be minimized.
  # If final = TRUE, passing the ML parameter estimates returns a
  #   list with the sigma estimates and the maximized log-likelihood.
  # Requires a transformation function y = phi(x, theta), and a
  #   function phiprime(x, theta) for the derivative dy/dx, where
  #   theta is a list containing all the parameters.
  x, t, # data vectors
  beta0, beta1, eta, eta0, x0, t0, # SDE parameters
  lambda, # named list with additional parameter(s) for phi
  final = FALSE)
{
  theta <- c(list(beta0=beta0, beta1=beta1, eta=eta, eta0=eta0,
                 x0=x0, t0=t0), lambda)
  s <- order(t); t <- t[s]; x <- x[s] # ensure increasing t
  n <- length(x)
  y <- phi(x, theta)
  y0 <- phi(x0, theta)
  Dt <- diff(c(t0, t))
  if (beta1 != 0) {
    ex <- exp(beta1 * Dt)
    ex2 <- ex^2
    z <- y + beta0 / beta1 - ex * (c(y0, y[-n]) + beta0 / beta1)
    cdiag <- ex2 * eta + eta + (1 - eta) * (ex2 - 1) / (2 * beta1)
    cdiag[1] <- cdiag[1] - ex2[1] * (eta - eta0)
    csub <- -ex * eta
  } else { # beta1 == 0
    z <- y - c(y0, y[-n]) - beta0 * Dt
    cdiag <- 2 * eta + (1 - eta) * Dt
    cdiag[1] <- cdiag[1] - eta + eta0
    csub <- -rep(eta, n)
  }
  ld.v <- logdet.and.v(cdiag, csub, z)
  logJ <- sum(log(abs(hiprime(x, theta)))) - ld.v$logdet
  Jn <- exp(logJ / n) # J^(1/n)
  u <- ld.v$v / Jn
  if (!final) return (u) # "normal" exit
  # Else, at optimum, calculate sigma.p, sigma.m and sigma.0
  # estimates, and the log-likelihood:
  ms <- sum(u^2) / n # mean square
  sigma2 <- Jn^2 * ms # estimate for sigma^2
  list(sigma.p = sqrt((1 - eta) * sigma2), sigma.m = sqrt(eta *
    sigma2), sigma.0 = sqrt(eta0 * sigma2), loglikelihood =
    -(n / 2) * (log(ms) + log(2 * pi) + 1))
}
```

Appendix C. Example 2, Section 3.2.1

Extract tree 301:

```
> (lob301 <- Loblolly[Loblolly$Seed == 301, ])
  height age Seed
1    4.51  3  301
15   10.89  5  301
29   28.72 10  301
43   41.74 15  301
57   52.70 20  301
71   60.92 25  301
```

Appendix C.1. Additive process noise

The functions for the transformation H^c and its derivative cH^{c-1} are

```
> phi <- function(H, theta) H^theta$c
> phiprime <- function(H, theta) with(theta, c * H^(c - 1))
```

To facilitate convergence, start by fixing $\eta = 0.5$ (η is the relative measurement variance from eqns. (28)–(29)). Minimizing the sum of squares of \mathbf{u} using `nls`, with an initial guess $a = 70$, $b = 0.1$, $c = 1$,

```
> nls(~ uvector(x = height, t = age, beta0 = b * a^c, beta1 = -b, eta = 0.5,
+ eta0 = 0, x0 = 0, t0 = 0, lambda = list(c = c)), data = lob301,
+ start = list(a = 70, b = 0.1, c = 1))
Nonlinear regression model
model: 0 ~ uvector(x = height, t = age, beta0 = b * a^c, beta1 = -b,
eta = 0.5, eta0 = 0, x0 = 0, t0 = 0, lambda = list(c = c))
data: lob301
      a      b      c
71.96058 0.09947 0.49217
residual sum-of-squares: 1.829
```

Number of iterations to convergence: 6

Achieved convergence tolerance: 1.145e-06

The option `algorithm = "port"` allows upper and lower bounds on the parameters, and we use it now for $0 \leq \eta \leq 1$:

```
> nls(~ uvector(x = height, t = age, beta0 = b * a^c, beta1 = -b, eta = eta,
+ eta0 = 0, x0 = 0, t0 = 0, lambda = list(c = c)), data = lob301,
+ start = list(a = 70, b = 0.1, c = 0.5, eta = 0.5), algorithm = "port",
+ lower = c(0, 0, 0, 0), upper = c(100, 1, 2, 1))
Nonlinear regression model
model: 0 ~ uvector(x = height, t = age, beta0 = b * a^c, beta1 = -b,
eta = eta, eta0 = 0, x0 = 0, t0 = 0, lambda = list(c = c))
data: lob301
      a      b      c      eta
72.5459 0.0967 0.5024 1.0000
residual sum-of-squares: 1.327
```

Algorithm "port", convergence message: relative convergence (4)

Finally, calculate the variance estimates and the maximized log-likelihood:

```
> uvector(x = lob301$height, t = lob301$age, beta0 = 0.0967 * 72.5459^0.5024,
+ beta1 = -0.0967, eta = 1, eta0 = 0, x0 = 0, t0 = 0,
+ lambda = list(c = 0.5024), final = TRUE)
$sigma.p
[1] 0

$sigma.m
[1] 0.04865072

$sigma.0
[1] 0

$logLikelihood
[1] -3.9882
```

Appendix C.2. Multiplicative process noise

```
> phi <- function(H, theta)
+   with(theta, log(abs(a^c - H^c)))
> phiprime <- function(H, theta)
+   with(theta, - c * H^(c - 1) / (a^c - H^c))
> nls(~ uvector(x = height, t = age, beta0 = -b, beta1 = 0, eta = eta,
+ eta0 = 0, x0 = 0, t0 = 0, lambda = list(a = a, c = c)), data = lob301,
+ start = list(a = 72, b = 0.1, c = 0.5, eta = 0.5), algorithm = "port",
+ lower = c(0, 0, 0, 0), upper = c(100, 1, 2, 1))
Nonlinear regression model
  model: 0 ~ uvector(x = height, t = age, beta0 = -b, beta1 = 0, eta = eta,
    eta0 = 0, x0 = 0, t0 = 0, lambda = list(a = a, c = c))
  data: lob301
         a         b         c         eta
77.10687  0.08405  0.54946  1.00000
residual sum-of-squares: 1.154
```

Algorithm "port", convergence message: relative convergence (4)

```
> uvector(x = lob301$height, t = lob301$age, beta0 = -0.08405, beta1 = 0,
+ eta = 1, eta0 = 0, x0 = 0, t0 = 0,
+ lambda = list(a = 77.10687, c = 0.54946), final = TRUE)
$sigma.p
[1] 0

$sigma.m
[1] 0.01576683

$sigma.0
[1] 0

$logLikelihood
[1] -3.568224
```


Appendix D. ML estimation in hierarchical SDE models, Section 3.3

Extending the function `uvector` of Appendix B,

```
> uvector
function(
  # If final = FALSE (default), returns vector whose sum of squares
  #   should be minimized.
  # If final = TRUE, passing the ML parameter estimates returns a
  #   list with the sigma estimates and the maximized log-likelihood.
  # Requires a transformation function y = phi(x, theta), and a
  #   function phiprime(x, theta) for the derivative dy/dx, where
  #   theta is a list containing all the parameters.
  x, t, unit = NULL, # data and unit id vectors
  beta0, beta1, eta, eta0, x0, t0, # SDE parameters. Some of these
  #   may be local, given as a vector of values for each observation
  lambda, # named list of additional parameters(s) for phi,
  #   possibly local vectors
  mum = 1, mu0 = 1, mup = 1, # optional sigma multipliers
  sorted = FALSE, # data already ordered by increasing t?
  final = FALSE)
{
  if(is.null(unit)) unit <- rep(1, length(x)) # single unit
  if(length(unique(eta)) > 1 || length(unique(eta0)) > 1)
    stop("eta and eta0 must be global")
  theta <- data.frame(unit, beta0, beta1, eta, eta0, x0, t0,
    lambda, mum, mu0, mup)[!duplicated(unit), ] # one row per unit
  v <- c()
  n <- logJ <- 0
  for(id in theta$unit){
    theta.j <- theta[match(id, theta$unit), ]
    j <- unit == id
    x.j <- x[j]
    t.j <- t[j]
    if(!sorted){ # ensure increasing t
      s <- order(t.j)
      x.j <- x.j[s]
      t.j <- t.j[s]
    }
    n.j <- length(x.j)
    y <- phi(x.j, theta.j)
    y0 <- phi(theta.j$x0, theta.j)
    Dt <- diff(c(theta.j$t0, t.j))
    muetam <- theta.j$mum^2 * theta.j$eta
    mueta0 <- theta.j$mu0^2 * theta.j$eta0
    muetap <- theta.j$mup^2 * (1 - theta.j$eta)
    if (theta.j$beta1 != 0) {
      ex <- exp(theta.j$beta1 * Dt)
      ex2 <- ex^2
      z <- y + theta.j$beta0 / theta.j$beta1 - ex *
        (c(y0, y[-n.j]) + theta.j$beta0 / theta.j$beta1)
      cdiag <- (ex2 + 1) * muetam + muetap * (ex2 - 1) /
```

```

      (2 * theta.j$beta1)
      cdiag[1] <- cdiag[1] - ex2[1] * (muetam - mueta0)
      csub <- -ex * muetam
    } else { # beta1 == 0
      z <- y - c(y0, y[-n.j]) - theta.j$beta0 * Dt
      cdiag <- 2 * muetam + muetap * Dt
      cdiag[1] <- cdiag[1] - muetam + mueta0
      csub <- rep(- muetam, n.j)
    }
    ld.v <- logdet.and.v(cdiag, csub, z)
    v <- c(v, ld.v$v)
    logJ <- logJ + sum(log(abs(phpime(x.j, theta.j)))) -
      ld.v$logdet
    n <- n + n.j
  }
  if(n != length(x)) stop("Should not happen, something wrong!")
  Jn <- exp(logJ / n) # J^(1/n)
  u <- v / Jn
  if (!final) return (u) # "normal" exit
  # Else, at optimum, calculate sigma.P, sigma.M and sigma.Z
  # estimates, and the log-likelihood:
  ms <- sum(u^2) / n # mean square
  sigma2 <- Jn^2 * ms # estimate for sigma^2
  list(sigma.P = sqrt((1 - eta) * sigma2),
       sigma.M = sqrt(eta * sigma2),
       sigma.Z = sqrt(eta0 * sigma2),
       logLikelihood = -(n / 2) * (log(ms) + log(2 * pi) + 1))
}

```

For clarity, a simple for loop has been used rather than potentially more efficient or elegant *R*-specific shortcuts. This `uvector` is backward compatible with the previous version, giving the same results for single units.

Appendix E. Example 3, Section 3.3.1

The transformation functions are

```

> phi <- function(H, theta) with(theta,
+   ifelse(abs(rep_len(c, length(H))) < 1e-9,
+     log(H / a), ((H / a)^c - 1) / c))
> # (works for vectors and scalars)
> phpime <- function(H, theta) with(theta,
+   (H / a)^(c - 1) / a)

```

First, take *a* as local. Local parameters in `nls` are indexed by a factor identifying the units, `Seed` in this case. A vector with one element per observation is passed on to the model function. On the other hand, values for the factor levels, one element per unit, are used in the starting values and output. The ML estimates are:

```

> (alocal <- nls(~uvector(x = height, t = age, unit = Seed,
+ beta0 = 0, beta1 = -b, eta = 0, eta0 = 0, x0 = 0, t0 = 0,
+ lambda = list(a = a[Seed], c = c), mup = sqrt(abs(b))),

```

```

+ data = Loblolly, start = list(a = rep(72, 14), b = 0.1, c = 0.5)))
Nonlinear regression model
  model: 0 ~ uvector(x = height, t = age, unit = Seed, beta0 = 0,
    beta1 = -b, eta = 0, eta0 = 0, x0 = 0, t0 = 0, lambda =
    list(a = a[Seed], c = c), mup = sqrt(abs(b)))
  data: Loblolly
    a1      a2      a3      a4      a5      a6      a7      a8
68.36651 69.11596 71.87593 70.69002 70.44039 71.38285 72.90628 70.92199
    a9      a10     a11     a12     a13     a14      b      c
74.01902 74.77264 75.44943 76.41765 76.91871 78.84126 0.09472 0.49182
  residual sum-of-squares: 40.35

Number of iterations to convergence: 3
Achieved convergence tolerance: 5.947e-06
> p <- unname(coef(alocal)) # parameter estimates
> uvector(x = Loblolly$height, t = Loblolly$age, unit = Loblolly$Seed,
+ beta0 = 0, beta1 = - p[15], eta = 0, eta0 = 0, x0 = 0,
+ t0 = 0, lambda = list(a = p[Loblolly$Seed], c = p[16]), mup =
+ sqrt(abs(p[15])), final = TRUE)
$sigma.P
[1] 0.03358892

$sigma.M
[1] 0

$sigma.Z
[1] 0

$logLikelihood
[1] -88.39581

> logLik(alocal) # another way
'log Lik.' -88.39581 (df=17)
> AIC(alocal) # AIC or BIC for comparing models with different
[1] 210.7916
> BIC(alocal) # numbers of parameters (lower is better)
[1] 252.1155

```

Now try a global and b local:

```

> (blocal <- nls(~uvector(x = height, t = age, unit = Seed,
+ beta0 = 0, beta1 = -b[Seed], eta = 0, eta0 = 0, x0 = 0, t0 = 0,
+ lambda = list(a = a, c = c), mup = sqrt(abs(b[Seed]))),
+ data = Loblolly, start = list(a = 72, b = rep(0.1, 14), c = 0.5)))
Nonlinear regression model
  model: 0 ~ uvector(x = height, t = age, unit = Seed, beta0 = 0,
    beta1 = -b[Seed], eta = 0, eta0 = 0, x0 = 0, t0 = 0,
    lambda = list(a = a, c = c), mup = sqrt(abs(b[Seed]))))
  data: Loblolly
    a      b1      b2      b3      b4      b5      b6      b7
73.08143 0.08912 0.09082 0.09495 0.09053 0.08915 0.09111 0.09496
    b8      b9      b10     b11     b12     b13     b14      c

```

```
0.08957 0.09680 0.09819 0.09843 0.09984 0.09984 0.10313 0.49156
residual sum-of-squares: 37.35
```

```
Number of iterations to convergence: 4
```

```
Achieved convergence tolerance: 2.256e-06
```

```
> p <- unname(coef(blocal))
```

```
> uvector(x = Loblolly$height, t = Loblolly$age, unit = Loblolly$Seed,
```

```
+ beta0 = 0, beta1 = -(p[-1])[Loblolly$Seed], eta = 0, eta0 = 0,
```

```
+ x0 = 0, t0 = 0, lambda = list(a = p[1], c = p[16]),
```

```
+ mup = sqrt(abs((p[-1])[Loblolly$Seed])), final = TRUE)
```

```
$sigma.P
```

```
[1] 0.03231109
```

```
$sigma.M
```

```
[1] 0
```

```
$sigma.Z
```

```
[1] 0
```

```
$logLikelihood
```

```
[1] -85.15201
```

```
> c(AIC(blocal), BIC(blocal))
```

```
[1] 204.3040 245.6279
```

Finally, with both a and b locals,

```
> (ablocal <- nls(~uvector(x = height, t = age, unit = Seed,
```

```
+ beta0 = 0, beta1 = - b[Seed], eta = 0, eta0 = 0, x0 = 0, t0 = 0,
```

```
+ lambda = list(a = a[Seed], c = c), mup = sqrt(abs(b[Seed]))),
```

```
+ data = Loblolly, start = list(a = rep(72, 14), b = rep(0.1, 14),
```

```
+ c = 0.5)))
```

```
Nonlinear regression model
```

```
model: 0 ~ uvector(x = height, t = age, unit = Seed, beta0 = 0,
```

```
beta1 = -b[Seed], eta = 0, eta0 = 0, x0 = 0, t0 = 0,
```

```
lambda = list(a = a[Seed], c = c), mup = sqrt(abs(b[Seed])))
```

```
data: Loblolly
```

	a1	a2	a3	a4	a5	a6	a7
	68.31203	67.27849	68.61879	73.81971	75.66951	75.05226	72.12885
	a8	a9	a10	a11	a12	a13	a14
	76.93856	72.22306	72.02841	73.72472	74.01390	75.85147	76.05177
	b1	b2	b3	b4	b5	b6	b7
	0.09497	0.09821	0.10087	0.08990	0.08659	0.08913	0.09631
	b8	b9	b10	b11	b12	b13	b14
	0.08571	0.09807	0.09972	0.09784	0.09888	0.09668	0.09960

```
c
```

```
0.49060
```

```
residual sum-of-squares: 30.67
```

```

Number of iterations to convergence: 4
Achieved convergence tolerance: 6.068e-07
> logLik(ablocal)
'log Lik.' -76.87568 (df=30)
> c(AIC(ablocal), BIC(ablocal))
[1] 213.7514 286.6759

```

Appendix F. Example 4, Section 3.3.2

Fit the b -local version from Example 3 as a mixed-effects model:

```

> library(nlme)
> (blocal.nlme <- nlme(0 ~ uvector(x = height, t = age, unit = Seed,
+ beta0 = 0, beta1 = -b, eta = 0, eta0 = 0, x0 = 0, t0 = 0,
+ lambda = list(a = a, c = c), mup = sqrt(abs(b))), data = Loblolly,
+ fixed = a + b + c ~ 1, random = b ~ 1, groups = ~Seed,
+ start = c(a = 72, b = 0.1, c = 0.5),
+ control = nlmeControl(pnlsTol = 0.01)))
Nonlinear mixed-effects model fit by maximum likelihood
  Model: 0 ~ uvector(x = height, t = age, unit = Seed, beta0 = 0,
    beta1 = -b, eta = 0, eta0 = 0, x0 = 0, t0 = 0,
    lambda = list(a = a, c = c), mup = sqrt(abs(b)))
Data: Loblolly
Log-likelihood: -101.7804
Fixed: a + b + c ~ 1
      a          b          c
73.43276915  0.09381183  0.49381270

Random effects:
Formula: b ~ 1 | Seed
      b Residual
StdDev: 0.003814812 0.7307142

Number of Observations: 84
Number of Groups: 14

```

To achieve convergence it was necessary to increase the tolerance `pnlsTol` from the default 0.001. It may be wise to scale the variables, e. g., specifying `x = height / 10`, `t = age / 10` and adjusting the `start` values accordingly, to make quantities closer to 1 as recommended in the `nlme` documentation. That did not make much difference in this instance.

One can obtain *predictions* for the units (not *estimates*, since here the locals are random variables and not parameters):

```

> coef(blocal.nlme)
      a          b          c
329 73.43277 0.08972720 0.4938127
327 73.43277 0.09095265 0.4938127
325 73.43277 0.09400202 0.4938127
307 73.43277 0.09087038 0.4938127
331 73.43277 0.08972664 0.4938127
311 73.43277 0.09129559 0.4938127

```

```

315 73.43277 0.09407617 0.4938127
321 73.43277 0.09010012 0.4938127
319 73.43277 0.09536543 0.4938127
301 73.43277 0.09631403 0.4938127
323 73.43277 0.09647272 0.4938127
309 73.43277 0.09739088 0.4938127
303 73.43277 0.09743172 0.4938127
305 73.43277 0.09964012 0.4938127
>
> c(AIC(blocal.nlme), BIC(blocal.nlme))
[1] 213.5608 225.7149
> c(AIC(blocal), BIC(blocal))
[1] 204.3040 245.6279

```

References

- Allen E (2007) Modeling with Itô Stochastic Differential Equations. Springer Netherlands, 228 p.
- Arnold L (1974) Stochastic Differential Equations: Theory and Applications. Wiley, New York, 228 p.
- Artzrouni M, Reneke J (1990) Stochastic differential equations in mathematical demography: a review. *Appl Math Comput* 39(3):139–153, DOI 10.1016/0096-3003(90)90078-h
- Batho A, García O (2014) A site index model for lodgepole pine in British Columbia. *Forest Science* 60(5):982–987, DOI 10.5849/forsci.13-509
- von Bertalanffy L (1949) Problems of organic growth. *Nature* 163:156–158, DOI 10.1038/163156a0
- Bishwal JP (2008) Parameter Estimation in Stochastic Differential Equations. Springer-Verlag, Berlin, Heidelberg, New York, DOI 10.1007/978-3-540-74448-1, 264 p.
- Box GEP, Cox DR (1964) An analysis of transformations. *Journal of the Royal Statistical Society, B* 26(2):211–252
- Broad L, Lynch T (2006) Growth models for Sitka spruce in Ireland. *Irish Forestry* 63(1 & 2):53–79
- Brouste A, Fukasawa M, Hino H, Iacus S, Kamatani K, Koike Y, Masuda H, Nomura R, Ogihara T, Shimuzu Y, Uchida M, Yoshida N (2014) The yuima project: A computational framework for simulation and inference of stochastic differential equations. *Journal of Statistical Software* 57(1):1–51, DOI 10.18637/jss.v057.i04
- Bu R, Giet L, Hadri K, Lubrano M (2010) Modeling multivariate interest rates using time-varying copulas and reducible nonlinear stochastic differential equations. *Journal of Financial Econometrics* 9(1):198–236, DOI 10.1093/jjfinec/nbq022
- Bu R, Cheng J, Hadri K (2016) Reducible diffusions with time-varying transformations with application to short-term interest rates. *Economic Modelling* 52, Part A:266–277, DOI 10.1016/j.econmod.2014.10.039
- Cleu EM (2000) Maximum likelihood estimates of a class of one dimensional stochastic differential equation models from discrete data. *Journal of Time Series Analysis* 22(5):505–515, DOI 10.1111/1467-9892.00238
- CRAN (2017) The Comprehensive R Archive Network. <https://cran.r-project.org/>, accessed 2017-07-24
- Cuenod CA, Favetto B, Genon-Catalot V, Rozenholc Y, Samson A (2011) Parameter estimation and change-point detection from Dynamic Contrast Enhanced MRI data using stochastic differential equations. *Math Biosci* 233(1):68 – 76, DOI 10.1016/j.mbs.2011.06.006
- Davidian M, Giltinan DM (2003) Nonlinear models for repeated measurement data: An overview and update. *Journal of Agricultural, Biological, and Environmental Statistics* 8(4):387–419, DOI 10.1198/1085711032697
- Donnet S, Samson A (2008) Parametric inference for mixed models defined by stochastic differential equations. *ESAIM: Probability and Statistics* 12:196–218, DOI 10.1051/ps:2007045
- Donnet S, Samson A (2013) A review on estimation of stochastic differential equations for pharmacokinetic/pharmacodynamic models. *Adv Drug Delivery Rev* 65(7):929 – 939, DOI 10.1016/j.addr.2013.03.005
- Donnet S, Foulley JL, Samson A (2010) Bayesian analysis of growth curves using mixed models defined by stochastic differential equations. *Biometrics* 66(3):733–741, DOI 10.1111/j.1541-0420.2009.01342.x
- Driver CC, Oud JHL, Voelkle MC (2017) Continuous time structural equation modeling with R package ctsem. *Journal of Statistical Software* 77(5):1–35, DOI 10.18637/jss.v077.i05
- Favetto B, Samson A (2010) Parameter estimation for a bidimensional partially observed Ornstein-Uhlenbeck process with biological application. *Scand J Stat* 37:200–220, DOI 10.1111/j.1467-9469.2009.00679.x
- Filipe PA, Braumann CA, Brites NM, Roquete CJ (2010) Modelling animal growth in random environments: An application using nonparametric estimation. *Biometrical Journal* 52(5):653–666, DOI 10.1002/bimj.200900273

- Furnival GM (1961) An index for comparing equations used in constructing volume tables. *Forest Science* 7(4):337–341
- García O (1983) A stochastic differential equation model for the height growth of forest stands. *Biometrics* 39(4):1059–1072, DOI 10.2307/2531339
- García O (1984) New class of growth models for even-aged stands: *Pinus radiata* in Golden Downs forest. *N Z J For Sci* 14:65–88
- García O (2008) Visualization of a general family of growth functions and probability distributions — The Growth-curve Explorer. *Environmental Modelling & Software* 23(12):1474–1475, DOI 10.1016/j.envsoft.2008.04.005
- Gardiner CW (1985) *Handbook of Stochastic Methods for Physics, Chemistry and the Natural Sciences*, 4th edn. Springer-Verlag, 442 p.
- Griewank A, Toint PL (1982) Partitioned variable metric updates for large structured optimization problems. *Numerische Mathematik* 39:119–137, DOI 10.1007/bf01399316
- Henderson D, Plachko P (2006) *Stochastic Differential Equations in Science and Engineering*. World Scientific, DOI 10.1142/9789812774798, 215 p.
- Hotelling H (1927) Differential equations subject to error, and population estimates. *J Am Stat Assoc* 22:283–314, DOI 10.1080/01621459.1927.10502963
- King A, Nguyen D, Ionides E (2016) Statistical inference for partially observed Markov processes via the R package pomp. *Journal of Statistical Software* 69(1):1–43, DOI 10.18637/jss.v069.i12
- Klim S, Mortensen SB, Kristensen NR, Overgaard RV, Madsen H (2009) Population stochastic modelling (PSM) — An R package for mixed-effects models based on stochastic differential equations. *Comput Methods Programs Biomed* 94(3):279–289, DOI 10.1016/j.cmpb.2009.02.001
- Kloeden PE, Platen E (1992) *Numerical Solution of Stochastic Differential Equations*. Springer, 632 p.
- Kristensen NR, Madsen H, Ingwersen SH (2005) Using stochastic differential equations for PK/PD model development. *J Pharmacokinet Pharmacodyn* 32(1):109–141, DOI 10.1007/s10928-005-2105-9
- Kunita H (2010) Itô’s stochastic calculus: Its surprising power for applications. *Stochastic Processes Appl* 120(5):622–652, DOI 10.1016/j.spa.2010.01.013
- Lv Q, Pitchford JW (2007) Stochastic von Bertalanffy models, with applications to fish recruitment. *J Theor Biol* 244(4):640–655, DOI 10.1016/j.jtbi.2006.09.009
- Martínez AS, González RS, Terçariol CAS (2008) Continuous growth models in terms of generalized logarithm and exponential functions. *Physica A* 387(23):5679–5687, DOI 10.1016/j.physa.2008.06.015
- Moler C, Van Loan C (2003) Nineteen dubious ways to compute the exponential of a matrix, twenty-five years later. *SIAM Rev* 45(1):3–49, DOI 10.1137/S00361445024180
- Møller J, Madsen H (2010) From state dependent diffusion to constant diffusion in stochastic differential equations by the Lamperti transform. IMM-Technical Report-2010-16, Technical University of Denmark, DTU Informatics, Lyngby, Denmark, URL http://orbit.dtu.dk/files/5175845/tr10_16.pdf, 25 p.
- Nielsen JN, Madsen H, Young PC (2000) Parameter estimation in stochastic differential equations: An overview. *Annual Reviews in Control* 24:83–94, DOI 10.1016/S1367-5788(00)90017-8
- Øksendal B (2003) *Stochastic Differential Equations*, sixth edn. Springer, Berlin, Heidelberg, DOI 10.1007/978-3-642-14394-6_5, 360 p.
- Paige R, Allen E (2010) Closed-form likelihoods for stochastic differential equation growth models. *Canadian Journal of Statistics-revue Canadienne De Statistique* 38(3):474–487, DOI 10.1002/cjs.10071
- Picchini U, Ditlevsen S (2011) Practical estimation of high dimensional stochastic differential mixed-effects models. *Computational Statistics & Data Analysis* 55(3):1426–1444, DOI 10.1016/j.csda.2010.10.003
- Picchini U, de Gaetano A, Ditlevsen S (2010) Stochastic differential mixed-effects models. *Scand J Stat* 37(1):67–90, DOI 10.1111/j.1467-9469.2009.00665.x
- Pinheiro JC, Bates DM (2000) *Mixed-Effects Models in S and S-Plus*. Springer, New York, DOI 10.1007/b98882, 528 p.
- R Development Core Team (2009) *R: A Language and Environment for Statistical Computing*. R Foundation for Statistical Computing, Vienna, Austria, ISBN 3-900051-07-0 (<http://www.R-project.org>)
- Richards FJ (1959) A flexible growth function for empirical use. *J Exp Bot* 10(29):290–300, DOI 10.1093/jxb/10.2.290
- Seber GAF, Wild CJ (2003) *Nonlinear Regression*. Wiley-Interscience, New York, DOI 10.1002/0471725315, 768 p.
- Sen P (1989) A stochastic model for coupled enzyme system. *Biometrical Journal* 31(8):973–992, DOI 10.1002/bimj.4710310814
- Simoncini V (2016) Computational methods for linear matrix equations. *SIAM Review* 58(3):377–441, DOI 10.1137/130912839
- Snijders TAB (2003) Multilevel analysis. In: Lewis-Beck M, Bryman A, Liao T (eds) *The SAGE Encyclopedia of*

- Social Science Research Methods, vol II, Sage, Thousand Oaks, CA, pp 673–677
- Soo YW, Bates DM (1992) Loosely coupled nonlinear least squares. *Computational Statistics & Data Analysis* 14(2):249–259, DOI 10.1016/0167-9473(92)90177-h
- Strathe AB, Sorensen H, Danfaer A (2009) A new mathematical model for combining growth and energy intake in animals: The case of the growing pig. *J Theor Biol* 261(2):165–175, DOI 10.1016/j.jtbi.2009.07.039
- Venables WN, Ripley BD (2002) *Modern Applied Statistics with S*, 4th edn. Springer, New York, DOI 10.1007/978-0-387-21706-2, 495 p.
- Whitaker GA, Golightly A, Boys RJ, Sherlock C (2017) Bayesian inference for diffusion-driven mixed-effects models. *Bayesian Analysis* 12(2):435–463, DOI 10.1214/16-BA1009
- Wilks SS (1962) *Mathematical Statistics*, 2nd edn. Wiley, 644 p.
- Zacks S (1971) *The Theory of Statistical Inference*. Wiley, New York, 609 p.