

UC Irvine

UC Irvine Previously Published Works

Title

Estimation for General Birth-Death Processes

Permalink

<https://escholarship.org/uc/item/7m81c13b>

Journal

Journal of the American Statistical Association, 109(506)

ISSN

0162-1459

Authors

Crawford, Forrest W
Minin, Vladimir N
Suchard, Marc A

Publication Date

2014-04-03

DOI

10.1080/01621459.2013.866565

Copyright Information

This work is made available under the terms of a Creative Commons Attribution License, available at <https://creativecommons.org/licenses/by/4.0/>

Peer reviewed

Estimation for general birth-death processes

Forrest W. Crawford¹, Vladimir N. Minin², and Marc A. Suchard^{3,4,5}

Typeset November 22, 2011

1. Department of Biomathematics
University of California Los Angeles
CHS AV-611
Los Angeles, CA 90095-1766 USA
fcrawford@ucla.edu

2. Department of Statistics
University of Washington
Padelford Hall C-315, Box 354322
Seattle, WA 98195-4322 USA
vminin@u.washington.edu

3. Departments of Biomathematics,
4. Human Genetics, and
5. Biostatistics
University of California Los Angeles
6558 Gonda Building,
Los Angeles, CA 90095-1766 USA
msuchard@ucla.edu

Abstract

Birth-death processes (BDPs) are continuous-time Markov chains that track the number of “particles” in a system over time. While widely used in population biology, genetics and ecology, statistical inference of the instantaneous particle birth and death rates remains largely limited to restrictive linear BDPs in which per-particle birth and death rates are constant. Researchers often observe the number of particles at discrete times, necessitating data augmentation procedures such as expectation-maximization (EM) to find maximum likelihood estimates. The E-step in the EM algorithm is available in closed-form for some linear BDPs, but otherwise previous work has resorted to approximation or simulation. Remarkably, the E-step conditional expectations can also be expressed as convolutions of computable transition probabilities for any general BDP with arbitrary rates. This important observation, along with a convenient continued fraction representation of the Laplace transforms of the transition probabilities, allows novel and efficient computation of the conditional expectations for all BDPs, eliminating the need for approximation or costly simulation. We use this insight to derive EM algorithms that yield maximum likelihood estimation for general BDPs characterized by various rate models, including generalized linear models. We show that our Laplace convolution technique outperforms competing methods when available and demonstrate a technique to accelerate EM algorithm convergence. Finally, we validate our approach using synthetic data and then apply our methods to estimation of mutation parameters in microsatellite evolution.

Keywords: Birth-death process, EM algorithm, MM algorithm, maximum likelihood estimation, continuous-time Markov chain, microsatellite evolution

1 Introduction

A birth-death process (BDP) is a continuous-time Markov chain that models a non-negative integer number of particles in a system (Feller, 1971). The state of the system at a given time is the number of particles in existence. At any moment in time, one of the particles may “give birth” to a new particle, increasing the count by one, or one particle may “die”, decreasing the count by one. BDPs are popular modeling tools in a wide variety of quantitative disciplines, such as population biology, genetics, and ecology (Thorne et al, 1991; Krone and Neuhauser, 1997; Novozhilov et al, 2006). For example, BDPs can characterize epidemic dynamics, (Bailey, 1964; Andersson and Britton, 2000), speciation and extinction (Nee et al, 1994; Nee, 2006), evolution of gene families (Cotton and Page, 2005; Demuth et al, 2006), and the insertion and deletion events for probabilistic alignment of DNA sequences (Thorne et al, 1991; Holmes and Bruno, 2001).

Traditionally, most modeling applications have used the “simple linear” BDP with constant per-particle birth and death rates, which arises from an assumption of independence among particles and no background birth and death rates. When individual birth and death rates instead depend on the size of the population as a whole, the model is called a “general” BDP. Previous statistical estimation in BDPs has focused mainly on estimating the constant per-particle birth and death rates of the simple linear BDP based on observations of the number of particles over time. However, the simple linear BDP is often unrealistic, and nonlinear dependence of the birth and death rates on the current number of particles provides the means to model more sophisticated and realistic patterns of stochastic population dynamics in a wide variety of biological disciplines. For example, populations sometimes exhibit logistic-like growth as their number approaches the carrying capacity of their environment (Tan and Piantadosi, 1991). In genetic models, the rate of new offspring carrying an allele often depends on the proportions of both individuals already carrying the allele and those who do not (Moran, 1958). In coalescent theory, the rate of coalescence changes with the square of the number of lineages (Kingman, 1982). In addition, researchers may wish to assess the influence of covariates on birth and death rates by fitting a regression model (Kalbfleisch and Lawless, 1985; Liu et al, 2007).

Progress in estimating birth and death rates in BDPs has also typically been limited to continuous observation of the process (Moran, 1951, 1953; Anscombe, 1953; Darwin, 1956; Wolff, 1965;

Reynolds, 1973; Keiding, 1975). However, in practice researchers may observe data from BDPs only at discrete times through longitudinal observations. Estimating transition rates in continuous-time Markov processes using discrete observations is difficult since the state path between observations is not observed. Furthermore, direct analytic maximization of the likelihood for general BDPs remains infeasible for partially observed samples since the likelihood usually cannot be written in closed-form. Despite these challenges, several researchers have made progress in estimating parameters of the simple linear BDP under discrete observation (Keiding, 1974; Thorne et al, 1991; Holmes and Bruno, 2001; Rosenberg et al, 2003; Dauxois, 2004). However, none of these developments provides a robust method to find exact maximum likelihood estimates (MLEs) of parameters in discretely observed general BDPs with arbitrary birth and death rates.

A major insight comes from the fact that the likelihood of the continuously observed process has a simple form which easily yields expressions for estimation of rate parameters. This fact is the basis for expectation-maximization (EM) algorithms for maximum likelihood estimation in missing data problems (Dempster et al, 1977). In finite state-space Markov chains, the relevant conditional expectations (the E-step of the EM algorithm) can often be computed efficiently, and several researchers have derived EM algorithms for estimating transition rates in this context (Lange, 1995a; Holmes and Rubin, 2002; Hobolth and Jensen, 2005; Bladt and Sorensen, 2005; Metzner et al, 2007). Unfortunately, finding these conditional expectations for general BDPs poses challenges since the joint distribution of the states and waiting times (or its generating function) is usually not available in closed-form. Notably, Holmes and Bruno (2001); Holmes and Rubin (2002) and Doss et al (2010) are able to find analytic expressions or numerical approximations for these expectations in EM algorithms for certain BDPs whose rates depend linearly on the current number of particles. While these developments are promising, there remains a great need for estimation techniques that can be applied to more sophisticated BDPs under a variety of sampling scenarios. Indeed, more complex and realistic models like those reviewed by Novozhilov et al (2006) may be of little use to applied researchers if no practical method exists to estimate their parameters.

Here we seek to fill this apparent void by providing a framework for deriving EM algorithms for estimating rate parameters of a general BDP. We first formally define the general BDP and give an exact expression for the Laplace transform of the transition probabilities in the form of a continued fraction. We then give the likelihood for continuously-observed BDPs and outline

the EM algorithm. Next, we describe a novel method to efficiently compute the expectations of the E-step for BDPs with arbitrary rates. Since these expectations are convolutions of transition probabilities, we perform the convolution in the Laplace domain, and then invert the Laplace transformed expressions to obtain the desired conditional expectation. This technique obviates the costly numerical integration or repeated simulation that has plagued previous approaches. We provide examples of the maximization step for several different classes of BDPs and demonstrate a technique for accelerating convergence of the EM algorithm. We show that our method is faster than competing techniques and validate it using simulated data. Finally, we conclude with an application that analyzes microsatellite evolution and answers an open question in evolutionary genomics.

2 General BDPs and their EM algorithms

2.1 Formal description and transition probabilities

Consider a general BDP $X(\tau)$ counting the number of particles k in existence at times $\tau \geq 0$. From state $X(\tau) = k$, transitions to state $k + 1$ happen with instantaneous rate λ_k , and transitions to state $k - 1$ happen with instantaneous rate μ_k . The transition rates λ_k and μ_k may depend on k but are time-homogeneous. As we show below, it is often necessary to evaluate finite-time transition probabilities to derive efficient EM algorithms for estimation of arbitrary birth and death rates in general BDPs. This proves useful both in completing the E-step of the EM algorithm and in computing incomplete data likelihoods for validation of our EM estimates. For a starting state $i \geq 0$, the finite-time transition probabilities $P_{i,j}(\tau) = \Pr(X(\tau) = j \mid X(0) = i)$ obey the system of ordinary differential equations

$$\begin{aligned} \frac{dP_{i,0}(\tau)}{d\tau} &= \mu_1 P_{i,1}(\tau) - \lambda_0 P_{i,0}(\tau), \text{ and} \\ \frac{dP_{i,j}(\tau)}{d\tau} &= \lambda_{j-1} P_{i,j-1}(\tau) + \mu_{j+1} P_{i,j+1}(\tau) - (\lambda_j + \mu_j) P_{i,j}(\tau), \end{aligned} \tag{1}$$

for $j \geq 1$ with $P_{i,i}(0) = 1$ and $P_{i,j}(0) = 0$ for $i \neq j$ (Feller, 1971).

For some simple parameterizations of λ_k and μ_k , closed-form solutions exist for the transition probabilities $P_{i,j}(\tau)$, but this is not possible for most models. Karlin and McGregor (1957) show

that for any parameterization of λ_k and μ_k , it is possible to express the transition probabilities in terms of orthogonal polynomials. However, in practice these special polynomials are difficult to find, and even when they are available, they rarely yield solutions in closed-form or expressions that are amenable to computation (Novozhilov et al, 2006; Renshaw, 2011). In contrast, the continued fraction method we outline below does not require additional model-specific insight beyond specification of λ_k and μ_k .

To solve for the transition probabilities, it is advantageous to work in the Laplace domain (Karlin and McGregor, 1957). This transformation also proves essential in maintaining numerical stability of transition probabilities in general BDPs and in computing the conditional expectations necessary for the EM algorithm derived in a subsequent section. Laplace transforming equation (1) yields

$$\begin{aligned} s f_{i,0}(s) - \delta_{i0} &= \mu_1 f_{i,1}(s) - \lambda_0 f_{i,0}(s), \\ s f_{i,j}(s) - \delta_{ij} &= \lambda_{j-1} f_{i,j-1}(s) + \mu_{j+1} f_{i,j+1}(s) - (\lambda_j + \mu_j) f_{i,j}(s), \end{aligned} \tag{2}$$

where $f_{i,j}(s)$ is the Laplace transform of $P_{i,j}(\tau)$ and $\delta_{ij} = 1$ if $i = j$ and zero otherwise. Letting $i = 0$ and rearranging (2), we obtain the recurrence relations

$$\begin{aligned} f_{0,0}(s) &= \frac{1}{s + \lambda_0 - \mu_1 \left(\frac{f_{0,1}(s)}{f_{0,0}(s)} \right)}, \text{ and} \\ \frac{f_{0,j}(s)}{f_{0,j-1}(s)} &= \frac{\lambda_{j-1}}{s + \mu_j + \lambda_j - \mu_{j+1} \left(\frac{f_{0,j+1}(s)}{f_{0,j}(s)} \right)}. \end{aligned} \tag{3}$$

We can inductively combine these expressions for $j = 1, 2, 3, \dots$ to arrive at the well-known generalized continued fraction

$$f_{0,0}(s) = \frac{1}{s + \lambda_0 - \frac{\lambda_0 \mu_1}{s + \lambda_1 + \mu_1 - \frac{\lambda_1 \mu_2}{s + \lambda_2 + \mu_2 - \dots}}}. \tag{4}$$

This is an exact expression for the Laplace transform of the transition probability $P_{0,0}(\tau)$. In (4), let $a_1 = 1$ and $a_j = -\lambda_{j-2} \mu_{j-1}$, and let $b_1 = s + \lambda_0$ and $b_j = s + \lambda_{j-1} + \mu_{j-1}$ for $j \geq 2$. Then (4)

becomes

$$f_{0,0}(s) = \frac{a_1}{b_1 + \frac{a_2}{b_2 + \frac{a_3}{b_3 + \dots}}} \quad (5)$$

We can write this more compactly as

$$f_{0,0}(s) = \frac{a_1}{b_1 +} \frac{a_2}{b_2 +} \frac{a_3}{b_3 +} \dots \quad (6)$$

The k th convergent of $f_{0,0}(s)$ is

$$f_{0,0}^{(k)}(s) = \frac{a_1}{b_1 +} \frac{a_2}{b_2 +} \dots \frac{a_k}{b_k} = \frac{A_k(s)}{B_k(s)}, \quad (7)$$

where $A_k(s)$ and $B_k(s)$ are the numerator and denominator of the rational function $f_{0,0}^{(k)}$. The transition probabilities $P_{i,j}(\tau)$ for $i, j > 0$ can be derived in continued fraction form by combining (2) and (4) to obtain

$$f_{i,j}(s) = \begin{cases} \left(\prod_{k=j+1}^i \mu_k \right) \frac{B_j(s)}{B_{i+1}(s)+} \frac{B_i(s)a_{i+2}}{b_{i+2}+} \frac{a_{i+3}}{b_{i+3}+} \dots & \text{for } j \leq i, \\ \left(\prod_{k=i}^{j-1} \lambda_k \right) \frac{B_i(s)}{B_{j+1}(s)+} \frac{B_j(s)a_{j+2}}{b_{j+2}+} \frac{a_{j+3}}{b_{j+3}+} \dots & \text{for } i \leq j, \end{cases} \quad (8)$$

(Murphy and O'Donohoe, 1975; Crawford and Suchard, 2011).

Although the Laplace transforms of the transition probabilities are generally still not available in closed-form, a continued fraction representation is desirable for several reasons: 1) continued fraction representations of functions often converge much faster than equivalent power series; 2) there are efficient algorithms for evaluating them to a finite depth; and 3) there exist methods for bounding the error of truncated continued fractions (Bankier and Leighton, 1942; Wall, 1948; Blanch, 1964; Lorentzen and Waadeland, 1992; Craviotto et al, 1993; Abate and Whitt, 1999; Cuyt et al, 2008). For an arbitrary BDP, we recover the transition probabilities through numerical inversion of the Laplace-transformed expressions. We evaluate the continued fraction to a moni-

tored depth that controls the overall error and generates stable approximations to the transition probabilities unattainable by previous methods (Murphy and O’Donohoe, 1975; Parthasarathy and Sudhesh, 2006; Crawford and Suchard, 2011).

The ability to compute transition probabilities for general BDPs with arbitrary rate parameterizations proves useful in two ways. First, if we interpret finite-time transition probabilities as functions of an unknown parameter vector $\boldsymbol{\theta}$, then $P_{a,b}(t)$ given $\boldsymbol{\theta}$ returns the *likelihood* of a discrete observation from a BDP such that $X(0) = a$ and $X(t) = b$, where the trajectory in time t between a and b is unobserved. Second, transition probabilities play an important role in computing conditional expectations of sufficient statistics, as we shall see below.

2.2 Likelihood expressions and surrogate functions

[Figure 1 about here.]

With a formal description of a general BDP and the finite-time transition probabilities in hand, we now proceed with our task of estimating the parameters of a general BDP using discrete observations. Given one or more independent observations of the form $\mathbf{Y} = (X(0) = a, X(t) = b)$ from a general BDP, we wish to find maximum likelihood estimates of the rate parameters λ_k and μ_k for $k = 0, 1, 2, \dots$. We will assume that the birth and death rates at state k depend on both k and a finite-dimensional parameter vector $\boldsymbol{\theta}$, so that the form of $\lambda_k(\boldsymbol{\theta})$ and $\mu_k(\boldsymbol{\theta})$ is known for all k .

For a single realization of the process starting at $X(0) = a$ and ending at $X(t) = b$, let T_k be the total time spent in state k . Let U_k be the number of “up” steps (births) from state k , and let D_k be the number of “down” steps (deaths) from state k . Let the total number of up and down steps in a realization of the process be denoted by $U = \sum_{k=0}^{\infty} U_k$ and $D = \sum_{k=0}^{\infty} D_k$ respectively. We also define the total particle time,

$$T_{\text{particle}} = \int_0^t X(\tau) \, d\tau = \sum_{k=0}^{\infty} kT_k, \quad (9)$$

that counts the amount of time lived by each particle since time $\tau = 0$. Of course, the total elapsed time is $t = \sum_{k=0}^{\infty} T_k$. We demonstrate these concepts schematically in Figure 1.

The log-likelihood for a continuously observed process takes a simple form when we sum over all possible states k (Wolff, 1965):

$$\ell(\boldsymbol{\theta}) = \sum_{k=0}^{\infty} U_k \log [\lambda_k(\boldsymbol{\theta})] + D_k \log [\mu_k(\boldsymbol{\theta})] - [\lambda_k(\boldsymbol{\theta}) + \mu_k(\boldsymbol{\theta})] T_k. \quad (10)$$

However, when a BDP is sampled discretely such that only $X(0) = a$ and $X(t) = b$ are observed, the quantities U_k , D_k , and T_k are unknown for every state k , and we cannot maximize the log-likelihood (10) without them.

We therefore appeal to the EM algorithm for iterative maximum likelihood estimation with missing data (Dempster et al, 1977). In the EM algorithm, we define a surrogate objective function Q by taking the expectation of the complete data log-likelihood (10), conditional on the observed data \mathbf{Y} and the parameter values $\boldsymbol{\theta}^{(m)}$ from the previous iteration of the EM algorithm (the E-step). Then we find the parameter values $\boldsymbol{\theta}^{(m+1)}$ that maximize this surrogate function (the M-step). This two-step process is repeated until convergence to the maximum likelihood estimate of $\boldsymbol{\theta}$. Taking the expectation of (10) conditional on \mathbf{Y} and $\boldsymbol{\theta}^{(m)}$, we form the surrogate function Q :

$$\begin{aligned} Q(\boldsymbol{\theta} \mid \boldsymbol{\theta}^{(m)}) &= \mathbb{E}[\ell(\boldsymbol{\theta}) \mid \mathbf{Y}, \boldsymbol{\theta}^{(m)}] \\ &= \sum_{k=0}^{\infty} \mathbb{E}(U_k \mid \mathbf{Y}) \log [\lambda_k(\boldsymbol{\theta})] + \mathbb{E}(D_k \mid \mathbf{Y}) \log [\mu_k(\boldsymbol{\theta})] - \mathbb{E}(T_k \mid \mathbf{Y}) [\lambda_k(\boldsymbol{\theta}) + \mu_k(\boldsymbol{\theta})], \end{aligned} \quad (11)$$

where for clarity we have omitted the dependence of the expectations on the parameter value $\boldsymbol{\theta}^{(m)}$ from the m th iterate. In general, we assume that the maximum likelihood estimator exists; see Bladt and Sorensen (2005) for a discussion of the issues of identifiability, existence, and uniqueness.

2.3 Computing the expectations of the E-step

Computing the expectations of U_k , D_k , and T_k in the E-step is difficult in birth-death estimation since the unobserved state path and waiting times are not independent conditional on the observed data \mathbf{Y} . Doss et al (2010) adopt an approach for linear BDPs that combines analytic results with simulations. For some models, these authors are able to derive the generating function for the joint distribution of U , D , T_{particle} , and the state path conditional on $X(0) = a$ and can manipulate this generating function to complete the E-step. For a more complicated linear model, Doss et al resort

to approximating the relevant conditional expectations by simulating sample paths, conditional on \mathbf{Y} (Hobolth, 2008).

Our solution is to recognize that we do not need to know very much about the missing data to find the conditional expectations used in the sufficient statistics above. In fact, the transition probabilities are all that we require. The following integral representations of the conditional expectations in the EM algorithm will prove useful:

$$\mathbb{E}(U_k|\mathbf{Y}) = \frac{\int_0^t P_{a,k}(\tau)\lambda_k P_{k+1,b}(t-\tau) d\tau}{P_{a,b}(t)}, \quad (12a)$$

$$\mathbb{E}(D_k|\mathbf{Y}) = \frac{\int_0^t P_{a,k}(\tau)\mu_k P_{k-1,b}(t-\tau) d\tau}{P_{a,b}(t)}, \quad \text{and} \quad (12b)$$

$$\mathbb{E}(T_k|\mathbf{Y}) = \frac{\int_0^t P_{a,k}(\tau)P_{k,b}(t-\tau) d\tau}{P_{a,b}(t)}. \quad (12c)$$

These formulas have appeared in many types of studies related to EM estimation for continuous-time Markov chains (Lange, 1995a; Holmes and Rubin, 2002; Bladt and Sorensen, 2005; Hobolth and Jensen, 2005; Metzner et al, 2007). For general BDPs whose transition probabilities must be computed numerically, numerical integration over the product of the densities can be computationally prohibitive.

However, the numerators in (12) a-c are convolutions of integrable time-domain functions. Since the Laplace transforms $f_{a,b}(s)$ of these transition probabilities are available and easy to compute, we take advantage of the Laplace convolution property, arriving at the representations

$$\mathbb{E}(U_k|\mathbf{Y}) = \lambda_k \frac{\mathcal{L}^{-1}\left[f_{a,k}(s) f_{k+1,b}(s)\right](t)}{P_{a,b}(t)}, \quad (13a)$$

$$\mathbb{E}(D_k|\mathbf{Y}) = \mu_k \frac{\mathcal{L}^{-1}\left[f_{a,k}(s) f_{k-1,b}(s)\right](t)}{P_{a,b}(t)}, \quad \text{and} \quad (13b)$$

$$\mathbb{E}(T_k|\mathbf{Y}) = \frac{\mathcal{L}^{-1}\left[f_{a,k}(s) f_{k,b}(s)\right](t)}{P_{a,b}(t)}. \quad (13c)$$

where \mathcal{L}^{-1} denotes inverse Laplace transformation. Although these formulas are equivalent to (12), they offer substantial time savings over computing the integral directly, and render tractable the

computation of expectations in the EM algorithm for arbitrary general BDPs.

To calculate the numerators of (13), we use the Laplace inversion method popularized by Abate and Whitt (1992, 1995). This involves a Riemann sum approximation of the inverse transform that stabilizes the discretization error and is amenable to series acceleration methods (Abate and Whitt, 1999; Press, 2007). To evaluate the continued fraction Laplace transforms $f_{a,b}(s)$, we use the modified Lentz method (Lentz, 1976; Thompson and Barnett, 1986; Press, 2007).

2.4 Maximization techniques for various BDPs

In contrast to the generic technique outlined above for computing the expectations of the E-step, the M-step depends explicitly on the functional form of the birth and death rates $\lambda_k(\boldsymbol{\theta})$ and $\mu_k(\boldsymbol{\theta})$. Here we give several representative examples of BDPs and techniques for completing the M-step of the EM algorithm, such as analytic maximization, minorize-maximize (MM), and Newton’s method.

2.4.1 Simple linear BDP

In the simple linear BDP, births and deaths happen at constant per-capita rates, so $\lambda_k = k\lambda$ and $\mu_k = k\mu$. The unknown parameter vector is $\boldsymbol{\theta} = (\lambda, \mu)$, and the surrogate function becomes

$$Q(\boldsymbol{\theta}) = \sum_{k=0}^{\infty} \mathbb{E}(U_k|\mathbf{Y}) \log[k\lambda] + \mathbb{E}(D_k|\mathbf{Y}) \log[k\mu] - \mathbb{E}(T_k|\mathbf{Y})k(\lambda + \mu). \quad (14)$$

Taking the derivative of (14) with respect to the unknown parameters, setting the result to zero, and solving for λ and μ gives the M-step updates

$$\lambda^{(m+1)} = \frac{\mathbb{E}(U|\mathbf{Y})}{\mathbb{E}(T_{\text{particle}}|\mathbf{Y})}, \text{ and} \quad (15a)$$

$$\mu^{(m+1)} = \frac{\mathbb{E}(D|\mathbf{Y})}{\mathbb{E}(T_{\text{particle}}|\mathbf{Y})}. \quad (15b)$$

These updates correspond to the usual maximum likelihood estimators in the continuously observed process (Reynolds, 1973). Note that the transition probabilities $P_{a,b}(t)$ in the denominators of the expectations in (12) cancel out in (15a) and (15b). When this is the case, transition probabilities are not necessary to derive an EM algorithm.

2.4.2 Linear BDP with immigration

Sometimes populations are not closed, and new individuals can enter; we call this action “immigration.” Another interpretation arises in models of point mutations in DNA sequences. Suppose new mutations arise in a DNA sequence via two distinct processes: one inserts new mutants at a rate proportional to the number already present, and the other creates new mutations at a constant rate, regardless of how many already exist. To model this behavior, we augment the simple linear BDP above with a constant term ν representing immigration, so that $\lambda_k = k\lambda + \nu$ and $\mu_k = k\mu$. The log-likelihood becomes

$$\ell(\boldsymbol{\theta}) = \sum_{k=0}^{\infty} U_k \log(k\lambda + \nu) + D_k \log(\mu) - T_k [k(\lambda + \mu) + \nu]. \quad (16)$$

Unfortunately, if we take the derivative of the log-likelihood with respect to λ or ν , the unknown appears in the denominator of the terms of the infinite sum. However, since each summand is a concave function of the unknown parameters, we can separate them in a minorizing function H such that for all $\boldsymbol{\theta}$, $H(\boldsymbol{\theta}|\boldsymbol{\theta}^{(m)}) \leq \ell(\boldsymbol{\theta})$ and $H(\boldsymbol{\theta}^{(m)}|\boldsymbol{\theta}^{(m)}) = \ell(\boldsymbol{\theta}^{(m)})$ as follows:

$$\begin{aligned} \ell(\boldsymbol{\theta}) &\geq H(\boldsymbol{\theta}|\boldsymbol{\theta}^{(m)}) \\ &= \sum_{k=0}^{\infty} U_k [p_k \log(p_k \lambda) + (1 - p_k) \log((1 - p_k)\nu)] + D_k \log(\mu) - [k(\lambda + \mu) + \nu] T_k, \end{aligned} \quad (17)$$

where

$$p_k = \frac{k\lambda^{(m)}}{k\lambda^{(m)} + \nu^{(m)}}. \quad (18)$$

Then letting $Q(\boldsymbol{\theta} | \boldsymbol{\theta}^{(m)}) = \mathbb{E} \left(H(\boldsymbol{\theta}) | \mathbf{Y}, \boldsymbol{\theta}^{(m)} \right)$ be the surrogate function, this minorization forms the basis for an EM algorithm in which a step of the minorize-maximize (MM) algorithm takes the place of the M-step, and the ascent property of the EM algorithm is preserved (Lange, 2010).

Maximizing Q with respect to λ and ν yields the updates

$$\lambda^{(m+1)} = \frac{\sum_{k=0}^{\infty} p_k \mathbb{E}(U_k | \mathbf{Y})}{\mathbb{E}(T_{\text{particle}} | \mathbf{Y})}, \text{ and} \quad (19a)$$

$$\nu^{(m+1)} = \frac{\sum_{k=0}^{\infty} (1 - p_k) \mathbb{E}(U_k | \mathbf{Y})}{t}. \quad (19b)$$

Expression (19a) is similar to (15a), the update for λ in the simple BDP. The difference lies in that each $\mathbb{E}(U_k | \mathbf{Y})$ in this case is weighted by the proportion of additions at state k due to births, not immigrations. The update for μ is the same as (15b).

2.4.3 Logistic/restricted growth

To illustrate an EM algorithm for more complicated rate specifications in which no MM update is evident and the rates no longer depend on the current state k in a linear way, we examine a model for restricted population growth. Typical *deterministic* population models often incorporate limitations on population size due to the carrying capacity K of the environment. One famous example is the logistic model of population growth (Murray, 2002). Continuous-time stochastic analogs have previously required a finite cap on population size (Tan and Piantadosi, 1991). These stochastic models roughly mimic the behavior of the deterministic model for population sizes below K , but are limited because they do not allow growth beyond K . Here we present a model which supports transient growth beyond the carrying capacity, but where the population size tends to a balance between restricted growth and death.

Suppose births are cooperative, requiring two parents, but fecundity decays as the number of extant particles increases, and death remains an independent process such that $\lambda_k = \lambda k^2 e^{-\beta k}$ and $\mu_k = k\mu$. Here, we can interpret the carrying capacity roughly as the population size $k > 0$ at which $\lambda_k \approx \mu_k$. Ignoring irrelevant terms, the surrogate function becomes

$$Q(\boldsymbol{\theta} | \boldsymbol{\theta}^{(m)}) = \sum_{k=0}^{\infty} \mathbb{E}(U_k | \mathbf{Y}) [\log(\lambda) - \beta k] + \mathbb{E}(D_k | \mathbf{Y}) \log(\mu) - \mathbb{E}(T_k | \mathbf{Y}) [\lambda k^2 e^{-\beta k} + k\mu]. \quad (20)$$

Since λ and β appear together, we opt for a numerical Newton step. The gradient of Q with respect

to these parameters is

$$F = \begin{pmatrix} \frac{\mathbb{E}(U|\mathbf{Y})}{\lambda} - \sum_{k=0}^{\infty} k^2 e^{-\beta k} \mathbb{E}(T_k|\mathbf{Y}) \\ - \sum_{k=0}^{\infty} [k \mathbb{E}(U_k|\mathbf{Y}) + \lambda k^3 e^{-\beta k} \mathbb{E}(T_k|\mathbf{Y})] \end{pmatrix}, \quad (21)$$

and the Hessian is

$$H = \begin{pmatrix} -\frac{\mathbb{E}(U|\mathbf{Y})}{\lambda^2} & -\sum_{k=0}^{\infty} k^3 e^{-\beta k} \mathbb{E}(T_k|\mathbf{Y}) \\ -\sum_{k=0}^{\infty} k^3 e^{-\beta k} \mathbb{E}(T_k|\mathbf{Y}) & \lambda \sum_{k=0}^{\infty} k^4 e^{-\beta k} \mathbb{E}(T_k|\mathbf{Y}). \end{pmatrix}. \quad (22)$$

Then we update these parameters by

$$\begin{pmatrix} \lambda^{(m+1)} \\ \beta^{(m+1)} \end{pmatrix} = \begin{pmatrix} \lambda^{(m)} \\ \beta^{(m)} \end{pmatrix} - H^{-1} F. \quad (23)$$

The ascent property is preserved when a Newton step is used in place of an exact M-step (Lange, 1995a). The update for μ is the same as (15b).

2.4.4 SIS epidemic models

Under a very common epidemic model, members of a finite population of size N are classified as either “susceptible” to a given disease or “infected” (Bailey, 1964; Andersson and Britton, 2000). Susceptibles become infected in proportion to the number of currently infected in the population, and infecteds revert to susceptible status with a certain rate independent of how many infecteds there are. This idealized susceptible-infectious-susceptible (SIS) infectious disease model specifies a general birth-death process in which we track the number of infecteds. Let $\lambda_k = \beta k(N - k)/N$ be the rate of new infections when there are already k infected in the population. Let $\mu_k = \gamma k/N$ be the rate of recovery of infecteds to susceptibles. Then if $\boldsymbol{\theta} = (\beta, \gamma)$, we have

$$Q(\boldsymbol{\theta}|\boldsymbol{\theta}^{(m)}) = \sum_{k=0}^N \mathbb{E}(U_k|\mathbf{Y}) \log(\beta) + \mathbb{E}(D_k|\mathbf{Y}) \log(\gamma) - \mathbb{E}(T_k|\mathbf{Y})(k(N - k)\beta + k\gamma)/N, \quad (24)$$

and the update for β is

$$\beta^{(m+1)} = \frac{N\mathbb{E}(U|\mathbf{Y})}{\sum_{k=0}^N (N-k)k\mathbb{E}(T_k|\mathbf{Y})}. \quad (25)$$

The update for γ is

$$\gamma^{(m+1)} = \frac{N\mathbb{E}(D|\mathbf{Y})}{\mathbb{E}(T_{\text{particle}}|\mathbf{Y})}. \quad (26)$$

2.4.5 Generalized linear models

Our general framework allows assessment of the influence of covariates on the rates of a general BDP in a novel way. Suppose we sample observations from independent processes $X_i(\tau)$, $i = 1, \dots, N$ and observe $\mathbf{Y}_i = (X_i(0), X_i(t_i))$ associated with d covariates $\mathbf{z}_i = (z_{i1}, \dots, z_{id})^t$. These processes may represent different subjects in a study. We model the birth and death rates λ_{ik} and μ_{ik} for each process/subject X_i as functions of \mathbf{z}_i and unknown d -dimensional regression coefficients $\boldsymbol{\theta}_\lambda$ and $\boldsymbol{\theta}_\mu$ in a generalized linear model (GLM) framework. We link

$$\log(\lambda_{ik}) = g(k, \mathbf{z}_i^t \boldsymbol{\theta}_\lambda) \quad \text{and} \quad \log(\mu_{ik}) = h(k, \mathbf{z}_i^t \boldsymbol{\theta}_\mu), \quad (27)$$

where $g(\cdot)$ and $h(\cdot)$ are scalar-valued functions. We note the possibility that covariates may differ between $\boldsymbol{\theta}_\lambda$ and $\boldsymbol{\theta}_\mu$ through trivial modification; to ease notation, we do not explore this direction. Given N independent processes, we sum log-likelihoods to arrive at the multiple-subject surrogate function:

$$Q(\boldsymbol{\theta}|\boldsymbol{\theta}^{(m)}) = \sum_{i=1}^N \sum_{k=0}^{\infty} \left[\mathbb{E}(U_k|\mathbf{Y}_i)g(k, \mathbf{z}_i^t \boldsymbol{\theta}_\lambda) + \mathbb{E}(D_k|\mathbf{Y}_i)h(k, \mathbf{z}_i^t \boldsymbol{\theta}_\mu) - \mathbb{E}(T_k|\mathbf{Y}_i) \left(e^{g(k, \mathbf{z}_i^t \boldsymbol{\theta}_\lambda)} + e^{h(k, \mathbf{z}_i^t \boldsymbol{\theta}_\mu)} \right) \right]. \quad (28)$$

Although we cannot usually maximize this surrogate function for all elements of $(\boldsymbol{\theta}_\lambda, \boldsymbol{\theta}_\mu)$ simultaneously, a Newton step is often straightforward to derive.

As an example, consider generalized linear model extension of the simple linear BDP in which

$$\log(\lambda_{ik}) = \log(k) + \mathbf{z}_i^t \boldsymbol{\theta}_\lambda, \quad \text{and} \quad \log(\mu_{ik}) = \log(k) + \mathbf{z}_i^t \boldsymbol{\theta}_\mu. \quad (29)$$

Taking the gradient of the corresponding surrogate function Q with respect to the parameters $\boldsymbol{\theta}_\lambda$ yields

$$\nabla_{\boldsymbol{\theta}_\lambda} Q = \sum_{i=1}^N \mathbb{E}(U|\mathbf{Y}_i)\mathbf{z}_i - e^{\mathbf{z}_i^t \boldsymbol{\theta}_\lambda} \mathbb{E}(T_{\text{particle}}|\mathbf{Y}_i)\mathbf{z}_i \quad (30)$$

and the second differential (Hessian) of Q is

$$\mathbf{d}_{\boldsymbol{\theta}_\lambda}^2 Q = - \sum_{i=1}^N e^{\mathbf{z}_i^t \boldsymbol{\theta}_\lambda} \mathbb{E}(T_{\text{particle}}|\mathbf{Y}_i)\mathbf{z}_i \mathbf{z}_i^t. \quad (31)$$

Combining these, we arrive at the Newton step for the parameter vector $\boldsymbol{\theta}_\lambda$:

$$\boldsymbol{\theta}_\lambda^{(m+1)} = \boldsymbol{\theta}_\lambda^{(m)} - (\mathbf{d}_{\boldsymbol{\theta}_\lambda}^2 Q)^{-1} \nabla_{\boldsymbol{\theta}_\lambda} Q. \quad (32)$$

A similar update can be found for $\boldsymbol{\theta}_\mu$. These updates are examples of the gradient EM algorithm for regression in Markov processes described by Wanek et al (1993) and Lange (1995a). It is worth noting that the Hessian matrix $\mathbf{d}_{\boldsymbol{\theta}_\lambda}^2 Q$ can become ill-conditioned, making it difficult to invert for the Newton step in (32) for some problems. Unfortunately there is no quasi-Newton option since in general $\mathbb{E}(T_{\text{particle}}|\mathbf{Y})e^{\mathbf{z}_i^t \boldsymbol{\theta}_\lambda}$ is unbounded. An alternative to inversion of the Hessian matrix is cyclic coordinate descent in which a Newton step is performed for each coordinate $\boldsymbol{\theta}_j$ individually. This carries the advantage of avoiding matrix inversion, but convergence is slower and the ascent property must be checked at each Newton step.

2.5 Implementation

Before presenting simulation results and our application to microsatellite evolution, we briefly outline some implementation details that ease our subsequent analyses.

2.5.1 E-step acceleration

The E-step in these EM algorithms for BDP estimation usually involves infinite weighted sums of the conditional expectations $\mathbb{E}(U_k|\mathbf{Y})$, $\mathbb{E}(D_k|\mathbf{Y})$, and $\mathbb{E}(T_k|\mathbf{Y})$. For example, when estimating λ in

the simple linear BDP, we must evaluate

$$\mathbb{E}(U|\mathbf{Y}) = \sum_{k=0}^{\infty} \mathbb{E}(U_k|\mathbf{Y}) = \frac{\sum_{k=0}^{\infty} \lambda_k \mathcal{L}^{-1} \left[f_{a,k}(s) f_{k+1,b}(s) \right] (t)}{P_{a,b}(t)}. \quad (33)$$

Fortunately, the conditional expectations of U_k , D_k , and T_k are usually small for $k \ll \min(a, b)$ and $k \gg \max(a, b)$, so it is possible to replace the infinite sum in (33) by a finite one. We find an additional increase in computational efficiency by exchanging the order of Laplace inversion and summation. Then (33) becomes

$$\mathbb{E}(U|\mathbf{Y}) \approx \frac{\mathcal{L}^{-1} \left[\sum_{k=k_{\min}}^{k_{\max}} \lambda_k f_{a,k}(s) f_{k+1,b}(s) \right] (t)}{P_{a,b}(t)}, \quad (34)$$

where we choose k_{\min} to be the largest $k < \min(a, b)$ such that $\lambda_k |f_{a,k}(s) - f_{k+1,a}| < 10^{-8}$ and k_{\max} to be the first $k > \max(a, b)$ such that $\lambda_k |f_{a,k}(s) f_{k+1,b}(s)| < 10^{-8}$. In practice, we rarely need to compute expectations for k less than $\min(a, b) - 10$ or greater than $\max(a, b) + 10$.

2.5.2 Quasi-Newton acceleration of EM iterates

EM algorithms are notorious for slow convergence, especially near optima. When appropriate, we exploit the quasi-Newton acceleration method introduced by Lange (1995b) in our implementations. Other acceleration methods exist, and may give better results, depending on the problem (Lange, 1995a; Louis, 1982; Meilijson, 1989; Jamshidian and Jennrich, 1993). Figure 2 shows the log-likelihood function and iterates for the basic EM and accelerated EM methods in the simple linear model. Since the quasi-Newton acceleration method does not guarantee that the likelihood increases at each step, “step-halving” is occasionally necessary to achieve ascent. Note that this requires likelihood evaluation at least once per iteration. Our approach is advantageous in that we can efficiently calculate this likelihood (transition probability) for any general BDP (Crawford and Suchard, 2011).

[Figure 2 about here.]

2.5.3 Asymptotic variance of EM estimates

Finding the observed information matrix for an EM estimate can be challenging. Louis (1982) gives formulae for the observed information, which Doss et al (2010) use to derive analytic expressions for the observed information for very simple BDPs. However, analytic expressions for the asymptotic variance are generally hard to find for more complicated models. We instead turn to the supplemented EM (SEM) algorithm of Meng and Rubin (1991), which computes the information matrix of the EM estimate of θ after the MLE $\hat{\theta}$ has been found. The observed information is $\mathbf{I}(\hat{\theta}) = -d^2Q(\hat{\theta}|\hat{\theta})(\mathbf{I} - d\mathbf{M}(\hat{\theta}))$, where $\mathbf{M}(\theta)$ is the EM algorithm map such that $\theta^{(m+1)} = \mathbf{M}(\theta^{(m)})$. We numerically approximate the differential $d\mathbf{M}$ at the termination of the EM algorithm.

We note also that since we are able to calculate transition probabilities directly, the observed data log-likelihood is easily computed as

$$\ell(\theta) = \sum_{i=1}^N \log P_{a_i, b_i}(t_i), \tag{35}$$

where $a_i = X_i(0)$ and $b_i = X_i(t_i)$. As an alternative to the approaches outlined above, we can calculate the Hessian using purely numerical techniques. If $\mathbf{H}(\hat{\theta}) = d^2\ell(\hat{\theta})$ is the numerical Hessian evaluated at the estimated value $\hat{\theta}$, then $\hat{\mathbf{I}} \approx -\mathbf{H}(\hat{\theta})$.

3 Results

3.1 Laplace convolution E-step comparison

To illustrate the computational speedup that the Laplace convolution formulae (13) and their acceleration in section 2.5.1 achieve over existing methods, we calculate conditional expectations for various BDP models for performing the E-step and report computing times in Table 2. The first method in the table employs rejection sampling of trajectories where we condition on the starting state, and reject based on the ending state (Bladt and Sorensen, 2005). The second method adapts an endpoint-conditioned simulation algorithm (Hobolth, 2008; Hobolth and Stone, 2009). The third considers naïve time-domain convolution (Equation (12)) using the `integrate` function in `R`. Finally, we compute the same quantities via the Laplace-domain convolution method outlined in section 2.3. In our implementations, we have made every effort to reuse as much shared `R` code as possible,

with the aim of making the routines comparable. We consider four different BDPs. For a simple linear BDP and a linear BDP with immigration, we use the data $\mathbf{Y} = (X(0) = 19, X(2) = 27)$. Under a logistic model, the data are $\mathbf{Y} = (X(0) = 10, X(2) = 16)$, and for the SIS model the data are $\mathbf{Y} = (X(0) = 10, X(2) = 31)$. We list all model parameter values in Table 2.

As seen in Table 2, the Laplace convolution method is often more than 10 times faster than the other methods. In terms of time-performance, the endpoint-conditioned simulation stands as second best, achieving almost comparable speed in the logistic BDP. To interpret this finding, we recall that Hobolth (2008) constructs an endpoint-conditioned simulation for performing the E-step in finite state-space Markov chains. Therefore, to adapt this method we approximate each BDP by a Markov chain with a finite transition rate matrix. To choose the arbitrary dimension of this matrix we truncate the process at the first state $k > \max(a, b)$ such that $P_{a,k}(t) < 10^{-5}$, and the resulting estimates agree substantially with the other methods. We are aware that the size of the rate matrix affects the speed of the simulation routine, so we wish to keep the matrix as small as possible. On the other hand, the matrix must remain large enough to include states that may be visited with high probability in a path from a to b over time t . For the logistic model, such a stringent upper bound lies just above the relatively small carrying capacity. However, endpoint-conditioned simulation completely fails for the SIS model, an issue we discuss later. Finally, and quite naturally, the two convolution methods arrived at nearly the same answer for each model; the difference is largely due to very different sources of numerical error, but at disparate computational costs.

[Table 1 about here.]

3.2 Synthetic examples

To evaluate the performance of our EM algorithms, we simulate discrete observations from several of the BDPs outlined above. For each sample, we draw starting points $X_i(0)$ uniformly from the integers 0 to 20, and times t_i uniformly from 0.1 to 3. We then simulate a trajectory of the BDP and record the state $X_i(t_i)$. For the generalized linear model (GLM), we employ the simple linear parameterization with a log link with $d = 2$ covariates. We specify the covariates $\mathbf{z}_i = (z_{i,1}, z_{i,2})$ as follows: $z_{i,1} \sim N(1, \sigma^2)$, $z_{i,2} \sim N(2, \sigma^2)$ for $i = 1, \dots, N/2$, $z_{i,1} \sim N(2, \sigma^2)$ and $z_{i,2} \sim N(1, \sigma^2)$ for

$i = N/2 + 1, \dots, N$, where $\sigma^2 = 0.1$.

Table 3 reports the number of simulated observations, true parameter values, point-estimates, asymptotic standard error estimates for all model parameters. It is important to note that the MLEs can differ substantially from the parameter values used to perform the simulation, regardless of the algorithm used to find the MLEs. This is due to several factors, including: 1) missing state paths; 2) stochasticity of the BDP generating the state paths; 3) arbitrary choice of starting states $X_i(0)$; and 4) finite sample sizes. Despite these limitations inherent in learning from partially observed stochastic processes, the point-estimates match the true parameter values rather well.

[Table 2 about here.]

3.3 Application to microsatellite evolution

Microsatellites are short tandem repeats of characters in a DNA sequence (Schlötterer, 2000; Ellegren, 2004; Richard et al, 2008). The number of repeated “motifs” in a microsatellite often changes over evolutionary timescales. The molecular mechanism responsible for changes in repeat numbers is known as “polymerase slippage” (Schlötterer, 2000). Several researchers have proposed linear BDPs for use in analyzing evolution of microsatellite repeat numbers (Whittaker et al, 2003; Calabrese and Durrett, 2003; Sainudiin et al, 2004). However, many investigations demonstrate that microsatellite mutability depends on the number of repeats already present, motif size, and motif nucleotide composition (Chakraborty et al, 1997; Eckert and Hile, 2009; Kelkar et al, 2008; Amos, 2010). Exactly how these factors affect addition and deletion rates remains an open question (Bhargava and Fuentes, 2010).

To our knowledge, no previous study formulates or fits a general BDP in which motif size and composition are treated as a covariates in a generalized regression framework, despite the scientific interest in examining such effects on microsatellite evolution. Webster et al (2002) study the evolution of 2467 microsatellites common (orthologous) to both humans and chimpanzees, providing an ideal dataset for studying the influence of repeat number and motif size on addition and deletion rates. For each of these observed microsatellites, Webster et al (2002) record the motif nucleotide pattern and the number of repeats of this motif found in chimpanzees and humans, and estimate a mutability parameter that controls the rate of addition and deletion.

We now present an extended application of our BDP inference technique to chimpanzee-human microsatellite evolution, drawing on the data in Table 6 of the supplementary information in Webster et al (2002). We introduce several novel modeling and inferential techniques relevant to the study of microsatellites, and deduce the effect of motif size and composition on microsatellite addition and deletion rates. While the likelihood takes a slightly more complicated form, our BDP regression technique is straightforward to implement and yields insight into the complicated process of microsatellite evolution.

3.3.1 Evolutionary model

To analyze the data as realizations from a BDP, we must acknowledge the evolutionary relationship between chimpanzees and humans. Suppose the most recent common ancestor of chimpanzees and humans lived at time t in the past, so that an evolutionary time of $2t$ separates contemporary humans and chimpanzees. We note that under mild conditions, general BDPs are reversible Markov chains (Renshaw, 2011). Therefore, assuming stationarity of the chimpanzee microsatellite length distributions, we stand justified in reversing the evolutionary process from the ancestor to chimpanzee, so that for estimation purposes we may regard humans as direct descendants of modern chimpanzees (or vice-versa) over an evolutionary time of $2t$. If C is the number of repeats in a chimpanzee microsatellite and H is the number of repeats in the corresponding human microsatellite, then the likelihood of the observation $\mathbf{Y} = (C, H)$ is

$$\begin{aligned}
 \Pr(\mathbf{Y}) &= \sum_{k=0}^{\infty} \pi_k P_{k,C}(t) P_{k,H}(t) \\
 &= \pi_C \sum_{k=0}^{\infty} P_{C,k}(t) P_{k,H}(t) \\
 &= \pi_C P_{C,H}(2t),
 \end{aligned} \tag{36}$$

where π_k is the equilibrium probability of the microsatellite having k repeats. The second line follows by reversibility and the third by the Chapman-Kolmogorov equality. Therefore, the log-likelihood of the observation \mathbf{Y} is now $\log \pi_C + \ell(\boldsymbol{\theta}; \mathbf{Y})$. Figure 3 shows a schematic representation of this reversibility argument.

[Figure 3 about here.]

3.3.2 BDP rates and equilibrium distribution

The observed data for microsatellite i are $\mathbf{Y}_i = (X_i(0), X_i(1))$, where $X_i(0)$ is the number of repeats observed in chimpanzees, $X_i(1)$ is the number of repeats observed in humans, and the evolutionary time separating humans and chimpanzees is scaled to unity. In addition to the evolutionary relationship explained above, there are other complications: in the Webster et al (2002) dataset, it is evident that microsatellites with small numbers of repeats are not detected. Rose and Falush (1998) argue that there is a minimum number of repeats necessary for microsatellite mutation via polymerase slippage. Sainudiin et al (2004) interpret this finding as justification for truncating the state-space of BDP at x_{\min} , so that $X(\tau) \geq x_{\min}$. To avoid questions of ascertainment bias (see e.g. Vowles and Amos (2006)), and to make our results comparable to those of past researchers, we *define* a microsatellite to be a collection of more than x_{\min} repeated motifs, where x_{\min} is 9 for repeats of size 1, 5 for repeats of size 3 and 4, and 2 for repeats of size 5.

Researchers have also observed that microsatellites do not tend to grow indefinitely (Kruglyak et al, 1998). The maximum number of repeats in the Webster et al dataset is 47. This suggests a finite nonzero equilibrium distribution of microsatellite lengths. To achieve such an equilibrium distribution, we preliminarily view the evolution as a linear BDP with immigration on a state-space that is truncated below x_{\min} . It is reasonable to assume that rates of addition and deletion depend linearly on how many repeats are already present. Then for a microsatellite that currently has k repeats, the birth and death rates are

$$\lambda_k = \begin{cases} k\lambda + \lambda & k \geq x_{\min} \\ 0 & k < x_{\min} \end{cases} \quad \text{and} \quad \mu_k = \begin{cases} k\mu & k > x_{\min} \\ 0 & k \leq x_{\min}. \end{cases} \quad (37)$$

This gives a geometric equilibrium distribution for the number of repeats:

$$\pi_k = \begin{cases} \left(1 - \frac{\lambda}{\mu}\right) \left(\frac{\lambda}{\mu}\right)^{k-x_{\min}-1} & k \geq x_{\min} \\ 0 & k < x_{\min}, \end{cases} \quad (38)$$

when $\lambda < \mu$ (Renshaw, 2011). We choose this simple model so that the BDP has a simple closed-form nonzero equilibrium solution that is easy to incorporate into the log-likelihood. Note that the

constraint $\lambda < \mu$ does not mean that the rate of microsatellite repeat addition is always less than the rate of deletion, since it is possible that $\lambda_k > \mu_k$ for small k . Additionally, $\lambda < \mu$ does not mean that the number of repeats in a microsatellite tends to zero over long evolutionary times — the equilibrium distribution (38) assigns positive probability to all repeat numbers greater than or equal to x_{\min} .

3.3.3 Likelihood and surrogate function

Now we augment the log-likelihood with the log-equilibrium probability of observing $X_i(0)$ chimpanzee repeats

$$F(\boldsymbol{\theta}) = \sum_{i=1}^N \log \pi_{X_i(0)} + \ell(\boldsymbol{\theta}; \mathbf{Y}_i), \quad (39)$$

where $\ell(\boldsymbol{\theta}; \mathbf{Y}_i)$ is equivalent to (10). Including the influence of the equilibrium distribution is similar to imposing a prior distribution on λ and μ . To ensure the existence of the equilibrium distribution (38), we must also incorporate the constraint $\lambda < \mu$. To achieve maximization of the augmented log-likelihood (39) under this constraint, we impose a barrier term of the form $\gamma \log(\mu - \lambda)$. By iteratively maximizing and sending the barrier penalty $\gamma \rightarrow 0$, we can achieve maximization under the inequality constraint. More formally, if we let

$$H(\boldsymbol{\theta}) = \sum_{i=1}^N [\log \pi_{X_i(0)} + \ell(\boldsymbol{\theta}; \mathbf{Y}_i)] + \gamma \log(\mu - \lambda), \quad (40)$$

then

$$\operatorname{argmax}_{\boldsymbol{\theta}} H(\boldsymbol{\theta}) \rightarrow \operatorname{argmax}_{\boldsymbol{\theta}} F(\boldsymbol{\theta}) \quad (41)$$

under the constraint $\lambda < \mu$ as $\gamma \rightarrow 0$.

To incorporate and evaluate the influence of motif size and composition heterogeneity, we now treat λ and μ in the i th observation as functions of the covariate vector \mathbf{z}_i in a general BDP. Suppose microsatellite i has motif size r_i . We code the vectors \mathbf{z}_i as follows:

$$\mathbf{z}_i = \begin{cases} (1, 0, 0, p_a, p_c, p_t)^t & r_i = 1 \\ (1, 1, 0, p_a, p_c, p_t)^t & r_i = 2 \\ (1, 0, 1, p_a, p_c, p_t)^t & r_i \geq 3 \end{cases} \quad (42)$$

where p_x is the proportion of x nucleotides per repeat. We define a single parameter α that controls the difference between λ and μ . Then in the i th microsatellite, the complete model becomes

$$\log(\lambda_{k,i}) = \log(k + 1) + \alpha + \mathbf{z}_i^t \boldsymbol{\theta} \quad \text{and} \quad \log(\mu_{k,i}) = \log(k) + \mathbf{z}_i^t \boldsymbol{\theta}. \quad (43)$$

Therefore $(\alpha, \boldsymbol{\theta})^t$ is the 7×1 vector of unknown parameters. Putting all this together, the surrogate function becomes

$$Q(\boldsymbol{\theta} | \boldsymbol{\theta}^{(m)}) \propto \left(\sum_{i=1}^N X_i(0) \alpha + \log(1 - e^\alpha) + \left[\sum_{k=0}^{\infty} \mathbb{E}(U_k | \mathbf{Y}_i) (\alpha + \mathbf{z}_i^t \boldsymbol{\theta}) + \mathbb{E}(D_k | \mathbf{Y}_i) \mathbf{z}_i^t \boldsymbol{\theta} - \mathbb{E}(T_k | \mathbf{Y}_i) \left((k + 1) e^{\alpha + \mathbf{z}_i^t \boldsymbol{\theta}} + k e^{\mathbf{z}_i^t \boldsymbol{\theta}} \right) \right] \right) + \gamma \log(-\alpha), \quad (44)$$

where $\alpha < 0$ since $\lambda < \mu$, and we send the penalty $\gamma \rightarrow 0$ as the algorithm converges. We use a gradient EM algorithm to find the MLE of $(\alpha, \boldsymbol{\theta})$.

Table 4 reports the parameter estimates, along with asymptotic standard errors. From these results, we infer that motifs of different sizes and composition have different characteristics under our evolutionary model. Specifically, λ and μ are greatest for dinucleotide repeats, as compared to motifs with one or at least three repeats. Motifs consisting mostly of A and T nucleotides also give rise to higher λ and μ . Table 1 shows the estimated λ and μ for each unique motif pattern in the dataset. These conclusions are largely consistent with the descriptive results obtained by Webster et al (2002). Our analysis also provides a natural probabilistic justification for the existence of a finite nonzero equilibrium distribution of microsatellite repeat numbers and a formal statistical framework for deducing the effect of motif size and repeat number on mutation rates.

[Table 3 about here.]

4 Discussion

Application of stochastic models in statistics requires a flexible and general approach to parameter estimation, without which even the most realistic model becomes unappealing to researchers who wish to learn from the data they have collected. Estimation for continuously observed BDPs is

Motif	λ	μ	Motif	λ	μ	Motif	λ	μ	Motif	λ	μ	Motif	λ	μ
A	0.3605	0.3969	AGGA	0.025	0.0276	CCAT	0.0051	0.0056	GGGG	0.0133	0.0147	TCTT	0.0085	0.0094
AAAAC	0.0128	0.0141	AGGG	0.0266	0.0293	CCATC	0.004	0.0044	GGGGG	0.0155	0.0171	TCTTT	0.0096	0.0106
AAAAG	0.0233	0.0257	AGGA	0.0256	0.0282	CCT	0.0031	0.0035	GGG	0.0094	0.0093	TG	0.6094	0.6708
AAAAT	0.0207	0.0228	AGTG	0.0108	0.0119	CCTC	0.0026	0.0028	GGGAA	0.0256	0.0282	TGA	0.0214	0.0235
AAAC	0.0111	0.0123	AGTG	0.0229	0.0252	CCTCC	0.0023	0.0025	GGGGA	0.0269	0.0296	TGAA	0.0216	0.0237
AAACA	0.0128	0.0141	AT	0.5407	0.5952	CCTG	0.0054	0.006	GGT	0.0231	0.0255	TGAGT	0.0212	0.0233
AAACC	0.0074	0.0081	ATA	0.0197	0.0217	CCTT	0.0047	0.0052	GGTA	0.0229	0.0252	TGAT	0.0197	0.0217
AAAG	0.0236	0.026	ATAA	0.0203	0.0224	CCTTT	0.006	0.0066	GGTG	0.0243	0.0268	TGC	0.0085	0.0094
AAAGA	0.0233	0.0257	ATAAA	0.0207	0.0228	CG	0.1835	0.202	GGTGT	0.0222	0.0245	TGCC	0.0054	0.006
AAAGG	0.0244	0.0269	ATAC	0.0102	0.0112	GGC	0.0038	0.0042	GT	0.6094	0.6708	TGG	0.0231	0.0255
AAAT	0.0203	0.0224	ATAG	0.0216	0.0237	GGG	0.0104	0.0114	GTCAC	0.0125	0.0138	TGGA	0.0229	0.0252
AAATA	0.0207	0.0228	ATATG	0.0202	0.0222	CT	0.1362	0.1499	GTAAT	0.0197	0.0217	TGGG	0.0243	0.0268
AAATT	0.0217	0.0239	ATC	0.0079	0.0087	CTC	0.0031	0.0035	GTC	0.0231	0.0255	TGT	0.019	0.0209
AAAT	0.0193	0.0212	ATCT	0.0093	0.0103	CTCC	0.0026	0.0028	GTGAG	0.0239	0.0263	TGTA	0.0197	0.0217
AAC	0.0089	0.0098	ATG	0.0214	0.0235	CTCCT	0.0037	0.0041	GTTG	0.0243	0.0268	TGTC	0.0099	0.0109
AACA	0.0111	0.0123	ATGA	0.0216	0.0237	CTG	0.0085	0.0094	GTT	0.019	0.0209	TGTT	0.018	0.0199
AACAA	0.0128	0.0141	ATGAC	0.0125	0.0138	CTGG	0.0138	0.0151	GTTG	0.0209	0.0231	TGTTT	0.0175	0.0193
AACCA	0.0056	0.0062	ATGAT	0.0202	0.0222	CTT	0.007	0.0077	GTTT	0.018	0.0199	TTA	0.0175	0.0193
AACCT	0.0102	0.0112	ATT	0.0175	0.0193	CTTC	0.0047	0.0052	GTTTA	0.0188	0.0207	TTAA	0.0186	0.0205
AAG	0.0241	0.0265	ATTA	0.0186	0.0205	CTTT	0.0085	0.0094	GTTTG	0.0197	0.0217	TTAAT	0.0179	0.0197
AAGA	0.0236	0.026	ATTC	0.0093	0.0103	CTTTC	0.006	0.0066	GTTTT	0.0175	0.0193	TTAG	0.0197	0.0217
AAGC	0.0118	0.013	ATTG	0.0197	0.0217	CTTTT	0.0096	0.0106	T	0.2524	0.2778	TTAT	0.017	0.0187
AAGG	0.025	0.0276	ATTT	0.017	0.0187	G	0.4579	0.5041	TA	0.5407	0.5952	TTAAT	0.0167	0.0184
AAGGG	0.0256	0.0282	ATTTA	0.0179	0.0197	GA	0.7283	0.8017	TAA	0.0197	0.0217	TTT	0.007	0.0077
AAGT	0.0216	0.0237	ATTTG	0.0103	0.0114	GAA	0.0241	0.0265	TAAA	0.0188	0.0207	TTCA	0.0093	0.0103
AAGTG	0.0228	0.0251	ATTTG	0.0188	0.0207	GAAA	0.0236	0.026	TAAAA	0.0207	0.0228	TTCC	0.0047	0.0052
AAT	0.0197	0.0217	ATTTT	0.0167	0.0184	GAAAA	0.0233	0.0251	TAAAT	0.0193	0.0212	TTCC	0.0085	0.0094
AATA	0.0203	0.0224	C	0.0229	0.0252	GAAAG	0.0244	0.0269	TAAAT	0.0186	0.0212	TTCTC	0.006	0.0066
AATAA	0.0207	0.0228	CA	0.1628	0.1792	GAAAG	0.025	0.0276	TAAATG	0.0202	0.0222	TTCTG	0.0108	0.0119
AATAG	0.0217	0.0239	CAA	0.0089	0.0098	GAAAG	0.0256	0.0282	TAAAT	0.0179	0.0197	TTCTT	0.0096	0.0106
AATG	0.0216	0.0237	CAAA	0.0111	0.0123	GAAT	0.0216	0.0237	TAC	0.0079	0.0087	TTG	0.019	0.0209
AATT	0.0186	0.0205	CAAAA	0.0128	0.0141	GAATT	0.0202	0.0222	TACTA	0.0111	0.0122	TTGAA	0.0202	0.0222
AC	0.1628	0.1792	CAAAAC	0.0074	0.0081	GACAG	0.0141	0.0155	TAGA	0.0216	0.0237	TTGT	0.018	0.0199
ACA	0.0089	0.0098	CAAAAG	0.0134	0.0148	GAG	0.0261	0.0287	TAT	0.0175	0.0193	TTGTT	0.0175	0.0193
ACAA	0.0111	0.0123	CAC	0.0035	0.0039	GAGAA	0.0244	0.0269	TATC	0.0093	0.0103	TTTA	0.017	0.0187
ACAAA	0.0128	0.0141	CACAC	0.0042	0.0047	GAGG	0.0266	0.0293	TATG	0.0197	0.0217	TTTAA	0.0179	0.0197
ACAG	0.0118	0.013	CACC	0.0028	0.0031	GAT	0.0214	0.0235	TAT	0.017	0.0187	TTTAG	0.0188	0.0207
ACC	0.0035	0.0039	CACCA	0.0042	0.0047	GATA	0.0216	0.0237	TATTT	0.0167	0.0184	TTTAT	0.0167	0.0184
ACCA	0.0056	0.0062	CAG	0.0096	0.0106	GATT	0.0197	0.0217	TC	0.1362	0.1499	TTTT	0.0085	0.0094
AG	0.7283	0.8017	CAGA	0.0118	0.013	GCACA	0.0077	0.0085	TCAAA	0.0119	0.0131	TTTTCC	0.006	0.0066
AGAA	0.0236	0.026	CAGAG	0.0141	0.0155	GCCGC	0.0047	0.0051	TCAC	0.0051	0.0056	TTTTCT	0.0096	0.0106
AGAAA	0.0233	0.0257	CAGG	0.0126	0.0138	GCT	0.0085	0.0094	TCAT	0.0093	0.0103	TTTTG	0.018	0.0199
AGAC	0.0118	0.013	CAT	0.0079	0.0087	GCTGT	0.0122	0.0134	TCATT	0.0103	0.0114	TTTTGG	0.0197	0.0217
AGAGG	0.0256	0.0282	CATA	0.0102	0.0112	GGA	0.0261	0.0287	TCC	0.0031	0.0035	TTTTG	0.0175	0.0193
AGAT	0.0216	0.0237	CATC	0.0051	0.0056	GGAA	0.025	0.0276	TCCA	0.0051	0.0056	TTTTA	0.0167	0.0184
AGCAA	0.0134	0.0148	CATG	0.0108	0.0119	GGAAG	0.0256	0.0282	TCCC	0.0047	0.0052	TTTTC	0.0096	0.0106
AGCCA	0.0059	0.0065	CATT	0.0093	0.0103	GGAG	0.0266	0.0293	TCCCT	0.0047	0.0052	TTTTTG	0.0175	0.0193
AGCTC	0.0072	0.0079	CCA	0.0035	0.0039	GGAGG	0.0269	0.0296	TCT	0.007	0.0077			
AGG	0.0261	0.0287	CCAAC	0.0042	0.0047	GGCCA	0.0081	0.0089	TCTG	0.0089	0.0099			

Table 1: Estimates of birth and death rates for each unique motif in the human-chimpanzee dataset of Webster et al (2002). Under our model of microsatellite mutation, AT repeats have the highest associated rates and CTCG repeats have the lowest.

straightforward and well-established. For partially observed BDPs, our approach is unique because it requires only two simple ingredients: the functional form of the birth and death rates $\lambda_k(\boldsymbol{\theta})$ and $\mu_k(\boldsymbol{\theta})$ for all k , and an exact or approximate M-step. A third ingredient is optional: the Hessian of the surrogate function is useful when asymptotic standard errors are desired. However, this matrix can often be approximated numerically upon convergence of the EM algorithm, since the observed-data likelihood is available numerically via (35). With these ingredients in hand, even elusive general BDPs become tractable.

In previous work on estimation for BDPs, completion of the E-step typically relies on time-domain numerical integration or simulation of BDP trajectories. As we show in Table 2, both rejection sampling and endpoint-conditioned simulation can occasionally perform satisfactorily, especially in comparison to time-domain convolution. However, endpoint-conditioning is designed for finite state-space Markov chains, and it relies on a matrix eigendecomposition to calculate transition probabilities. As we show for the SIS model, this matrix becomes nearly singular, causing the simulation algorithm to fail, even when we choose parameter values that are not biologically unreasonable. The Laplace convolution in the E-step of our algorithm is more generic with equivalent or better performance. For this reason, a variation on our Laplace convolution method for computing the E-step may offer further use in estimation for non-BDP finite Markov chains as well, such as nucleotide or codon substitution models. For some linear BDPs, the availability of a generating function furnishes analytic E- and M-steps yielding very fast parameter updates in closed-form (Doss et al, 2010). For some models, these tools provide the asymptotic variance of the MLE in closed-form. However, for the majority of BDPs, we must return to the Laplace convolution method outlined in this paper.

If one cannot find analytic parameter updates in the M-step, several options remain available. With a minorizing function as in section 2.4.2, an EM-MM algorithm is viable. Further, one or more numerical Newton steps offers an alternative, as in sections 2.4.3 and 2.4.5. One may employ other gradient-based methods as well. Although the MM update derived for the BDP with immigration (section 2.4.2) is appealing in its simplicity, multiple minorizations of the likelihood can result in very slow convergence, since the surrogate function lies far from the true likelihood for most values of $\boldsymbol{\theta}$. In addition, Newton steps that require matrix inversion may suffer since the Hessian of the surrogate can become ill-conditioned.

Even with the substantial speedup offered by our Laplace convolution method for performing the E-step and quasi-Newton acceleration of the EM iterates, our algorithms can move slowly toward the MLE. Here, naïve numerical optimization of the incomplete data likelihood can sometimes run computationally faster. However, such techniques perform very poorly when the number of parameters increases and they often require specification of tuning constants in order to reach the global optimum. For BDP estimation problems, EM algorithms offer several other advantages over naïve numerical optimization, and these benefits are especially stark when the M-step is available in closed-form. First, when the log-likelihood is locally convex, the EM algorithm is robust with respect to the initial parameter values near the maximum, and EM algorithms generally do not need tuning parameters. Further, the ascent property ensures the iterates will approach a maximum. Perhaps the most important reason to consider EM algorithms is that they can accommodate high-dimensional parameter spaces without substantially increasing the computational complexity of the algorithm. This is especially useful in models with many unknown parameters when performing regression with covariates (section 2.4.5), or our microsatellite example. We also note the potential for substantial computational speedup by parallelizing the E-step. When discrete observations from a BDP are independent, the E-step may be performed in parallel for every observation. For example, $\mathbb{E}(U|\mathbf{Y}_i)$ can be computed simultaneously for $i = 1, \dots, N$. When speed is an issue, graphics processing units may prove useful in reducing the computational cost of EM algorithms (Zhou et al, 2010).

With regard to our example, we present a novel way of studying the evolution of microsatellite repeats using a generalized linear model. Previous efforts often ignore the evolutionary relationship between organisms, use incomplete or equilibrium models of repeat numbers, or fit separate models to motifs of different sizes. We treat motif size as a categorical variable and incorporate motif nucleotide composition, allowing us to fit a single model to all the microsatellite observations simultaneously. Though our rate specification (37) and resulting equilibrium distribution (38) are intended to be somewhat simplistic, more sophisticated models that are informed by biological considerations may be fruitful. The only requirement in our setup is that the gradient and Hessian of λ_k , μ_k , and π_k be available for any repeat number k . Although our microsatellite example is limited in scope, it is easy to imagine a more comprehensive study. For example, incorporating more sophisticated motif nucleotide composition covariates and location of the microsatellite on the

chromosome might provide additional insight into the evolutionary process. Our EM framework is nearly ideal for these types of studies, since the number of unknown parameters does not substantially increase the computational burden of the M-step, and the E-step is completely unaffected by the number of parameters.

Interestingly, we attempted to use the generic nonlinear regression R function `nlm` to validate the MLEs obtained by our EM algorithm for the microsatellite evolution problem, starting at a variety of initial values, including the MLE found by our EM algorithm. This naïve optimizer failed to converge in every case. We speculate that this is because the small numerical errors in the likelihood evaluation have similar order of magnitude as the curvature of the likelihood function near the maximum. Our EM algorithms take advantage of analytic derivatives of the surrogate function instead of the likelihood, and hence are less susceptible to small errors in the numerical gradient.

5 Conclusion

Previous work on parameter estimation in BDPs almost exclusively confines itself to inference of birth and death rates under the simple linear model. To rectify this situation, we present a flexible and robust framework for deriving EM algorithms to estimate parameters in any general BDP, using discrete observations. We hope that this contribution encourages development of more sophisticated and realistic birth-death models in applied work, since researchers can now estimate parameters using more complicated rate structures, even when the data are observed at discrete times.

Software

A software implementation of the EM algorithms for general BDPs used in this paper is currently available from FWC (by request through the Editor to maintain reviewer anonymity) and will be deposited in CRAN (2011) before publication.

Acknowledgements

We are grateful to Kenneth Lange, Hua Zhou, and Gabriela Cybis for helpful comments. This work was supported by NIH grants R01 GM086887, HG006139, T32GM008185, and NSF grant DMS-0856099.

References

- Abate J, Whitt W (1992) Numerical inversion of probability generating functions. *Oper Res Lett* 12:245–251
- Abate J, Whitt W (1995) Numerical inversion of Laplace transforms of probability distributions. *ORS J Comput* 7(1):36–43
- Abate J, Whitt W (1999) Computing Laplace transforms for numerical inversion via continued fractions. *INFORMS J Comput* 11(4):394–405
- Amos W (2010) Mutation biases and mutation rate variation around very short human microsatellites revealed by human-chimpanzee-orangutan genomic sequence alignments. *J Mol Evol* 71:192–201
- Andersson H, Britton T (2000) *Stochastic Epidemic Models and their Statistical Analysis*. Lecture notes in statistics, Springer New York
- Anscombe FJ (1953) Sequential estimation. *J Roy Stat Soc B* 15(1):1–29
- Bailey NTJ (1964) *The Elements of Stochastic Processes with Applications to the Natural Sciences*. Wiley New York
- Bankier JD, Leighton W (1942) Numerical continued fractions. *Am J Math* 64(1):653–668
- Bhargava A, Fuentes F (2010) Mutational dynamics of microsatellites. *Mol Biotechnol* 44:250–266
- Bladt M, Sorensen M (2005) Statistical inference for discretely observed Markov jump processes. *J Roy Stat Soc B Met* 67(3):395–410
- Blanch G (1964) Numerical evaluation of continued fractions. *SIAM Rev* 6(4):383–421

- Calabrese P, Durrett R (2003) Dinucleotide repeats in the drosophila and human genomes have complex, length-dependent mutation processes. *Mol Biol Evol* 20(5):715–725
- Chakraborty R, Kimmel M, Stivers D, Davison L, Deka R (1997) Relative mutation rates at di-, tri-, and tetranucleotide microsatellite loci. *P Natl Acad Sci USA* 94(3):1041–1046
- Cotton JA, Page RDM (2005) Rates and patterns of gene duplication and loss in the human genome. *Proc R Soc B* 272:277–283
- CRAN (2011) The comprehensive R archive network. URL <http://cran.r-project.org>
- Craviotto C, Jones WB, Thron WJ (1993) A survey of truncation error analysis for Padé and continued fraction approximants. *Acta Appl Math* 33:211–272
- Crawford FW, Suchard MA (2011) Transition probabilities for general birth-death processes with applications in ecology, genetics, and evolution. *J Math Biol*
- Cuyt A, Petersen V, Verdonk B, Waadeland H, Jones W (2008) *Handbook of Continued Fractions for Special Functions*. Springer Berlin / Heidelberg
- Darwin JH (1956) The behaviour of an estimator for a simple birth and death process. *Biometrika* 43(1):23–31
- Dauxois J (2004) Bayesian inference for linear growth birth and death processes. *J Stat Plan Infer* 121(1):1–19
- Dempster AP, Laird NM, Rubin DB (1977) Maximum likelihood from incomplete data via the EM algorithm. *J Roy Stat Soc B* 39(1):1–38
- Demuth JP, Bie TD, Stajich JE, Cristianini N, Hahn MW (2006) The evolution of mammalian gene families. *PLoS ONE* 1(1):e85
- Doss CR, Suchard MA, Holmes I, Kato-Maeda M, Minin VN (2010) Great expectations: EM algorithms for discretely observed linear birth-death-immigration processes. *ArXiv e-prints* 1009.0893
- Eckert KA, Hile SE (2009) Every microsatellite is different: Intrinsic DNA features dictate mutagenesis of common microsatellites present in the human genome. *Mol Carcinogen* 48(4):379–388

- Ellegren H (2004) Microsatellites: simple sequences with complex evolution. *Nat Rev Genet* 5(6):435–445
- Feller W (1971) *An Introduction to Probability Theory and its Applications*. Wiley series in probability and mathematical statistics, Wiley New York
- Hobolth A (2008) A Markov chain Monte Carlo expectation maximization algorithm for statistical analysis of DNA sequence evolution with neighbor-dependent substitution rates. *J Comput Graph Stat* 17(1):1–25
- Hobolth A, Jensen JL (2005) Statistical inference in evolutionary models of DNA sequences via the EM algorithm. *Stat Appl Genet Mol* 4(1):1–19
- Hobolth A, Stone EA (2009) Simulation from endpoint-conditioned, continuous-time Markov chains on a finite state space, with applications to molecular evolution. *Ann Appl Stat* 3(3):1024–1231
- Holmes I, Bruno WJ (2001) Evolutionary HMMs: a Bayesian approach to multiple alignment. *Bioinformatics* 17(9):803–820
- Holmes I, Rubin G (2002) An expectation maximization algorithm for training hidden substitution models. *J Mol Biol* 317(5):753–764
- Jamshidian M, Jennrich RI (1993) Conjugate gradient acceleration of the EM algorithm. *J Am Stat Assoc* 88(421):221–228
- Kalbfleisch JD, Lawless JF (1985) The analysis of panel data under a Markov assumption. *J Am Stat Assoc* 80(392):863–871
- Karlin S, McGregor J (1957) The differential equations of birth-and-death processes, and the Stieltjes moment problem. *Trans Am Math Soc* 85(2):589–646
- Keiding N (1974) Estimation in the birth process. *Biometrika* 61(1)
- Keiding N (1975) Maximum likelihood estimation in the birth-and-death process. *Ann Stat* 3(2):363–372
- Kelkar YD, Tyekucheva S, Chiaromonte F, Makova KD (2008) The genome-wide determinants of human and chimpanzee microsatellite evolution. *Genome Res* 18(1):30–38

- Kingman JFC (1982) On the genealogy of large populations. *J Appl Probab* 19:27–43
- Krone SM, Neuhauser C (1997) Ancestral processes with selection. *Theor Popul Biol* 51:210–237
- Kruglyak S, Durrett RT, Schug MD, Aquadro CF (1998) Equilibrium distributions of microsatellite repeat length resulting from a balance between slippage events and point mutations. *P Natl Acad Sci USA* 95(18):10,774–10,778
- Lange K (1995a) A gradient algorithm locally equivalent to the EM algorithm. *J Roy Stat Soc B Met* 57(2):425–437
- Lange K (1995b) A quasi-Newton acceleration of the EM algorithm. *Stat Sinica* 5:1–18
- Lange K (2010) *Numerical Analysis for Statisticians (Statistics and Computing)*, 2nd edn. Springer New York
- Lentz WJ (1976) Generating Bessel functions in Mie scattering calculations using continued fractions. *Appl Opt* 15(3):668–671
- Liu H, Beckett LA, DeNardo GL (2007) On the analysis of count data of birth-and-death process type: with application to molecularly targeted cancer therapy. *Statist Med* 26:11141135
- Lorentzen L, Waadeland H (1992) *Continued Fractions with Applications*. North-Holland, Amsterdam
- Louis TA (1982) Finding the observed information matrix when using the EM algorithm. *J Roy Stat Soc B Met* 44(2):226–233
- Meilijson I (1989) A fast improvement to the EM algorithm on its own terms. *J Roy Stat Soc B Met* 51(1):127–138
- Meng XL, Rubin DB (1991) Using EM to obtain asymptotic variance-covariance matrices: The SEM algorithm. *J Am Stat Assoc* 86(416):899–909
- Metzner P, Dittmer E, Jahnke T, Schütte C (2007) Generator estimation of Markov jump processes. *J Comput Phys* 227:353–375
- Moran PAP (1951) Estimation methods for evolutive processes. *J Roy Stat Soc B Met* 13(1):141–146

- Moran PAP (1953) The estimation of the parameters of a birth and death process. *J Roy Stat Soc B Met* 15(2):241–245
- Moran PAP (1958) Random processes in genetics. *Math Proc Cambridge* 54(01):60–71
- Murphy JA, O’Donohoe MR (1975) Some properties of continued fractions with applications in Markov processes. *IMA J Appl Math* 16(1):57–71
- Murray J (2002) *Mathematical Biology: An Introduction, Interdisciplinary applied mathematics*, vol 1. Springer, New York
- Nee S (2006) Birth-death models in macroevolution. *Annu Rev Ecol Evol S* 37:1–17
- Nee S, May RM, Harvey PH (1994) The reconstructed evolutionary process. *Philos T Roy Soc B* 344(1309):305–311
- Novozhilov AS, Karev GP, Koonin EV (2006) Biological applications of the theory of birth-and-death processes. *Brief Bioinform* 7(1):70–85
- Parthasarathy PR, Sudhesh R (2006) Exact transient solution of a state-dependent birth-death process. *J Appl Math Stoch Anal* 82(6):1–16
- Press WH (2007) *Numerical Recipes: the Art of Scientific Computing*. Cambridge University Press
New York
- Renshaw E (2011) *Stochastic Population Processes: Analysis, Approximations, Simulations*. Oxford University Press
- Reynolds JF (1973) On estimating the parameters of a birth-death process. *Aust J Stat* 15(1):35–43
- Richard GF, Kerrest A, Dujon B (2008) Comparative genomics and molecular dynamics of DNA repeats in eukaryotes. *Microbiol Mol Biol Rev* 72(4):686–727
- Rose O, Falush D (1998) A threshold size for microsatellite expansion. *Mol Biol Evol* 15(5):613–615
- Rosenberg NA, Tsolaki AG, Tanaka MM (2003) Estimating change rates of genetic markers using serial samples: applications to the transposon IS6110 in *Mycobacterium tuberculosis*. *Theor Popul Biol* 63(4):347–363

- Sainudiin R, Durrett RT, Aquadro CF, Nielsen R (2004) Microsatellite mutation models. *Genetics* 168(1):383–395
- Schlötterer C (2000) Evolutionary dynamics of microsatellite DNA. *Chromosoma* 109:365–371
- Tan WY, Piantadosi S (1991) On stochastic growth processes with application to stochastic logistic growth. *Stat Sinica* 1:527–540
- Thompson IJ, Barnett AR (1986) Coulomb and Bessel functions of complex arguments and order. *J Comput Phys* 64:490–509
- Thorne J, Kishino H, Felsenstein J (1991) An evolutionary model for maximum likelihood alignment of DNA sequences. *J Mol Evol* 33(2):114–124
- Vowles EJ, Amos W (2006) Quantifying ascertainment bias and species-specific length differences in human and chimpanzee microsatellites using genome sequences. *Mol Biol Evol* 23(3):598–607
- Wall HS (1948) *Analytic Theory of Continued Fractions*. University Series in Higher Mathematics, D. Van Nostrand Company, Inc. New York
- Wanek LA, Goradia TM, Elashoff RM, Morton DL (1993) Multi-stage Markov analysis of progressive disease applied to melanoma. *Biom J* 35(8):967–983
- Webster MT, Smith NGC, Ellegren H (2002) Microsatellite evolution inferred from human and chimpanzee genomic sequence alignments. *P Natl Acad Sci USA* 99(13):8748–8753
- Whittaker JC, Harbord RM, Boxall N, Mackay I, Dawson G, Sibly RM (2003) Likelihood-based estimation of microsatellite mutation rates. *Genetics* 164(2):781–787
- Wolff RW (1965) Problems of statistical inference for birth and death queuing models. *Oper Res* 13(3):343–357
- Zhou H, Lange K, Suchard M (2010) Graphics processing units and high-dimensional optimization. *Stat Sci* 25(3):311–324




Figure 1: A sample path from a birth-death process (BDP) $X(\tau)$. The process starts at state $X(0) = 1$ and is at state $X(t) = 4$ at time t . At right are schematic representations of the time spent in each state T_k , the number of up steps U_k , and the number of down steps D_k . These quantities are the sufficient statistics for estimators of rate parameters in general birth-death processes.




Figure 2: Effect of quasi-Newton acceleration on iterates of the expectation-maximization (EM) algorithm for a simple linear BDP with birth rate λ and death rate μ . Contour lines sketch the log-likelihood from $N = 50$ discrete samples. Iterates are shown with the “+” symbol. On the left, ordinary EM iterates converge very slowly in the neighborhood of the maximum, for a total of 36 iterations. On the right, EM iterates using quasi-Newton acceleration make large jumps and converge rapidly in 15 iterations.




Figure 3: Reversibility of the BDP implies that the evolutionary relationship between contemporary chimpanzees and the most recent common ancestor can be inverted. On the left, the most recent common ancestor of chimpanzees and humans lived at time T in the past. At a certain locus, chimpanzees have a microsatellite consisting of 2 repeats of the motif AAC , and at an orthologous locus, humans have 3 repeats of the motif. The number of repeats in the ancestor is unknown. On the right, using a probabilistic justification explained in the text, we may interpret the evolutionary relationship between chimpanzees and humans as unidirectional, while “integrating out” the number of repeats at the ancestral locus.

Model	Quantity	Rejection sampling	Endpoint-conditioned simulation	Time-convolution	Laplace-convolution
Simple linear (2.4.1) $\lambda = 0.5, \mu = 0.3$	$\mathbb{E}(U \mathbf{Y})$	1.449	0.741	19.606	0.084
	$\mathbb{E}(D \mathbf{Y})$	1.375	0.743	21.224	0.086
	$\mathbb{E}(T_{\text{particle}} \mathbf{Y})$	1.432	0.636	16.488	0.087
Immigration (2.4.2) $\lambda = 0.5, \nu = 0.2$ $\mu = 0.3$	$\sum_k p_k \mathbb{E}(U \mathbf{Y})$	1.192	0.697	15.669	0.085
	$\mathbb{E}(D \mathbf{Y})$	1.324	0.689	21.058	0.086
	$\mathbb{E}(T_{\text{particle}} \mathbf{Y})$	1.319	0.703	14.961	0.089
Logistic (2.4.3) $\lambda = 0.5, \alpha = 0.2$ $\mu = 0.3$	$\mathbb{E}(U \mathbf{Y})$	50.810	0.162	21.907	0.102
	$\mathbb{E}(D \mathbf{Y})$	56.957	0.180	20.851	0.102
	$\sum_k k^2 e^{-k\alpha} \mathbb{E}(T_k \mathbf{Y})$	50.764	0.168	21.623	0.107
SIS (2.4.4) $\beta = 0.5, \gamma = 0.3$	$\mathbb{E}(U \mathbf{Y})$	7.880	*	5.295	0.059
	$\mathbb{E}(D \mathbf{Y})$	8.886	*	2.749	0.048
	$\sum_k (N - k) k \mathbb{E}(T_k \mathbf{Y})$	8.456	*	4.269	0.053

Table 2: Compute times (seconds) to perform various E-steps for four different BDP models. We report text section numbers in which the models are described in parentheses. For each E-step, we consider several methods. In all cases, the Laplace method takes substantially less time. The endpoint-conditioned simulation method fails for the susceptible-infectious-susceptible (SIS) infectious disease model.

Model	Parameter	True	Estimate	SE
Simple linear ($N = 500$) (2.4.1)	λ	0.5	0.5039	0.0269
	μ	0.2	0.1981	0.0254
Immigration ($N = 800$) (2.4.2)	λ	0.2	0.2182	0.0129
	ν	0.1	0.1016	0.0213
	μ	0.25	0.2488	0.0231
Logistic ($N = 1500$) (2.4.3)	λ	0.3	0.2917	0.0035
	α	0.5	0.4942	0.0397
	μ	0.05	0.0456	0.0633
SIS ($N = 1000$) (2.4.4)	β	0.1	0.1025	0.0048
	γ	2.0	2.1374	0.0367
GLM ($N = 1000$) (2.4.5)	$\theta_{\lambda,1}$	0.25	0.2585	0.0393
	$\theta_{\lambda,2}$	0.1	0.1143	0.0402
	$\theta_{\mu,1}$	0.2	0.1973	0.0457
	$\theta_{\mu,2}$	0.05	0.0877	0.0457

Table 3: Point-estimates and their standard errors (SE) for simulated observations under various BDPs. We report the text section describing each of the models in parentheses. The method for generating the rates in the generalized linear model (GLM) BDP is described in the text.

Parameter	Covariate	Estimate	SE
θ_1	Intercept	-1.3105	0.1236
θ_2	$r_i = 2$	0.2854	0.0983
θ_3	$r_i \geq 3$	-1.5405	0.1079
θ_4	p_a	0.2207	0.1725
θ_5	p_c	-0.3822	0.0577
θ_6	p_t	0.0477	0.0002
α	birth	-0.0889	0.0039

Table 4: Maximum likelihood estimates of parameters in the microsatellite model and their asymptotic standard errors. The first three elements of θ correspond to the motif size r_i , and the last three correspond to the motif nucleotide composition. The parameter α controls the difference between the birth and death rates. The i th microsatellite birth rate is then $\lambda = \exp(\alpha + \mathbf{z}_i^t \theta)$ and the death rate is $\mu = \exp(z_i^t \theta)$. Estimated birth and death rates are higher for dinucleotide repeats than for mononucleotide repeats or microsatellites whose motifs have 3, 4, or 5 nucleotides. Microsatellites whose motif consists, for example, of A nucleotides have higher birth and death rates compared to G nucleotides.