

Unbiased likelihood-based estimation of Wright-Fisher diffusion processes

Celia García-Pareja and Fabio Nobile

Abstract In this paper we propose an unbiased Monte Carlo maximum likelihood estimator for discretely observed Wright-Fisher diffusions. Our approach is based on exact simulation techniques that are of special interest for diffusion processes defined on a bounded domain, where numerical methods typically fail to remain within the required boundaries. We start by building unbiased maximum likelihood estimators for scalar diffusions and later present an extension to the multidimensional case. Consistency results of our proposed estimator are also presented and the performance of our method is illustrated through a numerical example.

1 Introduction

Problems that can be modeled as continuous-time phenomena are ubiquitous in the sciences, and have a wide range of applicability. In this context, diffusion models have been extensively used in numerous areas, namely, in economics, physics, the life sciences and engineering. However, inference on diffusion models for discretely observed data is challenging because a closed-form expression of the transition density of the process, and thus, of the likelihood, is often unavailable, see [1] for a survey on available inference methods for diffusion models, or [2] for a more recent review.

The object of study in this paper is the Wright-Fisher diffusion process $X = \{X_t, 0 \leq t \leq T\}$ determined as the weakly unique solution of the scalar stochastic differential equation (SDE)

Celia García-Pareja, Fabio Nobile
Institute of Mathematics, EPF Lausanne, Station 8, CH-1015 Lausanne, Switzerland

Celia García-Pareja
Institute of Bioengineering, EPF Lausanne, Station 15, CH-1015 Lausanne, Switzerland, e-mail: celia.garciapareja@epfl.ch

Fabio Nobile, e-mail: fabio.nobile@epfl.ch

$$dX_t = \gamma(X_t, \vartheta)dt + \sqrt{X_t(1 - X_t)}dB_t, \quad t \in [0, T], \quad X_0 = x_0 \in [0, 1], \quad (1)$$

where $\gamma(\cdot, \vartheta) : [0, 1] \mapsto \mathbb{R}$ is the drift function that depends on an unknown scalar parameter $\vartheta \in \Theta$ and satisfies the usual regularity conditions (it is locally bounded and with a linear growth bound), and B_t denotes a one-dimensional Brownian motion.

In the following, we will refer to the transition density function of X as

$$p_\vartheta(x, y; t) = P_\vartheta(X_t \in dy | X_0 = x) / dy, \quad t > 0, \quad x, y \in [0, 1] \quad (2)$$

where t refers to the time increment between the instances x and y .

The Wright-Fisher diffusion process has been widely used in population genetics modeling, where it describes the evolution of the frequency of different genetic variants over time. The advent of whole genome sequencing, and the subsequent increased availability of observed genetic data, calls for the development of suitable inference methods.

In its simplest case, the Wright-Fisher diffusion model considers two genetic variants, namely, type a and type A . Thus, X describes the frequency of variant a over time $t \in [0, T]$, and the drift function $\gamma(X_t, \vartheta)$ incorporates different evolutionary forces, e.g., mutation or natural selection. In this paper, we consider drift functions $\gamma(x; \vartheta)$ admitting the general form

$$\gamma(x; \vartheta) = \alpha(x) + x(1 - x)\eta(x; \vartheta), \quad \vartheta \in \Theta,$$

where $\alpha(x) = \frac{1}{2}(\theta_a - (\theta_a + \theta_A)x)$ and $\eta(x; \vartheta)$ is continuously differentiable in $(0, 1)$.

In a population genetics context, $\alpha(x)$ describes the recurrent mutant behaviour of the process, with θ_a and θ_A strictly positive and denoting the rates of mutation towards type a and type A , respectively. Furthermore, $\eta(x; \vartheta)$ describes the natural selection pattern accounting for possible fitness differences between the types.

In case $\eta(x; \vartheta) \equiv 0$, the population is said to follow neutral evolution, where none of the types displays a selective advantage, that is, all types are equally likely to reproduce.

A question of interest in applications is whether a newly appeared mutation is likely to sweep over the population. Think, for instance, in seasonal flu vaccine planning, where interest lies in predicting the most prevalent genetic variant a year ahead of time, so that effective vaccines can be designed before the next season, see, e.g., [3], [4], [5], [6]. In order to study deviations from neutrality, accurate inferences on ϑ are needed.

Our estimating approach is based on the Simultaneous Acceptance Method (SAM) presented in [7]. In short, our method provides an unbiased estimator of the likelihood of model (1), based on Monte Carlo samples drawn from the recently developed exact simulation algorithms for Wright-Fisher diffusions, see [8], [9]. Uniform convergence results of our Monte Carlo estimator to the true likelihood function ensure convergence of the maximizers and are used to show consistency of our maximum likelihood estimator (MLE).

The main advantage of our method is that it is based on samples drawn from the exact probability distribution of model (1). Likelihood methods for discretely

observed diffusions are usually approximate because of the unavailability of the transition density function. However, in the context of diffusions with bounded state space such as the Wright-Fisher diffusion, approximate numerical methods fail to stay within the state space and perform poorly near the boundaries, see [10]. Thus, estimates based on numerical approximations might yield biologically meaningless and unrealistic results.

Inference methods for SDEs that use approximate likelihood-based approaches include computationally intensive imputation methods [11], methods which approximate analytically the transition density [12] or methods that use closed-form density expansions [13]. Moreover, approximate Monte Carlo maximum likelihood approaches have been proposed, see [14]. Methods based on non-likelihood approaches include estimating functions [15], efficient method of moments [16] and indirect inference [17]. More specifically related to our work are other approaches on inference for Wright-Fisher diffusions, see [18], [19].

It is also worth mentioning an entire line of work on inference for SDEs with continuous data, see, e.g., [20]. However, in this paper we focus our attention in discretely observed diffusions, where observations are assumed to be given at a given collection of time points and without error.

The rest of this paper is structured as follows. In Section 2 we present our general estimation approach, followed by a detailed construction of our proposed Monte Carlo estimator, which we expose in Section 3. Section 4 is devoted to show the performance of our proposed estimator in an illustrative numerical example, and in Section 5 we show the applicability of our method to the multivariate case. Finally, Section 6 concludes the paper.

2 Estimation approach

Before proceeding with the basic formulation of our inferential problem, we fix the following preliminary notation. Let $C \equiv C([0, T], [0, 1])$ be the set of continuous mappings from $[0, T]$ to $[0, 1]$ and let us denote ω a typical element of C . Consider the cylinder σ -algebra $\Sigma = \sigma(\{X_t, 0 \leq t \leq T\})$, where X_t are the coordinate mappings $X_t : C \rightarrow [0, 1]$ such that for every $t \in [0, T]$, $X_t(\omega) = \omega(t)$. In what follows, we also assume that Θ is a compact subset of \mathbb{R}^d that contains the MLE of ϑ .

Given n discrete observations of a path from process (1) observed without error at times $0 < t_1 < \dots < t_n = T$, $\{x_{t_i}\}_{i=1}^n$, the aim is to propose an unbiased MLE, ϑ^n , for the selection parameter ϑ . Let us denote $x_{t_i} \equiv x^i$. Following the definition in (2), the likelihood of the process reads

$$\mathcal{L}^n(\vartheta) = \prod_{i=1}^n L^i(\vartheta) = \prod_{i=1}^n p_{\vartheta}(x^{i-1}, x^i; \Delta t_i),$$

where $\Delta t_i = t_i - t_{i-1}$, for $i = 1, \dots, n$, and $t_0 = 0$.

Because of the Markov property of the Wright-Fisher diffusion, each contribution $L^i(\vartheta)$ can be estimated independently. Let $x = x^{i-1}$ and $y = x^i$ for $i = 1, \dots, n$. In what follows, we write $L(\vartheta)$ to refer to any contribution $L^i(\vartheta)$.

In this paper we exploit the exact simulation technique proposed in [8] to sample paths of X solution of (1). In order for our estimating strategy to be applicable, we require the following assumptions.

1. The drift function $\gamma(x; \vartheta)$ is such that

$$\gamma(x; \vartheta) = \alpha(x) + x(1-x)\eta(x; \vartheta),$$

with $\alpha(x) = \frac{1}{2}(\theta_a - (\theta_a + \theta_A)x)$ for given $\theta_a, \theta_A > 0$ and $\eta(\cdot; \vartheta)$ is continuously differentiable in $(0,1)$, for any $\vartheta \in \Theta$.

2. The function $\eta(x; \cdot)$ is continuous in Θ , for any $x \in [0, 1]$.
3. The function $\phi : [0, 1] \times \Theta \rightarrow \mathbb{R}$,

$$\phi(x; \vartheta) = \frac{1}{2}[x(1-x)(\eta^2(x; \vartheta) + \eta'(x; \vartheta)) + 2\eta(x; \vartheta)\alpha(x)],$$

is continuous in $[0, 1]$ and thus bounded, with $\phi^-(\vartheta) \leq \phi(x; \vartheta) \leq \phi^+(\vartheta)$, $\forall \vartheta \in \Theta$ and $x \in [0, 1]$, with $\phi^-(\cdot)$ and $\phi^+(\cdot)$ continuous in Θ .

4. The function $A : [0, 1] \times \Theta \rightarrow \mathbb{R}$,

$$A(x; \vartheta) = \int_0^x \eta(z; \vartheta) dz$$

is continuous in $[0, 1]$ and thus bounded, with $A^+(\vartheta) \geq A(x; \vartheta)$, $\forall \vartheta \in \Theta$ and $x \in [0, 1]$, with $A(x; \cdot)$ continuous in Θ .

Let $\mathbb{W}\mathbb{F}_\vartheta$ denote the probability measure induced by the solution X of (1) on (C, Σ) and $\mathbb{W}\mathbb{F}$ the probability measure induced by the corresponding process X^α , that is, the restriction of X to the neutral case where $\eta(x; \vartheta) \equiv 0$. Then, if assumption 1. holds, and $\phi(\cdot; \vartheta)$ and $A(\cdot; \vartheta)$ are defined as above, the Radon-Nykodým derivative of $\mathbb{W}\mathbb{F}_\vartheta$ w.r.t. $\mathbb{W}\mathbb{F}$ exists and follows from Girsanov's transformation of measures and Itô's formula, with

$$\frac{d\mathbb{W}\mathbb{F}_\vartheta}{d\mathbb{W}\mathbb{F}}(\omega) = \exp\{A(\omega_T; \vartheta) - A(\omega_0; \vartheta)\} \exp\left\{-\int_0^T \phi(\omega_s; \vartheta) ds\right\}. \quad (3)$$

The following lemma enables our proposed estimating approach.

Lemma 1 *Let $\mathbb{W}\mathbb{F}_\vartheta^{t,x,y}$ denote the probability measure induced by the solution X of (1) on (C, Σ) conditioned to start at x and to finish at y after time t . Moreover, consider $\mathbb{W}\mathbb{F}^{t,x,y}$ the probability measure induced by the corresponding process X^α , then*

$$p_\vartheta(x, y; t) = p(x, y, t) \exp\{A(y; \vartheta) - A(x; \vartheta) - t\phi^-(\vartheta)\} a(x, y; \vartheta),$$

where $p(x, \cdot, t)$ denotes the probability density of X^α and

$$a(x, y; \vartheta) = \mathbb{E}_{\mathbb{W}\mathbb{F}^{t,x,y}} \left[\exp \left\{ - \int_0^t \phi(\omega_s; \vartheta) - \phi^-(\vartheta) ds \right\} \right],$$

with $\omega \sim \mathbb{W}\mathbb{F}^{t,x,y}$, $\omega = \{\omega_s, s \in [0, t]\}$, such that $\omega_0 = x, \omega_t = y$ and $\phi^-(\vartheta)$ a lower bound of the mapping $s \mapsto \phi(\omega_s; \vartheta)$.

Proof By Assumption 1. and the definitions of $\phi(\cdot; \vartheta)$ and $A(\cdot; \vartheta)$ provided in 3. and 4., we can write the Radon-Nykodým derivative of $\mathbb{W}\mathbb{F}_\vartheta$ w.r.t. $\mathbb{W}\mathbb{F}$ as in (3). Using Bayes theorem, we have:

$$\begin{aligned} \frac{d\mathbb{W}\mathbb{F}_\vartheta^{t,x,y}}{d\mathbb{W}\mathbb{F}^{t,x,y}}(\omega) &= \frac{p(x, y, t)}{p_\vartheta(x, y, t)} \frac{d\mathbb{W}\mathbb{F}_\vartheta}{d\mathbb{W}\mathbb{F}}(\omega) \\ &= \frac{p(x, y, t)}{p_\vartheta(x, y, t)} \exp \{A(y; \vartheta) - A(x; \vartheta)\} \exp \left\{ - \int_0^t \phi(\omega_s; \vartheta) ds \right\}, \end{aligned}$$

Taking expectations at both sides w.r.t. $\mathbb{W}\mathbb{F}^{t,x,y}$ and multiplying by $p_\vartheta(x, y, t)$ yields

$$p_\vartheta(x, y, t) = p(x, y, t) \exp \{A(y; \vartheta) - A(x; \vartheta)\} \mathbb{E}_{\mathbb{W}\mathbb{F}^{t,x,y}} \left[\exp \left\{ - \int_0^t \phi(\omega_s; \vartheta) ds \right\} \right]$$

where we recall that

$$\mathbb{E}_{\mathbb{W}\mathbb{F}^{t,x,y}} \left[\frac{d\mathbb{W}\mathbb{F}_\vartheta^{t,x,y}}{d\mathbb{W}\mathbb{F}^{t,x,y}} \right] = \int_C \frac{d\mathbb{W}\mathbb{F}_\vartheta^{t,x,y}}{d\mathbb{W}\mathbb{F}^{t,x,y}} d\mathbb{W}\mathbb{F}^{t,x,y} = \mathbb{W}\mathbb{F}_\vartheta^{t,x,y}(C) = 1.$$

Finally, we have

$$p_\vartheta(x, y, t) = p(x, y, t) \exp \{A(y; \vartheta) - A(x; \vartheta) - t\phi^-(\vartheta)\} a(x, y, \vartheta),$$

which concludes the proof. \square

3 Unbiased Monte Carlo estimator

Following [7] and [21], we shall devise a random function such that the mapping $\vartheta \mapsto L(\Lambda, \vartheta)$ is a.s. continuous, the random element Λ is independent of ϑ , and such that for any fixed ϑ in the parameter space

$$L(\vartheta) = \mathbb{E}[L(\Lambda, \vartheta)],$$

where the expectation is taken w.r.t. the probability law of the element Λ . Note that $\mathbb{E}[L(\Lambda, \vartheta)]$ is amenable to Monte Carlo estimation by the functional averages

$$L_N(\vartheta) = \frac{1}{N} \sum_{j=1}^N L(\Lambda^j, \vartheta), \quad (4)$$

where $\Lambda^j, j = 1, \dots, N$ are independent Monte Carlo samples.

In view of Lemma 1, we need to define

$$L(\Lambda, \vartheta) = \exp \{A(y; \vartheta) - A(x; \vartheta) - t\phi^-(\vartheta)\} \Pi(\Lambda, x, y; \vartheta), \quad (5)$$

where $\Pi(\Lambda, x, y; \vartheta) = p(M, x, y, t)a(\Upsilon, \omega, x, y; \vartheta)$, with $\Lambda = (M, \Upsilon, \{\omega_{t_i}, 1 \leq i \leq K\})$ and M independent of Υ and ω , so that

$$\mathbb{E}[\Pi(\Lambda, x, y; \vartheta)] = \mathbb{E}[p(M, x, y, t)]\mathbb{E}[a(\Upsilon, \omega, x, y; \vartheta)],$$

with $\mathbb{E}[p(M, x, y, t)] = p(x, y, t)$ and $\mathbb{E}[a(\Upsilon, \omega, x, y; \vartheta)] = a(x, y; \vartheta)$.

In the next subsections, the functions $p(M, x, y, t)$ and $a(\Upsilon, \omega, x, y; \vartheta)$ are described.

3.1 Exact rejection sampling of Wright-Fisher diffusion bridges

In this subsection, we briefly summarize the exact rejection algorithm for simulating Wright-Fisher diffusion bridges and that is central for the construction of the estimator in (4).

The rejection scheme is based on Lemma 1 above and uses candidate paths distributed according to $\mathbb{W}\mathbb{F}^{t,x,y}$, for which an exact simulation procedure exists, see Algorithms 4 and 5 in [8] and Algorithms 5 and 6 in [9] for a multivariate generalization. Note that, if assumptions 3. and 4. are satisfied, then,

$$\frac{d\mathbb{W}\mathbb{F}_\vartheta^{t,x,y}}{d\mathbb{W}\mathbb{F}^{t,x,y}}(\omega) \propto \exp \left\{ - \int_0^t \phi(\omega_s; \vartheta) - \phi^-(\vartheta) ds \right\},$$

where we recognize the right hand side as

$$\Pr(N = 0 | \omega) = \exp \left\{ - \int_0^t \phi(\omega_s; \vartheta) - \phi^-(\vartheta) ds \right\} \leq 1, \quad (6)$$

where N is the number of points of a marked Poisson process Φ on $[0, t] \times [0, 1]$ with rate $\phi^+(\vartheta)$ that lie below the graph of

$$s \mapsto g(\omega_s; \vartheta) := \frac{\phi(\omega_s; \vartheta) - \phi^-(\vartheta)}{\phi^+(\vartheta) - \phi^-(\vartheta)}.$$

In the following, we formally describe how the expression in (6) can be rewritten as the expectation of a certain indicator, for which an unbiased estimator is available.

Consider $\Phi = \{\Upsilon, \Psi\}$ and $K \sim \text{Po}(\phi^+(\vartheta)t)$, with $\Upsilon = \{t_1, \dots, t_K\}$ the projection of Φ on the time-axis with time-ordered points $t_i, 1 \leq i \leq K$ that are uniformly distributed on $[0, t]$, and $\Psi = \{\psi_1, \dots, \psi_K\}$ their corresponding marks that are uniformly distributed on $[0, 1]$.

Then Algorithm 1 provides skeletons of paths $\omega \sim \mathbb{W}\mathbb{F}_\vartheta^{(t,x,y)}$, where

$$I(x, y, \vartheta, \Phi, \omega) = \prod_{i=1}^K \mathbb{I} \left[\frac{\phi(\omega_{t_i}; \vartheta) - \phi^-(\vartheta)}{\phi^+(\vartheta) - \phi^-(\vartheta)} \leq \psi_i \right]$$

is the acceptance indicator.

Algorithm 1 Exact rejection algorithm for simulating skeletons of paths $\omega \sim \mathbb{WF}_{\vartheta}^{(t,x,y)}$

- 1 Simulate Φ , a Poisson process on $[0, t] \times [0, 1]$ with rate $\phi^+(\vartheta)$
 - 2 Given $\Phi = \{(t_i, \psi_i) : i = 1, \dots, K\}$, simulate $\omega \sim \mathbb{WF}^{t,x,y}$ at times $\{t_1, \dots, t_K\}$.
 - 3 **if** $I(x, y, \vartheta, \Phi, \omega) = 1$ **then**
 - 4 **return** $S(\omega) = \{(0, x), (t_1, \omega_{t_1}), \dots, (t_K, \omega_{t_K}), (T, y)\}$
 - 5 **else**
 - 6 Go back to Step 1.
 - 7 **end if**
-

By definition of $a(x, y; \vartheta)$, we have

$$a(x, y; \vartheta) = \mathbb{E}[I(x, y, \vartheta, \Phi, \omega)], \quad (7)$$

where the expectation is taken w.r.t. the joint distribution of Φ and ω .

Because Ψ is uniformly distributed on $[0, 1]$, the probability of none of the sampled marks ψ_i , $1 \leq i \leq K$ to be below $g(\omega_{t_i}; \vartheta)$ is $\prod_{i=1}^K 1 - (\phi(\omega_{t_i}^j; \vartheta) - \phi^-(\vartheta)) / (\phi^+(\vartheta) - \phi^-(\vartheta))$ and for each $\vartheta \in \Theta$,

$$a(Y^j, \omega^j, x, y; \vartheta) = \prod_{i=1}^{K^j} 1 - \frac{\phi(\omega_{t_i}; \vartheta) - \phi^-(\vartheta)}{\phi^+(\vartheta) - \phi^-(\vartheta)} \quad (8)$$

is an unbiased Monte Carlo estimator of (7).

In order to make the estimator proposed in (8) suitable for any possible value of ϑ , we refer to what is coined as the Simultaneous Acceptance Method (SAM) in [7].

Let Φ_+ be a marked Poisson process on $[0, t] \times [0, 1]$ with rate $\phi^+ \geq \sup_{\vartheta} (\phi^+(\vartheta) - \phi^-(\vartheta))$, and $U = \{U_1, \dots, U_K\}$ a vector of i.i.d. uniform random variables on $[0, 1]$. Using the thinning property of the Poisson process, see, e.g., Section 5 of [22], we can recover the process Φ by deleting each point of Φ_+ with probability $1 - (\phi^+(\vartheta) - \phi^-(\vartheta)) / \phi^+$, yielding

$$I(x, y, \vartheta, \Phi_+, \omega, U) = \prod_{i=1}^K \mathbb{I} \left[1 - \mathbb{I} \left[U_i < \frac{\phi^+(\vartheta) - \phi^-(\vartheta)}{\phi^+} \right] \frac{\phi(\omega_{t_i}; \vartheta) - \phi^-(\vartheta)}{\phi^+(\vartheta) - \phi^-(\vartheta)} \leq \psi_i \right].$$

After taking the expectation w.r.t. U one obtains

$$a(x, y; \vartheta) = \mathbb{E}[\mathbb{E}_U[I(x, y, \vartheta, \Phi_+, \omega, U)]],$$

where, again, the outer expectation is taken w.r.t. the joint distribution of Φ_+ and ω . Thus, a (simultaneous) unbiased estimator of $a(x, y; \vartheta)$ is

$$a(\Upsilon_+, \omega, x, y; \vartheta) = \frac{1}{N} \sum_{j=1}^N a(\Upsilon_+^j, \omega^j, x, y; \vartheta), \quad (9)$$

where

$$a(\Upsilon_+^j, \omega^j, x, y; \vartheta) = \prod_{i=1}^{K^j} 1 - (\phi(\omega_{t_i}^j; \vartheta) - \phi^-(\vartheta)) / \phi^+.$$

3.2 Estimation of conditioned neutral Wright-Fisher diffusion densities

To complete our estimation approach, it only remains to devise a strategy for computing $p(M, x, y, t)$. From equation (14) in [8], we know that

$$p(x, y; t) = \sum_{m=0}^{\infty} q_m^\theta(t) \mathbb{E}[\mathcal{D}_{\theta, L}(y)], \quad (10)$$

where the expectation $\mathbb{E}[\cdot]$ is taken w.r.t. the random variable $L \equiv L_{m, x}$, which is distributed as a binomial random variable with parameters m and x .

$$\mathcal{D}_{\theta, l}(y) = \frac{1}{B(\theta_a + l, \theta_A + m - l)} y^{\theta_a + l - 1} (1 - y)^{\theta_A + m - l - 1}$$

is the probability density function of a beta random variable with parameters $\theta_a + l$ and $\theta_A + m - l$, for each realization l of $L_{m, x} \sim \text{Binomial}(m, x)$, and, for each t , $q_m^\theta(t)$ is the probability mass function of a certain discrete random variable M taking values in $\{m = 0, 1, \dots\}$. Then, by definition, we can rewrite (10) as

$$p(x, y; t) = \mathbb{E}_M [\mathbb{E}_L [\mathcal{D}_{\theta, L}(y)]] = \sum_{m=0}^{\infty} P(M = m) \sum_{l=0}^m P(L_{m, x} = l) \mathcal{D}_{\theta, l}(y), \quad (11)$$

where the expectations are taken with respect to the laws of M and $L_{m, x}$, respectively.

Were an analytic expression for $\mathbb{E}_M[\cdot]$ available, one could compute (11) exactly from its definition. Unfortunately, $q_m^\theta(t)$ is only known in infinite series form, see [23] and [24], that is,

$$q_m^\theta(t) = \sum_{k=m}^{\infty} (-1)^{k-m} b_k^{(t, \theta)}(m),$$

with

$$b_k^{(t, \theta)}(m) = \frac{(\theta + 2k - 1) \Gamma(\theta + m + k - 1)}{m!(k - m)! \Gamma(\theta + m)} e^{-k(k+\theta-1)t/2}$$

and $\theta = \theta_a + \theta_A$, rendering an exact computation of $\mathbb{E}_M[\cdot]$ impossible.

However, there exists an exact sampling strategy for drawing samples from M , see Algorithm 2 in [8] and Algorithm 3 in [9], and, given $m^j, j = 1, \dots, N$ independent Monte Carlo samples distributed according to $q_m^\theta(t)$, we obtain

$$p(M^j, x, y; t) = \sum_{l=0}^{m^j} \binom{m^j}{l} x^l (1-x)^{m^j-l} \mathcal{D}_{\theta, t}(y),$$

with

$$p(M, x, y; t) = \frac{1}{N} \sum_{j=1}^N p(M^j, x, y; t), \quad (12)$$

an unbiased estimator of $p(x, y; t)$.

Thus, combining (9) and (12) we obtain

$$\begin{aligned} \Pi(\Lambda, x, y; \vartheta) &= p(M, x, y; t) a(Y_+, \omega, x, y; \vartheta) \\ &= \frac{1}{N} \sum_{j=1}^N p(M^j, x, y; t) \frac{1}{N} \sum_{j=1}^N a(Y_+^j, \omega^j, x, y; \vartheta), \end{aligned}$$

with the desired properties.

3.3 Theoretical guarantees

Theoretical results on consistency of our proposed MLE, ϑ^n , follow from those shown in [21], as we detail in this section. So far, we have provided unbiased estimators for each independent contribution $L^i(\vartheta)$. Following the expression in (4), an estimator of the full likelihood function is

$$\mathcal{L}_N^n(\vartheta) = \prod_{i=1}^n \frac{1}{N} \sum_{j=1}^N L^i(\Lambda^j, \vartheta),$$

where the contributions are estimated independently for each observation $i = 1, \dots, n$ over N Monte Carlo samples. By Kolmogorov's Strong Law of Large Numbers we know that $\mathcal{L}_N^n(\vartheta) \rightarrow \mathcal{L}^n(\vartheta)$ a.s. as $N \rightarrow \infty$, which, however, does not guarantee convergence of the maximizers ϑ_N^n of the functional averages $\mathcal{L}_N^n(\vartheta)$ to ϑ^n . A sufficient condition is uniform convergence in ϑ , that is,

$$\lim_{N \rightarrow \infty} \sup_{\vartheta \in \Theta} |\mathcal{L}_N^n(\vartheta) - \mathcal{L}^n(\vartheta)| = 0, \text{ a.s.} \quad (13)$$

In order to prove (13), we revert to the following general result for random elements on a Banach space [25].

Theorem 1 [Theorem 2 in [21]] *Let $(\mathcal{X}, \|\cdot\|)$ be a separable Banach space, with random elements $\mathbf{X} \in \mathcal{X}$ with norm $\|\mathbf{X}\|$, such that*

$$\mathbb{E}[\|\mathbf{X}\|] < \infty \text{ and } \mathbb{E}[\mathbf{X}] = 0,$$

where recall that $\mathbb{E}[\mathbf{X}]$ is defined as the unique element $v \in \mathcal{X}$ such that $T(v) = \mathbb{E}[T(\mathbf{X})]$ for all linear functionals $T \in \mathcal{X}^*$. Then, if $\mathbf{X}^1, \mathbf{X}^2, \dots$ are independent copies of \mathbf{X} ,

$$\lim_{N \rightarrow \infty} \left\| \frac{1}{N} \sum_{j=1}^N \mathbf{X}^j \right\| = 0.$$

Then, the following corollary provides a proof for (13).

Corollary 1 *Let $L(\Lambda, \vartheta)$ be defined as in (5), and $\Lambda^1, \Lambda^2, \dots$, i.i.d. copies of the random element Λ . Then,*

$$\lim_{N \rightarrow \infty} \sup_{\vartheta \in \Theta} \left| \frac{1}{N} \sum_{j=1}^N L(\Lambda^j, \vartheta) - L(\vartheta) \right| = 0, \text{ a.s. .}$$

Proof Take $(\mathcal{X}, \|\cdot\|) = (F, \|\cdot\|)$ the set of continuous mappings on the compact set Θ , endowed with the supremum norm $\|f\| = \sup_{\vartheta \in \Theta} |f(\vartheta)|$, for any $f \in F$. It is well known that $(F, \|\cdot\|)$ is a separable Banach space, see, e.g., [26].

Assumption 2. implies that $\phi(x; \cdot)$ and $A(x; \cdot)$ are continuous in Θ . Then, a.s. continuity of $L(\Lambda, \cdot)$ follows from assumption 3, and we have $L(\Lambda, \cdot) \in F$.

Now, because of (6) we know that $a(Y, \omega, x, y; \vartheta) \leq 1$ a.s. . Moreover, assumption 4. implies that $|L(\Lambda, \vartheta)| \leq \kappa < \infty$, where

$$\kappa = \sup_{\vartheta \in \Theta} \exp \{A(y; \vartheta) - A(x; \vartheta) - t\phi^-(\vartheta)\}$$

is finite for being the supremum of a continuous function over the compact set Θ . Then, $\mathbb{E}[\|L(\Lambda, \cdot)\|] < \infty$. Using Corollary 1 in [21] we obtain $L(\cdot) \in F$, and by uniqueness of the expectation we have $\mathbb{E}[L(\Lambda, \cdot) - L(\cdot)] = 0$. Applying Theorem 1 to $L(\Lambda, \cdot) - L(\cdot)$ concludes the proof. \square

Finally, the uniform convergence in (13) and the compactness of Θ imply the consistency of our proposed estimator, which validates our estimating strategy. We formalize this statement in the following Corollary.

Corollary 2 *Let $\{\vartheta_N^n\}_N$ be any sequence of maximizers of $\{\mathcal{L}_N^n(\vartheta)\}_N$. If ϑ^n is the unique element of $\arg \max_{\vartheta \in \Theta} \mathcal{L}^n(\vartheta)$, then $\lim_{N \rightarrow \infty} \vartheta_N^n = \vartheta^n$.*

4 Numerical example

We exemplify the performance of our method through a particular numerical example that refers to a widely used instance of the Wright-Fisher diffusion model.

We consider (1) with $\gamma(x; \vartheta) = \alpha(x) + \vartheta x(1-x)[x+h(1-2x)]$, which corresponds to the Wright-Fisher diffusion with diploid selection. For illustrative purposes, we focus on the haploid case, that is, $h = 1/2$, for which the quantities of relevance are significantly simplified. In this case,

$$\eta(x; \vartheta) = \frac{\vartheta}{2}, \text{ so that } A(x; \vartheta) = \frac{\vartheta}{2}x \text{ and } \phi(x; \vartheta) = \frac{1}{2}[x(1-x)\frac{\vartheta^2}{4} + \vartheta\alpha(x)].$$

The bounds for $\phi(\cdot; \vartheta)$ are defined from the quantities

$$k_1 = \frac{1}{32} \left(\vartheta^2 + 4\vartheta(\theta_a - \theta_A) + 4(\theta_a + \theta_A)^2 \right), \quad k_2 = \frac{\vartheta}{4}\theta_a, \quad k_3 = -\frac{\vartheta}{4}\theta_A,$$

where k_1 is the vertex of the parabola defined by $\phi(\cdot; \vartheta)$, $k_2 = \phi(0; \vartheta)$ and $k_3 = \phi(1; \vartheta)$. Considering a range of mutation rates $0 < \theta_a, \theta_A < \frac{1}{2}$, $\phi^+(\vartheta) = k_1$ and $\phi^-(\vartheta) = \min(k_2, k_3)$.

In this example, the parameter ϑ quantifies the relative fitness between type a and type A in the time-scaled weak selection-weak mutation diffusion model, see, e.g., [27]. Relative fitness between type a and type A is often defined as $1 + s : 1$, where $s = (f_a - f_A)/f_A$, for f_a and f_A that quantify the natural selective advantage of types a and A , respectively. Then, s is referred to as selection coefficient. If f_a is interpreted as division rate, one can assume it to be as small as 0 (in which case type a would be lethal or such that completely prevents reproduction of the individual carrying it), yielding $s = -1$. On the opposite end, assuming type a to be highly beneficial, s could be, in principle, as large as necessary. However, experimental results on RNA viruses have shown selection coefficient values of individual mutations to be $|s| \approx 1$, see [28], which in the diffusion limit of large (but finite) populations yields $\vartheta \in \Theta = [-c, c]$, with c the population size.


Similarly, most experimental results on RNA viruses quantify mutation probabilities per gene per individual per generation of an order from $\approx 10^{-6}$ to $\approx 10^{-4}$, see [29], which, in the time-scaled diffusion limit, can reach mutation rates of an order from $\approx 10^{-3}$ to $\approx 10^{-1}$.

Numerical results showing the consistency of our proposed MLE are reported in Table 1. The benchmark dataset had a size of $n = 100$ with $\Delta t_i = 1$ for $i = 1, \dots, n$, and parameters $\theta_a = \theta_A = 0.02$, $\vartheta = 0.7$, corresponding to a weak selection-weak mutation regime, see Fig 1. The data were simulated using the exact algorithm (Algorithm 6) presented in [8], with $A^+(\vartheta) = \lfloor \frac{\vartheta}{2} \rfloor$. Neutral Wright-Fisher bridges' paths $\omega \sim \mathbb{WF}^{t, x, y}$ in Step 2 of Algorithm 1 above, were simulated using the EWF sampler presented in [30], wrapped in our Python 3.9.12 implementation. Maximum (log)-likelihood estimators were computed using Brent's optimization algorithm, see [31].

Table 1 Consistency of the MLE, $N \rightarrow \infty$. Results obtained for the MLE, ϑ_N^n , with $n = 100$ and $\Delta t_i = 1$ for $i = 1, \dots, n$, and parameters $\theta_a = \theta_A = 0.02$, $\vartheta = 0.7$. Standard errors based on 50 bootstrapped samples are shown in parenthesis.

N	ϑ_N^n
1	0.763 (0.1686)
10	0.738 (0.1061)
20	0.713 (0.0586)
50	0.709 (0.0276)
100	0.719 (0.0249)
200	0.695 (0.0148)
500	0.707 (0.0100)
Large $N (= 10^3)$	0.701 (0.0083)

Fig. 1 Generated benchmark data set with $n = 100$ points and $\Delta t_i = 1$ for $i = 1, \dots, n$, with $\theta_a = \theta_A = 0.02$ and $\vartheta = 0.7$.



5 Monte Carlo maximum likelihood estimator for the coupled Wright-Fisher diffusion

The approach presented so far can also be applied in a multidimensional setting. Let us now consider the evolution of the frequency of genetic variants across L different locations in the genome. As above, we focus on the simplest case, where only two different genetic variants (or types) are possible in each location. The interest now lies, not only on the selective advantage of one type over the other at each location, but also in considering pairwise selective interactions of genetic variants across

locations. The study of genetic interactions is a highly active field in biological evolution, where it is known that genetic effects seldom act individually, but they rather combine their influence resulting into complex networks of interactions, see, e.g., [32] for an overview.

In this context we consider the coupled Wright-Fisher diffusion model, where X is described as the weakly unique solution of the SDE

$$dX_t = [\alpha(X_t; \theta) + G(X_t; \vartheta)]dt + D^{\frac{1}{2}}(X_t)dB_t, \quad X_0 = x_0 \in [0, 1]^L, \quad t \in [0, T], \quad (14)$$

with X an L -dimensional vector of frequencies of allele types at locations $k = 1, \dots, L$, $\alpha(x^k; \theta)$ is defined as above, $G(x; \vartheta)$ is a coupling term that describes natural selection and pairwise selective interactions across locations and $D(x)$ is a diagonal matrix with entries $x^k(1 - x^k)$.

The coupling term $G : [0, 1]^L \times \Theta \rightarrow \mathbf{R}^L$ has entries with the general form

$$G^k(x; \vartheta) = x^k(1 - x^k) \left(s^{k1} - s^{k2} + \sum_{\substack{l=1 \\ l \neq k}}^L (h_{12}^{kl} - h_{22}^{kl} + (h_{11}^{kl} - h_{12}^{kl} - h_{21}^{kl} + h_{22}^{kl})x^l) \right),$$

where s^{kj} for $j = 1, 2$, is the selective advantage coefficient of type j at location k , and h_{jr}^{kl} is the pairwise selective interaction between type j at location k and type r at location l , for $r = 1, 2$.

Following the recent results on exact simulation for coupled Wright-Fisher diffusions [9], one can write an equivalent to Lemma 1 above, and build an unbiased estimator for each corresponding likelihood contribution as in (5). Given $i = 1, \dots, n$ discrete observations without error of paths $k = 1, \dots, L$ from the process (14), each likelihood contribution can be written as

$$L^i(\Lambda, \vartheta) = \exp \{ \tilde{A}(y; \vartheta) - \tilde{A}(x; \vartheta) - t\tilde{\phi}^-(\vartheta) \} \tilde{p}(M, x, y, t) \tilde{a}(Y, \tilde{\omega}, x, y; \vartheta), \quad (15)$$

with $\vartheta = (s, h)$ a multidimensional selection parameter, and

i) The function $\tilde{\phi} : [0, 1]^L \times \Theta \rightarrow \mathbb{R}$ is defined as

$$\tilde{\phi}(x; \vartheta) := \frac{1}{2} \left[(V(x; \vartheta))^T D(x) V(x; \vartheta) + 2(V(x; \vartheta))^T \alpha(x; \theta) \right],$$

where $V^k(x; \vartheta) = s^{k1} - s^{k2} + \sum_{\substack{l=1 \\ l \neq k}}^L (h_{12}^{kl} - h_{22}^{kl} + (h_{11}^{kl} - h_{12}^{kl} - h_{21}^{kl} + h_{22}^{kl})x^l)$, for

$k, l = 1, \dots, L$, and the corresponding bounds are $\tilde{\phi}^-(\vartheta) \leq \tilde{\phi}(x; \vartheta) \leq \tilde{\phi}^+(\vartheta)$.

ii) The function $\tilde{A} : [0, 1]^L \times \Theta \rightarrow \mathbb{R}$ is defined as

$$\tilde{A}(X_t; \vartheta) = \int_0^t V(X_s) dX_s = \sum_{k=1}^L \left(K_s^k X_t^k + \sum_{l=1}^L K_l^k X_t^k + \sum_{l=k+1}^L K_{lr}^{kj} X_t^k X_t^l \right),$$

with K_s^k, K_l^k and K_{lr}^{kj} appropriate constants.

iii) The function $\tilde{p}(x, y; t)$ is the joint distribution of L independent neutral Wright-Fisher bridges from x to y , that is,

$$\tilde{p}(x, y; t) = \prod_{k=1}^L p(x^k, y^k; t) = \prod_{k=1}^L \mathbb{E}_M [\mathbb{E}_L [\mathcal{D}_{\theta, L}(y^k)]].$$

Then, function $\tilde{p}(M, x, y; t) = \frac{1}{N} \sum_{j=1}^N \tilde{p}(M^j, x, y; t)$ is an unbiased estimator of $\tilde{p}(x, y; t)$, where

$$\tilde{p}(M^j, x, y; t) = \prod_{k=1}^L \left(\sum_{l=0}^{m^{jk}} \binom{m^{jk}}{l} (x^k)^l (1-x^k)^{m^{jk}-l} \mathcal{D}_{\theta, t}(y^k) \right),$$

with m^{jk} , $j = 1, \dots, N$, independent Monte Carlo samples from $q_m^\theta(t)$, which are drawn independently for every location $k = 1, \dots, L$.

iv) The function $\tilde{a}(Y^j, \tilde{\omega}, x, y; \vartheta)$ is defined as in (8), with

$$\tilde{a}(Y^j, \omega, x, y; \vartheta) = \prod_{i=1}^{\tilde{K}^j} 1 - \frac{\tilde{\phi}(\tilde{\omega}_i; \vartheta) - \tilde{\phi}^-(\vartheta)}{\tilde{\phi}^+(\vartheta) - \tilde{\phi}^-(\vartheta)},$$

where $\tilde{K}^j \sim \text{Po}(\tilde{\phi}^+(\vartheta)t)$ and $\tilde{\omega} \sim \text{WFL}^{t, x, y}$, where the latter denotes the joint law of L independent neutral Wright-Fisher bridges from x to y .

Clearly, all assumptions required in Section 3.3 are fulfilled, and thus, an unbiased estimator of (15) is available, providing an extension of our method to the multidimensional setting.

The analogous to Algorithm 1 in the coupled Wright-Fisher multidimensional setting can be derived from Algorithm 4 in [9].

6 Conclusion

In this paper we have presented an unbiased estimation technique for a class of Wright-Fisher diffusion processes that cover a wide range of applications of interest. The main advantage of our method is that it is based on exact simulation of Wright-Fisher diffusions, which circumvents any source of error due to numerical approximations and provides unbiased maximum likelihood estimators. We have illustrated the performance of our method in a numerical example showing promising results.

Future work includes exploring the joint estimation of mutation and selection. However, the rejection mechanism at the core of existing exact simulation algorithms for Wright-Fisher diffusions requires mutation parameters for candidate and target paths to be the same, and thus, this extension would require to develop new exact simulation techniques, which are outside of the scope of this paper.

Acknowledgements The authors wish to thank Anne-Florence Bitbol for useful discussions and Davide Pradovera for helpful suggestions on code implementation.

References

1. Helle Sørensen. Parametric inference for diffusion processes observed at discrete points in time: a survey. *International Statistical Review*, 72(3):337–354, December 2004.
2. Peter Craigmile, Radu Herbei, Ge Liu, and Grant Schneider. Statistical inference for stochastic differential equations. *WIREs Comput Stat*, n/a(n/a):e1585, April 2022.
3. Richard A. Neher, Colin A. Russell, and Boris I. Shraiman. Predicting evolution from the shape of genealogical trees. *eLife*, 3:e03568, November 2014.
4. Marta Łuksza and Michael Lässig. A predictive fitness model for influenza. *Nature*, 507(7490):57–61, March 2014.
5. John Huddleston, John R. Barnes, Thomas Rowe, Xiyan Xu, Rebecca Kondor, David E. Wentworth, Lynne Whittaker, Burcu Ermetal, Rodney Stuart Daniels, John W. McCauley, Seiichiro Fujisaki, Kazuya Nakamura, Noriko Kishida, Shinji Watanabe, Hideki Hasegawa, Ian Barr, Kanta Subbarao, Pierre Barrat-Charlaix, Richard A. Neher, and Trevor Bedford. Integrating genotypes and phenotypes improves long-term forecasts of seasonal influenza a/h3n2 evolution. *eLife*, 9:e60067, September 2020.
6. Pierre Barrat-Charlaix, John Huddleston, Trevor Bedford, and Richard A. Neher. Limited predictability of amino acid substitutions in seasonal influenza viruses. *Mol Biol Evol*, 38(7):2767–2777, July 2021.
7. Alexandros Beskos, Omiros Papaspiliopoulos, Gareth O. Roberts, and Paul Fearnhead. Exact and computationally efficient likelihood-based estimation for discretely observed diffusion processes. *Journal of the Royal Statistical Society. Series B (Statistical Methodology)*, 68(3):333–382, 2006.
8. Paul A. Jenkins and Dario Spanò. Exact simulation of the Wright-Fisher diffusion. *Annals of Applied Probability*, 27(3):1478–1509, 2017.
9. Celia García-Pareja, Henrik Hult, and Timo Koski. Exact simulation of coupled Wright-Fisher diffusions. *Advances in Applied Probability*, 53(4):923–950, 2021.
10. C. E. Dangerfield, D. Kay, S. MacNamara, and K. Burrage. A boundary preserving numerical algorithm for the Wright-Fisher model with mutation. *BIT Numerical Mathematics*, 52(2):283–304, June 2012.
11. G. O. Roberts and O. Stramer. On inference for partially observed nonlinear diffusion models using the Metropolis-Hastings algorithm. *Biometrika*, 88(3):603–621, 2001.
12. Yacine Aït-Sahalia. Maximum likelihood estimation of discretely sampled diffusions: A closed-form approximation approach. *Econometrica*, 70(1):223–262, January 2002.
13. Chenxu Li. Maximum-likelihood estimation for diffusion processes via closed-form density expansions. *The Annals of Statistics*, 41(3):1350–1380, June 2013.
14. Garland B. Durham and A. Ronald Gallant. Numerical techniques for maximum likelihood estimation of continuous-time diffusion processes. *Journal of Business & Economic Statistics*, 20(3):297–338, July 2002.
15. Bo Martin Bibby, Martin Jacobsen, and Michael Sørensen. Chapter 4 – estimating functions for discretely sampled diffusion-type models. 2010.
16. A. Ronald Gallant and George Tauchen. Which moments to match? *Econometric Theory*, 12(4):657–681, 1996.
17. C. Gourieroux, A. Monfort, and E. Renault. Indirect inference. *Journal of Applied Econometrics*, 8:S85–S118, 1993.
18. Joshua G. Schraiber, Robert C. Griffiths, and Steven N. Evans. Analysis and rejection sampling of Wright-Fisher diffusion bridges. *Theoretical Population Biology*, 89:64–74, November 2013.

19. Paula Tataru, Maria Simonsen, Thomas Bataillon, and Asger Hobolth. Statistical inference in the Wright–Fisher model using allele frequency data. *Systematic Biology*, 66(1):e30–e46, 2017.
20. Assyr Abdule, Giacomo Garegnani, Grigorios A. Pavliotis, Andrew M. Stuart, and Andrea Zanoni. Drift estimation of multiscale diffusions based on filtered data. *Found. Comput. Math.*, 2021.
21. Alexandros Beskos, Omiros Papaspiliopoulos, and Gareth Roberts. Monte Carlo maximum likelihood estimation for discretely observed diffusion processes. *The Annals of Statistics*, 37(1):223–245, February 2009.
22. John Frank Charles Kingman. *Poisson processes*, volume 3. Clarendon Press, 1992.
23. Robert C. Griffiths. Lines of descent in the diffusion approximation of neutral wright-fisher models. *Theoretical Population Biology*, 17(1):37 – 50, 1980.
24. Simon Tavaré. Line-of-descent and genealogical processes, and their applications in population genetics models. *Theoretical Population Biology*, 26(2):119–164, October 1984.
25. Daniel P. Giesy. Strong laws of large numbers for independent sequences of banach space-valued random variables. In *Probability in Banach Spaces*, pages 89–99, Berlin, Heidelberg, 1976. Springer Berlin Heidelberg.
26. Zbigniew Semadeni. Spaces of continuous functions on compact sets. *Advances in Mathematics*, 1(3):319–382, 1965.
27. Alison Etheridge. *Some Mathematical Models from Population Genetics: École D’Été de Probabilités de Saint-Flour XXXIX-2009*, volume 2012. Springer-VERLAG Berlin, Heidelberg, 2011.
28. Elisa Visher, Shawn E. Whitefield, John T. McCrone, William Fitzsimmons, and Adam S. Lauring. The mutational robustness of influenza a virus. *PLOS Pathogens*, 12(8):e1005856, August 2016.
29. Hongrui Cui, Guangsheng Che, Mart C. M. de Jong, Xuesong Li, Qinfang Liu, Jianmei Yang, Qiaoyang Teng, Zejun Li, and Nancy Beerens. The pb1 gene from h9n2 avian influenza virus showed high compatibility and increased mutation rate after reassorting with a human h1n1 influenza virus. *Virology Journal*, 19(1):20, January 2022.
30. Jaromir Sant, Paul A. Jenkins, Jere Koskela, and Dario Spanò. Ewf: simulating exact paths of the wright-fisher diffusion. *Bioinformatics*, 39(1):btad017, January 2023.
31. Richard P Brent. *Algorithms for minimization without derivatives*, chapter An algorithm with guaranteed convergence for finding a zero of a function. Englewood Cliffs, NJ: Prentice-Hall, 1973.
32. Júlia Domingo, Pablo Baeza-Centurion, and Ben Lehner. The causes and consequences of genetic interactions (epistasis). *Annu. Rev. Genom. Hum. Genet.*, 20(1):433–460, 2019.