Boosting Bayesian parameter inference of nonlinear stochastic differential equation models by Hamiltonian scale separation

Carlo Albert,¹ Simone Ulzega,¹ and Ruedi Stoop²

¹*Eawag, Swiss Federal Institute of Aquatic Science and Technology, 8600 Dübendorf, Switzerland*

²*Institute of Neuroinformatics and Institute of Computational Science UZH/ETHZ, Irchel Campus 8057 Zurich, Switzerland*

(Received 23 December 2015; published 15 April 2016)

Parameter inference is a fundamental problem in data-driven modeling. Given observed data that is believed to be a realization of some parameterized model, the aim is to find parameter values that are able to explain the observed data. In many situations, the dominant sources of uncertainty must be included into the model for making reliable predictions. This naturally leads to stochastic models. Stochastic models render parameter inference much harder, as the aim then is to find a distribution of likely parameter values. In Bayesian statistics, which is a consistent framework for data-driven learning, this so-called posterior distribution can be used to make probabilistic predictions. We propose a novel, exact, and very efficient approach for generating posterior parameter distributions for stochastic differential equation models calibrated to measured time series. The algorithm is inspired by reinterpreting the posterior distribution as a statistical mechanics partition function of an object akin to a polymer, where the measurements are mapped on heavier beads compared to those of the simulated data. To arrive at distribution samples, we employ a Hamiltonian Monte Carlo approach combined with a multiple time-scale integration. A separation of time scales naturally arises if either the number of measurement points or the number of simulation points becomes large. Furthermore, at least for one-dimensional problems, we can decouple the harmonic modes between measurement points and solve the fastest part of their dynamics analytically. Our approach is applicable to a wide range of inference problems and is highly parallelizable.

DOI: [10.1103/PhysRevE.93.043313](https://doi.org/10.1103/PhysRevE.93.043313)

I. INTRODUCTION

Modeling a dynamical process starts with a basic model that is usually obtained from a more or less deep insight into the nature of the process. The next step is the determination of the parameters of the model, based on observed data, which is generally a highly nontrivial task, in particular when complex behavior of such systems needs to be predicted or when the measurements are noisy. A minimal example is a perceptron [1,2], the basic element of a neuronal network, that predicts the double cosine value associated with the input of the corresponding sine function plus the sine's value at a fixed earlier time. While this task can easily be achieved for clean data using gradient descent learning, for noisy input data, this is largely impossible, as noise cannot be learned. The result is a distribution of potential parameter values (Fig. 1).

In Bayesian statistics, knowledge about parameters is expressed by probability distributions and learning is implemented as an update rule on these distributions (see, e.g., Ref. [3]). If a constant noise term is added to the output of a deterministic model, such as in the perceptron example above, then Bayesian inference is straightforward. If noise enters the formulation of the model equations, however, then Bayesian inference all of a sudden becomes computationally very expensive.

In our paper, we demonstrate the calibration of ordinary one-dimensional (1D) stochastic differential equation (SDE) models based on noisy time series and the quantification of the resulting parametric uncertainty. The generic approach that we use is exemplified by a simple SDE model from hydrology.

Problems of this kind are commonly solved by Monte Carlo (MC) methods that are based on simulating model realizations and comparing them to the data. Popular methods are particle filters [4,5], Metropolis-within-Gibbs algorithms

[6,7], or approximate Bayes computations [8–11]. A major problem with these simulation-based methods is, however, their inefficiency in the presence of many data points or high dimensions. One solution is to map the output space to a smaller dimensional space of summary statistics and accept or reject proposed model parameters depending on how well associated model runs conform with the data in terms of these summary statistics [12]. However, a method to choose the summary statistics to achieve a significant representation of the posterior parameter distribution is a largely unsolved problem.

These difficulties can be remedied with a reinterpretation of the Bayesian posterior distribution as the partition function of a statistical mechanics system and by simulating the dynamics of the latter. After discretizing the time of the original problem, we are led to a problem akin to the statistical mechanics of a polymer with harmonic bonds in an exterior potential [13]. In this framework, the measurements are interpreted as an additional exterior potential that acts only on the polymer's "measurement beads" and confines their dynamics within the measurement uncertainty. The model parameters are interpreted as additional degrees of freedom coupling to all the beads of the polymer. To simulate the dynamics of this system, we apply the Hamiltonian Monte Carlo (HMC) algorithm [14], which combines molecular dynamics [15,16] with the Metropolis algorithm [17]. Compared to traditional methods, the use of the Hamiltonian approach achieves much higher acceptance rates since data points are already used for the suggestion of new parameters, and thus model realizations incompatible with the data are never considered. The drawback is that the model equations need to be known and derivatives have to be calculated.

HMC requires two sets of parameters to be tuned: (i) the parameters that define the kinetic energy of the statistical

FIG. 1. General problem setting: Parameter estimation (synaptic weights w_1, w_2) of a perceptron from (a) noiseless and (b) noisy data (noise sampled from a flat distribution over the interval $[-0.2, 0.2]$). Whereas for noiseless data the estimates perfectly converge, for noisy data the estimates the system attempts to also include the noise, leading to a nontrivial distribution of the parameter estimates and rendering the extraction of the optimal parameters a nontrivial task. Open circles: location of the optimal parameters in the noiseless case.

mechanics system and (ii) the parameters that define the numerical integration scheme of Hamilton's equation in the molecular dynamics part of the HMC algorithm. Efficiency of HMC algorithms can be gained if the kinetic term is made dependent on the configuration geometry of the statistical mechanics system. If the Riemann geometry of the parameter space of statistical models is taken into account, then the simulated search of paths across this manifold samples the target density in a very efficient way [18]. Unfortunately, this procedure is both demanding and computationally costly, depending strongly on the quality of the space's extracted geometry.

Here we explore a computationally simpler approach that, to our knowledge, has never been applied in the context of Bayesian inference before. Depending on the number of discretization points needed to approximate the original SDE system and the number of measurement points, the dynamics of the statistical mechanics system happen on very different time scales. This suggests a multiple time-scale integration technique for the simulation of the statistical mechanics system [19]. We will show that for 1D SDE we can always find a parametrization that decouples the harmonic modes in between measurement points from both the measurement points and the model parameters and allows for an analytical time-saving solution of the fastest part of the dynamics. While for higher-dimensional problems it is not always possible to solve part of the dynamics analytically, we believe that scale separation alone will render many SDE amenable to a full-fledged Bayesian inference with time series. In fact, scale separation appears to be a generic feature if the dynamics of the SDE requires a large number of discretization points.

II. INFERENCE PROBLEM SETTING

Consider, for simplicity and concreteness, a reservoir dynamics $S(t)$ that on the observation time scale is linear, with other (inflow and outflow) processes happening at much shorter time scales, so they can be described by white noise. Furthermore, assume that this noise scales linearly with the system state $S(t)$. The model equation is thus given by the

SDE

$$\dot{S}(t) = r(t) - \frac{1}{K} \left(1 + \frac{\gamma}{2}\right) S(t) + \sqrt{\frac{\gamma}{K}} S(t) \eta(t), \quad (1)$$

where $r(t)$ denotes the time-varying rain input, K denotes the retention time, γ is the noise strength, and $\eta(t)$ indicates the white noise property, i.e.,

$$\langle \eta(t) \eta(t') \rangle = \delta(t - t'). \quad (2)$$

Equation (1) is to be understood in the Stratonovich sense [20].

Properties of a transformed version of (1) have been derived for constant input [21–23]. Here, suffice it to say that, for constant input $r(t) = r_0$, the equilibrium distribution, $P_{\text{eq}}(S)$, is an inverse Γ distribution with scale parameter $2Kr_0/\gamma$ and shape parameter $(2 + \gamma)/\gamma$ (see Sec. III A), i.e.,

$$P_{\text{eq}}(S) \propto S^{-2(1+\gamma)/\gamma} e^{-2Kr_0/(\gamma S)}. \quad (3)$$

The mean of this expression equals the equilibrium solution of the unperturbed system ($\gamma = 0$) $\langle S \rangle_{\text{eq}} = Kr_0$ and its variance, for $\gamma < 2$, is given by $\langle (S - \langle S \rangle_{\text{eq}})^2 \rangle_{\text{eq}} = K^2 r_0^2 \gamma / (2 - \gamma)$, which, for $\gamma \geq 2$, is seen to diverge. The power-law decay of the inverse Γ distribution is reminiscent of the invariance of Eq. (1) under rescaling of both $r(t)$ and $S(t)$. In real-world hydrology, for which Eq. (1) is a model, indeed often fat-tailed error distributions are observed [24].

While this equation was motivated by a popular hydrological model [25,26], it is by no means restricted to this context. By means of the transformation $S(t) = 1/n(t)$, Eq. (1) turns into a model that has been suggested, e.g., as a phenomenological description of the dynamics of the neutron density in nuclear reactors [21].

In our setting, the input $r(t)$ is a smooth and nowhere-vanishing function. We assume the observed time series, y_s , to be the outflow of the reservoir, $S(t)/K$, observed at times $0 = t_1 < t_2 < \dots < t_{n+1} = T$, with multiplicative independent log-normal errors,

$$\ln(y_s) = \ln \left[\frac{S(t_s)}{K} \right] + \sigma \epsilon_s, \quad s = 1, \dots, n+1, \quad (4)$$

where the ϵ_s are uncorrelated standard normal errors. For simplicity, we assume σ as well as the input $r(t)$ to be known, so we are left with the task of inferring parameter combinations (K, γ) that are compatible with the data given by Eq. (4) in the Bayesian sense, where knowledge about parameters is expressed in terms of probability distributions.

We start our treatise by assuming that we have prior knowledge about a parameter vector, θ , in the form of a probability distribution, $f_{\text{prior}}(\theta)$, and measured data, $\mathbf{y}$, believed to be a realization of the model. The posterior knowledge, combining prior knowledge with the one acquired from the data, is calculated by means of equation

$$f_{\text{post}}(\theta|\mathbf{y}) = \frac{f_{\text{prior}}(\theta) L(\mathbf{y}|\theta)}{\int f_{\text{prior}}(\theta') L(\mathbf{y}|\theta') d\theta'}, \quad (5)$$

where $L(\mathbf{y}|\theta)$ is the probability distribution for model outputs given model parameters, evaluated at the measured data (the infamous *likelihood function*).

Before we set out to derive from Eqs. (1), (2), and (4) the likelihood function, we express the parameters and state

variables by dimensionless quantities. Due to scale invariance of the noise term, γ is already dimensionless. State variable $S(t)$ and parameter K are replaced by the dimensionless quantities $q(t)$ and β , respectively, which are defined by the transformations $\beta = \sqrt{T\gamma/K}$ and $S(t) = [T\gamma r(t)/\beta^2]e^{\beta q(t)}$. In these new variables and parameters, Eq. (1) becomes the nonlinear SDE with constant noise

$$\dot{q}(t) = \frac{\beta}{T\gamma} e^{-\beta q(t)} - \frac{1}{T} \rho(t) + \frac{1}{\sqrt{T}} \eta(t), \quad (6)$$

with $\rho(t) = (T/\beta)(d/dt)\{\ln[r(t)]\} + (2 + \gamma)\beta/(2\gamma)$.

The probability $P(q_1, T|q_0, 0)$ of finding the system in a state q_1 at time $t = T$, if it was in an initial state q_0 at time $t = 0$, is expressed as a *path integral* as

$$P(q_1, T|q_0, 0) = \frac{1}{Z} \int e^{-S[q, \dot{q}]} \delta[q(T) - q_1] \delta[q(0) - q_0] \mathcal{D}q, \quad (7)$$

where the integral extends over all paths $q : [0, T] \rightarrow \mathbb{R}$ and where the path-measure $\mathcal{D}q$ is formally written as the infinite product $\mathcal{D}q = \prod_t dq(t)$. The *action* is a functional on the space of paths and reads [27]

$$S[q, \dot{q}] = \frac{1}{T} \int_0^T dt \left\{ \frac{1}{2} \left[T\dot{q}(t) + \rho(t) - \frac{\beta}{\gamma} e^{-\beta q(t)} \right]^2 - \frac{\beta^2}{2\gamma} e^{-\beta q(t)} \right\}. \quad (8)$$

This action includes the Jacobian that is introduced when changing coordinates from $\eta(t)$ to $q(t)$.

If we denote the parameter vector $\theta = (\beta, \gamma)^T$ and assume a flat prior, then the posterior (5) is, as a function of θ , proportional to the likelihood function

$$f_{\text{post}}(\theta|\mathbf{y}) \propto \int \exp \left(-\frac{1}{2} \sum_{s=1}^{n+1} \frac{\{\ln[y_s/r(t_s)] - \beta q(t_s)\}^2}{\sigma^2} - S[q, \dot{q}] \right) \mathcal{D}q. \quad (9)$$

Whereas the first term in the exponent describes the log-probability distribution of model outputs, for given model parameters, inputs, and a system realization $q(t_s)$, the second term is the log-probability of the associated system realization $q(t)$.

When applying this approach now to real-world problems, instead of undertaking a prohibitive numerical computation of the path integral, we apply HMC to sample parameter vectors from a joint distribution of system realizations and model parameters given by an appropriate discretization of the action of the path integral. By doing so, we observe that we obtain distinct regimes of time scales in the Hamiltonian that can be separated (see Sec. III B). This time-scale separation simplifies and boosts our algorithm; in many cases it even permits parts of the required integrations to be done analytically.

III. ALGORITHM

A. Inference algorithm

For the inference algorithm, it is necessary to rewrite action (8) with the help of the time-dependent potential $U(q, t) = \frac{1}{\gamma} e^{-\beta q} + q\rho(t)$ as

$$\begin{aligned} S[q, \dot{q}] &= \frac{1}{T} \int_0^T dt \left\{ \frac{1}{2} T^2 \dot{q}^2(t) + \frac{1}{2} \left[\rho(t) - \frac{\beta}{\gamma} e^{-\beta q(t)} \right]^2 \right. \\ &\quad \left. - T \frac{\partial U(q, t)}{\partial t} - \frac{\beta^2}{2\gamma} e^{-\beta q(t)} \right\} + U(q(T), T) - U(q(0), 0) \\ &= \frac{1}{T} \int_0^T dt \left\{ \frac{1}{2} T^2 \dot{q}^2(t) + \frac{1}{2} \left[\rho(t) - \frac{\beta}{\gamma} e^{-\beta q(t)} \right]^2 \right. \\ &\quad \left. - T q(t) \dot{\rho}(t) - \frac{\beta^2}{2\gamma} e^{-\beta q(t)} \right\} \\ &\quad + \frac{1}{\gamma} e^{-\beta q(T)} + q(T) \rho(T) - \frac{1}{\gamma} e^{-\beta q(0)} - q(0) \rho(0). \quad (10) \end{aligned}$$

With the action in this form, we easily derive the equilibrium distribution, for constant input $r(t) = r_0$, by plugging (7) and (10) into the detailed balance condition

$$P(q_1 t_1 | q_0 t_0) P_{\text{eq}}(q_0) = P(q_0 t_1 | q_1 t_0) P_{\text{eq}}(q_1), \quad (11)$$

and using the transformation $q(t) \rightarrow q(-t)$. We get, since $\dot{\rho}(t) = 0$,

$$P_{\text{eq}}(q) \propto e^{-2U(q)}.$$

Back-transformation to the original variables leads to Eq. (3).

For efficiently drawing parameter samples from (9), we interpret the latter as the partition function of a 1D statistical mechanics system and simulate its dynamics employing the HMC algorithm [14]. The model parameters θ are interpreted as additional dynamical degrees of freedom coupling to the system variables $q(t)$. Each degree of freedom, $q(t)$ and θ , is paired with a conjugate variable, $p(t)$ and π , respectively, so the system is defined by the Hamiltonian

$$\mathcal{H}_{\text{HMC}}(q, \theta; p, \pi) = K(p, \pi) + V(q, \theta), \quad (12)$$

where

$$K(p, \pi) = \int_0^T \frac{p^2(t)}{2m(t)} dt + \sum_{\alpha=1}^2 \frac{\pi_{\alpha}^2}{2m_{\alpha}}, \quad (13)$$

and $V(q, \theta)$ is the negative logarithm of the kernel of (9). The posterior (9) can then be expressed by the phase-space path integral

$$f_{\text{post}}(\theta|\mathbf{y}) \propto \int e^{-\mathcal{H}_{\text{HMC}}(q, \theta; p, \pi)} \mathcal{D}p \mathcal{D}q d\pi. \quad (14)$$

The HMC method, as a combination of the *Metropolis algorithm* [17] and *molecular dynamics* methods [15, 16], iterates the following steps:

(1) Momenta $p(t)$ and π are sampled from the Gaussian distributions defined by Eq. (13).

(2) The system is then allowed to evolve in $(q, \theta; p, \pi)$ phase space for an arbitrary time interval τ according to a volume-preserving and time-reversible solution of a discretized set of Hamilton equations.

(3) The discretization error on the energy preservation due to the previous step is corrected by a Metropolis acceptance or rejection step.

The last step is the standard Metropolis algorithm, while the first two steps permit arbitrarily large jumps in phase space, while maintaining an arbitrarily large acceptance rate. Each new phase-space configuration is associated with a combination of model parameters θ , which is compatible with the data in the Bayesian sense. Thus, omitting a possible burn-in, the parameter marginal of the simulated Markov chain of configurations represents a sample of the posterior probability distribution.

In order to simulate the dynamics of the Hamiltonian (12), we first need to discretize the path integral (14). Let us assume that the measurement time points $\{y_s\}_{s=1, \dots, n+1}$ of the time series (4) are equidistantly distributed on the time interval $[0, T]$, with $t_1 = 0$ and $t_{n+1} = T$. Each interval between two consecutive data points is further partitioned into j bins, such that we have a total of $nj + 1 = N \gg 1$ discretization points. The path integral (14) is then approximated by an ordinary integral, with the approximate path measure $\mathcal{D}p\mathcal{D}q \approx \prod_i dp_i dq_i$. The discretized versions of $K(p, \pi)$ and $V(q, \theta)$ are now given by

$$K(p, \pi) \approx \sum_{i=1}^N \frac{p_i^2}{2m_i} \Delta t + \sum_{\alpha=1}^2 \frac{\pi_\alpha^2}{2m_\alpha}, \quad (15)$$

$$V(q, \theta) \approx \frac{\Delta t}{T} \sum_{i=2}^N \left[\frac{1}{2} T^2 \dot{q}_i^2 + \frac{1}{2} \left(\rho_i - \frac{\beta}{\gamma} e^{-\beta q_i} \right)^2 - \frac{\beta^2}{2\gamma} e^{-\beta q_i} - T q_i \dot{\rho}_i \right] + \frac{1}{\gamma} e^{-\beta q_N} + q_N \rho_N - \frac{1}{\gamma} e^{-\beta q_1} - q_1 \rho_2 + \sum_{s=1}^{n+1} \frac{\{\ln[y_s/r_{(s-1)j+1}] - \beta q_{(s-1)j+1}\}^2}{2\sigma^2}, \quad (16)$$

with $\dot{q}_i = (q_i - q_{i-1})/\Delta t$, $\rho_i = T \ln[r(t_i)/r(t_{i-1})]/(\beta \Delta t) + (2 + \gamma)\beta/(2\gamma)$, and $\dot{\rho}_i = (\rho_i - \rho_{i-1})/\Delta t$ and where terms of order $\mathcal{O}(N^{-1/2})$ were neglected. Note that we did not apply the midpoint discretization that is associated with the Stratonovich convention. In Eq. (15) this leads to a different dynamics that does not, however, alter the posterior we are interested in, and in Eq. (16) we produce errors of the order $\mathcal{O}(N^{-1/2})$ that we neglect.

Physically, the discretized Hamiltonian can be identified with a classical polymer chain of N beads with harmonic bonds between neighboring beads in an external field [13]. The latter consists of two parts, a field that results from the measurements and is felt by the measurement beads only [last term on the right-hand side of Eq. (16)] and a field that results from the dynamics of the original Eq. (1) and is felt by all the beads. The masses m_i and m_α are tunable parameters of the algorithm. Since measurement beads are constrained more than intermediate beads, we will assign larger masses to the former. Figure 2 shows a typical realization of the dynamics of the polymer. Measurement beads only move

FIG. 2. Simulated polymer chain dynamics, with $n + 1 = 11$ data points (large circles) and $j - 1 = 9$ intermediate beads (small circles). For other parameters, see Secs. IV and V. Bottom: Initial state, where intermediate beads are on a linear interpolation between data points. Top: Polymer after $N_{MC} = 1000$ iterations of the propagation algorithm (dotted line: initial configuration). Clearly, the new configuration is mostly determined by the dynamics of the light-mass intermediate beads, while the heavy-mass data points move to a much lesser extent.

within the measurement uncertainty, while the intermediate beads explore much larger regions of phase space.

We have thus reduced the original Bayesian inference problem to simulating the dynamics of a linear polymer (cf. Fig. 2). Each state of this fictitious molecule corresponds to a well-defined configuration in the original phase space, characterized by a set of system variables $\{q_i\}_{i=1, \dots, N}$ and a parameter vector θ . It is now essential to note that potential (16) contains terms of distinct scaling in the potentially large numbers N and n , which refer to dynamics on distinct time scales. In particular, for large N , $V(q, \theta)$ is dominated by its harmonic part, and to resolve its dynamics, brute-force numerical integration of Hamilton's equations in step 2 of the HMC algorithm would require a very small discretization time step.

B. Time-scale separation

Whereas an interesting *approximate* approach would be to employ a partial averaging of the fast Fourier modes [28], we will use an *exact multiple time-scale integration* based on Trotter's formula [19]. For this, we introduce so-called *staging variables*, and diagonalize the harmonic part between the measurement points. To this end, we rewrite the discretized harmonic part of the action as

$$\sum_{i=2}^N \frac{T}{2\Delta t} (q_i - q_{i-1})^2 = \frac{T}{2} \sum_{s=1}^n \left\{ \frac{[q_{(s-1)j+1} - q_{sj+1}]^2}{j\Delta t} + \sum_{k=2}^j \frac{k}{(k-1)\Delta t} [q_{(s-1)j+k} - q_{(s-1)j+k}^*]^2 \right\}, \quad (17)$$

with $q_{(s-1)j+k}^* = [(k-1)q_{(s-1)j+k+1} + q_{(s-1)j+1}]/k$. The boundary beads, corresponding to the original measurement

points, are not transformed, $u_{sj+1} = q_{sj+1}$, $s = 0, \dots, n$, while, for the intermediate staging beads, we apply the coordinate transformations $u_{sj+k} = q_{sj+k} - q_{sj+k}^*$, $s = 0, \dots, n-1$, $k = 2, \dots, j$. Their inverse transformations are given by

$$q_{sj+1} = u_{sj+1}, \quad (18)$$

$$q_{sj+k} = \sum_{l=k}^{j+1} \frac{k-1}{l-1} u_{sj+l} + \frac{j-k+1}{j} u_{sj+1}, \quad (19)$$

which can be captured by the recursive relation

$$q_{sj+k} = u_{sj+k} + \frac{k-1}{k} q_{sj+k+1} + \frac{1}{k} u_{sj+1}. \quad (20)$$

The momenta are not transformed, which means we are using a noncanonical transformation. This alters only the dynamics of the system, not the posterior we are interested in.

We now split the Hamiltonian $\mathcal{H}_{\text{HMC}}$ into components according to their scaling behavior in n and N and write

$$\mathcal{H}_{\text{HMC}} = \mathcal{H}_N + \mathcal{H}_n + \mathcal{H}_1, \quad (21)$$

where

$$\mathcal{H}_N = \frac{1}{2} \sum_{s=1}^n \sum_{k=2}^j \left[\frac{\Delta t}{m'} p_{(s-1)j+k}^2 + \frac{Tk}{\Delta t(k-1)} u_{(s-1)j+k}^2 \right], \quad (22)$$

$$\begin{aligned} \mathcal{H}_n = & \frac{1}{2} \sum_{s=1}^{n+1} \left(\frac{\Delta t}{M} p_{(s-1)j+1}^2 \right. \\ & \left. + \frac{\{\ln[y_s/r_{(s-1)j+1}] - \beta u_{(s-1)j+1}\}^2}{\sigma^2} \right) \\ & + \frac{T}{2j\Delta t} \sum_{s=1}^n [u_{(s-1)j+1} - u_{sj+1}]^2, \end{aligned} \quad (23)$$

$$\begin{aligned} \mathcal{H}_1 = & \sum_{\alpha=1}^2 \frac{\pi_\alpha^2}{2m_\alpha} + \frac{\Delta t}{T} \sum_{i=2}^N \left\{ \frac{1}{2} \left(\rho_i - \frac{\beta}{\gamma} e^{-\beta q_i} \right)^2 \right. \\ & \left. - \frac{\beta^2}{2\gamma} e^{-\beta q_i} - T q_i \dot{\rho}_i \right\} + \frac{1}{\gamma} e^{-\beta q_N} + q_N \rho_N \\ & - \frac{1}{\gamma} e^{-\beta q_1} - q_1 \rho_2. \end{aligned} \quad (24)$$

Here we have introduced two masses, M and m' , for the boundary and staging beads, respectively. The scaling of these Hamiltonians can be derived from basic properties of discretized SDEs, from which we conclude that $u_i \sim \sqrt{\Delta t}$. Furthermore, in agreement with the *equipartition law*, we find that $p_i \sim 1/\sqrt{\Delta t}$. Accordingly, we find that the harmonic part (22), for the staging beads, scales linearly with N . The terms of Eq. (23), including both the harmonic part for the boundary beads and the measurement term, scale linearly with n . Finally, Eq. (24) scales with neither n nor N . Thanks to the staging variables, $\mathcal{H}_N$ and $\mathcal{H}_n$ have become fully decoupled. We use Trotter's formula [29] in order to design a reversible molecular dynamics integrator that takes the presence of the different

time scales into account. For an appropriate partition of the Hamiltonian, three distinct regimes can be distinguished:

- (i) $\mathcal{H}_N \sim \mathcal{H}_n \gg \mathcal{H}_1$,
- (ii) $\mathcal{H}_N \gg \mathcal{H}_n \sim \mathcal{H}_1$,
- (iii) $\mathcal{H}_N \gg \mathcal{H}_n \gg \mathcal{H}_1$.

In the following, we restrict ourselves to regime (ii), where the number of measurements n is assumed to be not too large and/or the measurement error σ to be not too small (the generalization of the method to the other regimes would, however, be straightforward). In this regime we may simply separate the harmonic part of the action for the staging beads from the rest and write

$$\mathcal{H}_{\text{HMC}} = \mathcal{H}_N + \mathcal{H}'. \quad (25)$$

For obtaining reversible integrators, we define the Liouville operators $iL_N = \{\cdot, \mathcal{H}_N\}$, $iL' = \{\cdot, \mathcal{H}'\}$, where $\{\cdot, \cdot\}$ denote the Poisson brackets that apply to functions on the phase space. Trotter's formula [30] allows us to write the Hamiltonian propagator as

$$e^{i(L_N+L')\tau} = [e^{iL_N(\Delta\tau/2)} e^{iL'\Delta\tau} e^{iL_N(\Delta\tau/2)}]^P + \mathcal{O}(\tau^3/P^2) \quad (26)$$

for $\tau = P\Delta\tau$. Here the outer propagator $\exp[iL_N(\Delta\tau/2)]$ reflects much faster dynamics than the inner one. However, thanks to our reparametrization, it describes the dynamics of uncoupled harmonic oscillators, which we can readily solve. Masses and frequencies of the oscillators are derived from (22) as

$$m = m'/\Delta t, \quad \omega_k = \sqrt{\frac{Nk}{(k-1)m}}. \quad (27)$$

The outer propagator becomes

$$\begin{aligned} u_{(s-1)j+k}(\Delta\tau/2) = & u_{(s-1)j+k}(0) \cos(\omega_k \Delta\tau/2) \\ & + \frac{p_{(s-1)j+k}(0)}{m\omega_k} \sin(\omega_k \Delta\tau/2), \end{aligned} \quad (28)$$

$$\begin{aligned} p_{(s-1)j+k}(\Delta\tau/2) = & p_{(s-1)j+k}(0) \cos(\omega_k \Delta\tau/2) \\ & - m\omega_k u_{(s-1)j+k}(0) \sin(\omega_k \Delta\tau/2), \end{aligned} \quad (29)$$

for $s = 1, \dots, n$ and $k = 2, \dots, j$. For the inner propagator, we employ the time-reversible and volume-preserving velocity Verlet algorithm [31], which leads for the boundary beads to

$$\begin{aligned} u_{(s-1)j+1}(\Delta\tau) = & u_{(s-1)j+1}(0) + \frac{\Delta\tau}{M} p_{(s-1)j+1}(0) \\ & + \frac{\Delta\tau^2}{2M} F_{(s-1)j+1}[\mathbf{u}(0), \boldsymbol{\theta}(0)], \end{aligned} \quad (30)$$

$$\begin{aligned} p_{(s-1)j+1}(\Delta\tau) = & p_{(s-1)j+1}(0) + \frac{\Delta\tau}{2} (F_{(s-1)j+1}[\mathbf{u}(0), \boldsymbol{\theta}(0)] \\ & + F_{(s-1)j+1}[\mathbf{u}(\Delta\tau), \boldsymbol{\theta}(\Delta\tau)]), \end{aligned} \quad (31)$$

with $s = 1, \dots, n+1$ and where $F_i[\mathbf{u}, \boldsymbol{\theta}]$ denotes the partial derivative of $\mathcal{H}'[\mathbf{u}, \boldsymbol{\theta}]$ with respect to u_i . Analogous equations emerge for the model parameters $\boldsymbol{\theta}$ and their momenta $\boldsymbol{\pi}$, by exchanging u by θ and p by π , along with the corresponding masses, respectively. For the staging beads, only the momenta need to be updated (because the associated kinetic term is

FIG. 3. System realization (solid line) with synthetic observations (filled circles, with error bars indicating the assumed measurement uncertainty). The system response closely follows the oscillations of the sinusoidal input (dashed) in a time-delayed manner. Parameters for this figure and all following figures: $K_{\text{true}} = 50$ and $\gamma_{\text{true}} = 0.2$.

not part of $\mathcal{H}'$ but of $\mathcal{H}_N$). Thus, with $s = 1, \dots, n$ and $k = 2, \dots, j$ as before,

$$p_{(s-1)j+k}(\Delta\tau) = p_{(s-1)j+k}(0) + \frac{\Delta\tau}{2} (F_{(s-1)j+k} \times [\mathbf{u}(0), \boldsymbol{\theta}(0)] + F_{(s-1)j+k}[\mathbf{u}(\Delta\tau), \boldsymbol{\theta}(\Delta\tau)]). \quad (32)$$

The propagators (28) through (32) are applied sequentially P times to calculate the system evolution over time τ . The proposed configuration, $(\mathbf{u}', \boldsymbol{\theta}'; \mathbf{p}', \boldsymbol{\pi}')$, is accepted with Metropolis probability $\min(1, e^{\mathcal{H}_{\text{HMC}}(\mathbf{u}, \boldsymbol{\theta}; \mathbf{p}, \boldsymbol{\pi}) - \mathcal{H}_{\text{HMC}}(\mathbf{u}', \boldsymbol{\theta}'; \mathbf{p}', \boldsymbol{\pi}')})$. The next iteration then starts with sampling a new momentum vector $(\mathbf{p}, \boldsymbol{\pi})$.

The analytical solution (28) and (29) is one main boosting part of our algorithm. To find such a solution, it is important to arrive at model equations of the form (6), where the noise term depends neither on the state variables nor on the parameters to be inferred. In a one-dimensional model this can always be achieved through reparametrization (see, e.g., chapter 5 in Ref. [32]). In higher dimensions, this will not always be possible. But even in such cases, we will be able to boost our

FIG. 4. Simulated system realizations associated with synthetic data, based on $n + 1 = 11$ measurement points and $N = 301$ discretization points.

FIG. 5. Markov chain evolution of the inferred parameter K .

algorithm through assigning smaller time intervals $\Delta\tau$ to the fast dynamics and larger ones to the slow dynamics.

IV. RESULTS

For our toy system, we have considered a simple sinusoidal input $r(t) = \sin^2(0.01t) + 0.1$. A system realization was first obtained from Eq. (1) using $K_{\text{true}} = 50$ (in arbitrary units of time) and $\gamma_{\text{true}} = 0.2$. Such system realization was then used to generate a synthetic time series of observed data according to Eq. (4). The error σ was set to 0.1. The input signal, the “true” system realization, and the corresponding data time series are shown in Fig. 3.

A set of 200 system realizations sampled from the integrand of Eq. (9), based on $n + 1 = 11$ measurement points and $N = 301$ discretization points, is shown in Fig. 4, together with the generated synthetic data. These samples were generated with a HMC algorithm, where the masses were set (in arbitrary units) to $M = 720$ for the measurement beads, to $m' = 130$ for the intermediate discretization beads, and to $m_\alpha = 150$ for both the dimensionless parameters β and γ . The different dynamics of the heavy measurement beads and the light discretization beads can be appreciated in Fig. 4.

The Markov chains for parameters K and γ , obtained after $N_{\text{MC}} = 50\,000$ iterations of the HMC algorithm, are shown in Figs. 5 and 6, respectively.

FIG. 6. Markov chain evolution of the inferred parameter γ .

FIG. 7. System dynamics in the phase space $K - \gamma$. The circle represents the initial state, while the square corresponds to the true parameter values used to generate the data.

The efficiency of the algorithm can be appreciated best by inspecting the system evolution in the phase space $K - \gamma$ (Fig. 7). The starting point of the algorithm was set to a linear interpolation of the data points, together with values for the model parameters that were deliberately chosen far off the truth (open circle in Fig. 7). Nevertheless, the very first step of the algorithm already takes the system to the vicinity of the true parameter values, where most of its dynamics then occurs. Few excursions lead far away from the true parameter values. These explore the heavy tails of the posterior parameter distribution.

The Markov chains (Figs. 5 and 6) determine the probability density functions (PDF) for K and γ , respectively. The results obtained by using the built-in kernel density estimator provided by MATHEMATICA (version 10) are exhibited in Fig. 8; they are fully compatible with the true, to-be-inferred, parameter values K_{true} and γ_{true} .

V. IMPLEMENTATION

The algorithm was implemented in C++ (version C++11) using the open-source ADEPT library (version 1.1; [33]), which provides a powerful tool for fast reverse-mode automated differentiation (AD). Our algorithm benefits greatly from the use of AD. This gives us the possibility to modify Eq. (1) and therefore the action (10), while leaving the implementation of

the algorithm unaltered. This makes our program extremely flexible and suitable for a much broader range of applications than the simple exemplary SDE model described here.

The simulations were run on both serial and parallel implementations of the algorithm on a 64-bit Linux system equipped with two 12-core 2.7-GHz processors (Intel Xeon E5-2697v2) and 64 GB of memory clocked at 1866 MHz. We used $n + 1 = 11$ measurement points and $N = 101, 201, \dots, 501$ discretization points. In the Hamiltonian propagator (26), we set $\Delta\tau = 0.25$ and $P = 3$, with a constant total observation time $T = 833$ (arbitrary units of time). The initial values of the parameters were set to $K = 200$ and $\gamma = 0.5$.

For example, a complete run with $N_{\text{MC}} = 50\,000$ iterations with $n = 10$ and $N = 301$ required about 43 s with the serial implementation of the algorithm. In the case of our toy system the burn-in phase is extremely short and can be safely ignored. Under these conditions the algorithm can be parallelized in a straightforward manner simply by breaking up the Markov chain into several smaller independent chains. The execution time with a fixed-size problem scales in a reasonably linear way with the number of processes (strong scaling). Our example could be therefore run in only about 3 s using 16 processors. An alternative strategy, suitable for long time series, would be to parallelize the updates of the polymer beads in each step of a single MC chain.

VI. CONCLUSIONS

We presented an extremely efficient and versatile approach for data-based SDE parameter estimation. Our algorithm obtains its strength from translating the problem of generating posterior parameter samples into the problem of simulating the dynamics of a statistical mechanics system; the main novelty in our algorithm is the exploitation of the fact that this dynamics generically happens on very different time scales. Furthermore, at least for 1D systems, our approach also allows for an analytical, and therefore computationally efficient, integration of the fastest part of the dynamics.

In most application cases, our choice of a fixed diagonal mass matrix, for reasonable choices of masses—heavier for the measurement beads and lighter for the discretization beads—can be expected to work well. Nonetheless, if the curvature of the potential varies strongly, then it might be beneficial to adapt the mass matrix to the local curvature as suggested in Ref. [18]. For such cases, a combination of the scale separation

FIG. 8. Probability density functions for the inferred parameters (a) K and (b) γ . The true values used to generate the data are represented by the dashed vertical lines.

method proposed in this paper with the local mass matrix adaptation of Ref. [18] might be the most efficient solution. This extension, however, comes at the price of a computational overhead (second derivatives have to be calculated and implicit equations have to be solved), and we will no longer be able to solve part of the dynamics analytically. On a more general level, given the wide field of very distinct applications, optimal parameter inference for SDE models will not be provided by one single approach but will require a set of tools to make the optimal choice from. We expect statistical physics to continue making strong contributions toward this aim.

The structure of our algorithm is well suited to parallelization, which is important in particular if we deal with a high number of measurements. The algorithm can easily be adapted to other inference problems, such as higher-dimensional SDE and SDE coupled to ordinary differential equations. These adaptations and extensions will, however, be addressed in future works.

ACKNOWLEDGMENT

This work was partly financed by the Eawag Discretionary Fund.

-
- [1] F. Rosenblatt, *Psychol. Rev.* **65**, 386 (1958).
 - [2] W.-H. Steeb, Y. Hardy, A. Hardy, and R. Stoop, *Problems & Solutions in Scientific Computing with C++ and Java Simulations* (World Scientific, Singapore, 2004).
 - [3] G. E. Box and G. C. Tiao, *Bayesian Inference in Statistical Analysis*, Vol. 40 (John Wiley & Sons, New York, 2011).
 - [4] N. Chopin, P. E. Jacob, and O. Papaspiliopoulos, *J. Roy. Stat. Soc. B* **75**, 397 (2013).
 - [5] X. Liu and M. Niranjana, *Bioinformatics* **28**, 1501 (2012).
 - [6] L. Tomassini, P. Reichert, H. R. Künsch, C. Buser, R. Knutti, and M. E. Borsuk, *J. Roy. Stat. Soc. C* **58**, 679 (2009).
 - [7] P. Reichert and J. Mieleitner, *Water Resour. Res.* **45**, W10402 (2009).
 - [8] J. Marin, P. Pudlo, C. Robert, and R. Ryder, *Stat. Comput.* **22**, 1167 (2012).
 - [9] C. Albert, H. R. Künsch, and A. Scheidegger, *Stat. Comput.* **25**, 1217 (2015).
 - [10] T. Toni, D. Welch, N. Strelkowa, A. Ipsen, and M. Stumpf, *J. Roy. Soc. Interface* **6**, 187 (2009).
 - [11] T. Toni and M. Stumpf, *Bioinformatics* **26**, 104 (2010).
 - [12] P. Fearnhead and D. Prangle, *J. Roy. Stat. Soc. B* **74**, 419 (2012).
 - [13] D. Chandler and P. G. Wolynes, *J. Chem. Phys.* **74**, 4078 (1981).
 - [14] S. Duane, A. D. Kennedy, B. J. Pendleton, and D. Roweth, *Phys. Lett. B* **195**, 216 (1987).
 - [15] B. J. Alder and T. E. Wainwright, *J. Chem. Phys.* **31**, 459 (1959).
 - [16] A. Rahman, *Phys. Rev.* **136**, A405 (1964).
 - [17] N. Metropolis, A. W. Rosenbluth, M. N. Rosenbluth, A. H. Teller, and E. Teller, *J. Chem. Phys.* **21**, 1087 (1953).
 - [18] M. Girolami and B. Calderhead, *J. Roy. Stat. Soc. B* **73**, 123 (2011).
 - [19] M. E. Tuckerman, B. J. Berne, G. J. Martyna, and M. L. Klein, *J. Chem. Phys.* **99**, 2796 (1993).
 - [20] R. L. Stratonovich, *Conditional Markov Processes and Their Application to the Theory of Optimal Control* (Elsevier, New York, 1968).
 - [21] W. L. Dutré and A. F. Debusscher, *Nucl. Sci. Eng.* **62**, 355 (1977).
 - [22] A. Schenzle and H. Brand, *Phys. Rev. A* **20**, 1628 (1979).
 - [23] H. Fujisaka, H. Ishii, M. Inoue, and T. Yamada, *Prog. Theor. Phys.* **76**, 1198 (1986).
 - [24] M. Thyer, B. Renard, D. Kavetski, G. Kuczera, S. W. Franks, and S. Srikanthan, *Water Resour. Res.* **45**, W00B14 (2009).
 - [25] A. Breinholt, F. O. Thordarson, J. K. Møller, M. Grum, P. S. Mikkelsen, and H. Madsen, *Environmetrics* **22**, 946 (2011).
 - [26] V. Livina, Y. Ashkenazy, Z. Kizner, V. Strygin, A. Bunde, and S. Havlin, *Physica A* **330**, 283 (2003).
 - [27] A. W. C. Lau and T. C. Lubensky, *Phys. Rev. E* **76**, 011123 (2007).
 - [28] J. D. Doll, R. D. Coalson, and D. L. Freeman, *Phys. Rev. Lett.* **55**, 1 (1985).
 - [29] M. Tuckerman, B. J. Berne, and G. J. Martyna, *J. Chem. Phys.* **97**, 1990 (1992).
 - [30] H. F. Trotter, *Proc. Am. Math. Soc.* **10**, 545 (1959).
 - [31] W. C. Swope, H. C. Andersen, P. H. Berens, and K. R. Wilson, *J. Chem. Phys.* **76**, 637 (1982).
 - [32] H. Risken, *The Fokker-Planck Equation: Methods of Solution and Applications*, 2nd ed. (Springer, Berlin, 1989).
 - [33] R. J. Hogan, *ACM Trans. Math. Softw.* **40**, 1 (2014).