

Article

Interacting Particle Solutions of Fokker–Planck Equations Through Gradient–Log–Density Estimation

Dimitra Maoutsa ^{1,*} , Sebastian Reich ²  and Manfred Opper ^{1,*} ¹ Artificial Intelligence Group, Technische Universität Berlin, Marchstraße 23, 10587 Berlin, Germany² Institute of Mathematics, University of Potsdam, Karl-Liebknecht-Str. 24/25, 14476 Potsdam, Germany; sereich@uni-potsdam.de

* Correspondence: dimitra.maoutsa@tu-berlin.de (D.M.); manfred.opper@tu-berlin.de (M.O.)

Received: 7 June 2020; Accepted: 15 July 2020; Published: 22 July 2020



Abstract: Fokker–Planck equations are extensively employed in various scientific fields as they characterise the behaviour of stochastic systems at the level of probability density functions. Although broadly used, they allow for analytical treatment only in limited settings, and often it is inevitable to resort to numerical solutions. Here, we develop a computational approach for simulating the time evolution of Fokker–Planck solutions in terms of a mean field limit of an interacting particle system. The interactions between particles are determined by the gradient of the logarithm of the particle density, approximated here by a novel statistical estimator. The performance of our method shows promising results, with more accurate and less fluctuating statistics compared to direct stochastic simulations of comparable particle number. Taken together, our framework allows for effortless and reliable particle-based simulations of Fokker–Planck equations in low and moderate dimensions. The proposed gradient–log–density estimator is also of independent interest, for example, in the context of optimal control.

Keywords: stochastic systems; Fokker–Planck equation; interacting particles; multiplicative noise; gradient flow; stochastic differential equations

1. Introduction

Many biological and physical systems are characterised by the presence of stochastic forces that influence their dynamics. These forces may be attributed either to intrinsic or extrinsic sources [1–3], i.e. arising either from random fluctuations of constituent subsystems [4,5], or from fluctuating interactions with the environment [6,7].

Often, deterministic analysis of these systems suffices to describe their macroscopic behaviour, and the fluctuations contribute only negligible perturbations around the deterministic dynamics. In systems biology, for example, rate equations describing mean concentrations of considered system's species have provided a useful description of chemical reaction networks' dynamics, and have enabled answering invaluable questions pertaining chemical systems comprising large numbers of molecules [8].

However, in several settings, the effect of stochastic forces influences considerably the resulting system's behaviour by qualitatively altering its evolution. In those settings, random fluctuations have to be accounted for, and thus stochastic analysis becomes essential [9,10]. Phenomena such as stochastic resonance [11,12], noise induced transitions [13–15], and stochastic synchronisation, to name a few, are prevalent in many physical systems, and highlight the importance of considering fluctuating forces in the analysis of a system's behaviour. Manifestations of these phenomena abound in nature and have been encountered in genetics [16], neuroscience [17,18], climate science [12,19,20] and other fields.

Indispensable tools for the analysis of stochastic systems are Kolmogorov equations, and in particular the Fokker–Planck equation (FPE) [21,22]. FPEs characterise the evolution of the probability density functions (PDF) for the state variables of dynamical systems described by stochastic differential equations (SDE). SDEs commonly arise in modeling random effects in systems with continuous state variables [23], or after diffusion approximation of Master equations for systems involving discrete state transitions in time [8,24,25]. The associated FPEs have been widely used for modelling stochastic phenomena in various fields, such as, for example, in physics, finance, biology, neuroscience, traffic flow [26].

Yet, although commonly used, explicit closed-form solutions of FPEs are rarely available [27], especially in settings where the underlying dynamics is nonlinear. In particular, exact analytic solutions may be obtained only for a restricted class of systems following linear dynamics perturbed by white Gaussian noise, and for some nonlinear Hamiltonian systems [21,28]. Further systems that admit analytical treatment independent of system dimension are those with discrete state transitions approximated via van Kampen expansion (linear noise approximation), resulting thus in linear SDEs with time dependent coefficients [9].

Existing numerical approaches for computing Fokker–Planck solutions may be grouped into three broad categories: grid-based, semi-analytical, and sample-based methods. The first category comprises mainly finite difference and finite element methods [29–32]. These frameworks, based on integration of FPEs employing numerical solvers for partial differential equations, entail computationally demanding calculations with inherent finite spatial resolution [33].

Conversely, semi-analytical approaches try to reduce the number of required computations by assuming conditional Gaussian structures [34], or by employing cumulant neglect closures [35], statistical linearisation [36,37], or stochastic averaging [38]. Although efficient in the settings they are devised for, their applicability is limited since the resulting solutions are imprecise or unstable in certain settings.

On the other hand, in the sample-based category, Monte Carlo methods resort to stochastic integration of a large number of *independent* stochastic trajectories that as an ensemble represent the probability density [39,40]. These methods are appropriate for computing *unbiased* estimates of exact expectations from empirical averages. Nevertheless, as we show in the following, cumulants of resulting distributions exhibit strong temporal fluctuations when the number of simulated trajectories is not sufficiently large.

Surprisingly, there is an alternative sample-based approach built on *deterministic* particle dynamics. In this setting, the particles are not independent, but they rather *interact* via an (approximated) probability density, and the FPE describes the mean field limit when their number grows to infinity. This approach introduces a bias in the approximated expectations, but significantly reduces the variance for a given particle number.

Recent research, see e.g., [41–44], has focused on particle methods for models of thermal equilibrium, where the stationary density is known analytically. For these models, interacting particle methods have found interesting new applications in the field of probabilistic Bayesian inference: by treating the Bayesian posterior probability density as the stationary density of a FPE, the particle dynamics provides posterior samples in the long time limit. For this approach, the particle dynamics are constructed by exploiting the gradient structure of the probability flow of the FPE. This involves the relative entropy distance to the equilibrium density as a Lyapunov function. Unfortunately, this structure does not apply to general FPEs in *non-equilibrium* settings, where the stationary density is usually unknown.

In this article, we introduce a framework for interacting particle systems that may be applied to general types of Fokker–Planck equations. Our approach is based on the fact that the instantaneous effective force on a particle due to diffusion is proportional to the *gradient* of the *logarithm* of the exact probability density (GLD). Rather than computing a differentiable estimate of this density (say by a kernel density estimator), we estimate the GLD directly without requiring knowledge of a stationary

density. Therefore, we introduce an approximation to the effective force acting on each particle, which becomes exact in the large particle number limit given the consistency of the estimator.

Our approach is motivated by recent developments in the field of machine learning, where GLD estimators have been studied independently and are used to fit probabilistic models to data. An application of these techniques to particle approximations for FPE is, to our knowledge, new. (The approach in [45] uses a GLD estimator different from ours for particle dynamics but with a probability flow towards equilibrium which is not given by a standard FPE.) Furthermore, our method provides straightforward approximations of entropy production rates, which are of primary importance in non-equilibrium statistical physics [46].

This article is organised as follows: Section 2 describes the deterministic particle formulation of the Fokker–Planck equation. Section 3 shows how a gradient of the logarithm of a density may be represented as the solution of a variational problem, while in Section 4 we discuss an empirical approximation of the gradient-log-density. In Section 5, we introduce function classes for which the variational problem may be solved explicitly, while in Section 6 we compare the temporal derivative of empirical expectations based on the particle dynamics with exact results derived from the Fokker–Planck equation. Section 7 is devoted to the class of equilibrium Fokker–Planck equations, where we discuss relations to Stein Variational Gradient Descent and other particle approximations of Fokker–Planck solutions. In Section 8, we show how our method may be extended to general diffusion processes with state-dependent diffusion, while Section 9 discusses how our framework may be employed to simulate second order Langevin dynamics. In Section 10 we demonstrate various aspects of our method by simulating Fokker–Planck solutions for different dynamical models. Finally, we conclude with a discussion and an outlook in Section 11.

2. Deterministic Particle Dynamics for Fokker–Planck Equations

We consider Fokker–Planck equations of the type

$$\frac{\partial p_t(x)}{\partial t} = -\nabla \cdot \left[f(x)p_t(x) - \frac{\sigma^2}{2} \nabla p_t(x) \right]. \quad (1)$$

Given an initial condition $p_0(x)$, Equation (1) describes the temporal development of the density $p_t(x)$ for the random variable $X(t) \in \mathbb{R}^d$ following the stochastic differential equation

$$dX(t) = f(X(t))dt + \sigma dB(t). \quad (2)$$

In Equation (2), $f(x) \in \mathbb{R}^d$ denotes the drift function characterising the deterministic part of the driving force, while $dB(t) \in \mathbb{R}^d$ represents the differential of a vector of independent Wiener processes capturing stochastic, Gaussian white noise excitations. For the moment, we restrict ourselves to state independent and diagonal *diffusion matrices*, i.e. diffusion matrices independent of $X(t)$ (additive noise) with diagonal elements σ^2 characterising the noise amplitude in each dimension. Extensions to more general settings are deferred to Section 8.

We may rewrite the FPE Equation (1) in the form of a *Liouville* equation

$$\frac{\partial p_t(x)}{\partial t} = -\nabla \cdot [g(x, t) p_t(x)] \quad (3)$$

for the *deterministic* dynamical system

$$\frac{dX}{dt} = g(X, t), \quad X(0) \sim p_0(x), \quad (4)$$

(dropping the time argument in $X(t)$ for simplicity) with velocity field

$$g(x, t) = f(x) - \frac{\sigma^2}{2} \nabla \ln p_t(x). \quad (5)$$

Hence, by evolving an ensemble of N independent realisations of Equation (4) (to be called ‘particles’ in the following) according to

$$\frac{dX_i}{dt} = g(X_i, t), \quad i = 1, \dots, N \quad X_i(0) \sim p_0(x), \quad (6)$$

we obtain an empirical approximation to the density $p_t(x)$.

Since the only source of randomness in Equation (4) can be attributed to the initial conditions $X_i(0)$, averages computed from the particle approximation (Equation (6)) are expected to have smaller variance compared to N independent realisations of the SDE (Equation (2)). Unfortunately, this approach requires perfect knowledge of the unknown instantaneous density $p_t(x)$ (Equation (5)), which is actually the unknown quantity of interest.

Here, we circumvent this issue by introducing *statistical estimators* for the term $\nabla \ln p_t(x)$, computed from the entire ensemble $(X_1(t), \dots, X_N(t))$ of particles at time t . Although this additional approximation introduces interactions among the particles via the estimator, for sufficiently large particle number N , fluctuations of the estimator are expected to be negligible and the limiting dynamics should converge to its mean field limit (Equation (4)) provided the estimator is asymptotically consistent. Thus, rather than computing a differentiable approximation to $p_t(x)$ from the particles, e.g. by a kernel density estimator, we show in the following section, how the function $\nabla \ln p_t(x)$ may be directly estimated from samples of $p_t(x)$.

3. Variational Representation of Gradient–Log–Densities

To construct a gradient–log–density (GLD) estimator we rely on a variational representation introduced by Hyvärinen in his *score–matching* approach for the estimation of non–normalised statistical models [47]. We favoured this approach over other estimators [48,49] due to its flexibility to adapt to different function classes chosen to approximate the GLD.

Here, we use a slightly more general representation compared to [47] allowing for an extra arbitrary reference function $r(x) = (r^{(1)}(x), \dots, r^{(d)}(x))$ such that the component α of the gradient is represented as

$$\partial_\alpha \ln p(x) = r^{(\alpha)}(x) + \arg \min_{\phi} \mathcal{L}_\alpha^r[\phi, p](x), \quad (7)$$

where $\partial_\alpha \doteq \frac{\partial}{\partial x^{(\alpha)}}$ stands for the partial derivative with respect to coordinate α of the vector $x \equiv (x^{(1)}, \dots, x^{(d)})$.

The cost function is defined as an expectation with respect to the density $p(x)$ by

$$\mathcal{L}_\alpha^r[\phi, p] = \int p(x) \left(\phi^2(x) + 2r^{(\alpha)}(x)\phi(x) + 2\partial_\alpha \phi(x) \right) dx, \quad (8)$$

with dx representing the volume element in R^d . To obtain this relation, we use integration by parts (assuming appropriate behaviour of densities and ϕ at boundaries), and get

$$\mathcal{L}_\alpha^r[\phi, p] = \int p(x) \left(\phi(x) + r^{(\alpha)}(x) - \partial_\alpha \ln p(x) \right)^2 dx - \int p(x) \left(\partial_\alpha \ln p(x) - r^{(\alpha)}(x) \right)^2 dx. \quad (9)$$

Minimisation with respect to ϕ yields Equation (7).

4. Gradient–Log–Density Estimator

To transform the variational formulation into a GLD estimator based on N sample points $(X_1, \dots, X_N)$, we replace the density $p(x)$ in Equation (8) by the empirical distribution $\hat{p}_t(x) = \frac{1}{N} \sum_{i=1}^N \delta(x - X_i(t))$, i.e.,

$$\mathcal{L}_\alpha^r[\phi, p_t] \approx \mathcal{L}_\alpha^r[\phi, \hat{p}_t] = \frac{1}{N} \sum_{i=1}^N \left(\phi^2(X_i) + 2r^{(\alpha)}(X_i)\phi(X_i) + 2\partial_\alpha \phi(X_i) \right), \quad (10)$$

and

$$\partial_\alpha \ln p_t(x) \approx r^{(\alpha)}(x) + \arg \min_{\phi \in \mathcal{F}} \mathcal{L}_\alpha^r[\phi, \hat{p}_t](x), \quad (11)$$

where $\mathcal{F}$ is an appropriately chosen family of functions with controllable complexity. By introducing the estimator of Equation (11) in Equation (6), we obtain a particle representation for the Fokker–Planck equation

$$\frac{dX_i^{(\alpha)}}{dt} = f^{(\alpha)}(X_i) - \frac{\sigma^2}{2} \left(r^{(\alpha)}(X_i) + \arg \min_{\phi \in \mathcal{F}} \mathcal{L}_\alpha^r[\phi, \hat{p}_t](X_i) \right), \quad (12)$$

for $i = 1, \dots, N$ and $\alpha = 1, \dots, d$, with

$$\hat{p}_t(x) = \frac{1}{N} \sum_{i=1}^N \delta(x - X_i), \quad X_i(0) \sim p_0(x).$$

Although in this article, we use reference functions $r(\cdot) \equiv 0$ for all simulated examples, the choice $r(x) = \frac{2}{\sigma^2} f(x)$, which cancels the first two terms in Equation (12), leads to interesting relations with other particle approaches for simulating Fokker–Planck solutions for equilibrium systems (c.f. Section 7) and impacts on the numerical approximation of ϕ (c.f. Section 5). In particular, when considering Brownian dynamics, where $f(x) = -\nabla U(x)$, the choice $r(x) = -\frac{2}{\sigma^2} \nabla U(x)$ leads to $\phi(x) = 0$ once the system has reached thermal equilibrium. More generally speaking, one may choose $r(x) = \frac{2}{\sigma^2} \nabla \ln p_*$, where p_* is an appropriate reference measure such as the equilibrium measure of the underlying stochastic process.

Estimating the Entropy Rate

Interestingly, the variational approach provides us with a simple, built in method for computing the entropy rate (temporal change of entropy) of the stochastic process (Equation (2)).

Using the FPE (1) and integration by parts, one can derive the well-known relation, see e.g., [50],

$$-\frac{d}{dt} \int p_t(x) \ln p_t(x) dx = \frac{\sigma^2}{2} \sum_{\alpha=1}^d \int p_t(x) (\partial_\alpha \ln p_t(x))^2 dx + \int p_t(x) \nabla \cdot f(x) dx. \quad (13)$$

The first term on the right hand side is usually called entropy production, whereas the second term corresponds to the entropy flux. In the stationary state, the total entropy rate vanishes. For equilibrium dynamics, both terms vanish individually at stationarity. This should be compared to the minimum of the cost function (Equation (9)), which for $r \equiv 0$, equals

$$\min_{\phi} \mathcal{L}_\alpha^0[\phi, p_t] = - \int p_t(x) (\partial_\alpha \ln p_t(x))^2 dx. \quad (14)$$

Thus, we obtain the estimator

$$-\frac{d}{dt} \int p_t(x) \ln p_t(x) dx \approx -\frac{\sigma^2}{2} \sum_{\alpha=1}^d \min_{\phi} \mathcal{L}_\alpha^0[\phi, \hat{p}_t] + \frac{1}{N} \sum_{i=1}^N \nabla \cdot f(X_i). \quad (15)$$

We will later see for the case of equilibrium dynamics that a similar method may be employed to approximate the relative entropy distance to the equilibrium density.

5. Function Classes

In the following, we discuss choices for families of functions $\mathcal{F}$ leading to explicit, closed form solutions for estimators.

5.1. Linear Models

A simple possibility is to choose linearly parametrised functions of the form

$$\phi(x) = \sum_{k=1}^m a_k \phi_k(x), \quad (16)$$

where the $\phi_k(x)$ are appropriate basis functions, e.g., polynomials, radial basis functions or trigonometric functions. For this linear parametrisation, the empirical cost (Equation (10)) is quadratic in the parameters a_k and can be minimised explicitly. A straightforward computation from Equation (12) shows that

$$\frac{dX_i}{dt} = f(X_i) - \frac{\sigma^2}{2} r(X_i) + \frac{\sigma^2}{2} \sum_{k,j=1}^m (C^{-1})_{kj} \phi_k(X_i) \sum_{l=1}^N \{ \nabla \phi_j(X_l) + \phi_j(X_l) r(X_l) \}, \quad (17)$$

with $C_{kl} = \sum_{i=1}^N \phi_k(X_i) \phi_l(X_i)$.

Here, we require the number of samples to be greater than the number of employed basis functions, i.e., $N \geq m + 1$, to have a non-singular matrix C . This restriction may be lifted by adding a penalty to the empirical cost (Equation (10)) for regularisation, similar to ridge regression estimators. Equation (17) is independent of the reference function r , when r belongs to the linear span of the selected basis functions. However, this model class with a finite parameter number has limited complexity.

When the particle number N grows large, we expect convergence of the dynamics to a mean field limit which would be given by

$$\frac{dX_i}{dt} = f(X_i) - \frac{\sigma^2}{2} r(X_i) + \frac{\sigma^2}{2} \sum_{k,j=1}^m (C^{-1})_{kj} \phi_k(X_i) \langle \nabla \phi_j(x) + \phi_j(x) r(x) \rangle_{q_t},$$

where the brackets $\langle \cdot \rangle_{q_t}$ denote expectation with respect to the limiting density $q_t(x)$ and $C_{kl} = \langle \phi_k(X_i) \phi_l(X_i) \rangle_{q_t}$. Since the linear model class (Equation (16)) exhibits limited complexity for fixed m , we do not expect the approximated solution $q_t(x)$ to equal the exact solution $p_t(x)$ of the FPE. Nevertheless, for rare cases, where both $\nabla \ln p_t(x)$ and also $r(x)$ are linear combinations of the employed basis functions for all times t , $q_t(x)$ would provide an exact solution. For example, in a setting with linear drift function $f(x) = -\gamma x$, reference function $r(x) = 0$, and dimensionality $d = 1$, a basis consisting of a constant $\phi_1(x) = 1$ and a linear function $\phi_2(x) = x$, would be able to perfectly represent the GLD of the Gaussian density $p_t(x)$.

5.2. Kernel Approaches

Here, we consider a family $\mathcal{F}$ of functions for which the effective number of parameters to be computed is not fixed beforehand, but rather increases with the sample number N : a *reproducing kernel Hilbert space* (RKHS) of functions defined by a positive definite (Mercer) kernel $K(\cdot, \cdot)$. Statistical models based on such function spaces have played a prominent role in the field of machine learning in recent years [51].

A common, kernel-based approach to regularise the minimisation of empirical cost functions is via penalisation using the RKHS norm $\|\cdot\|_{\text{RKHS}}$ of functions in $\mathcal{F}$. This can also be understood as a penalised version of a linear model (Equation (16)) with infinitely many feature functions ϕ_k . For so-called universal kernels [52], this unbounded complexity suggests that we could expect asymptotic convergence of the GLD estimator (see [53] for related results) and a corresponding convergence of the particle model to the FPE as its mean field limit. However, a rigorous proof may not be trivial, since particles in our setting are not independent.

The explicit form of the kernel-based approximation is given by

$$\partial_\alpha \ln p(x) \approx r^{(\alpha)}(x) + \arg \min_{\phi \in \mathcal{F}} \left\{ \mathcal{L}_\alpha^r[\phi, \hat{p}] + \frac{\lambda}{N} \|\phi\|_{\text{RKHS}}^2 \right\}(x), \quad (18)$$

where the parameter λ controls the strength of the penalisation. Again, this optimisation problem can be solved in closed form in terms of matrix inverses. One can prove a *representer theorem* which states that the minimiser $\phi(x)$ in Equation (18) is a linear combination of kernel functions evaluated at the sample points X_i , i.e.,

$$\phi(x) = \sum_{i=1}^N a_i K(x, X_i). \quad (19)$$

For such functions, the RKHS norm is given by

$$\|\phi\|_{\text{RKHS}}^2 = \sum_{i,j=1}^N a_i a_j K(X_i, X_j). \quad (20)$$

Hence, this representation leads again to a quadratic form in the N coefficients.

The solution of the minimisation problem is given by

$$a_j = - \sum_{k=1}^N \left((K^2 + \lambda K)^{-1} \right)_{jk} \sum_{l=1}^N \left\{ \partial_{\alpha_l} K(X_l, X_k) + K(X_l, X_k) r^{(\alpha)}(X_l) \right\}, \quad (21)$$

where $K_{ij} \doteq K(X_i, X_j)$. Similar approaches for kernel-based GLD estimators have been discussed in [48,49]. For $r(\cdot) = 0$, Equation (21) agrees with the GLD estimator of [48] derived by inverting *Stein's* equation, or by minimising the *Kernelised Stein discrepancy*.

The resulting particle dynamics is given by

$$\frac{dX_i}{dt} = f(X_i) - \frac{\sigma^2}{2} r(X_i) + \frac{\sigma^2}{2} \sum_{k=1}^N \left((K + \lambda I)^{-1} \right)_{ik} \sum_{l=1}^N \left\{ \nabla_l K(X_l, X_k) + K(X_l, X_k) r(X_l) \right\}. \quad (22)$$

Please note that here the inverse matrix also depends on the particles X_k .

We may simplify Equation (22) by adding and subtracting a term $\lambda \delta_{kl} r(X_l)$ in the summation over l , with δ_{kl} denoting the Kronecker delta. This yields

$$\frac{dX_i}{dt} = f(X_i) + \frac{\sigma^2}{2} \sum_{k=1}^N \left((K + \lambda I)^{-1} \right)_{ik} \left\{ \sum_{l=1}^N \nabla_l K(X_l, X_k) - \lambda r(X_k) \right\}. \quad (23)$$

In the limit of small λ , the right hand side becomes independent of the reference function r .

In the present article, we employ Gaussian radial basis function (RBF) kernels given by

$$K(x, x') = \exp \left[-\frac{1}{2l^2} \|x - x'\|^2 \right], \quad (24)$$

with a length scale l . A different possibility would be given by kernels with a *finite dimensional* feature representation

$$K(x, x') = \sum_{j=1}^m \phi_j(x) \phi_j(x'), \quad (25)$$

which may also be interpreted as a *linear model* as in Equation (16) with a L_2 penalty on the unknown coefficients.

5.3. A Sparse Kernel Approximation

Since the inversions of the $N \times N$ matrices in Equation (22) have to be performed at each step of a time discretised ODE system (Equation (22)), for large particle number N , the cubic complexity becomes too computationally demanding. Hence, here, we resort to a well established approximation in machine learning to overcome this issue, by applying a sparse approximation to the optimisation problem of Equation (18), see e.g., [54]. In particular, we introduce a smaller set of $M \ll N$ *inducing points* $\{z_k\}_{k=1}^M$ that need not necessarily be a subset of the N particles. We then minimise the penalised cost function (Equation (18)) in the finite dimensional family of functions

$$\phi(x) = \sum_{i=1}^M a_i K(x, z_i). \quad (26)$$

This may also be understood as a special linear parametric approximation. To keep matrices well conditioned, in practice we add a small ‘jitter’ term to Equation (18), i.e. we use

$$\lambda \|\phi\|_{\text{RKHS}}^2 + \epsilon \|\phi\|_2^2, \quad (27)$$

as the total penalty. In the limit $\lambda, \epsilon \rightarrow 0$, this representation reduces to an approximation of the form of Equation (16) with M basis functions $K(\cdot, z_l)$ for $l = 1, \dots, M$.

By introducing the matrices

$$K_{kl}^{zz} \doteq K(z_k, z_l) + \epsilon \delta_{kl}, \quad K_{ij}^{xz} \doteq K(X_i, z_j), \quad (28)$$

and

$$A \doteq K^{xz} \left[(\lambda + \epsilon) I + (K^{zz})^{-1} (K^{xz})^\top (K^{xz}) \right]^{-1} (K^{zz})^{-1}, \quad (29)$$

we replace the particle dynamics of Equation (22) by

$$\frac{dX_i}{dt} = f(X_i) - \frac{\sigma^2}{2} r(X_i) + \frac{\sigma^2}{2} \sum_k A_{ik} \sum_l \{ \nabla_l K(X_l, z_k) + K(X_l, z_k) r(X_l) \}. \quad (30)$$

Hence, for this approximation we have to invert only $M \times M$ matrices. For fixed M , the complexity of the GLD estimator is limited. Results for log-density estimators in machine learning (obtained for independent data) indicate that for a moderate growth of the number of inducing points M with the number of particles N , similar approximation rates may be obtained as for full kernel approaches.

6. A Note on Expectations

In this section, we present a preliminary discussion of the quality of the particle method to approximate expectations of scalar functions h of the random variable $X(t)$. We concentrate on the temporal development of $h(X(t))$. While it would be important to obtain an estimate of the approximation error over time, we will defer such an analysis to future publications and only concentrate on a result for the first time derivative of expectations, i.e., the evolution over infinitesimal times.

Using the FPE (Equation (1)) and integrations by parts, one derives the exact result

$$\frac{d\langle h(X) \rangle}{dt} = \langle L_x h(X) \rangle, \quad (31)$$

where $\langle \cdot \rangle$ denotes the expectation with respect to $p_t(x)$ and the operator L_x is defined as

$$L_x \doteq f(X) \cdot \nabla + \frac{\sigma^2}{2} \nabla^2. \quad (32)$$

In Equation (32), L_x denotes the *generator* of the diffusion process defined by the corresponding stochastic differential equation (Equation (2)). To obtain a related result for the particle dynamics and the empirical expectations denoted by $\langle \cdot \rangle_{\hat{p}_t}$, we employ the relation

$$\frac{d\langle h(X) \rangle_{\hat{p}_t}}{dt} = \langle \nabla h(X) \cdot \frac{dX}{dt} \rangle_{\hat{p}_t} = \frac{1}{N} \sum_{i=1}^N \nabla h(X_i) \cdot \frac{dX_i}{dt}, \quad (33)$$

where we have used the chain rule for the time derivative.

We will focus initially on the dynamics based on basis functions. Expressing the sum over X_i in Equation (17) as an expectation, we obtain

$$\frac{dX_i}{dt} = f(X_i) - \frac{\sigma^2}{2} r(X_i) + \frac{N\sigma^2}{2} \sum_{k,j=1}^m (C^{-1})_{kj} \phi_k(X_i) \langle (r(X) + \nabla) \phi_j(X) \rangle_{\hat{p}_t}. \quad (34)$$

Hence, by inserting Equation (34) into Equation (33) and adding and subtracting the term $\frac{\sigma^2}{2} \langle \nabla^2 h(X) \rangle_{\hat{p}_t}$ we obtain

$$\frac{d\langle h(X) \rangle_{\hat{p}_t}}{dt} = \langle L_x h(X) \rangle_{\hat{p}_t} + \Delta, \quad (35)$$

where the remainder term is given by

$$\Delta = \frac{\sigma^2}{2} \langle (r(X) + \nabla) \cdot (\hat{\nabla} h(x) - \nabla h(x)) \rangle_{\hat{p}_t}. \quad (36)$$

For the approximation of the vectorial function $\nabla h(x)$ we have

$$\hat{\nabla} h(x) = \sum_{j,k=1}^m \phi_j(x) \left(C^{-1} \right)_{jk} \sum_{i=1}^N \phi_k(X_i) \nabla h(X_i). \quad (37)$$

A simple comparison shows that each component of the vector $\hat{\nabla} h(x)$ can be written as the minimiser of $\sum_{l=1}^N (\Phi(X_l) - \partial_\alpha h(X_l))^2$ where $\Phi(x)$ is a linear combination of basis functions. Hence, $\hat{\nabla} h(x)$ equals the best approximation of the vectorial function $\nabla h(x)$ based on the 'data' $\nabla h(X_l)$ using regression with basis functions. Thus, if $\nabla h(x)$ is well approximated by basis functions, the remainder Δ is small. If indeed $\nabla h(x) = \sum_{n=1}^M c_n \phi_n(x)$, for some $c_n \in \mathbb{R}^d$, the remainder term vanishes, $\Delta = 0$. By its similarity to the finite basis function model, this result should also be valid for the sparse kernel dynamics of Equation (30), when the penalty λ is small. One might conjecture that the temporal development of expectations for reasonably smooth functions might be faithfully represented by the particle dynamics. This conjecture is supported by our numerical results.

A similar result also holds for the dynamics of Equation (22). In this case, the function $\hat{\nabla} h(x)$ in the remainder Δ (Equation (36)) is given by

$$\hat{\nabla} h(x) = \sum_{j,k=1}^N K(x, X_j) \left((K + \lambda I)^{-1} \right)_{jk} \nabla h(X_k), \quad (38)$$

which has also an interpretation as an approximation of $\nabla h(x)$ by regularised kernel regression.

7. Equilibrium Dynamics

An important class of stochastic dynamical systems describe *thermal equilibrium*, for which the drift function f is the negative gradient of a potential U , while the limiting equilibrium density p_∞ is explicitly given by a Gibbs distribution:

$$f(x) = -\nabla U(x) \quad (39)$$

$$\nabla \ln p_\infty(x) = \frac{2}{\sigma^2} f(x). \quad (40)$$

For this class of models, our method provides a simple and built-in estimator for the relative entropy between the instantaneous and the equilibrium density, p_t and p_∞ respectively. As we discuss here, our framework may also be related to two other particle approaches that converge to the (approximate) equilibrium density.

7.1. Relative Entropy

The relative entropy or *Kullback–Leibler divergence* is defined as

$$D(p_t|p_\infty) \doteq \int p_t(x) \ln \frac{p_t(x)}{p_\infty(x)} dx. \quad (41)$$

Following a similar calculation that led to Equation (13), we obtain

$$\frac{d}{dt} D(p_t|p_\infty) = -\frac{\sigma^2}{2} \int p_t(x) \|\nabla \ln p_t(x) - \nabla \ln p_\infty(x)\|^2 dx = -\frac{2}{\sigma^2} \int p_t(x) \|g(x, t)\|^2 dx, \quad (42)$$

where $g(x, t)$ indicates the velocity field of the particle system defined in Equation (4). The first equality holds for arbitrary drift functions. To obtain the second equality, we have inserted the explicit result for p_∞ .

Hence, we may compute the relative entropy at any time T as a time integral

$$D(p_T|p_\infty) = D(p_0|p_\infty) - \frac{2}{\sigma^2} \int_0^T \left\{ \int p_t(x) \|g(x, t)\|^2 dx \right\} dt, \quad (43)$$

where the inner expectation is easily approximated by our particle algorithm. This result shows that the exact velocity field $g(x, t)$ converges to 0 for $t \rightarrow \infty$, and one expects particles to also converge to fixed points. For other non-equilibrium systems, asymptotic fixed points are, however, the exception.

7.2. Relation to Stein Variational Gradient Descent

Recently, *Stein variational gradient descent* (SVGD), a kernel-based particle algorithm, has attracted considerable attention in the machine learning community [55,56]. The algorithm is designed to provide approximate samples from a given density p_∞ as the asymptotic fixed points of a deterministic particle system. Setting $-\ln p_\infty(x) = U(x) + \text{const}$, SVGD is based on the dynamics

$$\frac{dX_i}{dt} = \sum_l \{ -K(X_i, X_l) \nabla U(X_l) + \nabla_l K(X_i, X_l) \}. \quad (44)$$

This may be compared to our approximate FPE dynamics (Equation (22)) for the equilibrium case by setting $\sigma^2 = 2$ and $r(x) = f(x) = -\nabla U(x)$. For this setting, both algorithms have in fact, the same conditions

$$\sum_l \{ -K(X_i, X_l) \nabla U(X_l) + \nabla_l K(X_i, X_l) \} = 0, \quad (45)$$

for the ‘equilibrium’ fixed points. See [44] for a discussion of these fixed points for different kernel functions. However, both dynamics differ for finite times t , where a single time step of SVGD is computationally simpler, being free of the matrix inversion required by our framework. The mean field limit $N \rightarrow \infty$ of Equation (44) differs from the FPE limit, and the resulting partial differential equation is nonlinear [57]. Nevertheless, it is possible to interpolate between the two particle dynamics. In fact, in the limit of a large regularisation parameter $\lambda \rightarrow \infty$, the inverse matrix in Equation (22) becomes diagonal, i.e., $(K + \lambda I)^{-1} \simeq \frac{1}{\lambda} I$, and we recover SVGD (Equation (44)) by introducing a rescaled time $\tau \doteq t/\lambda$. This result could be of practical importance when the goal is to approximate the stationary distribution, irrespective of the finite-time dynamics. The SVGD combines faster matrix operations with slower relaxation times to equilibrium compared to the FPE dynamics. It would be interesting to see if an optimal computational speed of a particle algorithm might be achieved at some intermediate regularisation parameter λ .

7.3. Relation to Geometric Formulation of FPE Flow

Following Otto [58] and Villani [59], the FPE for the equilibrium case can be viewed as a gradient flow on the manifold of probability densities with respect to the Wasserstein metric. This formulation can be used to define an implicit Euler time discretisation method for the dynamics of the density p_t . For small times δt (and $\sigma^2 = 2$) this is given by the variational problem

$$p_{t+\delta t} = \arg \inf_p \left(W_2^2(p, p_t) + \delta t D(p \| p_\infty) \right) \quad (46)$$

in terms of the Kullback–Leibler divergence and the L_2 Wasserstein distance W_2 . The latter gives the minimum of $\langle \|X - X(t)\|^2 \rangle$ for two random variables $X(t)$ and X , where the expectation is over the joint distribution with fixed marginals p_t and p . Using the dual formulation for a regularised Wasserstein distance, approximate numerical algorithms for solving Equation (46) have been developed by [60] and by [61] with applications to simulations of FPE.

We show in the following that Equation (46) may be cast into a form closely related to our variational formulation (Equation (7)) for $r(x) = -\nabla U(x)$. Assuming that X and $X(t)$ are related through deterministic (transport) mappings of the form

$$X = X(t) + \delta t \nabla \psi(X(t)), \quad (47)$$

we may represent the Wasserstein distance in terms of ψ and the variational problem in Equation (46) may be reformulated as

$$\psi^* = \arg \min_{\nabla \psi} \frac{\delta t^2}{2} \int \|\nabla \psi(x)\|^2 p_t(x) dx + \delta t D(p_{t+\delta t} \| p_\infty). \quad (48)$$

where

$$p_{t+\delta t}(x) = p_t(x) - \delta t \nabla \cdot (p_t(x) \nabla \psi(x)) + O(\delta t^2). \quad (49)$$

To proceed, we expand the relative entropy in Equation (48) to first order in δt , inserting the representation of Equation (49) for $p_{t+\delta t}(x)$, thereby obtaining

$$\frac{\delta t}{2} \int \|\nabla \psi(x)\|^2 p_t(x) dx + D(p_{t+\delta t} \| p_\infty) = D(p_t \| p_\infty) + \quad (50a)$$

$$+ \frac{\delta t}{2} \left(\int p_t(x) \left\{ \|\nabla \psi(x)\|^2 - 2 \nabla^2 \psi(x) + 2 \nabla U(x) \cdot \nabla \psi(x) \right\} dx \right) + O(\delta t^2). \quad (50b)$$

Minimisation ignoring the $O(\delta t^2)$ terms (employing integration by parts) yields

$$\nabla \psi^*(x) = -\nabla U(x) - \nabla \ln p_t(x), \quad (51)$$

which is related to our cost function (Equation (8)) if we formally identify $\phi(x) = -\partial_\alpha \psi(x)$. More precisely, by replacing p_t by samples, the empirical cost function may be regularised with a RKHS norm penalty resulting in a nonparametric estimator for unnormalised log-density $\psi^*(x) = -\ln p_t(x) - U(x) + \text{const}$, as shown in [62]. One could use this estimator as an alternative to our approach. This would lead to a simultaneous estimate of all components of the GLD. In our approach, each of the d components of the gradient is computed individually. In this way, we avoid additional second derivatives of kernels, which would increase the dimensionality of the resulting matrices.

8. Extension to General Diffusion Processes

The Fokker–Planck equations for an SDE with arbitrary drift $f(x)$ and general, state-dependent diffusion matrix $D(x)$ is given by

$$\frac{\partial p_t(x)}{\partial t} = \nabla \cdot \left[-f(x)p_t(x) + \frac{1}{2} \nabla \cdot (D(x)p_t(x)) \right]. \quad (52)$$

This may again be written in the form of a Liouville equation (Equation (3)) where the effective force term equals

$$g(x, t) = f(x) - \frac{1}{2} \nabla \cdot D(x) - \frac{1}{2} D(x) \nabla \ln p_t(x). \quad (53)$$

9. Second Order Langevin Dynamics (Kramer's Equation)

For second-order Langevin equations, the system state comprises positions $X \in \mathbb{R}^d$ and velocities $V \in \mathbb{R}^d$ following the coupled SDE

$$dX = V dt \quad (54)$$

$$dV = (-\gamma V + f(X)) dt + \sigma dB_t. \quad (55)$$

In Equation (54), the dynamics describe the effect of a friction force, γV , an external force, $f(X)$, and a fluctuating force, where γ denotes the dissipation constant. In this setting, the effective *deterministic* ODE system is given by

$$\begin{aligned} \frac{dX}{dt} &= V \\ \frac{dV}{dt} &= -\gamma V + f(X) - \frac{\sigma^2}{2} \nabla_v \ln p_t(X, V). \end{aligned} \quad (56)$$

Considering here the equilibrium case, we set $f(x) = -\nabla U(x)$ for which the stationary density equals

$$\ln p_\infty(X, V) = -\beta \left(\frac{\|V\|^2}{2} + U(X) \right) \equiv -\beta H(X, V), \quad (57)$$

where $\beta = \frac{2\gamma}{\sigma^2}$ and $H(x, v) = \frac{\|v\|^2}{2} + U(x)$ denotes the *Hamiltonian* function. Inserting p_∞ into Equation (56), we find that for $t \rightarrow \infty$, the damping and the density-dependent part of the force cancel and we are left with pure Hamiltonian dynamics

$$\begin{aligned} \frac{dX}{dt} &= V \\ \frac{dV}{dt} &= -\nabla U(X), \end{aligned} \quad (58)$$

for which all particles become completely *decoupled*, with each one conserving energy separately. Of course, this result also precludes fixed point solutions to the particle dynamics. However,

this limiting dynamics captured by Equation (58) assumes the mean field limit $N \rightarrow \infty$ together with a consistent estimate of the GLD before taking the limit $t \rightarrow \infty$. For GLD estimators at finite N , we expect reasonably small stationary fluctuations of individual particle energies, which were also evident in our numerical experiments.

The exact asymptotic behaviour is also reflected in the expression for the change of the relative entropy for Kramer's equation. Similar to Equation (42) we obtain

$$\begin{aligned} \frac{d}{dt} D(p_t | p_\infty) &= -\frac{\sigma^2}{2} \int p_t(x, v) \|\nabla_v \ln p_t(x, v) - \nabla_v \ln p_\infty(x, v)\|^2 dx dv \\ &= -\frac{2}{\sigma^2} \int p_t(x, v) \|\gamma v + \frac{\sigma^2}{2} \nabla_v \ln p(x, v)\|^2 dx dv. \end{aligned} \quad (59)$$

When the system approaches equilibrium, both terms in the norm cancel out and the entropy production rate converges to 0.

10. Simulating Accurate Fokker–Planck Solutions for Model Systems

To demonstrate the accuracy of our approach, we simulated solutions of FPEs for a range of model systems and compared the results with those obtained from direct stochastic simulations (Monte Carlo sampling) with the same particle number, as well as with analytic solutions where relevant. We tested our framework on systems with diverse degrees of nonlinearity and dimensionality, as well as with various types of noise (additive/multiplicative). We quantified the accuracy of transient and steady state solutions resulting from our method in terms of 1-Wasserstein distance [59] and Kullback–Leibler (KL) divergence (Appendices C and D), along with the squared error of distances between distribution cumulants. For evaluating particle solutions for nonlinear processes, where analytical solutions of the Fokker–Planck equation are intractable, we simulated a very large number (N^∞) of stochastic trajectories that we considered to be ground truth Fokker–Planck solutions. We employed an Euler–Maruyama and forward Euler integration scheme of constant step size $dt = 10^{-3}$ for stochastic and deterministic simulations respectively. We provide a description of the employed algorithm along with analysis of its computational complexity in Appendix H, while further numerical experiments on the influence of hyperparameter values on the performance of the estimator are provided in Appendices G and F.

10.1. Linear Conservative System with Additive Noise

For a two dimensional Ornstein-Uhlenbeck process (Appendix A.1) transient and stationary densities evolved through deterministic particle simulations (D) consistently outperformed their stochastic counterparts (S) comprising the same number of particles in terms accuracy in approximating the underlying density (Figure 1). In particular, comparing the 1-Wasserstein distance between samples from analytically derived densities (P_t^A) (Appendix B)—considered here to reflect the ground truth—and the deterministically (D) or stochastically (S) evolved densities (P_t^N), $\mathcal{W}_1(P_t^A, P_t^N)$, we observed smaller Wasserstein distances to ground truth for densities evolved according to our deterministic particle dynamics, both for transient (Figure 1a) and stationary (Figure 1c) solutions. Specifically, we quantified the transient deviation of simulated densities from ground truth by the average temporal 1-Wasserstein distance, $\langle \mathcal{W}_1(P_t^A, P_t^N) \rangle_t$ (Appendix D). For small particle number, deterministically evolved interacting particle trajectories represented more reliably the evolution of the true probability density compared to independent stochastic ones, as portrayed by smaller average Wasserstein distances. For increasing particle number, the accuracy of the simulated solutions with the two approaches converged. Yet, while for $N = 2500$ particles the stochastically evolved densities suggest *on average* (over trials) comparable approximation precision with their deterministic counterparts, the deterministically evolved densities more reliably delivered densities of a certain accuracy, as proclaimed by the smaller dispersion of Wasserstein distances among different realisations (Figure 1a,c).




Figure 1. Accuracy of Fokker–Planck solutions for two dimensional Ornstein Uhlenbeck process.

(a) Mean, $\langle \mathcal{W}_1(P_t^A, P_t^N) \rangle_t$, and (c) stationary $\mathcal{W}_1(P_\infty^A, P_\infty^N)$, 1-Wasserstein distance, between analytic solution and deterministic(D)/stochastic(S) simulations of N particles (for different inducing point number M). (b) Average temporal deviations from analytic mean m_t and (d) covariance matrix C_t for deterministic and stochastic system for increasing particle number N . Deterministic particle simulations consistently outperformed stochastic ones in approximating the temporal evolution of the mean and covariance of the distribution for all examined particle number settings. (Further parameter values: regularisation constant $\lambda = 0.001$, Euler integration time step $dt = 10^{-3}$, and RBF kernel length scale l estimated at every time point as two times the standard deviation of the state vector. Inducing point locations were selected randomly at each time step from a uniform distribution spanning the state space volume covered by the state vector.)

Likewise, we observed similar results when comparing only the stationary distributions, $\mathcal{W}_1(P_\infty^A, P_\infty^N)$ (Figure 1c). While for small particle number, the interacting particle system more accurately captured the underlying limiting distribution, for increasing particle number the accuracy of both approaches converged, with our method consistently delivering more reliable approximations among individual repetitions.

Moreover, densities evolved with our deterministic framework exhibited less fluctuating cumulant trajectories in time, compared to their stochastic counterparts (Figure 2c). In particular, even for limited particle number, cumulants calculated over deterministically evolved particles progressed smoothly in time, while substantially more particles for the stochastic simulations were required for the same temporal cumulant smoothness. To further quantify the transient accuracy of Fokker–Planck solutions computed with our method, we compared the average transient discrepancy between the first two analytic cumulants (m_t and C_t) to those estimated from the particles ($\hat{m}_t$ and $\hat{C}_t$), $\langle \|\hat{m}_t - m_t\|_2 \rangle_t$ (Figure 1b) and $\langle \|\hat{C}_t - C_t\|_F \rangle_t$ (Figure 1d), where $\|\cdot\|_F$ stands for the Frobenious norm (Appendix E).

In line with our previous results, our deterministic framework delivered considerably more accurate transient cumulants when compared to stochastic simulations, with more consistent results among individual realisations, denoted by smaller dispersion of average cumulant differences. (Notice the logarithmic y-axis scale in Figure 1b,d. Error bars for the stochastic solutions were in fact larger than those for the deterministic solutions on a linear scale.)




Figure 2. Stationary and transient Fokker–Planck solutions computed with deterministic (green) and stochastic (brown) particle dynamics for a two dimensional Ornstein Uhlenbeck process. (a,b) Estimated stationary PDFs arising from deterministic ($N = 1000$) (green), and stochastic ($N = 1000$) (brown) particle dynamics. Purple contours denote analytically calculated stationary distributions, while top and side histograms display marginal distributions for each dimension. (c) Temporal evolution of marginal statistics, mean $\langle x \rangle$, standard deviation σ_x , skewness s_x , and kurtosis k_x , for analytic solution (A), and for stochastic (S) and deterministic (D) particle systems comprising $N = 1000$, with initial state distribution $\mathcal{N}\left(\begin{bmatrix} 0.5 \\ 0.5 \end{bmatrix}, \begin{bmatrix} 0.05^2 & 0 \\ 0 & 0.05^2 \end{bmatrix}\right)$, for $M = 100$ randomly selected inducing points employed in the gradient–log–density estimation. Deterministic particle simulations deliver smooth cumulant trajectories, as opposed to highly fluctuating stochastic particle cumulants. (Further parameter values: regularisation constant $\lambda = 0.001$, and RBF kernel length scale l estimated at every time point as two times the standard deviation of the state vector. Inducing point locations were selected randomly at each time step from a uniform distribution spanning the state space volume covered by the state vector.)

Interestingly, the number of sparse points M employed in the gradient–log–density estimation had only minor influence on the quality of the solution (Figure 1a,c). This hints to substantially low computational demands for obtaining accurate Fokker–Planck solutions, since our method is computationally limited by the inversion of the $M \times M$ matrix in Equation (29).

10.2. Bi-Stable Nonlinear System with Additive Noise

For nonlinear processes, since the transient solution of the FPE is analytically intractable, we compared the transient and stationary densities estimated by our method with those returned from stochastic simulations of $N^\infty = 26,000$ particles, and contrasted them against stochastic simulations with the same particle number.

For a system with bi-modal stationary distribution (Appendix A.2), the resulting particle densities from our deterministic framework closely agreed with those arising from the stochastic system with $N^\infty = 26,000$ particles (Figure 3a). In particular, deterministically evolved distributions respected the symmetry of the underlying double–well potential, while the stochastic system failed to accurately capture the potential symmetric structure Figure 3a (iii).

Systematic comparisons of the 1-Wasserstein distance between deterministic and stochastic N particle simulations with the “ N^∞ ” stochastic simulation comprising $N^\infty = 26,000$ particles revealed that our approach efficiently captured the underlying PDF already with $N = 500$ particles (Figure 3c,d). For increasing particle number, the stationary solutions of both systems converged to the “ N^∞ ” one. However, we observed a systematically increasing approximation accuracy delivered from the deterministic simulations compared to their stochastic counterparts.

It is noteworthy that, on average, deterministic simulations of $N = 500$ particles conveyed a better approximation of the underlying transient PDF compared to stochastic simulations of $N = 2500$ particles (Figure 3c).

Interestingly, for small particle number, the number of employed inducing points M did not significantly influence the accuracy of the approximated solution. However for increasing particle number, enlarging the set of inducing points contributed to more accurate approximation of Fokker–Planck equation solutions (Figure 3c), with the trade off of additional computational cost.

Similar to the Ornstein Uhlenbeck process (Section 10.1), comparing cumulant trajectories computed from both the deterministic and stochastic particle systems revealed less fluctuating cumulant evolution for densities evolved with our deterministic framework also in this nonlinear setting (Figure 3b).

10.3. Nonlinear System Perturbed by Multiplicative Noise

To assess the accuracy of our framework on general diffusion processes perturbed by state-dependent (multiplicative) noise, we simulated a bi-stable system with dynamics governed by Equation (A3) with diffusion function $D(x) = \sin^2(x)$ according to Equation (53). Similarly, in this setting, deterministic particle distributions delivered a closer approximation of the underlying density when compared to direct stochastic simulations. In particular, we found that in this setting, deterministically evolved distributions more accurately captured the tails of the underlying distribution, mediated here by stochastic simulations of $N^\infty = 35,000$ particles (Figure 4a,b).

Similar to the previously examined settings, the deterministic framework delivered more reliable and smooth trajectories for the marginal statistics of the underlying distribution (Figure 4c).

Comparing the temporal average and stationary 1-Wasserstein distance (Figure 4d,f) between the optimal stochastic distributions and the deterministic and stochastic particle distributions of size N , we found that the deterministic system delivered consistently more accurate approximations, as portrayed by smaller 1-Wasserstein distances.

Interestingly, we found that for deterministic particle simulations, the number of employed sparse points in the gradient–log–density estimation mediated a moderate approximation improvement for

small system sizes, while for systems comprising more than $N = 2000$ particles, the number of sparse points had minimal or no influence on the accuracy of the resulting distribution (Figure 4e,g).




Figure 3. Performance of deterministic (green) and stochastic (brown) N particle solutions compared to N^∞ (grey) stochastic particle densities for a nonlinear bi-stable process. (a) Instances of estimated pdfs arising from (left) stochastic ($N^\infty = 26,000$) (grey) and deterministic ($N = 1000$) (green), and (right) stochastic ($N^\infty = 26,000$) (grey) and stochastic ($N = 1000$) (brown) particle dynamics at times (i) $t = 0.005$, (ii) $t = 0.231$, and (iii) $t = 1.244$. (b) Temporal evolution of first four distribution cumulants, mean $\langle x \rangle$, standard deviation σ_x , skewness s_x , and kurtosis k_x , for stochastic (S^∞ and S) and deterministic (D) systems comprising $N^\infty = 26,000$, $N = 1000$, with initial state distribution $\mathcal{N}(0, 0.05^2)$, by employing $M = 150$ inducing points in the gradient–log–density estimation. (c) Mean, $\langle W_1(P_t^{N^\infty}, P_t^N) \rangle_t$, and (d) stationary, $W_1(P_\infty^A, P_\infty^N)$, 1-Wasserstein distance, between $N^\infty = 26,000$ stochastic, and deterministic (D)/stochastic (S) simulations of N particles (for different inducing point number M). (Further parameter values: regularisation constant $\lambda = 0.001$, Euler integration time step $dt = 10^{-3}$, and RBF kernel length scale $l = 0.5$. Inducing point locations were selected randomly at each time step from a uniform distribution spanning the state space volume covered by the state vector.)




Figure 4. Accuracy of Fokker-Planck solutions for a nonlinear system perturbed with state-dependent noise. (a) Instances of $N = 1000$ particle distributions resulting from deterministic (green) and (b) stochastic (brown) simulations against stochastic particle distributions comprising $N^\infty = 35,000$ particles (grey) for (i) $t = 0.1$, (ii) $t = 3.2$, and (iii) $t = 4.4$. Insets provide a closer view of details of distribution for visual clarity. Distributions resulting from deterministic particle simulations closer agree with underlying distribution for all three instances. (c) Temporal evolution of first four cumulants for the three particle systems (grey: S_∞ —stochastic with $N^\infty = 35000$ particles, brown: S —stochastic with $N = 1000$ particles, and green: D —deterministic with $N = 1000$ particles). Deterministically evolved distributions result in smooth cumulant trajectories. (d,e) Temporal average and (f,g) stationary 1-Wasserstein distance between distributions mediated through stochastic simulations of $N^\infty = 35000$, and through deterministic (green) and stochastic (brown) simulations of N particles against particle number N and inducing point number M . Shaded regions and error bars denote one standard deviation among 20 independent repetitions. Different green hues designate different inducing point number M employed in the gradient-log-density estimation. (Further parameter values: regularisation constant $\lambda = 0.001$, Euler integration time step $dt = 10^{-3}$, and RBF kernel length scale $l = 0.25$. Inducing points were arranged on a regular grid spanning the instantaneous state space volume captured by the state vector.)

10.4. Performance in Higher Dimensions

To quantify the scaling and performance of the proposed framework for increasing system dimension, we systematically compared simulated densities with analytically calculated ones for Ornstein–Uhlenbeck processes of dimension $D = \{2, 3, 4, 5\}$ following the dynamics of Equation (A4) for inducing point number $M = 100$ (Figure 5) and $M = 200$ (Figure 6). To evaluate simulated Fokker–Planck solutions, we calculated the Kullback–Leibler divergence between analytically evolved densities (Appendix B) and particle densities. We employed the closed-form equation for estimating KL divergence between two Gaussian distributions (Appendix C) for empirically estimated mean, $\hat{m}_t$, and covariance, $\hat{C}_t$, for particle distributions.




Figure 5. Accuracy of Fokker–Planck solutions for multi-dimensional Ornstein–Uhlenbeck processes for $M = 100$ inducing points. Comparison of deterministic particle Fokker–Planck solutions with stochastic particle systems and analytic solutions for multi-dimensional Ornstein–Uhlenbeck process of $D = \{2, 3, 4, 5\}$ dimensions. (a) Time-averaged and (d) stationary Kullback–Leibler (KL) divergence between simulated particle solutions (green: deterministic, brown: stochastic) and analytic solutions for different dimensions. Deterministic particle simulations outperform stochastic particle solutions even for increasing system dimensionality. (b) Time averaged and (e) stationary error between analytic, m_t , and sample mean, $\hat{m}_t$, for increasing particle number. (c) Time averaged and (f) stationary discrepancy between simulated, $\hat{C}_t$, and analytic covariances, C_t , as captured by the Frobenius norm of the relevant covariance matrices difference. The accuracy of the estimated covariance decreases for increasing dimensionality. (Further parameter values: regularisation constant $\lambda = 0.001$, Euler integration time step $dt = 10^{-3}$, and adaptive RBF kernel length scale l calculated at every time step as two times the standard deviation of the state vector. Inducing point locations were selected randomly at each time step from a uniform distribution spanning the state space volume covered by the state vector.)




Figure 6. Accuracy of Fokker-Planck solutions for multi-dimensional Ornstein-Uhlenbeck processes for $M = 200$ inducing points. Comparison of deterministic particle Fokker-Planck solutions with stochastic particle systems and analytic solutions for multi-dimensional Ornstein-Uhlenbeck process of $D = \{2,3,4,5\}$ dimensions. (a) Time-averaged and (d) stationary Kullback-Leibler (KL) divergence between simulated particle solutions (green: deterministic, brown: stochastic) and analytic solutions for different dimensions. Deterministic particle simulations outperform stochastic particle solutions even for increasing system dimensionality. (b) Time averaged and (e) stationary error between analytic, m_t , and sample mean, $\hat{m}_t$, for increasing particle number. (c) Time averaged and (f) stationary discrepancy between simulated, $\hat{C}_t$, and analytic covariances, C_t , as captured by the Frobenius norm of the relevant covariance matrices difference. The accuracy of the estimated covariance decreases for increasing dimensionality. (Further parameter values: regularisation constant $\lambda = 0.001$, Euler integration time step $dt = 10^{-3}$, and adaptive RBF kernel length scale l calculated at every time step as two times the standard deviation of the state vector. Inducing point locations were selected randomly at each time step from a uniform distribution spanning the state space volume covered by the state vector.)

For all dimensionalities, the deterministic particle solutions approximated transient and stationary densities remarkably accurately with Kullback-Leibler divergence between the simulated and analytically derived densities below 10^{-2} for all dimensions, both for transient and stationary solutions (Figures 5a,d and 6a,d). In fact, the deterministic particle solutions delivered more precise approximations of the underlying densities compared to direct stochastic simulations of the same particle number. Remarkably, even for processes of dimension $D = 5$, deterministically evolved solutions mediated through $N = 500$ particles resulted in approximately the same KL divergence of stochastic particle solutions of $N = 6500$ particles.

Our deterministic particle method delivered consistently better approximations of the mean of the underlying densities compared to stochastic particle simulations (Figure 5b,e). Specifically, estimations of the stationary mean of the underlying distributions were more than two orders of magnitude accurate that their stochastically approximated counterparts already for small particle number (Figure 5e).

Yet, the accuracy of our deterministic framework deteriorated for increasing dimension (Figure 5a,d). More precisely, while for low dimensionalities the covariance matrices of the underlying densities were accurately captured by deterministically evolved particles, for increasing system dimension approximations of covariance matrices became progressively worse. Yet, even for systems of dimension $D = 5$, covariance matrices computed from deterministically simulated solutions of $N = 500$ particles were at the same order of magnitude as accurate as covariances delivered by stochastic particle simulations of size $N = 6500$.

However, comparing the resulting performance of solutions delivered by employing different number of inducing points (Figures 5 and 6) reveals that for increasing dimension more inducing points are required to attain accurate FPE solutions. In particular, both transient and stationary KL divergences to ground truth improved remarkably for dimensions $D = 4$ and $D = 5$ by employing $M = 200$ inducing points in the gradient–log–density estimation (Figures 5a,d and 6a,d). In more detail, for nearly all dimensions, the estimation of the covariance of the underlying distribution improved considerably, both for transients and stationary solutions (Figures 5c,f and 6c,f), while only the stationary mean of dimension $D = 4$ showed significant improvement (Figures 5b,e and 6b,e). Figure 6c,f reveals that by increasing the number of inducing points our framework is able to capture more effectively the spread of the underlying distribution, clearly surpassing in approximation accuracy solutions mediated by stochastic particle simulations.

10.5. Second order Langevin Systems

To demonstrate the performance of our framework for simulating solutions of the FPEs for second order Langevin systems as described in Section 9, we incorporated our method in a symplectic Verlet integrator (Equations (A10)–(A12)) simulating the second-order dynamics captured by Equation (56) for a linear $f(x) = -4x$ and a nonlinear, $f(x) = -4x^3 + 4x$, drift function (Equation (A10)), and compared the results with stochastic simulations integrated by a semi-symplectic framework [63]. In agreement with previous results, cumulant trajectories evolved smoother in time for deterministic particle simulations when compared to their stochastic counterparts (Figures 7a and 8c). Stationary densities closely matched analytically derived ones (see Equation (A7)) (purple contour lines in Figures 7b and 8b), while transient densities captured the fine details of simulated stochastic particle densities comprising $N^\infty = 20,000$ (Figure 8a).

Furthermore, the symplectic integration contributed to the preservation of energy levels for each particle after the system reached equilibrium (Figures 7e and 8f), which was also evident when observing individual particle trajectories in the state space (Figures 7c,d and 8d,e).

As already conveyed in Section 9, the velocity term and the gradient–log–density term canceled out in the long time limit (Figures 7f and 8g) for each particle individually, while the average kinetic energy in equilibrium exactly resorted to the value dictated by the fluctuation–dissipation relation and the equipartition of energy property, i.e., $\langle \mathcal{K}^{(i)} \rangle_N = \frac{\sigma^2}{2\gamma}$ (Figures 7g and 8h).




Figure 7. Energy preservation for second order Langevin dynamics in a quadratic potential.

Comparison of deterministic particle Fokker–Planck solutions with stochastic particle systems for a harmonic oscillator (a) First four cumulant temporal evolution for deterministic (green) and stochastic (brown) system. (b) Stationary joint and marginal distributions for deterministic (green) and stochastic (brown) systems. Purple lines denote analytically derived stationary distributions. (c,d) State space trajectory of a single particle for deterministic (green) and stochastic (brown) system. Color gradients denote time. (e) Temporal evolution of individual particle energy $E_t^{(i)}$ for deterministic system for 5 particles. (f) Difference between velocity and gradient–log–density term for individual particles. After the system reaches stationary state the particle velocity and GLD term cancel out. (g) Ensemble average kinetic energy through time resorts to $\frac{\sigma^2}{2\gamma}$ (grey dashed line) after equilibrium is reached. (Further parameter values: regularisation constant $\lambda = 0.001$, integration time step $dt = 2 \times 10^{-3}$, and adaptive RBF kernel length scale l calculated at every time step as two times the standard deviation of the state vector. Number of inducing points $M = 300$. Inducing point locations were selected randomly at each time step from a uniform distribution spanning the state space volume covered by the state vector.)




Figure 8. Energy preservation for second order Langevin dynamics in a double well potential. Comparison of deterministic particle Fokker–Planck solutions with stochastic particle systems for a bistable process. (a,b) Joint and marginal distributions of system states mediated by $N = 8000$ particles evolved with our framework (green) and with direct stochastic simulations comprising $N^\infty = 20,000$ (grey) and $N = 8000$ (brown) particles at (a) $t = 0.6$, and (b) $t = 10$. Purple lines denote the analytically derived stationary density. (c) First four cumulants temporal evolution for deterministic (green) and stochastic (brown) system. (d) State space trajectory of a single particle for deterministic and (e) stochastic system. Color gradients denote time. (f) Temporal evolution of individual particle energy $E_t^{(i)}$ for deterministic system for 5 particles. (g) Temporal evolution of distribution of particle energies $E_t^{(i)}$ for deterministic (green) and stochastic (brown) system. (h) Difference between velocity and gradient log density term for individual particles. (i) Ensemble average kinetic energy through time resorts to $\frac{\sigma_v^2}{2\gamma}$ (grey dashed line) after equilibrium is reached. (Further parameter values: regularisation constant $\lambda = 0.001$, integration time step $dt = 2 \times 10^{-3}$ and adaptive RBF kernel length scale l calculated at every time step as two times the standard deviation of the state vector. Number of inducing points $M = 300$. Inducing point locations were selected randomly at each time step from a uniform distribution spanning the state space volume covered by the state vector.)

10.6. Nonconservative Chaotic System with Additive Noise (Lorenz Attractor)

As a final assessment of our framework for simulating accurate solutions of Fokker–Planck equations, we employed a Lorenz attractor model with parameters rendering the dynamics chaotic, perturbed by moderate additive Gaussian noise (Equation (A13)). By comparing stochastic simulations of $N^\infty = 150,000$ particles and deterministic and stochastic simulations of $N = 4000$ particles (Figure 9), we observed that the deterministic framework more precisely captured finer details of the underlying

distribution (Figure 9a), represented here by the N^∞ stochastic simulation. While both stochastic and deterministic simulations capture the overall butterfly profile of the Lorenz attractor, the deterministic system indeed delivered a closer match to the underlying distribution.




Figure 9. Deterministic (green) and stochastic (brown) Fokker–Planck particle solutions for a three dimensional Lorenz attractor system in the chaotic regime perturbed by additive Gaussian noise. (a) Joint and marginal distributions of system states mediated by $N = 4000$ particles evolved with our framework (green) and with direct stochastic simulations comprising $N = 150,000$ (grey) and $N = 4000$ (brown) particles at $t = 0.4$. (b) Cumulant trajectories for the three particle systems. Cumulants derived from deterministic particle simulations (green) closer match cumulant evolution of the underlying distribution (grey) compared to stochastic simulations (brown). (Further parameter values: regularisation constant $\lambda = 0.001$, Euler integration time step $dt = 10^{-3}$, adaptive RBF kernel length scale l calculated at every time step as two times the standard deviation of the state vector. Number of inducing points: $M = 200$. Inducing point locations were selected randomly at each time step from a uniform distribution spanning the state space volume covered by the state vector.)

Similar to the previously examined models, cumulant trajectories computed from deterministically evolved particles show closer agreement with those computed from the N^∞ stochastic system, compared to the stochastic system comprising N particles (Figure 9b). In particular, cumulants for the x and y states exhibited high temporal fluctuations when computed from stochastically evolved distributions, while our framework conveyed more accurate cumulant trajectories, closer to those delivered by the N^∞ stochastic system.

11. Discussion and Outlook

We presented a particle method for simulating solutions of FPEs governing the temporal evolution of the probability density for stochastic dynamical systems of the diffusion type. By reframing the FPE in a Liouville form, we obtained an effective dynamics in terms of independent deterministic particle trajectories. Unfortunately, this formulation requires the knowledge of the gradient of the logarithm of the instantaneous probability density of the system state, which is actually the unknown quantity of interest. We circumvented this complication by introducing statistical estimators for the gradient–log–density based on a variational formulation. To combine high flexibility of estimators with computational efficiency, we employed kernel-based estimation together with an additional sparse approximation. For the case of equilibrium systems, we related our framework to Stein Variational Gradient Descent, a particle-based dynamics to approximate the stationary density, and to a geometric formulation of Fokker–Planck dynamics. We further discussed extensions of our method to settings with multiplicative noise and to second order Langevin dynamics.

To demonstrate the performance of our framework, we provided detailed tests and comparisons with stochastic simulations and analytic solutions (when possible). We demonstrated the accuracy of our method on conservative and non-conservative model systems with different dimensionalities. In particular, we found that our framework outperforms stochastic simulations both in linear and nonlinear settings by delivering more accurate densities for small particle number when the dimensionality is small enough. For increasing particle number, the accuracy of both approaches converges. Yet, our deterministic framework *consistently* delivered results with smaller variance among individual repetitions. Furthermore, we showed that our method, even for small particle numbers, exhibits low-order cumulant trajectories with significantly less temporal fluctuations when compared against to stochastic simulations of the same particle number.

We envisage several ways to improve and extend our method. There is room for enhancement by optimising hyperparameters of our algorithm such as inducing point position and kernel length scale. Current grid-based and uniform random selection of inducing point position may contribute to the deterioration of solution accuracy in higher dimensions. Other methods, such as subsampling or clustering of particle positions may lead to further improvements. On the other hand, a hyperparameter update may not be at all necessary at each time step in certain settings, such that a further speedup of our algorithm could be achieved.

The implementation of our method depends on the function class chosen to represent the estimator. In this paper we have focused on linear representations, leading to simple closed form expressions. It would be interesting to see if other, nonlinear parametric models, such as neural networks, (see e.g., [64]) could be employed to represent estimators. While in this setting, there would be no closed-form solutions, the small changes in estimates between successive time steps suggest that only a few updates of numerical optimisation may be necessary at each step. Moreover, the ability of neural networks to automatically learn relevant features from data might help to improve performance for higher dimensional problems when particle motion is typically restricted on lower dimensional submanifolds.

From a theoretical point of view, rigorous results on the accuracy of the particle approximation would be important. These would depend on the speed of convergence of estimators towards exact gradients of log–densities. However, to obtain such results may not be easy. While rates of convergence

for kernel-based estimators have been studied in the literature, the methods for proofs usually rely on the independence of samples and would not necessarily apply to the case of interacting particles.

We have so far addressed only the forward simulation of FPEs. However, preliminary results indicate that related techniques may be applied to particle based simulations for smoothing (forward–backward) and related control problems for diffusion processes [65]. Such problems involve computations of an effective, controlled drift function in terms of gradient–log–densities. We defer further details and discussions on subsequent publications on the topic.

Taken together, the main advantage of our framework is its minimal requirement in simulated particle trajectories for attaining reliable Fokker–Planck solutions with smoothly evolving transient statistics. Moreover, our proposed method is nearly effortless to set up when compared to classical grid-based FPE solvers, while it delivers more reliable results than direct stochastic simulations.

Author Contributions: Conceptualization, S.R. and M.O.; methodology, D.M. and M.O.; software, D.M.; validation, D.M. and M.O.; formal analysis, D.M. and M.O.; investigation, D.M.; resources, M.O.; data curation, D.M.; writing—original draft preparation, D.M. and M.O.; writing—review and editing, D.M., S.R. and M.O.; visualization, D.M.; supervision, M.O.; project administration, M.O.; funding acquisition, S.R. and M.O. All authors have read and agreed to the published version of the manuscript.

Funding: This research has been partially funded by Deutsche Forschungsgemeinschaft (DFG)-SFB1294/1-318763901.

Acknowledgments: We would like to thank Paul Rozdeba for constructive criticism on the manuscript.

Conflicts of Interest: The authors declare no conflict of interest. The funders had no role in the design of the study; in the collection, analyses, or interpretation of data; in the writing of the manuscript, or in the decision to publish the results.

Appendix A. Simulated Systems

Appendix A.1. Two Dimensional Ornstein-Uhlenbeck Process

For comparing Fokker–Planck solutions computed with our approach with solutions derived from stochastic simulations, we considered the two dimensional Ornstein-Uhlenbeck process captured by the following equations

$$dX_t = (-4X_t + Y_t) dt + \sigma dB_1 \quad (A1)$$

$$dY_t = (-4Y_t + X_t) dt + \sigma dB_2, \quad (A2)$$

where the related potential is $U(x, y) = 2x^2 - xy + 2y^2$. The simulation time was set to $T = 3$ with Euler–Maruyama integration step $dt = 10^{-3}$. For estimating the instantaneous gradient log density we employed $M = \{50, 100, 150, 200\}$ inducing points, randomly selected at every time point from a uniform distribution spanning the state space volume covered by the particles at the current time point.

Appendix A.2. Bistable Nonlinear System

For testing our framework on nonlinear settings, we simulated

$$dX_t = (-4X_t^3 + 4X_t) dt + D(X_t)^{\frac{1}{2}} dB_t, \quad (A3)$$

with $D(x) = 1$ for evaluating solutions with additive Gaussian noise, and with $D(x) = \sin^2(x)$ for multiplicative noise FP solutions. The associated potential is $U(x) = x^4 - 2x^2$.

Appendix A.3. Multi-Dimensional Ornstein-Uhlenbeck Processes

For quantifying the scaling of our method for increasing system dimension, we simulated systems of dimensionality $D = \{2, 3, 4, 5\}$ according to the following equation

$$dX_{(i)t} = \left(-4X_{(i)t} + \sum_{j=1, i \neq j}^d \frac{1}{2} X_{(j)t} \right) dt + dB_{(i)}, \quad (\text{A4})$$

for $d \in D$. Simulation time was determined by the time required for analytic mean m_t to converge to its stationary solution within $\tilde{\epsilon}$ precision $\tilde{\epsilon} = 10^{-5}$, while the integration step was set to $dt = 10^{-3}$.

Appendix A.4. Second Order Langevin Dynamics

For demonstrating the energy preservation properties of our method for second-order Langevin dynamics, we incorporated our framework into a Verlet symplectic integration scheme (Equation (A10)), and compared the results with stochastic simulations integrated according to a semi-symplectic scheme [63].

We consider a system with dynamics for positions X and velocities V captured by

$$dX = Vdt \quad (\text{A5})$$

$$dV = (-\gamma V + f(X))dt + \sigma dB_t, \quad (\text{A6})$$

where the velocity change (acceleration) is the sum of a deterministic drift f , a velocity-dependent damping $-\gamma V$, and a stochastic noise term σdB_t .

In conservative settings, the drift is the gradient of a potential $f(x) = -\nabla U(x)$. Here we used a quadratic (harmonic) potential $U(x) = 2x^2$ and a double-well potential $U(x) = x^4 - 2x^2$.

In equilibrium, the Fokker–Planck solution is the Maxwell–Boltzmann distribution, i.e.,

$$p_\infty(X, V) = \frac{1}{Z} e^{-\beta H(X, V)} = \frac{1}{Z} e^{-\beta \left(\frac{\|V\|^2}{2} + U(X) \right)}, \quad (\text{A7})$$

with partition function $Z = \int e^{-\beta \left(\frac{\|V\|^2}{2} + U(X) \right)} dx dv$.

We may compute the energy of each particle at each time point as the sum of its kinetic and potential energies

$$E_t^{(i)} = \frac{1}{2} V^{(i)2} + U(X^{(i)}). \quad (\text{A8})$$

Here the superscripts denote individual particles. After the system has reached equilibrium, energy levels per particle are expected to remain constant.

From the equipartition of energy and the fluctuation–dissipation relation, in the long time limit the average kinetic energy of the system is expected to resort to

$$\lim_{t \rightarrow \infty} \langle \mathcal{K} \rangle = \lim_{t \rightarrow \infty} \frac{1}{2} \langle V^2 \rangle = \frac{\sigma^2}{2\gamma}. \quad (\text{A9})$$

Symplectic integration [33] of Equation (56) follows the equations

$$V_{n+\frac{1}{2}} = V_n + \frac{dt}{2} \left(-\gamma V_n + f(X_n) - \frac{\sigma^2}{2} \nabla_v \ln p_t(X_n, V_n) \right) \quad (\text{A10})$$

$$X_{n+1} = X_n + dt V_{n+\frac{1}{2}} \quad (\text{A11})$$

$$V_{n+1} = V_{n+\frac{1}{2}} + \frac{dt}{2} \left(-\gamma V_{n+\frac{1}{2}} + f(X_{n+1}) - \frac{\sigma^2}{2} \nabla_v \ln p_t(X_{n+1}, V_{n+\frac{1}{2}}) \right), \quad (\text{A12})$$

where n denotes a single integration step.

Appendix A.5. Lorenz attractor

For simulating trajectories of the noisy Lorenz system, we employed the following equations

$$dx_t = \sigma(y - x)dt + \sigma dW_x \quad (A13)$$

$$dy_t = (x(\rho - z) - y)dt + \sigma dB_y \quad (A14)$$

$$dz_t = (xy - \beta z)dt + \sigma dB_z, \quad (A15)$$

with parameters $\sigma = 10$, $\rho = 28$, and $\beta = \frac{8}{3}$, that render the deterministic dynamics chaotic [66], employing moderate additive Gaussian noise.

Appendix B. Computing Central Moment Trajectories for Linear Processes

For a linear process

$$dX_t = A X_t dt + \sigma dB, \quad (A16)$$

the joint density of the state vector X remains Gaussian for all times when the initial density is Gaussian. The mean vector m and covariance matrix C may be computed by solving the ODE system

$$\frac{dm}{dt} = Am \quad (A17)$$

$$\frac{dC}{dt} = AC + CA^\top + \sigma^2 I. \quad (A18)$$

Appendix C. Kullback–Leibler Divergence for Gaussian Distributions

We calculated the KL divergence between the theoretical and simulated distributions with

$$KL(P_1||P_2) = \frac{1}{2} \left(\log(|\Sigma_2|/|\Sigma_1|) - d + \text{Tr}(\Sigma_2^{-1}\Sigma_1) + (\mu_2 - \mu_1)^T \Sigma_2^{-1}(\mu_2 - \mu_1) \right), \quad (A19)$$

where $P_x \sim \mathcal{N}(\mu_x, \Sigma_x)$.

Appendix D. Wasserstein Distance

We employed the 1-Wasserstein distance [59] as a distance metric for comparing pairs of empirical distributions.

For two distributions P and Q , we denote with $\mathcal{J}(P, Q)$ all joint distributions J for a pair of random variables (X, Y) with marginals P and Q . Then the *Wasserstein distance* between these distributions is

$$\mathcal{W}_p(P, Q) = \left(\inf_{J \in \mathcal{J}(P, Q)} \int \|x - y\|^p dJ(x, y) \right)^{\frac{1}{p}}, \quad (A20)$$

where for the 1-Wasserstein distances (used in the present manuscript) $p = 1$.

Interestingly, the Wasserstein distance between two one dimensional distributions P and Q obtains a closed form solution

$$\mathcal{W}_p(P, Q) = \left(\int_0^1 \|F_P^{-1}(\tau) - F_Q^{-1}(\tau)\|^p d\tau \right)^{1/p}, \quad (A21)$$

with F_P and F_Q indicating the cumulative distribution functions of P and Q .

Moreover, for one dimensional empirical distributions P and Q with samples of same size $\{X_i\}_{i=1}^n$ and $\{Y_i\}_{i=1}^n$, the Wasserstein distance simplifies into computation of differences of order statistics

$$\mathcal{W}_p(P, Q) = \left(\sum_{i=1}^n \|X_{(i)} - Y_{(i)}\|^p \right)^{\frac{1}{p}}, \quad (A22)$$

where $X_{(i)}$ and $Y_{(i)}$ indicates the i -th order statistic of the sample $\{X_i\}_{i=1}^n$ and $\{Y_i\}_{i=1}^n$, i.e. $X_{(1)} \leq X_{(2)} \leq \dots \leq X_{(n)}$ and $Y_{(1)} \leq Y_{(2)} \leq \dots \leq Y_{(n)}$ [67].

We calculated the average temporal 1-Wasserstein distance as the time average of instantaneous 1-Wasserstein distances between the two distributions under comparison, P_t^A and P_t^B , i.e.,

$$\left\langle \mathcal{W}_1(P_t^A, P_t^B) \right\rangle_t = \frac{1}{T/dt} \sum_{t=1}^{T/dt} \mathcal{W}_1(P_t^A, P_t^B), \quad (\text{A23})$$

where as before T stands for the duration of the simulation and dt for the time discretisation step.

Appendix E. Frobenious Norm

For comparing covariance matrices of the simulated particle systems with ground truth we employed the Frobenious norm of the relevant matrices difference. The Frobenious norm of a $m \times n$ matrix A may be calculated from

$$\|A\|_F = \sqrt{\sum_{i=1}^m \sum_{j=1}^n |a_{i,j}|^2} = \sqrt{\text{Tr}(A^* A)}, \quad (\text{A24})$$

where $a_{i,j}$ denote the entries of matrix A , and A^* stands for the conjugate transpose of A .

Appendix F. Influence of Hyperparameter Values on the Performance of the Gradient-Log-Density Estimator

To determine the influence of the hyperparameter values on the performance of the gradient-log-density estimator, we systematically evaluated the approximation error of our estimator for $N = 1000$ samples of a one dimensional log-normal distribution with mean $\mu = 0$ and standard deviation $\sigma = 0.5$ for 20 independent realisations.

We quantified the approximation error as the average error between the analytically calculated and predicted gradient-log-density on each sample, i.e.,

$$\text{Approximation error} = \frac{1}{N} \sum_{i=1}^N \|\nabla \ln p(x_i) - (\widehat{\nabla \ln p(x_i)})\|, \quad (\text{A25})$$

where the analytically calculated gradient-log-density was determined as $\nabla \ln p(x) = \frac{\mu - \sigma^2 - \ln(x)}{\sigma^2 x}$.

By systematically varying the regularisation parameter λ , the kernel length scale l , and the inducing point number M we observed the following:

- The hyperparameter that strongly influences the approximation accuracy is the kernel length scale l (Figure A1).
- Underestimation of kernel length scale l has stronger impact on approximation accuracy than overestimation (Figure A1).
- For increasing regularisation parameter value λ , underestimation of l has less impact on the approximation accuracy (Figures A1 and A2).
- For overestimation of the kernel length scale l , the regularisation parameter λ and inducing point number M have nearly no effect on the resulting approximation error (Figure A1).
- For underestimation of kernel length scale l , increasing the number of inducing points M in the estimator results in larger approximation errors (Figure A2 (upper left)).




Figure A1. Approximation error for increasing kernel length scale l for different regularisation parameter values λ and inducing point number M .




Figure A2. Approximation error for increasing regularisation parameter value λ for different kernel length scale l and inducing point number M .

Appendix G. Required Number of Particles for Accurate Fokker–Planck Solutions

To compare the computational demands of the deterministic and stochastic particle systems, we determined the required particle number each system needed to attain a specified accuracy compared to ground truth transient solutions. In particular, for a two dimensional Ornstein–Uhlenbeck process, we identified the minimal number of particles N_{KL}^* both systems required to achieve a certain time-averaged Kullback–Leibler distance to ground truth transient solutions, $\langle \text{KL}(P_t^A, P_t^N) \rangle_t$. As already indicated in the previous sections, the stochastic system required considerably larger particle number to achieve the same time averaged KL distances to ground truth when compared to our proposed framework. In fact, for the entire range of examined KL distances, our method consistently required at least one order of magnitude less particles compared to the its stochastic counterpart.




Figure A3. Required particle number, N_{KL}^* , to attain time averaged Kullback–Leibler divergence to ground truth, $\langle \text{KL}(P_t^A, P_t^N) \rangle_t$, for deterministic (green) and stochastic (brown) particle systems for a two dimensional Ornstein–Uhlenbeck process. Markers indicate mean required particle number, while error bars denote one standard deviation over 20 independent realisations. Grey circles indicate required particle number for each individual realisation. The deterministic particle system consistently required at least one order of magnitude less particles compared to its stochastic counterpart. (Further parameter values: regularisation constant $\lambda = 0.001$, inducing point number $M = 100$, and RBF kernel length scale l estimated at every time point as two times the standard deviation of the state vector. Inducing point locations were selected randomly at each time step from a uniform distribution spanning the state space volume covered by the state vector.)

Appendix H. Algorithm for Simulating Deterministic Particle System

Here we provide the algorithm for simulating deterministic particle trajectories according to our proposed framework (Algorithms A1 and A2). In the comments, we denote the computational complexity of each operation in the gradient–log–density estimation in terms of big- $\mathcal{O}$ notation. Since the inducing point number M employed in the gradient–log–density estimation is considerably smaller than sample number N , i.e., $M \ll N$, the overall computational complexity of a *single* gradient–log–density evaluation amounts to $\mathcal{O}(NM^2)$.

Algorithm A1: Gradient Log Density Estimator

Input: X : $N \times D$ state vector

Z : $M \times D$ inducing points vector

d : dimension for gradient

l : RBF Kernel length scale

Output: G : $N \times 1$ vector for gradient–log–density at each position X in d dimension

- 1 $K^{xz} \leftarrow K(X, Z; l)$ // $N \times M$ $\mathcal{O}(NM)$
 - 2 $K^{zz} \leftarrow K(Z, Z; l)$ // $M \times M$ $\mathcal{O}(M^2)$
 - 3 $I_{K^{zz}} \leftarrow (K^{zz} + 10^{-3} I)^{-1}$ // $M \times M$ $\mathcal{O}(M^3)$
 - 4 $\text{grad}_K \leftarrow \nabla_{X^{(d)}} K(X, Z; l)$ // $N \times M$ $\mathcal{O}(NM)$
 - 5 $\text{sgrad}_K \leftarrow \sum_{X_i} \text{grad}_K$ // $1 \times M$
 - 6 $G \leftarrow K^{xz} (\lambda I + I_{K^{zz}} (K^{xz})^\top K^{xz} + 10^{-3} I)^{-1} I_{K^{zz}} \text{sgrad}_K^\top$ // $N \times 1$
// $\mathcal{O}(NM^2) + \mathcal{O}(M^3)$
-

Algorithm A2: Deterministic Particle Simulation

Input: x_0 : $1 \times D$ initial condition
 s_0 : variance of initial condition
 N : particle number
 M : inducing point number
 T : duration of simulation
 dt : integration time step
 $f(\cdot)$: drift function
 $D(\cdot)$: diffusion function
 l_0 : $1 \times D$ Kernel length scale or False // FALSE for adaptive length scale selection
 $random_M$: boolean variable // TRUE for selecting inducing points from random uniform distribution; FALSE for arranging them on a regular grid

Output: $\{X_t\}_{t=0}^{T/dt}$: $N \times D \times \left\lceil \frac{T}{dt} \right\rceil$ particle trajectories

```

1 Initialization:  $X_0 \leftarrow$  Draw  $N$  samples from Gaussian  $\mathcal{N}(x_0, s_0)$ 
2 for  $t \leftarrow 1$  to  $\frac{T}{dt}$  do
3   if  $random\_M$  is TRUE then                                     // Select inducing points
4     for  $d \leftarrow 1$  to  $D$  do
5        $Z^{(d)} \leftarrow$  Draw  $M$  samples from Uniform( $\min(X_{t-1}^{(d)}), \max(X_{t-1}^{(d)})$ )
6     end
7   else
8     for  $d \leftarrow 1$  to  $D$  do
9        $Z^{(d)} \leftarrow$  CreateRegularGrid( $\min(X_{t-1}^{(d)}), \max(X_{t-1}^{(d)}), M$ )
10    end
11  end
12  if  $l\_0$  is FALSE then                                           // Set length scale
13     $l \leftarrow 2 \text{std}(X_{t-1})$ 
14  else
15     $l \leftarrow l\_0$ 
16  end
17  for  $d \leftarrow 1$  to  $D$  do
18     $G^{(d)} \leftarrow$  GradientLogDensityEstimation( $X_{t-1}, Z, d, l$ )
19  end
20   $X_t \leftarrow X_{t-1} + \left( f(X_{t-1}) - \frac{1}{2} D(X_{t-1}) \circ G - \frac{1}{2} \nabla D(X_{t-1}) \right) dt$ 
21 end

```

References

- Swain, P.S.; Elowitz, M.B.; Siggia, E.D. Intrinsic and extrinsic contributions to stochasticity in gene expression. *Proc. Natl. Acad. Sci. USA* **2002**, *99*, 12795–12800.
- Pimentel, J.A.; Aldana, M.; Huepe, C.; Larralde, H. Intrinsic and extrinsic noise effects on phase transitions of network models with applications to swarming systems. *Phys. Rev. E* **2008**, *77*, 061138.
- Hilfinger, A.; Paulsson, J. Separating intrinsic from extrinsic fluctuations in dynamic biological systems. *Proc. Natl. Acad. Sci. USA* **2011**, *108*, 12167–12172.
- Elowitz, M.B.; Levine, A.J.; Siggia, E.D.; Swain, P.S. Stochastic gene expression in a single cell. *Science* **2002**, *297*, 1183–1186.
- White, J.A.; Rubinstein, J.T.; Kay, A.R. Channel noise in neurons. *Trends Neurosci.* **2000**, *23*, 131–137.
- Volfson, D.; Marciniak, J.; Blake, W.J.; Ostroff, N.; Tsimring, L.S.; Hasty, J. Origins of extrinsic variability in eukaryotic gene expression. *Nature* **2006**, *439*, 861–864.

7. Fellous, J.M.; Rudolph, M.; Destexhe, A.; Sejnowski, T.J. Synaptic background noise controls the input/output characteristics of single cells in an in vitro model of in vivo activity. *Neuroscience* **2003**, *122*, 811–829.
8. Schnoerr, D.; Sanguinetti, G.; Grima, R. Approximation and inference methods for stochastic biochemical kinetics—A tutorial review. *J. Phys. Math. Theor.* **2017**, *50*, 093001.
9. Van Kampen, N.G. The expansion of the master equation. *Adv. Chem. Phys.* **1976**, *34*, 245–309.
10. Suzuki, M. Passage from an initial unstable state to a final stable state. *Adv. Chem. Phys.* **1981**, *46*, 195–278.
11. Gammaitoni, L.; Hänggi, P.; Jung, P.; Marchesoni, F. Stochastic resonance. *Rev. Mod. Phys.* **1998**, *70*, 223.
12. Benzi, R.; Parisi, G.; Sutera, A.; Vulpiani, A. Stochastic resonance in climatic change. *Tellus* **1982**, *34*, 10–16.
13. Horsthemke, W. Noise induced transitions. In *Non-Equilibrium Dynamics in Chemical Systems*; Springer: Berlin/Heidelberg, Germany, 1984; pp. 150–160.
14. Adorno, D.P.; Pizzolato, N.; Valenti, D.; Spagnolo, B. External noise effects in doped semiconductors operating under sub-THz signals. *Rep. Math. Phys.* **2012**, *70*, 171–179.
15. Assaf, M.; Roberts, E.; Luthey-Schulten, Z.; Goldenfeld, N. Extrinsic noise driven phenotype switching in a self-regulating gene. *Phys. Rev. Lett.* **2013**, *111*, 058102.
16. Balázs, G.; van Oudenaarden, A.; Collins, J.J. Cellular decision making and biological noise: From microbes to mammals. *Cell* **2011**, *144*, 910–925.
17. Hughes, S.W.; Cope, D.W.; Toth, T.I.; Williams, S.R.; Crunelli, V. All thalamocortical neurones possess a T-type Ca^{2+} ‘window’ current that enables the expression of bistability-mediated activities. *J. Physiol.* **1999**, *517*, 805–815.
18. Rose, J.E.; Brugge, J.F.; Anderson, D.J.; Hind, J.E. Phase-locked response to low-frequency tones in single auditory nerve fibers of the squirrel monkey. *J. Neurophysiol.* **1967**, *30*, 769–793.
19. Nicolis, C. Solar variability and stochastic effects on climate. *Sol. Phys.* **1981**, *74*, 473–478.
20. Nicolis, C. Stochastic aspects of climatic transitions—response to a periodic forcing. *Tellus* **1982**, *34*, 1–9.
21. Risken, H. Fokker-Planck equation. In *The Fokker-Planck Equation*; Springer-Verlag: Berlin/Heidelberg, Germany, 1996; pp. 63–95.
22. Wang, M.C.; Uhlenbeck, G.E. On the theory of the Brownian motion II. *Rev. Mod. Phys.* **1945**, *17*, 323.
23. Särkkä, S.; Solin, A. *Applied Stochastic Differential Equations*; Cambridge University Press: Cambridge, UK, 2019; Volume 10.
24. Gillespie, D.T. The chemical Langevin equation. *J. Chem. Phys.* **2000**, *113*, 297–306.
25. Melykoti, B.; Burrage, K.; Zygalakis, K.C. Fast stochastic simulation of biochemical reaction systems by alternative formulations of the chemical Langevin equation. *J. Chem. Phys.* **2010**, *132*, 164109.
26. Schadschneider, A.; Chowdhury, D.; Nishinari, K. *Stochastic Transport in Complex Systems: From Molecules to Vehicles*; Elsevier: Amsterdam, The Netherlands, 2010.
27. Kumar, P.; Narayanan, S. Solution of Fokker-Planck equation by finite element and finite difference methods for nonlinear systems. *Sadhana* **2006**, *31*, 445–461.
28. Brics, M.; Kaupuzs, J.; Mahnke, R. How to solve Fokker-Planck equation treating mixed eigenvalue spectrum? *arXiv* **2013**, arXiv:1303.5211.
29. Chang, J.; Cooper, G. A practical difference scheme for Fokker-Planck equations. *J. Comput. Phys.* **1970**, *6*, 1–16.
30. Pichler, L.; Masud, A.; Bergman, L.A. Numerical solution of the Fokker-Planck equation by finite difference and finite element methods—A comparative study. In *Computational Methods in Stochastic Dynamics*; Springer-Verlag: Berlin/Heidelberg, Germany, 2013; pp. 69–85.
31. Harrison, G.W. Numerical solution of the Fokker Planck equation using moving finite elements. *Numer. Methods Partial Differ. Equ.* **1988**, *4*, 219–232.
32. Epperlein, E. Implicit and conservative difference scheme for the Fokker-Planck equation. *J. Comput. Phys.* **1994**, *112*, 291–297.
33. Leimkuhler, B.; Reich, S. *Simulating Hamiltonian Dynamics*; Cambridge University Press: Cambridge, UK, 2004; Volume 14.
34. Chen, N.; Majda, A.J. Efficient statistically accurate algorithms for the Fokker-Planck equation in large dimensions. *J. Comput. Phys.* **2018**, *354*, 242–268.
35. Lin, Y.; Cai, G. *Probabilistic Structural Dynamics: Advanced Theory and Applications*; McGraw-Hill: New York, NY, USA, 1995.

36. Roberts, J.B.; Spanos, P.D. *Random Vibration and Statistical Linearization*; Courier Corporation: Chelmsford, MA, USA, 2003.
37. Proppe, C.; Pradlwarter, H.; Schuëller, G. Equivalent linearization and Monte Carlo simulation in stochastic dynamics. *Probab. Eng. Mech.* **2003**, *18*, 1–15.
38. Grigoriu, M. *Stochastic Calculus: Applications in Science and Engineering*; Springer Science & Business Media: Berlin/Heidelberg, Germany, 2013.
39. Øksendal, B. *Stochastic Differential Equations*; Springer-Verlag: Berlin/Heidelberg, Germany, 2003.
40. Kroese, D.P.; Taimre, T.; Botev, Z.I. *Handbook of Monte Carlo Methods*; John Wiley & Sons: Hoboken, NJ, USA, 2013; Volume 706.
41. Carrillo, J.; Craig, K.; Patacchini, F. A blob method for diffusion. *Calc. Var. Partial Differ. Equ.* **2019**, *58*, 1–53.
42. Pathiraja, S.; Reich, S. Discrete gradients for computational Bayesian inference. *J. Comp. Dyn.* **2019**, *6*, 236–251.
43. Reich, S.; Weissmann, S. Fokker-Planck particle systems for Bayesian inference: Computational approaches. *arXiv* **2019**, arXiv:1911.10832.
44. Liu, Q.; Lee, J.; Jordan, M. A kernelized Stein discrepancy for goodness-of-fit tests. In Proceedings of the International Conference on Machine Learning, New York, NY, USA, 19–24 June 2016; pp. 276–284.
45. Taghvaei, A.; Mehta, P.G. Accelerated flow for probability distributions. *arXiv* **2019**, arXiv:1901.03317.
46. Velasco, R.M.; Scherer García-Colín, L.; Uribe, F.J. Entropy production: Its role in non-equilibrium thermodynamics. *Entropy* **2011**, *13*, 82–116.
47. Hyvärinen, A. Estimation of non-normalized statistical models by score matching. *J. Mach. Learn. Res.* **2005**, *6*, 695–709.
48. Li, Y.; Turner, R.E. Gradient estimators for implicit models. *arXiv* **2017**, arXiv:1705.07107.
49. Shi, J.; Sun, S.; Zhu, J. A spectral approach to gradient estimation for implicit distributions. *arXiv* **2018**, arXiv:1806.02925.
50. Tomé, T.; De Oliveira, M.J. *Stochastic Dynamics and Irreversibility*; Springer: Berlin/Heidelberg, Germany, 2015.
51. Shawe-Taylor, J.; Cristianini, N.; et al.. *Kernel Methods for Pattern Analysis*; Cambridge University Press: Cambridge, UK, 2004.
52. Scholkopf, B.; Smola, A.J. *Learning with Kernels: Support Vector Machines, Regularization, Optimization, and Beyond*; MIT Press: Cambridge, MA, USA, 2001.
53. Sutherland, D.J.; Strathmann, H.; Arbel, M.; Gretton, A. Efficient and principled score estimation with Nyström kernel exponential families. *arXiv* **2017**, arXiv:1705.08360.
54. Rasmussen, C.E. *Gaussian Processes in Machine Learning*; Summer School on Machine Learning; Springer-Verlag: Berlin/Heidelberg, Germany, 2003; pp. 63–71.
55. Liu, Q.; Wang, D. Stein variational gradient descent: A general purpose Bayesian inference algorithm. In Proceedings of the Advances in Neural Information Processing Systems, Barcelona, Spain, 5–10 December 2016; pp. 2378–2386.
56. Liu, Q. Stein variational gradient descent as gradient flow. In Proceedings of the Advances in Neural Information Processing Systems, Long Beach, CA, USA, 4–9 December 2017; pp. 3115–3123.
57. Garbuno-Inigo, A.; Nüsken, N.; Reich, S. Affine invariant interacting Langevin dynamics for Bayesian inference. *arXiv* **2019**, arXiv:1912.02859.
58. Otto, F. The geometry of dissipative evolution equations: The porous medium equation. *Commun. Partial Differ. Equ.* **2001**, *26*, 101–174.
59. Villani, C. *Optimal Transport: Old and New*; Springer Science & Business Media: Berlin/Heidelberg, Germany, 2008.
60. Frogner, C.; Poggio, T. Approximate inference with Wasserstein gradient flows. *arXiv* **2018**, arXiv:1806.04542.
61. Caluya, K.; Halder, A. Gradient flow algorithms for density propagation in stochastic systems. *IEEE Trans. Autom. Control* **2019**. doi:10.1109/TAC.2019.2951348.
62. Batz, P.; Ruttor, A.; Oppen, M. Variational estimation of the drift for stochastic differential equations from the empirical density. *J. Stat. Mech. Theory Exp.* **2016**, *2016*, 083404.
63. Milstein, G.; Tretyakov, M. Computing ergodic limits for Langevin equations. *Phys. D Nonlinear Phenom.* **2007**, *229*, 81–95.
64. Saremi, S.; Mehrjou, A.; Schölkopf, B.; Hyvärinen, A. Deep energy estimator networks. *arXiv* **2018**, arXiv:1805.08306.

65. Reich, S.; Cotter, C. *Probabilistic Forecasting and Bayesian Data Assimilation*; Cambridge University Press: Cambridge, UK, 2015.
66. Lorenz, E.N. Deterministic nonperiodic flow. *J. Atmos. Sci.* **1963**, *20*, 130–141.
67. Bobkov, S.; Ledoux, M. *One-Dimensional Empirical Measures, Order Statistics and Kantorovich Transport Distances*; Amer Mathematical Society: Providence, RI, USA, 2014.



© 2020 by the authors. Licensee MDPI, Basel, Switzerland. This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution (CC BY) license (<http://creativecommons.org/licenses/by/4.0/>).