
Learning effective stochastic differential equations from microscopic simulations: combining stochastic numerics and deep learning

Felix Dietrich *

Technical University of Munich
Munich, Germany
felix.dietrich@tum.de

Alexei Makeev

Moscow State University
Moscow, Russia
amak@cs.msu.ru

George Kevrekidis

University of Massachusetts
Amherst, MA, USA
georgek@math.umass.edu

Nikolaos Evangelou

Johns Hopkins University
Baltimore, MD, USA
nevange2@jh.edu

Tom Bertalan

Johns Hopkins University
Baltimore, MD, USA
tom@tombertalan.com

Sebastian Reich

University of Potsdam
Potsdam, Germany
sereich@uni-potsdam.de

Ioannis G. Kevrekidis

Johns Hopkins University
Baltimore, MD, USA
yannisk@jhu.edu

Abstract

We identify effective stochastic differential equations (SDE) for coarse observables of fine-grained particle- or agent-based simulations; these SDE then provide coarse surrogate models of the fine scale dynamics. We approximate the drift and diffusivity functions in these effective SDE through neural networks, which can be thought of as effective stochastic ResNets. The loss function is inspired by, and embodies, the structure of established stochastic numerical integrators (here, Euler-Maruyama and Milstein); our approximations can thus benefit from error analysis of these underlying numerical schemes. They also lend themselves naturally to “physics-informed” gray-box identification when approximate coarse models, such as mean field equations, are available. Our approach does not require long trajectories, works on scattered snapshot data, and is designed to naturally handle different time steps per snapshot. We consider both the case where the coarse collective observables are known in advance, as well as the case where they must be found in a data-driven manner.

1 Introduction

Residual networks (ResNets [22], but see also [46]) successfully employ a forward-Euler integrator based approach to create very deep architectures. This has inspired followup work to include other integrator schemes for deterministic, ordinary differential equations (ODE) such as symplectic integrators [6, 21, 53]. In this work, we propose to also include stochastic integrators in the search for proper loss functions for network training.

Identification of dynamical systems from data is an established systems task [45, 39, 43, 7, 9, 29]. Machine learning both in its “second spring” in the the 80-90 years, and its current “third spring” has

*Open source code available at <https://gitlab.com/felix.dietrich/sde-identification>.

made a major impact in the approaches and algorithms for this task. From ResNets [22] and Neural ODEs [11, 46] to “neural PDEs” [18] and DeepONets [35], the identification of dynamical systems from (discrete) spatio-temporal data is a booming business because it is now comparatively easy to program and train neural networks. Learning SDEs and SPDEs is a natural “next step” and there is also an array of excellent contributions to this as we outline in the related work section below. The time-honored SDE estimation techniques for local drift and diffusivity can now be synthesized in a (more or less global) surrogate model.

Given all these recently available options, what is the point of our work? On the application side, we were motivated by the identification of coarse, or *effective* SDEs: atomistic/stochastic/agent-based physics simulations codes whose macroscopic statistics (e.g. leading moment evolution) have long been modelled as effective SDEs, whether with mathematical proofs or heuristically. This is the reason for part of our title, and for some of the examples: lattice simulations of epidemic models.

There is also, however, a methodological motivation whose roots are in the numerical analysis of integration schemes for SDEs/SPDEs. These established algorithms (starting with Euler-Maruyama [23]) iteratively produce discrete time data; it makes sense to “mathematically inform” the neural network architecture used for identification with the structure of the numerical integrator one would use on an explicitly available, closed form coarse stochastic model. This way, our work naturally links with, and can take advantage of, the extensive body of work in what is called “backward error analysis” of these integration schemes [54].

A third motivation for our work: beyond its links to backward error analysis, the approach naturally lends itself to the what one might call “physics-informed” or “gray-box” identification algorithms: neural network architectures that seamlessly incorporate approximate or partially known analytical models if these are available. We developed such an approach for deterministic ODE identification in the early 90s, devising recurrent Runge-Kutta ResNet architectures for this purpose, and extending them to gray boxes and implicit ResNets [46, 44, 18]. It is a natural next step to extend these techniques so that they can be applied to the identification of effective stochastic models, allowing the possibility of physics-informed gray boxes, and exploiting mathematical support via numerical backward error analysis. Learned, black-box models for agent-based systems may be used in malicious or questionable ways in technology and society. An example could be careless, uncontrolled generalization in predicting the behavior of physical or societal systems using learned models. Using gray-box, physics-informed machine learning methods that are amenable to thorough mathematical analyses and interpretability is important to mitigate such risks, and our optimization approach is specifically designed for this type of knowledge integration and analysis.

Our contributions are as follows:

1. We introduce a neural network training loss function that allows us to approximate the state dependent drift and diffusivity of an SDE from discrete observation data.
2. The approach does not require long time series: paired snapshots (x_t, x_{t+h}) , scattered over the domain of interest, are sufficient. Step sizes h can vary for each snapshot.
3. The approach works point-wise: No distributions need to be approximated from data, nor do we need to compute distances between distributions.
4. We formulate explicit loss functions for Euler-Maruyama- and Milstein-type integration schemes using Gaussian noise.
5. We demonstrate that we can employ the approach to coarse-grain particle/lattice simulations; in particular, we approximate the drift and diffusivity of a two-dimensional mean-field SDE for a Gillespie-based [16, 17] for a kinetic Monte-Carlo lattice simulation of an epidemiological SIR model.

2 Related work

Loss functions and network architectures informed by numerical integration methods were devised several decades ago, e.g. [46, 44]. Recently, these ideas have been rediscovered and used in deep architectures such as ResNets [22], with symplectic schemes for Hamiltonian [6, 21, 53] and Poisson networks [25]. Scalable gradient approaches were developed for SDE approximation recently [33], which allow backpropagation through a long series of SDE iterations. If only a single time step for each data point is available, as in our work, this type of scalability is not an immediate training issue.

The theory of SDE is well-developed; We refer to standard books on stochastic processes [41], stochastic calculus [26], and stochastic algorithms [42]. In the broader machine learning community, SDE are used mostly for generative modeling, meaning that the actual drift and diffusivity functions learned from the data are not relevant, as long as the generated distribution is “acceptable”. A classical example are generative-adversarial networks [19, 24]. In recent years, learning random ODEs (deterministic ODE with random parameters) has been done using an approach with generative-adversarial networks [34]. Long-Short-Term-Memory networks were used in a stochastic setting for the generation of hand-written text [20]. The authors of [49] present a framework that learns an SDE for training and sampling from score-based models, allowing sample generation.

In a classical dynamical systems context, the SDE itself is the desired outcome of approximation and thus, typically, different methods are necessary. Gaussian Processes have been used to identify SDE from observations [52], with a target likelihood based on the Euler-Maruyama scheme (not explicitly mentioned by the authors). Our setting is very different from identifying ODE from noisy observations, for example through Gaussian Processes [51]. Generative stochastic modeling of strongly nonlinear flows with non-Gaussian statistics is also possible [3]. Density and ensemble approximation techniques [50] provide a complementary view to the particle-like approach with SDE. Stochastic dynamical systems can also be described through the associated Fokker-Planck operator. See, e.g. [5], for a non-parametric forecasting method using the “shift-map”, a numerical approximation of the operator projected on the eigenfunctions of the Laplace operator.

3 Mathematical problem setting

We discuss the SDE

$$dx_t = f(x_t)dt + \sigma(x_t)dW_t, \quad (1)$$

where $f : \mathbb{R}^n \rightarrow \mathbb{R}^n$ and $\sigma : \mathbb{R}^n \rightarrow \mathbb{R}^{n \times n}$ are smooth, possibly nonlinear functions; every $\sigma(x)$ is positive and bounded away from zero (and a positive-definite matrix if $n > 1$), and W_t a Wiener process such that for $t > s$, $W_t - W_s \sim \mathcal{N}(0, t - s)$. See [41] for an introduction to stochastic processes. We assume we have access to a set of N snapshots $D = \{(x_1^{(k)}, x_0^{(k)}, h^{(k)})\}_{k=1}^N$, where $x_0^{(k)}$ are points scattered in the state space of (1) and the value of $x_1^{(k)}$ results from the evolution of (1) under a small time-step $h^{(k)} > 0$, starting at $x_0^{(k)}$. The snapshots are samples from a distribution, $x_0 \sim p_0$, and the transition densities $x_1 \sim p_1(\cdot|x_0, h)$ are associated with (1) for a given time-step $h > 0$, which in turn is chosen from some distribution p_h . The joint data-generating distribution is therefore given by $p(x_0, x_1, h) = p_1(x_1|x_0, h)p_0(x_0)p_h(h)$. Alternatively, the data could be collected along a long trajectory $\{x_{t_i}\}$ of (1) with sample frequency $h_i > 0$, that is, $t_{i+1} = t_i + h_i$. The problem is to identify the drift f and diffusivity σ through two neural networks $f_\theta : \mathbb{R}^n \rightarrow \mathbb{R}^n$ and $\sigma_\theta : \mathbb{R}^n \rightarrow \mathbb{R}^{n \times n}$, parameterized by their weights θ , only from the data in D . We assume the points in D are sampled sufficiently densely in the region of interest.

As we will see, each choice of SDE integrator will give a different training algorithm, with a loss based on the particular discretization scheme; training algorithms based on integrators of different accuracy, integrators with different rates and types of convergence (strong or weak), even integrators based on different stochastic calculus type (Ito vs Stratonovich) fall under this umbrella. We will demonstrate these ideas here through two Ito-calculus based training algorithms: the Euler-Maruyama and the Milstein schemes.

3.1 Identification of drift and diffusivity with the Euler-Maruyama scheme

We now formulate and rationalize the loss term that we use to train the neural networks f_θ and σ_θ . The Euler-Maruyama scheme is a simple method to integrate (1) over a small time $h > 0$:

$$x_1 = x_0 + hf(x_0) + \sigma(x_0)\delta W_0, \quad (2)$$

where $h > 0$ is small and δW_0 is a vector of n random variables, all i.i.d. and normally distributed around zero with variance h . The convergence of (2) for $h \rightarrow 0$ has been studied at length; We refer to standard literature [41]. We can use this idea to construct a loss function for training the networks σ_θ and f_θ simultaneously. We initially restrict the discussion to the case $n = 1$ for simplicity. Essentially, conditioned on x_0 and h , we can think of x_1 as a point drawn from a multivariate normal distribution

$$x_1 \sim \mathcal{N}(x_0 + hf(x_0), h\sigma(x_0)^2). \quad (3)$$

In the training data set D , we only have access to triples $(x_0^{(k)}, x_1^{(k)}, h^{(k)})$, and not the drift f and diffusivity σ . To approximate them, we define the probability density p_θ of the normal distribution (3) and then, given the neural networks f_θ and σ_θ , ask that the log-likelihood of the data D under the assumption in equation (3) is high:

$$\theta := \arg \max_{\hat{\theta}} \mathbb{E} [\log p_{\hat{\theta}}(x_1|x_0, h)] \approx \arg \max_{\hat{\theta}} \frac{1}{N} \sum_{k=1}^N \log p_{\hat{\theta}}(x_1^{(k)} | x_0^{(k)}, h^{(k)}). \quad (4)$$

We can now formulate the loss function that will be minimized to obtain the neural network weights θ . The logarithm of the well-known probability density function of the normal distribution, together with the mean and variance from (3), yields the loss to minimize during training,

$$\mathcal{L}(\theta|x_0, x_1, h) := \frac{(x_1 - x_0 - hf_\theta(x_0))^2}{h \sigma_\theta(x_0)^2} + \log |h \sigma_\theta(x_0)^2| + \log(2\pi). \quad (5)$$

This formula can easily be generalized to higher dimensions (see supplement), and we use such generalizations for examples in more than one dimension. Minimizing \mathcal{L} in (5) over the data D implies maximization of the log marginal likelihood (4) with the constant terms removed (as they do not influence the minimization) [41]. Likelihood estimation in combination with the normal distribution is used in many variational and generative approaches [19, 28, 33, 52, 51, 49]. Note that here, the step size $h^{(k)}$ is defined *per snapshot*, so it is possible that it has different values for every index k . This will be especially useful in coarse-graining Gillespie simulations, where the time step is determined as part of the scheme.

We note that the specific form (3) of the transition probabilities p_θ induced by the Euler-Maruyama method implies that it will, in general, be impossible to exactly reproduce the data-generation density $p(x_0, x_1, h)$ for any finite step-size h . This motivates the application of more refined numerical methods, such as the Milstein scheme considered in section 3.2 below. Alternatively, one can think of p_θ as being induced by an SDE (1) with modified drift f and diffusion σ . Modified equations have been derived for a variety of numerical methods, such as the Euler or the Milstein method [54]. In general, there is no second-order modified equation for the Euler-Maruyama method, while one exists for the Milstein scheme. Modified equation analysis in our setting is typically performed with equal step-sizes $h^{(k)}$ for all data points.

3.2 Identification of drift and diffusivity with the Milstein scheme

We now discuss how the loss function can be based on the Milstein integration scheme [40], to illustrate complexities that arise when deviating from simple Euler schemes, even in one dimension. For $n = 1$, the Milstein scheme is a more accurate method to integrate (1), and is given through

$$x_1 = x_0 + hf(x_0) + \sigma(x_0) \delta W_0 + \frac{1}{2} \sigma(x_0) \frac{d}{dx} \sigma(x_0) ((\delta W_0)^2 - h), \quad (6)$$

where $h > 0$ is small and δW_0 is a random variable, normally distributed around zero with variance h . For the scheme (6), it is also possible to write down the distribution of x_1 , similar to equation (3). It is no longer normal, however, because of the square of the variable δW_0 . To determine the distribution of x_1 , we define

$$A := x_0 + hf(x_0) - h \frac{1}{2} \sigma(x_0) \frac{d}{dx} \sigma(x_0); \quad B := \sqrt{h} \sigma(x_0); \quad C := h \frac{1}{2} \sigma(x_0) \frac{d}{dx} \sigma(x_0).$$

Then, for a random variable $z \sim \mathcal{N}(0, 1)$, we have $x_1 = A + Bz + Cz^2$. Assuming $C \neq 0$,

$$\frac{x_{k+1} - A}{C} = \frac{B}{C} z + z^2 = \left(z + \frac{B}{2C} \right)^2 - \frac{B^2}{4C^2} \iff \frac{x_{k+1} - A + \frac{B^2}{4C}}{C} = \left(z + \frac{B}{2C} \right)^2 =: \xi.$$

The variable ξ has a non-central χ^2 distribution with one degree of freedom and non-centrality parameter $\lambda = \left(\frac{B}{2C} \right)^2$. With this derivation, we can see that x_1 is distributed according to a linear transformation of a non-central χ^2 distribution. The probability density of ξ is

$$p_\xi \left(y; k = 1, \lambda = \frac{B^2}{4C^2} \right) = \frac{1}{2} \exp(-(y + \lambda)/2) \left(\frac{y}{\lambda} \right)^{k/4 - 1/2} I_{k/2 - 1}(\sqrt{\lambda y}),$$

where $I_{k/2-1}$ is a modified Bessel function of the first kind. We can now use the function p_ξ to define the probability density p_θ of x_1 , given x_0 :

$$p_{\text{Milstein},\theta}(x_1|x_0,h) := \frac{1}{C} p_\xi \left(\frac{x_1 - A + \frac{B^2}{4C}}{C}; k=1, \lambda = \frac{B^2}{4C^2} \right). \quad (7)$$

The corresponding point-wise loss function then is very similar to (5),

$$\mathcal{L}(\theta|x_0, x_1, h) := -\log p_{\text{Milstein},\theta}(x_1|x_0, h). \quad (8)$$

Equation (7) defines a probability density that can vanish over large portions of its domain, which leads to issues both with log-marginal likelihood estimation and with gradient computations. These issues can be mitigated through regularization, but also through kernel density approximations, which we outline in Sec. 3.3. To illustrate the difference of the probability distributions for the Euler-Maruyama and the Milstein templates, we plot them for $h = 0.15$, $f(x) = x + 1$, $\sigma(x) = \frac{1}{4}x^2 + x + \frac{1}{2}$ in Fig. 1 against 10^5 sampled points using the two templates with the starting point $x_0 = 0$.




Figure 1: Probability density functions and sampled data from the two numerical integrator templates.

3.3 Approximation of the probability density of numerical integration schemes

Inspired by the iteration procedures use to discover kernels for deep Gaussian processes [31] and kernel density estimates [23, 47], we now describe an iterative method to approximate the probability density of numerical integration schemes. This will be useful in cases for which the density cannot be derived in analytical, closed form, or for schemes with potentially locally vanishing probability density, like the Milstein method.

Consider an arbitrary (but convergent and accurate) numerical integration scheme ϕ for (1), so that for a fixed step size $h > 0$ and a noise process δW_0 , a single time-step becomes

$$x_1 = \phi(x_0; h, \delta W_0). \quad (9)$$

Similar to Sections 3.1 and 3.2, given the scheme ϕ as well as the approximate drift f_θ and diffusivity σ_θ parametrized by θ , we need to approximate the probability density $p_\theta(x_1|x_0, h)$. For schemes that rely on noise processes δW_0 that can be sampled easily (e.g. Gaussian noise as in Euler-Maruyama and Milstein schemes), the key idea for our numerical approximation is the following. We begin by sampling a set of M points $\delta W_0^{(i)}$ from the distribution of δW_0 . Then, we use the SDE integration scheme ϕ to estimate candidates $\hat{x}_1^{(i)}$ for the next point in time, given x_0 and h . Finally, we describe the density of these candidates through a mixture of Gaussians with a small standard deviation $\sigma_r > 0$, so that the actual probability distribution p_ϕ is approximated:

$$p_\phi(x_1|x_0, h) \approx \frac{1}{M} \sum_{i=1}^M p_{\mathcal{N}(\hat{x}_1^{(i)}, \sigma_r^2)}(x_1). \quad (10)$$

The approximation accuracy of (10) improves as $M \rightarrow \infty$ and $\sigma_r \rightarrow 0$; for a discussion of this, we refer to standard literature on kernel density estimation [23, 47] and viscosity solutions for PDE [14, 13]. Fig. 2 demonstrates this convergence for varying σ , with fixed $M = 10^4$, $x_0 = 0.01$, $h = 1.0$, $f(x) = 2x$, $\sigma(x) = x$, and ϕ the Milstein scheme with the exact form for p_ϕ given through equation (7). Many other density estimation methods are applicable here as well, such as generative-adversarial networks [19] or normalizing flows [30].




Figure 2: Approximations of (7) at varying scales σ_r , in standard and log-scale.

3.4 Exploiting latent spaces

While we initially consider cases where the coarse variables are already known (e.g. their mean-field or quasi-chemical approximation variables), many observations in the real world do not immediately reveal the smallest dimensionality of the space in which the dynamics evolve. Latent space identification and related unsupervised learning are possible with a wide variety of techniques. We also learned such latent spaces with Diffusion Maps [12] as well as an SDE-integrator-informed auto-encoder with the SDE loss (5) augmented by the log mean squared loss: $\mathcal{L}(\theta|D) = \log(\frac{1}{N} \sum_{k=1}^N (\text{decoder}(\text{encoder}(x_0^{(k)})) - x_0^{(k)})^2) + \mathcal{L}_{\text{Euler-Maruyama}}$. The Diffusion Map results can be interpreted as a latent space; Learning an SDE on these new variables, and then mapping back to the original data is thus similar to latent SDE identification [33]. A description of the methods to obtain these latent spaces and the identification of SDE on them for kMC lattices can be found in the supplemental material.

4 Computational experiments

We first illustrate the neural network identification techniques on toy examples, and then show an application to learning effective SDEs from particle dynamics simulations. Tab. 1 lists important training parameters. All computational experiments were performed on a single desktop computer with an AMD Ryzen 7 3700X 8-Core Processor, 32GB RAM, and a NVIDIA GeForce GTX 1070 GPU. Training did not take longer than 30 minutes for each case, data generation did not take longer than two hours. The neural networks and loss functions were defined in TensorFlow V2.4.1 [1] (Apache 2.0 license). For the toy examples, we created snapshots by integrating the SDEs with 10 steps of Eq. (2) from $t = 0$ to $t = h$, only storing the initial and final result. The given network topology was used twice, separately for drift and diffusivity. If not otherwise stated, we use a test/validation split of 90%/10%, the ADAM optimizer with standard parameters $\text{learning_rate}=0.001$, $\beta_1 = 0.9$, $\beta_2 = 0.999$, $\epsilon = 10^{-7}$, batch size of 32, 100 epochs, and exponential linear unit activation functions.

Table 1: Parameters, equations for loss, training and validation errors for the experiments.

Example	h	# points	Network topology	Loss eq.	Train	Val.
Sec. 4.1: cubic, 1D	0.01	10,000	1 50 50 1	(5)	-1.59	-1.57
Sec. 4.1: cubic, 1D	0.01	10,000	1 50 50 1	(7)	-1.60	-1.61
Sec. 4.1: linear, 3D	0.25	15,000	3 50 50 3/6 drift/diff.	(5)	-7.60	-7.63
Sec. 4.1: 1D–20D	0.01	10,000	$n 50 50 n$	(5)	varies	varies
Sec. 4.2: Gillespie	variable	90,000	2 50 50 2	(5)	-10.9	-11.0
Sec. 4.2: kMC lattice	variable	90,000	2 50 50 2	(5)	-11.7	-11.8

4.1 Toy examples: cubic drift, linear diffusivity, higher dimensions

We tested our implementation with examples in one dimension first, with drift f and diffusivity σ defined through $f(x_t) = -2x_t^3 - 4x_t + 1.5$, $\sigma(x_t) = 0.05x_t + 0.5$. A comparison between the learned and true functions is shown in Fig. 3, along with densities obtained from sampling the true and approximate SDE. We trained networks with both Euler-Maruyama and Milstein loss functions, with very comparable results—even the training and validation curves were very similar. When increasing

the dimension from $n = 1$ to $n = 20$, approximation quality decreases after $n = 6$. Non-diagonal diffusivity matrices were approximated accurately for $n = 3$ (see supplemental material).




Figure 3: Approximation vs. true values of f and σ in example 4.1. Densities over time, from points sampled uniformly in $[-2, 2]$ and then integrated using the true SDE and using the approximation of drift and diffusivity with the network.

4.2 Learning coarse, effective SDE from fine-grained (particle) simulations

We validate our approach through a well-studied model for which analytical approximations at several levels of coarse-graining are available. The problem is the dynamics of epidemics on a lattice, which we briefly outline below. We consider two different versions of the model, and “learn” the corresponding effective SDEs. The first version at the finest scale involves Monte Carlo lattice simulations; the second version involves Gillespie simulation of the mechanism implied by the mean field model. Note that we discover (a) an approximation to the mean field SDE if we would use mean field SDE data, and (b) an approximation to dynamics in the mean field variables which is different than the mean field SDE, if we get data from Gillespie/lattice simulations and coarse-grain it in terms of mean field variables [36, 37].

Epidemic SIRS model: We consider a hierarchical sequence of three stochastic models describing the dynamics of the classical SIR-type (susceptible-infected-recovered) epidemic process [27, 48, 2]. The first, most detailed model is represented by kinetic Monte Carlo (kMC) simulations on a lattice. The kMC model considers a two-dimensional square lattice of size $N = N_x \times N_y$ with periodic boundary conditions, where each site contains a single individual. In accordance to the health state, each individual can be in one of three states: Susceptible (S), Infected (I) or Recovered (R). The elementary reactions (independent Poisson processes) occurring on a lattice are as follows: 1) $I + S \rightarrow I + I$ (rate constant k_1 ; infection), 2) $I \rightarrow R$ (rate constant k_2 ; recovery with immunity), 3) $R \rightarrow S$ (rate constant k_3 ; loss of immunity).

The infection spreads through nearest neighbor bonds only; thus the first reaction may occur only when I and S are the first neighbors on a lattice. Elementary events take place with transition probabilities k_p ($p = 1, 2, 3$). When $k_3 > 0$, the model is called as SIRS, while the SIR model is a particular case of the SIRS model with $k_3 = 0$. In addition to steps 1)-3), migration of individuals occurring via the exchange mechanism is considered: 4) $S_i + I_j \rightarrow I_i + S_j$ (rate constant d_1), 5) $S_i + R_j \rightarrow R_i + S_j$ (rate constant d_2). Here, the sites i and j are first neighbors on a lattice. The time evolution of the probability distribution for the populations of individuals on a lattice can be described by the chemical master equation, derived from the Markov property of the underlying stochastic process. To generate realizations of this process, we applied the so-called rejection-free direct kMC algorithm [8, 16] which produces a sequence of configurations and the times when they change from one to the other. It is interesting to mention that the same lattice kMC model can be formulated in a mathematically equivalent form using the following three types of states: Excited (infected), Refractory (recovered) or Quiescent (susceptible). In this case the model may reproduce the typical patterns of excitable media such as traveling pulses and spiral waves [38]. Furthermore, we assume that $d = d_1 = d_2$. In essence, the migration rate d may take into account quarantine restrictions and other regulations in place to keep the infection from spreading. Small d values correspond to strong restrictions. If a spatially homogeneous, random distribution of all individuals on a lattice is assumed, the problem can be significantly simplified. Such a homogeneous distribution can be achieved in the “fast migration” limit ($d \rightarrow \infty$).

The second stochastic model we study is represented by the Gillespie's stochastic simulation algorithm (SSA, [16]) that is used to simulate a well-stirred, spatially homogeneous system. For the SIRS model, SSA operates in terms of three integer stochastic variables: the number of susceptible individuals, n_0 , the number of infected individuals, n_1 , and the number of recovered individuals, n_2 ; their concentrations are calculated as: $y_p = n_p/N$. The mass balance implies that $y_0 + y_1 + y_2 = 1$. At each time step of the SSA algorithm the following three rates are calculated through $r_1 = 4k_1y_0y_1$, $r_2 = k_2y_1$, $r_3 = k_3y_2$. Then, an event to be performed is selected with a probability proportional to its rate and the numbers n_p are updated accordingly.

The third stochastic model is represented by a set of stochastic differential equations (SDE) of the Langevin type which, under certain assumptions, approximates the results of the SSA algorithm [17]. The SDE model is as follows:

$$\begin{aligned}\frac{dy_0}{dt} &= -r_1 + r_3 - \sqrt{r_1/N}dW_1(t) + \sqrt{r_3/N}dW_3(t), \\ \frac{dy_1}{dt} &= (r_1 - r_2) + \sqrt{r_1/N}dW_1(t) - \sqrt{r_2/N}dW_2(t), \\ \frac{dy_2}{dt} &= r_2 - r_3 + \sqrt{r_2/N}dW_2(t) - \sqrt{r_3/N}dW_3(t).\end{aligned}\tag{11}$$

Here, $W_p(t)$ are independent Wiener processes such that $dW_p(t) \sim \mathcal{N}(0, dt)$. Only two of three SDEs must be considered, since $y_0 + y_1 + y_2 = 1$. In the limit $N \rightarrow \infty$, the SDE system transforms to the standard mean-field SIRS model. In particular, for the SIR model this SDE system reduces to having a diagonal diffusivity matrix:

$$dy_0 = -r_1 dt + \sqrt{r_1/N}dW_1(t), \quad dy_2 = r_2 dt + \sqrt{r_2/N}dW_2(t).\tag{12}$$

Eq. (11) or (12) can be integrated using the Euler-Maruyama method with fixed time step h . The SDE models approximate the solution of both SSA and kMC models. The following parameters and initial conditions are used in the simulations: $k_1 = 1$, $k_2 = 1$, $k_3 = 0$; $y_0(0) = 0.9$, $y_1(0) = 0.1$.

We assume that SSA and kMC simulations generate trajectories which can be approximated by the solution of two latent SDEs with diagonal diffusion matrix, similar to Eq. (12). Under this assumption, the neural network was used to learn the unknown drift and diffusion terms at the fixed values of the model parameters k_1, k_2, k_3 and N . Training was performed using a set of N_{tr} short-time stochastic simulations on a $N_1 \times N_2$ lattice with the random, but physically-relevant initial conditions for the average concentrations $y_m \in (0, 1)$, $m \in \{0, 1, 2\}$. The initial species distribution on a lattice was also random. During each independent SSA and kMC simulation run, the data points were evaluated at the time moments $t_0 \in \{0, t_1, \dots, t_{\max}\}$, where $h_i = t_{i+1} - t_i \approx h$. The values of t_{\max} and h are the fixed parameters of the algorithm. Note that a time step in the SSA/kMC algorithm is a random variable. The time is advanced via an increment given from the exponential distribution: $dt = -\log(\xi)/R$, where R is the total reaction rate and $\xi \in (0, 1)$ is a uniformly distributed random number. Therefore, h_i are close to h , but they are not constant. Each data point used for training contains the following five values: $h_i, y_{(0,i)}, y_{(2,i)}, y_{(0,i+1)}, y_{(2,i+1)}$. The total number of data points is $N_{dp} = N_{tr} ([t_{\max}/h] - 1)$. The trained network approximates the unknown drift and diffusion terms at any given values of $y_m \in (0, 1)$. Thus, the SDE approximation in combination with the Euler-Maruyama method provides a time-stepper (decoder) which can reproduce the dynamics.

Gillespie simulations: Fig. 4 shows that a network can be trained to match the results of SSA simulations. A distribution function (histogram) of the time steps h_i is plotted in Fig. 4(a). It demonstrates that our approach works with variable time steps in the training data set. Three sample paths of y_0 and y_1 and the average of 200 paths are shown in Figs. 4(b) and 4(c).

Kinetic Monte-Carlo lattice simulations: Fig. 5 shows that our approach successfully reproduces the results of kMC simulations on a 32×32 lattice. The Euler-Maruyama method with a fixed time step $h_{EM} = 0.002$ was used in numerical integration of the network SDE approximation. Fig. 5(a) demonstrates three sample paths obtained using the neural network (blue) and kMC (red, "true solution"). The averaged paths are presented in Fig. 5(b). A sufficiently high accuracy network approximation is demonstrated in Figs. 5(c) and 5(d), where the probability density functions for y_0 and y_1 at $t = 2$ are shown. They were estimated by NN and kMC, and are visually identical.

The described examples consider the lattice SIR model which has only the trivial steady states with $y_1 = 0$. When $k_3 > 0$, the system exhibits richer dynamic behavior. Simulations with $k_3 = 0.2$ were performed as well, and the network model is capable of reproducing the correct steady state and damped oscillations (the mean-field SIRS model has a focus-type steady state at these values of




Figure 4: Histogram of the time steps h_i used in the training data set (a). Three sample paths of y_0 and y_1 (b); Average of 200 paths with deviation of $\max |y_{i,NN} - y_{i,SSA}| < 0.02$ (c). Blue (red) lines correspond to NN (SSA) results. Parameters: $N = 400$; $N_{tr} = 10000$, $t_{max} = 0.1$, $h = 0.01$.

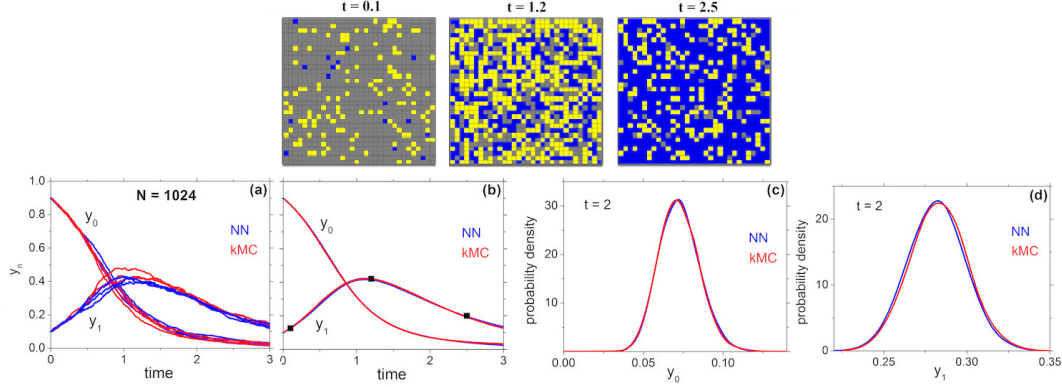


Figure 5: Illustration of the kMC lattice (top row), physically relevant values y_0, y_1 measured over time resp. simulated with the identified SDE from the network (a), averaged paths over 200 simulations with a deviation $\max |y_{i,NN} - y_{i,kMC}| < 0.01$ (b), and propagated densities from the initial condition until $t = 2$ (c,d). Lattice snapshots show S, I and R-type species as grey, yellow and blue squares, respectively. Parameters: 32×32 lattice, $d = 50$; $N_{tr} = 4000$, $t_{max} = 0.05$, $h = 0.01$.

parameters). This is actually surprising: When the SIRS model is considered, the diffusion matrix is not diagonal, because there are three different reactions. Nevertheless, the trained network is capable here of approximating the kMC results, even when only trained with a diagonal diffusivity matrix.

5 Discussion

Training neural networks with loss functions based on numerical integrators such as (5) and (8) has several limitations. If the time step between many samples is too small, we cannot accurately identify the drift, because the diffusivity term will dominate. This could be mitigated by starting with small h , estimating the diffusivity, and then estimating the drift with subsampled trajectories. Even with infinite data the drift is difficult to estimate, because the time step also has to go to zero [41]. Conversely, if the time step is too big, we cannot accurately identify the diffusivity. If the dynamics of the coarse-grained observables include rare events, then learning the corresponding SDE is a challenge, because these events will not be present in a lot of snapshots, and hence not a lot of data is available to learn them. Using a loss function based on Lévy noise SDE integrators might provide a useful surrogate for SDEs with rare-events.

There are many possible extensions and applications of our work. An important task is to find explicit probability density functions for integrators using other noise types, such as Lévy or Poisson noise. Analyzing integrators based on other types of calculus such as Heun's [10] method may help impose more variegated types of priors during training. It is a predictor-corrector method, $\tilde{x}_1 = x_0 + f(x_0)h + \sigma(x_0)z$; $x_1 = x_0 + 1/2[f(x_0) + f(\tilde{x})]h + 1/2[\sigma(x_0) + \sigma(\tilde{x}_1)]z$; $z \sim \mathcal{N}(0, h)$, based on Stratonovich calculus, and will give rise to implicit neural networks. Using the adjoint method [34] can allow us to propagate the loss gradients back through longer time series. The choice

of numerical approximation (9) entails a systematic error for the data generating distribution which can impact adversely on the capability to generalize beyond the training data. This effect can be minimized by higher-order numerical methods and analysed through modified equation analysis [54].

In terms of applications, many more particle- and agent-based models can be coarse grained. For many of these, it is important to find coarse variables, and latent space techniques such as VAE-type architectures and Diffusion Maps discussed in Sec. 3.4 will help to identify them. These techniques can help to construct local SDE models in the process of larger-scale simulations to guide sampling and exploration. Another clear next step based on numerical analysis is the identification of stochastic partial differential equations (SPDEs).

Acknowledgments and Disclosure of Funding

The research of S.R. has been partially funded by the Deutsche Forschungsgemeinschaft (DFG) - Project-ID 318763901 - SFB1294. The work of I.G.K was partially supported by the US Department of Energy and the DARPA ATLAS program.

References

- [1] M. Abadi, A. Agarwal, P. Barham, E. Brevdo, Z. Chen, C. Citro, G. S. Corrado, A. Davis, J. Dean, M. Devin, S. Ghemawat, I. Goodfellow, A. Harp, G. Irving, M. Isard, Y. Jia, R. Jozefowicz, L. Kaiser, M. Kudlur, J. Levenberg, D. Mané, R. Monga, S. Moore, D. Murray, C. Olah, M. Schuster, J. Shlens, B. Steiner, I. Sutskever, K. Talwar, P. Tucker, V. Vanhoucke, V. Vasudevan, F. Viégas, O. Vinyals, P. Warden, M. Wattenberg, M. Wicke, Y. Yu, and X. Zheng. TensorFlow: Large-scale machine learning on heterogeneous systems, 2015. Software available from tensorflow.org.
- [2] L. J. Allen. A primer on stochastic epidemic models: Formulation, numerical simulation, and analysis. *Infectious Disease Modelling*, 2(2):128–142, May 2017.
- [3] H. Arbabi and T. Sapsis. Generative stochastic modeling of strongly nonlinear flows with non-gaussian statistics. *arXiv:1908.08941*, 2019.
- [4] T. Berry, J. R. Cressman, Z. Gregurić-Ferenček, and T. Sauer. Time-scale separation from diffusion-mapped delay coordinates. *SIAM Journal on Applied Dynamical Systems*, 12(2):618–649, Jan. 2013.
- [5] T. Berry, D. Giannakis, and J. Harlim. Nonparametric forecasting of low-dimensional dynamical systems. *Physical Review E*, 91(3), 3 2015.
- [6] T. Bertalan, F. Dietrich, I. Mezić, and I. G. Kevrekidis. On learning hamiltonian systems from data. *Chaos: An Interdisciplinary Journal of Nonlinear Science*, 29(12):121107, Dec. 2019.
- [7] J. Bongard and H. Lipson. Automated reverse engineering of nonlinear dynamical systems. *Proceedings of the National Academy of Sciences*, 104(24):9943–9948, 2007.
- [8] A. B. Bortz, M. H. Kalos, and J. L. Lebowitz. A new algorithm for Monte Carlo simulation of Ising spin systems. *Journal of Computational Physics*, 17(1):10–18, 1975.
- [9] S. L. Brunton, J. L. Proctor, and J. N. Kutz. Discovering governing equations from data by sparse identification of nonlinear dynamical systems. *Proceedings of the National Academy of Sciences of the United States of America*, 113(15):3932–3937, 2016.
- [10] K. Burrage, P. M. Burrage, and T. Tian. Numerical methods for strong solutions of stochastic differential equations: an overview. *Proceedings of the Royal Society of London. Series A: Mathematical, Physical and Engineering Sciences*, 460(2041):373–402, 2004.
- [11] R. T. Q. Chen, Y. Rubanova, J. Bettencourt, and D. Duvenaud. Neural ordinary differential equations. In *NeurIPS conference 2018*, 2018.
- [12] R. R. Coifman and S. Lafon. Diffusion maps. *Applied and Computational Harmonic Analysis*, 21(1):5–30, July 2006.
- [13] M. G. Crandall, H. Ishii, and P.-L. Lions. User’s guide to viscosity solutions of second order partial differential equations. *Bull. Amer. Math. Soc.*, 27:1–67, 1992.
- [14] M. G. Crandall and P.-L. Lions. Viscosity solutions of hamilton-jacobi equations. *Transactions of the American Mathematical Society*, 277(1):1–42, 1983.

- [15] C. J. Dsilva, R. Talmon, R. R. Coifman, and I. G. Kevrekidis. Parsimonious representation of nonlinear dynamical systems through manifold learning: A chemotaxis case study. *Applied and Computational Harmonic Analysis*, 44(3):759–773, May 2018.
- [16] D. T. Gillespie. A general method for numerically simulating the stochastic time evolution of coupled chemical reactions. *Journal of Computational Physics*, 22(4):403–434, 1976.
- [17] D. T. Gillespie. The chemical langevin equation. *The Journal of Chemical Physics*, 113(1):297–306, July 2000.
- [18] R. González-García, R. Rico-Martínez, and I. Kevrekidis. Identification of distributed parameter systems: A neural net based approach. *Computers & Chemical Engineering*, 22:S965–S968, Mar. 1998.
- [19] I. Goodfellow, J. Pouget-Abadie, M. Mirza, B. Xu, D. Warde-Farley, S. Ozair, A. Courville, and Y. Bengio. Generative adversarial nets. In Z. Ghahramani, M. Welling, C. Cortes, N. D. Lawrence, and K. Q. Weinberger, editors, *Advances in Neural Information Processing Systems 27*, pages 2672–2680. Curran Associates, Inc., 2014.
- [20] A. Graves. Generating sequences with recurrent neural networks. *arXiv:1308.0850*, 2013.
- [21] S. Greydanus, M. Dzamba, and J. Yosinski. Hamiltonian neural networks. In *Proceedings of the 33rd International Conference on Neural Information Processing Systems*, 2019.
- [22] K. He, X. Zhang, S. Ren, and J. Sun. Deep residual learning for image recognition. In *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. IEEE, June 2016.
- [23] D. J. Higham. An algorithmic introduction to numerical simulation of stochastic differential equations. *SIAM Review*, 43(3):525–546, Jan. 2001.
- [24] A. Jalal, A. Ilyas, C. Daskalakis, and A. G. Dimakis. The robust manifold defense: Adversarial training using generative models. *arXiv:1712.09196v5*, 2017.
- [25] P. Jin, Z. Zhang, I. G. Kevrekidis, and G. E. Karniadakis. Learning poisson systems and trajectories of autonomous systems via poisson neural networks. *arXiv:2012.03133*, 2020.
- [26] I. Karatzas and S. E. Shreve. *Brownian Motion and Stochastic Calculus*. Springer New York, 1998.
- [27] W. O. Kermack and A. G. McKendrick. A contribution to the mathematical theory of epidemics. *Proceedings of the Royal Society of London. Series A, Containing Papers of a Mathematical and Physical Character*, 115(772):700–721, Aug. 1927.
- [28] D. P. Kingma and M. Welling. Auto-encoding variational bayes. In *Proceedings of the 2nd International Conference on Learning Representations (ICLR)*, Dec. 2014.
- [29] S. Klus, F. Nüske, S. Peitz, J.-H. Niemann, C. Clementi, and C. Schütte. Data-driven approximation of the koopman generator: Model reduction, system identification, and control. *Physica D: Nonlinear Phenomena*, 406:132416, May 2020.
- [30] I. Kobyzev, S. J. D. Prince, and M. A. Brubaker. Normalizing Flows: An Introduction and Review of Current Methods. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, Aug. 2019.
- [31] J. Lee, J. Sohl-dickstein, J. Pennington, R. Novak, S. Schoenholz, and Y. Bahri. Deep neural networks as gaussian processes. In *International Conference on Learning Representations*, 2018.
- [32] D. Lehmborg, F. Dietrich, G. Köster, and H.-J. Bungartz. datafold: data-driven models for point clouds and time series on manifolds. *Journal of Open Source Software*, 5(51):2283, July 2020.
- [33] X. Li, T.-K. L. Wong, R. T. Chen, and D. Duvenaud. Scalable gradients for stochastic differential equations. *arXiv:2001.01328*, 2020.
- [34] J. Liu, Z. Long, R. Wang, J. Sun, and B. Dong. Rode-net: Learning ordinary differential equations with randomness from data. *arXiv:2006.02377*, 2020.
- [35] L. Lu, P. Jin, G. Pang, Z. Zhang, and G. E. Karniadakis. Learning nonlinear operators via DeepONet based on the universal approximation theorem of operators. *Nature Machine Intelligence*, 3(3):218–229, 3 2021.
- [36] A. G. Makeev, D. Maroudas, and I. G. Kevrekidis. Coarse stability and bifurcation analysis using stochastic simulators: Kinetic Monte Carlo examples. *The Journal of Chemical Physics*, 116(23):10083–10091, June 2002.

- [37] A. G. Makeev, D. Maroudas, A. Z. Panagiotopoulos, and I. G. Kevrekidis. Coarse bifurcation analysis of kinetic monte carlo simulations: A lattice-gas model with lateral interactions. *The Journal of Chemical Physics*, 117(18):8229–8240, Nov. 2002.
- [38] A. G. Makeev and N. L. Semendyaeva. A basic lattice model of an excitable medium: Kinetic monte carlo simulations. *Mathematical Models and Computer Simulations*, 9(5):636–648, 2017-09.
- [39] I. Mezić. Spectral properties of dynamical systems, model reduction and decompositions. *Nonlinear Dynamics*, 41(1):309–325, Aug. 2005.
- [40] G. N. Mil’shtejn. Approximate integration of stochastic differential equations. *Theory of Probability & Its Applications*, 19(3):557–562, June 1975.
- [41] G. A. Pavliotis. *Stochastic Processes and Applications*. Springer New York, 2014.
- [42] R. Perez-Carrasco and J. Sancho. Stochastic algorithms for discontinuous multiplicative white noise. *Physical Review E*, 81(3):032104, 2010.
- [43] A. Rahimi and B. Recht. Unsupervised regression with applications to nonlinear system identification. In *Proceedings of the 19th International Conference on Neural Information Processing Systems, NIPS’06*, pages 1113–1120, Cambridge, MA, USA, 2006. MIT Press.
- [44] R. Rico-Martinez, J. Anderson, and I. Kevrekidis. Continuous-time nonlinear signal processing: a neural network based approach for gray box identification. In *Proceedings of IEEE Workshop on Neural Networks for Signal Processing*. IEEE, 1994.
- [45] R. Rico-Martínez and I. Kevrekidis. *Nonlinear system identification using neural networks: dynamics and instabilities*, chapter 16, pages 409–442. Elsevier Science, 1995.
- [46] R. Rico-Martínez, K. Krischer, I. Kevrekidis, M. Kube, and J. Hudson. Discrete- vs. continuous-time nonlinear signal processing of Cu electrodisolution data. *Chemical Engineering Communications*, 118(1):25–48, Nov. 1992.
- [47] B. Schölkopf and A. J. Smola. *Learning with Kernels: Support Vector Machines, Regularization, Optimization, and Beyond*. The MIT Press, 2018.
- [48] C. Siettos and L. Russo. Mathematical modeling of infectious disease dynamics. *Virulence*, 4:295 – 306, 2013.
- [49] Y. Song, J. Sohl-Dickstein, D. P. Kingma, A. Kumar, S. Ermon, and B. Poole. Score-based generative modeling through stochastic differential equations. In *International Conference on Learning Representations*, 2021.
- [50] L. Yang, C. Daskalakis, and G. E. Karniadakis. Generative ensemble-regression: Learning stochastic dynamics from discrete particle ensemble observations. *arXiv:2008.01915v1*, 2020.
- [51] S. Yang, S. W. K. Wong, and S. C. Kou. Inference of dynamic systems from noisy and sparse data via manifold-constrained gaussian processes. *Proceedings of the National Academy of Sciences*, 118(15):e2020397118, 2021.
- [52] C. Yildiz, M. Heinonen, J. Intosalmi, H. Mannerstrom, and H. Lahdesmaki. Learning stochastic differential equations with Gaussian Processes without gradient matching. In *2018 IEEE 28th International Workshop on Machine Learning for Signal Processing (MLSP)*. IEEE, Sept. 2018.
- [53] A. Zhu, P. Jin, and Y. Tang. Deep hamiltonian networks based on symplectic integrators. *arXiv:2004.13830*, 2020.
- [54] K. C. Zygalakis. On the existence and the applications of modified equations for stochastic differential equations. *SIAM Journal on Scientific Computing*, 33(1):102–130, Jan. 2011.

Table 2: Software we used to test that the code in the supplement can be compiled and run.

Software	Version
Microsoft Windows	10.0.19042
Python	3.8
Microsoft Visual Studio Community 2019	16.9.0
Microsoft Visual C++	2019

A Instructions for the software

The code to reproduce the results of the computational experiments is attached, compressed into a *.zip file. We will make the code available publicly once the paper is accepted. Tab. 2 lists the software we used to test if the code can be compiled and actually produces the experimental results.

All required python packages can be installed by running “`pip install -r requirements.txt`” in the folder where the `requirements.txt` file is placed. The toy examples and lattice computations are all available in separate files, the mapping is listed in Tab. 3. To compile the `kMCLattice.dll`, the Microsoft Visual Studio solution file in `kmc/kMCLattice_dll` can be used. Make sure to use an x64 setup (for both Python and the .dll file). The file needs to be placed in the same folder as the Python script `kmc/KMC_SIRS3.py`. The SSA results can be obtained by setting both `Dif1` and `Dif2` to 1000 in the file `kmc/KMC_SIRS3.py`.

Table 3: Mapping of computational experiments to Jupyter notebooks (.ipynb) / python (.py) files.

Experiment	File
Toy example: 1D, cubic drift, linear diff. (E.-M.)	toy - 1d sde-cubic.ipynb
Toy example: 1D, cubic drift, linear diff. (Milstein)	toy - 1d sde-cubic-with-milstein.ipynb
Toy example: 3D, linear drift, lower-triangular diff.	toy - 3d sde-nondiagonal.ipynb
Toy example: 3D, linear drift, SPD diff.	toy - 3d sde-spd.ipynb
Toy example: nD (1-20D), cubic drift, linear diff.	toy - nd sde-cubic.ipynb
Kinetic Monte-Carlo simulations	kmc/KMC_SIRS3.py

B Additional information regarding the computational experiments

The following subsections contain additional descriptions and results for the computational experiments. They are not necessary to follow the main results described in the paper.

B.1 Milstein loss

The TensorFlow implementation for the Milstein loss uses a slightly different Bessel function, with integer-valued parameter $\nu = 0$ instead of the required $\nu = -1/2$ (no other Bessel function was available). Figure 6 demonstrates the difference is quite small.

B.2 Toy example: cubic drift, linear diffusivity

The loss and validation loss curves for Euler-Maruyama and Milstein schemes are shown in figure 7.

B.3 Toy example: non-diagonal diffusivity

We now demonstrate that our approach can learn non-diagonal diffusivity matrices. We sample initial points $x_0 \in [-.3, .3]^3$, randomly sample a lower-triangular diffusivity matrix Σ with positive eigenvalues, which we use as the (constant) diffusivity $\sigma(x) = \Sigma$, and set the drift to be $f(x) = -x$. The absolute error between original and (the average over) identified matrix entries is smaller than 0.01, and the standard deviation of the diffusivity values over the entire data set is smaller than 0.003 for all elements of the matrix. Note that the network σ_θ is still a nonlinear function over the state space, not a constant. To ensure the matrix has positive eigenvalues, we pass the real-valued output




Figure 6: Probability density functions and sampled data from the two numerical integrator templates. The top panel shows the PDFs with the correct Bessel function with real-valued parameter $\nu = -1/2$, using `scipy.stats.ncx2` for the PDF of the non-centered χ^2 distribution. The bottom panel shows the PDF with an integer valued parameter $\nu = 0$ (the only one available in TensorFlow).




Figure 7: Milstein vs. Euler-Maruyama loss curves.

of the neurons that encode the diagonal of the matrix through a soft-plus activation function. Figure 8 shows that we learn a good approximation for Σ .

B.4 Toy example: increasing dimension from 1D to 20D

We define $f(x) = -(4x^3 - 8x + 3)/2$ and $\sigma(x) = 5e - 2x + 0.5$, where all operations are meant coordinate-wise (e.g. x^3 computes the third power of each individual coordinate of the vector). Figure 9 (left) illustrates how the training and validation losses change

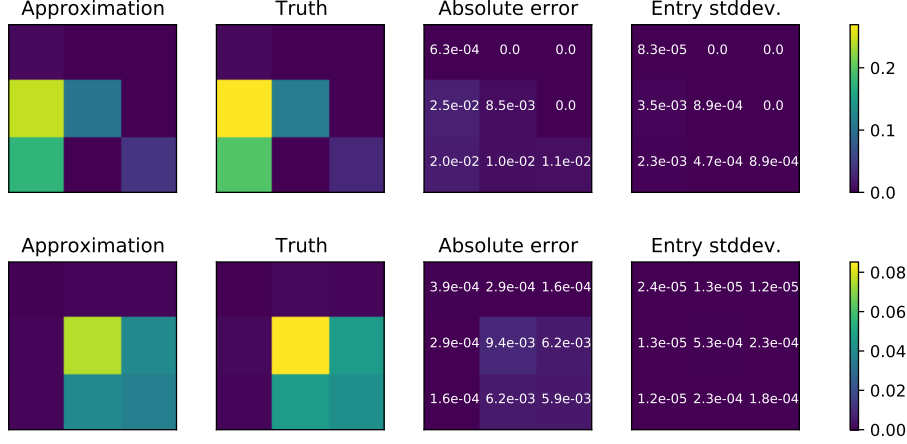


Figure 8: Lower-triangular (top row) and full symmetric, positive definite (bottom row) diffusivity matrix approximation. The panels show the averaged network output over the uniformly sampled data inside the training region, the true diffusivity matrices, the absolute error between the results shown in the first and second columns, and the deviation of the network matrix output over the training region. The color range is the same for all four plots.

when increasing n from 1 to 20. The loss ($\mathcal{L}_{\text{Euler-Maruyama}}$) is adapted to the changing dimensionality by adding $\log(2\pi)n$, as described by Eq. (14). This normalization is necessary because (apparently) TensorFlow 2.4.1 does not add this constant when computing the log probability (See https://www.tensorflow.org/probability/api_docs/python/tfp/distributions/MultivariateNormalDiag#log_prob). The increase in loss can be explained by the constant number of points ($N = 10,000$) we used: Increasing the intrinsic dimension of the problem by sampling the input data in the n -dimensional cube $[-2, 2]^n$ causes the data sampling to get sparser and sparser. By increasing the number of training data points linearly with the dimension (while keeping the number of training iterations per dimension constant), we can see that the training loss is relatively small even for $n = 12$ (Fig. 9, right).




Figure 9: Normalized training and validation loss when increasing the intrinsic dimension of the problem from $n = 1$ to $n = 20$. The left panel shows results when keeping the number of samples constant at $N = 10,000$, the right panel shows what happens when increasing $N = 10,000 \times n$, but training for the same number of iterations per dimension (Epochs = $1000/n$).

B.5 Kinetic Monte-Carlo lattice simulations

Results on a smaller 20×20 lattice are shown in Fig. 10.

Fig. 11 shows that the algorithm works at different system sizes, even when N is large—that is, even when the noise level is low. For $N \rightarrow \infty$, the system becomes deterministic and can be described by an ODE.

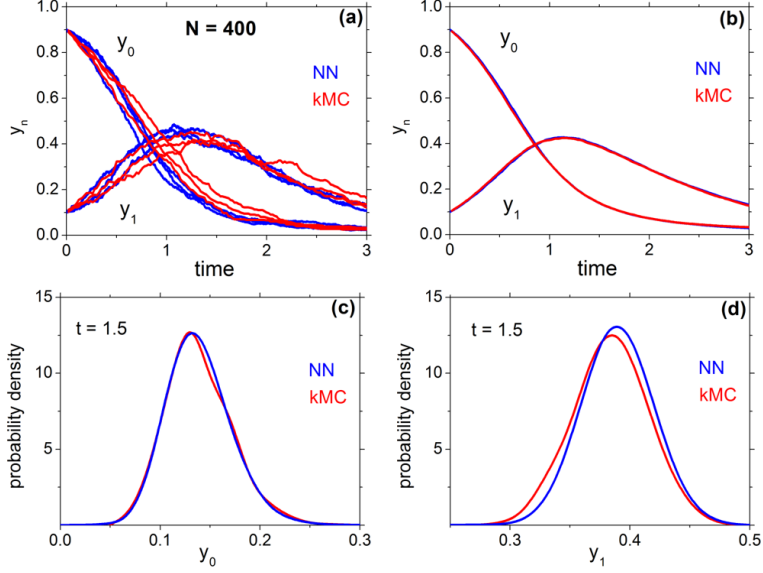


Figure 10: Blue (red) lines correspond to NN (kMC) results. (a) Three sample paths of y_0 and y_1 ; (b) Average of 200 paths; (c),(d) Probability density function for y_0 and y_1 at $t = 1.5$ computed from 2×10^4 sample paths. Parameters: 20×20 lattice, $d = 50$; $N_{tr} = 10000$, $t_{max} = 0.05$, $h = 0.01$.




Figure 11: Influence of the system size on the neural network-estimated probability density function at $t = 2$: 20×20 (black), 32×32 (navy), 100×100 (blue).

The data presented in Fig. 11 correspond to a relatively fast migration rate $d = d_1 = d_2 = 50$. In this case, a well-mixed state is achieved and the kMC results are close to that obtained with the SSA algorithm. However, the network can also be trained to match the kMC results at lower migration rate values. Fig. 12 shows the simulation results for $d = 5$. Since the initial distribution on a lattice is random, we omit the first few time steps and select the data points for training only when a “matured” state is achieved. For the results presented in Fig. 12, the training process was started at $t_0 = 0.1$. At $d = 5$, the accuracy of the network is acceptable, but at lower d values, the coincidence of the network and kMC results decreased considerably. This suggests that, around that value of d the dynamics may not effectively close in terms of the mean field variables, and that higher dimensional approximations (possibly including pair probabilities as additional independent variables) may become necessary.

Similar results can be obtained on a smaller 20×20 lattice. In this case, as expected, the influence of the stochasticity becomes more pronounced. In a separate experiment, we also confirmed the algorithm works at different system sizes, even when N is large—that is, even when the noise level is low. For $N \rightarrow \infty$, the system becomes deterministic and can be described by an ODE. At low migration rate values ($d < 5$), the approximation quality of the network decreased considerably as expected: The system can hardly be described in terms of two averaged concentrations because the species distribution on a lattice becomes inhomogeneous: Clustering occurs, and a higher dimensional SDE, possibly in terms of densities *as well as* pair probabilities, becomes necessary.

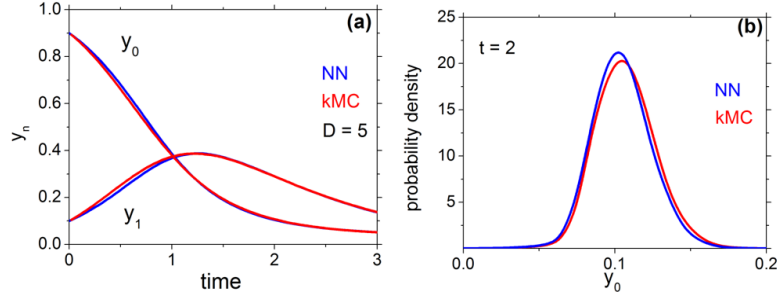


Figure 12: Relatively low migration rate. Parameters: 32×32 lattice, $d = 5$; $N_{tr} = 4000$, $t_{\max} = 0.4$, $h = 0.02$.

When $k_3 > 0$, the system exhibits richer dynamic behavior. An example of simulations with $k_3 = 0.2$ is shown in Fig. 13. The network model is capable of reproducing the correct steady state and damped oscillations (the mean-field SIRS model has a focus-type steady state at these values of parameters).




Figure 13: SIRS model with $k_3 = 0.2$, 60×60 lattice, $d = 50$; $N_{tr} = 5000$, $t_{\max} = 0.1$, $h = 0.01$. The left panel shows 200 averaged paths, the right panel shows the propagated densities until $t = 2$.

B.6 Learning latent spaces for the SIR model

We identify latent spaces of six-dimensional data obtained from the kMC lattice simulations. Fig. 14 illustrates that the latent spaces are both two-dimensional and triangular. This shape is due to the conservation law of the underlying model (when written in terms of fractions of the total population that are susceptible (S), infectious (I), and removed (R), the law would be $S + I + R = 1$). Sec. B.6.1 outlines the Diffusion Maps algorithm, Sec. B.6.2 describes how the auto-encoder was set up and trained.




Figure 14: Latent spaces identified by SDE-integrator-informed auto-encoder (a), and with Diffusion Maps (b). Both show a triangular structure, as expected from a system with three intrinsic variables with a single conservation law ($y_0 + y_1 + y_2 = \text{const}$).

B.6.1 Diffusion Maps

The Diffusion Maps algorithm by Coifman and Lafon [12] follows three main steps:

1. Construct the kernel matrix $K_{i,j} = \exp(-d(x_i, x_j)^2/\epsilon)$.
2. Normalize the matrix to mitigate effects of data density.
3. Compute eigenvectors to approximate eigenfunctions of Δ evaluated on the data.

These steps are expanded upon in algorithm 1. We use the implementation from [32].

Algorithm 1: Diffusion maps algorithm from [12], adapted from [4], implemented in [32].

Input: $n \in \mathbb{N}$: Dimension of points in the input. $\ell \in \mathbb{N}$: Number of eigenvectors to compute.

$X \in \mathbb{R}^{N \times n}$: input data set, with N points and n dimensions per point.

1. Compute kernel matrix $K \in \mathbb{R}^{N \times N}$ with $K_{ij} = \exp(-\|X_i - X_j\|^2/\epsilon)$ and kernel bandwidth ϵ through 5% of the median of squared distances between all points. Note that sparse computations are also possible in the software [32].
2. Normalize kernel matrix to make it invariant to sampling density:
 - (a) Form the diagonal normalization matrix $P_{ii} = \sum_{j=1}^N K_{ij}$.
 - (b) Form the matrix $W = P^{-1} K P^{-1}$.
 - (c) Form the diagonal normalization matrix $Q_{ii} = \sum_{j=1}^N W_{ij}$.
 - (d) Form the matrix $\hat{T} = Q^{-1/2} W Q^{-1/2}$.
3. Solve the eigensystem and construct eigenvectors:
 - (a) Find the ℓ largest eigenvalues a_p associated to non-harmonic eigenvectors v_p of \hat{T} . See [15] for a discussion and method to remove harmonic eigenvectors.
 - (b) Compute the eigenvalues of $\hat{T}^{1/\epsilon}$ by $\lambda_p = a_p^{1/(2\epsilon)}$.
 - (c) Compute the eigenvectors ϕ_p of the matrix $T = Q^{-1} W$ through $\phi_p = Q^{-1/2} v_p$.

Output: ℓ eigenvectors ϕ_p , each with N entries, and with corresponding eigenvalues λ_p .

B.6.2 SDE-informed auto-encoder

The SDE-informed auto-encoder has a single encoder network (from observations to latent space) and two additional networks: one decoder from latent space back to the observations, and one network that defines drift and diffusivity functions *on the latent space*. Note that the drift and diffusivity we find can still be considered functions of the original observations, because $f(x) \approx f_\theta(\text{encoder}(x))$ and $\sigma(x) \approx \sigma_\theta(\text{encoder}(x))$. In the computational example that lead to Fig. 14, we used an encoder architecture of 6|25|25|25|2, a decoder architecture 2|25|25|25|6, and an SDE approximation network with diagonal diffusivity matrix and architecture 2|25|25|25|6, all with `tf.nn.relu` activations. The loss function for training the auto-encoder is the following, where $\mathcal{L}_{\text{Euler-Maruyama}}$ is defined in the paper (Eq. (5)):

$$\mathcal{L}(\theta|D) = \log \left(\frac{1}{B} \sum_{k=1}^B \left(\text{decoder} \left(\text{encoder} \left(x_0^{(k)} \right) \right) - x_0^{(k)} \right)^2 \right) + \mathcal{L}_{\text{Euler-Maruyama}}(\theta|D), \quad (13)$$

where B is the size of the current mini-batch (here: $B = 32$). The Python class implementing this auto-encoder loss is called `AEModel`, and can be found in `code/notebooks/sde/sde_learning_network.py`

C Additional proofs

The logarithm of the probability density p of a multivariate normal distribution in n dimensions is

$$\log(p(x)) = -\frac{n}{2} \log(2\pi) - \frac{1}{2} \log |\det \Sigma| - \frac{1}{2} (x - \mu)^T \Sigma^{-1} (x - \mu). \quad (14)$$

Proof. This follows directly from

$$p(x) = \frac{1}{\sqrt{(2\pi)^n |\det \Sigma|}} \exp \left[-\frac{1}{2} (x - \mu)^T \Sigma^{-1} (x - \mu) \right].$$

□