

RESEARCH ARTICLE | JUNE 24 2025

## Improving the noise estimation of latent neural stochastic differential equations

L. Heck ; M. Gelbrecht ; M. T. Schaub ; N. Boers 



*Chaos* 35, 063139 (2025)

<https://doi.org/10.1063/5.0257224>



### Articles You May Be Interested In

Fast and reliable prediction of scramjet flowfields via Gaussian process latent variable model and deep learning

*Physics of Fluids* (April 2023)

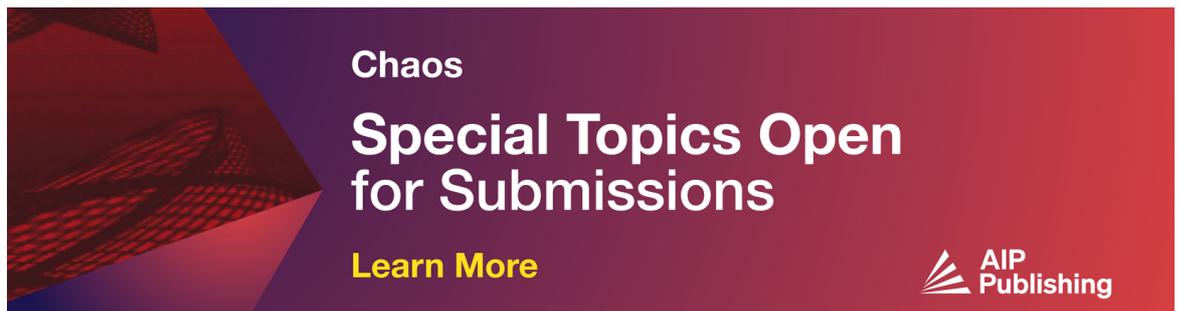
New measures of multimodality for the detection of a ghost stochastic resonance

*Chaos* (December 2009)

Data augmentation-based statistical inference of diffusion processes

*Chaos* (March 2023)

08 October 2025 11:46:11



**Chaos**  
**Special Topics Open  
for Submissions**  
[Learn More](#)



# Improving the noise estimation of latent neural stochastic differential equations

Cite as: Chaos 35, 063139 (2025); doi: 10.1063/5.0257224

Submitted: 9 January 2025 · Accepted: 8 June 2025 ·

Published Online: 24 June 2025



View Online



Export Citation



CrossMark

L. Heck,<sup>1,a)</sup> M. Gelbrecht,<sup>2,b)</sup> M. T. Schaub,<sup>3</sup> and N. Boers<sup>2,c)</sup>

## AFFILIATIONS

<sup>1</sup>Institute for Computing and Information Sciences, Radboud University, Nijmegen, The Netherlands

<sup>2</sup>Earth System Modelling, School of Engineering and Design, TU Munich, Munich, Germany

<sup>3</sup>Computational Network Science, Department of Computer Science, RWTH Aachen University, Aachen, Germany

<sup>a)</sup>Also at: Department of Computer Science, RWTH Aachen University, Germany and Complexity Science, Potsdam Institute for Climate Impact Research, Potsdam, Germany. **Author to whom correspondence should be addressed:** [linus.heck@ru.nl](mailto:linus.heck@ru.nl).

<sup>b)</sup>Also at: Complexity Science, Potsdam Institute for Climate Impact Research, Potsdam, Germany.

<sup>c)</sup>Also at: Complexity Science, Potsdam Institute for Climate Impact Research, Potsdam, Germany and Department of Mathematics and Statistics, University of Exeter, Exeter, United Kingdom.

## ABSTRACT

Latent neural stochastic differential equations (SDEs) have recently emerged as a promising approach for learning generative models from stochastic time series data. However, they systematically underestimate the noise level inherent in such data, limiting their ability to capture stochastic dynamics accurately. We investigate this underestimation in detail and propose a straightforward solution; by including an explicit additional noise regularization in the loss function, we are able to learn a model that accurately captures the diffusion component of the data. We demonstrate our results on a conceptual model system that highlights the improved latent neural SDE's capability to model stochastic bistable dynamics.

© 2025 Author(s). All article content, except where otherwise noted, is licensed under a Creative Commons Attribution-NonCommercial-NoDerivs 4.0 International (CC BY-NC-ND) license (<https://creativecommons.org/licenses/by-nc-nd/4.0/>). <https://doi.org/10.1063/5.0257224>

The goal of fitting models to stochastic time series data is to match the data's probability distribution. While latent neural stochastic differential equations (SDEs) are a promising tool for this task, it has been observed that these models underestimate the noise level in the data. On complex stochastic systems such as multistable systems that are forced by noise, this underestimation can lead to a complete mismatch of behavior between the model and the data. In this work, we analyze this underestimation and propose augmenting the training loss with an additional noise penalty. This penalty, adjustable via a hyperparameter, enables the model to align with user-defined target statistics. We demonstrate the effectiveness of our approach on a conceptual model system that exhibits bistable dynamics and is forced by noise. Our results show that the proposed noise penalty allows the latent neural SDE to accurately match the diffusion component of the data, improving the model's ability to fit, data distribution. The code for our experiments is publicly available.

## I. INTRODUCTION

Neural stochastic differential equations combine stochastic differential equations (SDEs) with a deep learning component. They have recently emerged as a powerful framework to model continuous data in an unsupervised, generative fashion. Compared to standard (stochastic) recurrent neural networks, the underlying process is not a discrete-time process but a continuous SDE. The model can, thus, be arbitrarily discretized and is amenable to direct analysis, thereby offering increased interpretability. Potential applications of neural SDEs cover a wide range of fields, including all instances in which SDEs are already used to investigate dynamical data, such as paleoclimatology and present-day climate,<sup>1,2</sup> solar and wind energy forecasts,<sup>3,4</sup> astrophysics,<sup>5</sup> or finance.<sup>6</sup>

In general terms, such a neural SDE is defined by

$$du(t) = h(u(t), t)dt + g(u(t), t)dB_t, \quad (1)$$

where  $u(t)$  denotes the hidden state (vector) of the system,  $B_t$  is the standard Brownian motion, and  $h(\cdot)$  and  $g(\cdot)$  denote the so-called

*drift* and the *diffusion* term, respectively, which are parametrized via deep neural networks. Two main methods have been proposed to fit the neural SDE to data: latent neural SDEs and SDE-Generative Adversarial Networks (GANs). First, *latent neural SDEs*<sup>7</sup> may be derived by minimizing the Kullback–Leibler divergence between the neural SDE and the empirically observed data  $u_{\text{obs}}$ ; these can also be considered a form of variational autoencoder.<sup>8</sup> Second, *SDE-GANs*<sup>9</sup> emerge from minimizing a Wasserstein type loss between  $u$  and  $u_{\text{obs}}$ , which leads to a generative adversarial neural network architecture. In this paper, we focus on latent neural SDE architectures, which are easier to train in practice, whereas SDE-GANs suffer both from instability generally present in GANs and from a lack of known in-practice tricks to alleviate this instability (Ref. 8, p. 84).

Our main contributions are as follows. First, we provide a detailed investigation of the previously observed effect (Refs. 7 and 8, p. 84) that latent neural SDE schemes are unable to accurately match the diffusion component of the data, resulting in failure to match the distribution of the data. To this end, we provide a detailed experimental evaluation of latent neural SDEs regarding this mismatch of the diffusion term and investigate the reasons for this effect. Even though there have been experiments on latent neural SDEs in the literature [e.g., Refs. 7, 8 (pp. 89–92), 10, and 11], a detailed empirical investigation of this phenomenon is still missing. These experiments aim at some form of *uncertainty quantification*, where noise underestimation is more acceptable, while we aim at a model that can accurately reproduce the data distribution.

Using the newly gained understanding of latent neural SDEs, multiple solutions can be proposed to accurately match the diffusion component. As our second contribution, we provide a possible remedy in terms of a new noise penalty that we add to the loss function to be optimized. We show that this enables the latent neural SDE to fit transition rates in multistable models and match the data's Kramers–Moyal (KM) coefficients, assuming constant diffusion in the data. While we choose these metrics, any metric can be used to tune the noise penalty's hyperparameter.

Our experiments and all underlying data of the presented results are publicly available.<sup>12</sup>

In the following, we first discuss related work in Sec. II and review latent neural SDEs in Sec. III. We then introduce the chosen experimental setups in Sec. IV, before investigating the noise underestimation in Sec. V. In Sec. VI, we present some approaches to fix the observed noise underestimation problem, before concluding with a short discussion of our results and possible future work.

## II. RELATED WORK

Neural SDEs can be used for uncertainty quantification within probabilistic machine learning, similar to Gaussian processes,<sup>13</sup> which also induce probability distributions on time series. In contrast to Gaussian processes, which are usually trained explicitly using a specified parametric kernel, we train neural SDEs using an unsupervised scheme. The reason is that, in general, there is no way to write down the probability distribution that is induced by a neural SDE. The advantage of using SDEs instead of Gaussian processes is that an SDE can express a much larger set of probability distributions than a Gaussian process can.

Direct applications of latent neural SDEs include the following works from different disciplines. Fagin *et al.*<sup>5</sup> apply latent neural SDEs to data from astrophysics. They use them as a drop-in replacement for Gaussian processes, so the noise is interpreted as uncertainty. Xu *et al.*<sup>14</sup> consider latent neural SDEs with a fixed diffusion, a continuous analog of Bayesian recurrent neural networks. Kidger *et al.*<sup>10</sup> introduce a reversible Heun solver for differential equations and train a latent neural SDE on an air quality data set from Beijing. Hasan *et al.*<sup>15</sup> propose a different latent neural SDE framework to recover a mapping from one high-dimensional probabilistic space to another. Furthermore, the framework has also been extended to include discontinuous (“jump”) processes by Jia and Benson,<sup>16</sup> which has been applied by Herrera *et al.*<sup>17</sup> to predict mortality rate numbers in intensive care units. Djeumou *et al.*<sup>18</sup> learn an SDE diffusion term that is penalized on the distance, and in contrast to our work, they use a custom training scheme tailored to their problem instead of one based on GANs or VAEs.

Aside from neural SDEs, there are many other approaches to data-driven modeling of stochastic time series. These include other generative models, such as denoising diffusion models<sup>19</sup> and GANs.<sup>20</sup> Reservoir computing, a deterministic, recurrent neural network, has been adapted as well to stochastic time series. However, these approaches usually have to include additional steps, such as filtering<sup>21</sup> or past noise forecasting.<sup>15</sup>

## III. LATENT NEURAL SDEs

In this section, we review latent neural SDEs. Latent neural SDEs, as introduced in Li *et al.*,<sup>7</sup> are generative models that aim to reproduce the probability distribution of a data set consisting of time series. A latent neural SDE may be seen as consisting of two neural SDEs: a so-called posterior and a prior SDE. The names “prior” and “posterior” stem from the variational inference framework. In this context, the posterior is approximate. Each of these SDEs consists of a drift and a diffusion term as in Eq. (1). We first provide an intuitive overview and then explain the components that form the latent neural SDE.

### A. Intuitive overview

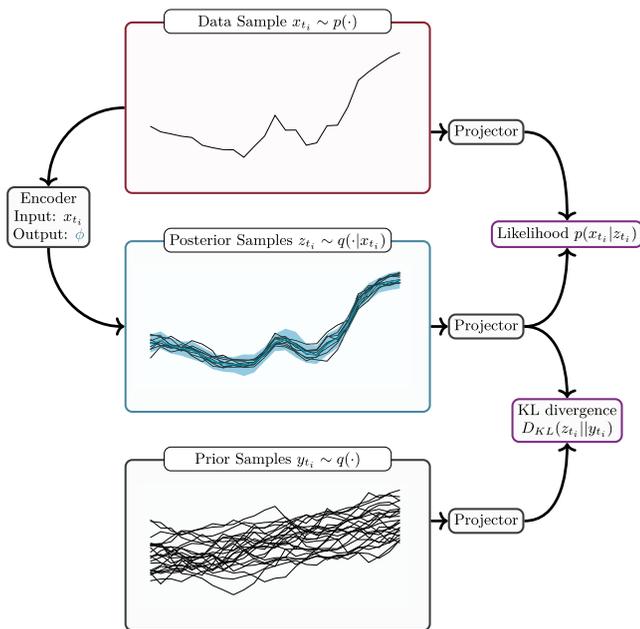
Latent neural SDEs aim to match the distribution of the *prior SDE* to the data distribution. Because the distribution of the prior SDE and the data set are not easily comparable, this comparison is mediated via a second SDE, which is the *posterior SDE* (cf. Fig. 1).

Latent neural SDEs aim to align the prior SDE's distribution to the posterior SDE's distribution and, at the same time, the posterior SDE's distribution to the data distribution. The prior SDE's distribution is aligned with the posterior SDE's distribution by minimizing the *Kullback–Leibler (KL) divergence* between the two SDEs. The posterior SDE's distribution is aligned with the data distribution by feeding a *context*, which is derived from samples of the data, into the posterior SDE. The posterior SDE is then trained to reconstruct these samples.

### B. Prior SDE and posterior SDE

The posterior SDE is a neural SDE of the form

$$du(t) = h_{\theta}(u(t), t, \phi)dt + g_{\theta}(u(t), t)dB_t, \quad (\text{posterior})$$



**FIG. 1.** Sketch of the latent neural SDE approach. The data sample is encoded and given as the context  $\phi$  to the posterior SDE, which yields a probability distribution. The likelihood of the data given realizations of the distribution is maximized during training. The distance between the posterior SDE and the prior SDE is quantified by the KL divergence, which is minimized during training.

where  $\theta$  are the neural SDE’s parameters and  $\phi$  is the context, i.e., time-dependent information that is computed by another neural network, using an input trajectory from the data set. Given  $\phi$ , the posterior SDE induces a probability distribution that aims to reconstruct the data distribution.

This reconstruction of the data by the posterior SDE is scored by an *observation model*. The observation model computes the likelihood that the data trajectories occur around the posterior SDE’s trajectories, which is maximized during training. In latent neural SDEs, we manually define a Gaussian prior with a fixed variance around the posterior SDE’s points and quantify the likelihood over this prior. Following Li *et al.*,<sup>7</sup> we write in integral form

$$\mathcal{L}_E = \int_0^{t_{\text{train}}} \log p(x_t | z_t) dt,$$

where  $x_t$  denote data samples and  $z_t$  sampled from integrating the posterior SDE. We illustrate the reconstruction that the posterior performs in Fig. 1. In training, we draw samples from the posterior SDE’s distribution and minimize their distance to the input data using the likelihood.

The context  $\phi$  can be computed by a recurrent neural network (RNN), which we call the *encoder*. The RNN operates on the reversed input trajectory, so its output at time  $t$  contains information from the input trajectory  $t$  until  $t_{\text{train}}$ . This output at time  $t$  is an input to the posterior SDE’s drift at time  $t$ . This way, the posterior

SDE has access to relevant information about the input trajectory’s future.

The posterior alone cannot model the data’s probability distribution: it always needs data as input (through the context) and thus does not provide a generative model on its own. Additionally, if we were to train the posterior alone, its diffusion would eventually be zero: The observation model would attempt to find a drift term that is close to the data by repeating the next value at each time step, observed through the context.

The prior SDE aims to model the data distribution without further input. Its general form is

$$d\tilde{u}(t) = f_\theta(\tilde{u}(t), t)dt + g_\theta(\tilde{u}(t), t)dB_t, \quad (\text{prior})$$

which lacks the context  $\phi$ . Note that while its drift  $f$  is different from the posterior drift  $h$ , the prior and posterior share the same diffusion  $g$ . Because of this, we can quantify the KL divergence between both SDEs.<sup>7,22</sup> Importantly, this is the KL divergence of the distributions over the path space, not the KL divergence over the marginals.

### C. Encoder and projector

The encoder transforms a given path from data into the context  $\phi$ . Theoretically, the encoder is not necessary, as we could directly set  $\phi = x_{t_i}$ . However, choosing the encoder as a recurrent neural network that has seen future observations in the data has proven helpful.<sup>7,8</sup> The intuitive advantage of an encoder is that the posterior SDE can look as far into the future as needed.

The latent space (containing  $y_{t_i}$  and  $z_{t_i}$ ) and data space (containing  $x_{t_i}$ ) might differ. For instance, the latent space could have more dimensions than the data space. In this case, we need a *projector* in the architecture, which would usually be a simple feed-forward neural network transforming a point in the latent space to a point in the data space.

### D. Training objective

With prior and posterior as above, the KL divergence between prior and posterior is

$$\begin{aligned} \mathcal{L}_{KL} &= D_{KL}(\mu_q || \mu_p) \\ &= \mathbb{E}_{B_t} \left[ \int_0^{t_{\text{train}}} \frac{1}{2} \left\| \frac{h_\theta(u(t), t, \phi) - f_\theta(u(t), t)}{g_\theta(u(t), t)} \right\|_2^2 dt \right]. \end{aligned}$$

Here,  $\mu_q$  is the probability distribution of the posterior SDE, and  $\mu_p$  of the prior SDE, and we integrate along a trajectory  $u(t)$  of the posterior SDE along Brownian motion  $B_t$ .<sup>7</sup>

We jointly minimize the distance between the data distribution and the posterior SDE’s distribution, as well as the distance between the prior SDE’s and the posterior SDE’s distribution, yielding the objective

$$\max_{\theta, \phi} \mathbb{E}_{x_{\text{data}}} (\mathcal{L}_E - \mathcal{L}_{KL}).$$

This objective can be interpreted as a so-called evidence lower bound (ELBO) from variational inference, making a latent neural SDE a Bayesian variational autoencoder, as first proposed by Kingma and Welling.<sup>23</sup>

To have more control over the model, it is desirable to weigh the KL divergence using a tunable factor  $\beta$ ,<sup>7,24</sup>

$$\max_{\theta, \phi} \mathbb{E}_{x_{\text{data}}} (\mathcal{L}_E - \beta \mathcal{L}_{KL}).$$

We explore the selection of this hyperparameter in Sec. IV.

#### IV. EXPERIMENTAL SETUP

In the following, we will consider two experimental systems. First, we set up a simple bistable stochastic system inspired by energy balance models (EBMs) in climate dynamics. The drift of this model exhibits two stable fixed points and the diffusion enables transitions between these two fixed points and their basins of attraction. Its governing equation is given by

$$dT(t) = (a_1 + a_2 \tanh(T(t) - T_0) - a_3 T(t)^4) dt + \sigma dB_t,$$

where  $a_1 = 235.175$ ,  $a_2 = 81.8$ ,  $a_3 = 3.402$ ,  $T_0 = 273$ ,  $\sigma = 40$ . These constants are similar to those found in the literature.<sup>25–27</sup> The EBM enables us to analyze the noise estimation of latent neural SDEs with a simple intuitive observable: the transition rate, which we define as the frequency at which trajectories cross the unstable fixed point between the two stable fixed points of the drift of the EBM. In a one-dimensional SDE, back-and-forth transitions over the unstable fixed point can only be caused by diffusion. Thus, assuming the drift is estimated correctly, the transition rate is a measure of whether the diffusion size is estimated correctly.

In the EBM, we have slightly increased values  $\alpha_1$  and  $\alpha_2$  from a realistic model to yield closer fixed points for more frequent transition behavior. Using the EBM, we intend to test whether latent neural SDEs can (a) model multistable systems and (b) accurately fit the diffusion size and, therefore, the rate of transition of such systems. We generally follow Li *et al.*<sup>7</sup> for our explicit latent neural SDE setup, except for the choices below. The differences are summarized in Table I.

##### A. Latent neural SDE architecture

We use gradients computed by directly autodifferentiating through the SDE solver, as opposed to solving an adjoint SDE, and choose an Euler–Maruyama SDE solver.

Li *et al.*<sup>7</sup> add a sigmoid activation after the final layer of the neural network of the diffusion term, and thus the value of the diffusion term is limited to the range  $[0, 1]$ . In our experiments, the diffusion did often train to the value 1 in this setup. To remove the limit on the diffusion value while keeping the sigmoid's limiting effect, we

TABLE I. Differences between Li *et al.*'s setup and ours.

| Aspect               | Li <i>et al.</i> | Our setup            |
|----------------------|------------------|----------------------|
| Diffusion activation | Sigmoid          | Sigmoid and softplus |
| Latent space dims.   | More than data   | Matches data         |
| Projector            | Trainable        | None                 |
| Training             | Adjoint SDE      | Autodifferentiation  |

add another layer with a final softplus activation to keep the output positive. We also attempted to leave out the sigmoid activation altogether, but this led to instabilities while training.

As in Li *et al.*,<sup>7</sup> the prior and posterior SDEs are time-independent, so they do not receive the current time  $t$  as input, only the current position in state space  $u(t)$ . The encoder is also time-independent, so it does not receive time information alongside the data.

We use a one-dimensional latent space for data generated from a model with a one-dimensional state space. In contrast, Li *et al.*<sup>7</sup> use a four-dimensional latent space for data generated from a model with a two-dimensional state space. They use a trainable projector, whereas we use no projector to make the model more transparent, which in turn enables our analysis of noise underestimation of the unmodified latent neural SDE.

##### B. Training and test set

The model trains on the timespan  $[0, t_{\text{train}})$ , which can be different for each model. We set  $t_{\text{train}} = 4$  time units for the EBM. The integration step size for the EBM is  $\Delta t = 0.01$ . To evaluate the statistics of the trained model, we use  $[0.5t_{\text{train}}, t_{\text{train}})$  as the training set and  $[t_{\text{train}}, 5t_{\text{train}})$  as the test set because we are less interested in the evolution from the Gaussian initial state and more interested in the long-run behavior of the system.

##### C. Data preprocessing and fixed hyperparameters

We normalize the data by subtracting the mean and dividing by the standard deviation. We do not corrupt or perturb the data after generation. We use an observation model based on a log-likelihood and we choose a variance of 0.01 following Li *et al.*<sup>7</sup> Changing the variance will impact the optimal values for  $\beta$ . We always use the following hyperparameters: we train for 10 000 epochs and use a linear annealing schedule for the KL divergence<sup>28</sup> that ramps up beta from zero to its final value in 1000 epochs; training without such a schedule is less stable in our experience. We use ADAM<sup>29</sup> as the optimizer with a learning rate that starts at 0.01 and decays by a factor of 0.997 each iteration. Our batch size is 1024, and we use two hidden layers with 100 neurons in all networks.

##### D. Analysis of results

We compute the marginals as histograms for the prior SDEs and the data. To quantify how close these are numerically, we use the *Wasserstein distance*, which intuitively represents the cost of transforming one distribution into another<sup>30</sup> and the *transition rate*, which is the rate at which trajectories leave the area of influence of one attractor and tip into the other. Additionally, we directly compare the drift and diffusion of some trained latent neural SDEs with the EBM. Not only do we directly compare the drift and diffusion terms of both the ground truth system and the trained prior SDE, but we also compute and compare the first two Kramers–Moyal coefficients using the implementation of Gorjão and Meirinhos.<sup>31</sup> The first two Kramers–Moyal coefficients are quantification of an SDE's drift and diffusion terms given data output by the system. We compute these factors even though we know the real drift and diffusion; our purpose is that different SDEs (i.e., the original EBM and the prior

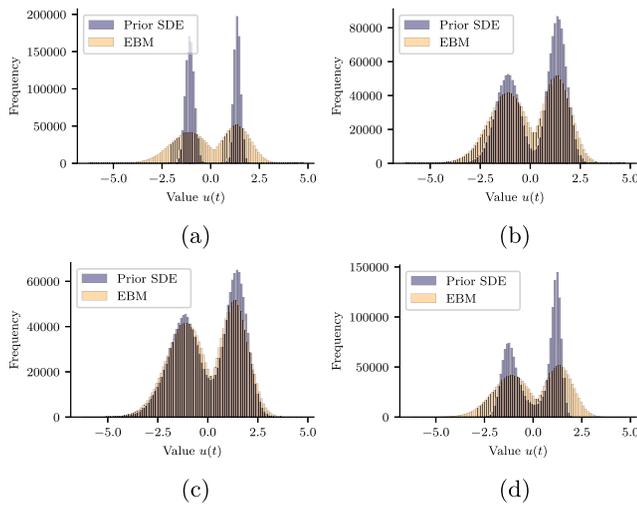


FIG. 2. Marginals for different values of hyperparameter  $\beta$  on test set. (a)  $\beta = 0.01$ , (b)  $\beta = 1$ , (c)  $\beta = 100$ , (d)  $\beta = 10000$ .

SDE) may produce similar behavior represented in different ways, resulting in more similar KM coefficients than their actual drift and diffusion terms.

E. Selection of KL weighing factor

The main hyperparameter to vary in a latent neural SDE is the KL weighting factor  $\beta$ . We sweep  $\beta$  from 0.01 to 10000 in increments of a factor 10.

We train on the EBM data and plot the marginals in Fig. 2. There seems to be a sweet spot for  $\beta$  between  $10^1$  and  $10^3$ , where the Wasserstein distances of the marginals between data and prior SDE are much lower. For  $\beta$  smaller or larger, the histograms of the latent neural SDEs’s prior become more concentrated on the attractors and less broadly around them.

The Wasserstein distances of the train and test sets on the EBM data set, depending on the hyperparameter  $\beta$ , are shown in Fig. 3(a). The Wasserstein distances of the test set are close to those of the training set because we chose a one-dimensional latent neural SDE with the knowledge that the underlying model is one-dimensional.

Next, we investigate the transition rates in Fig. 3(b). We can see that the transition rates of all latent neural SDEs are too low, less than half of the EBM’s transition rates. With values of  $\beta$  that are too high or too low, the transition rates decrease alongside the size of the diffusion. Underestimation of the diffusion size is a known problem of latent neural SDEs, mentioned by Li *et al.*<sup>7</sup> and subsequently discussed in GitHub issues of packages that implement them. However, there is, as of now, no answer to why the model behaves this way. In Sec. V, we investigate this effect.

V. INVESTIGATION OF NOISE UNDERESTIMATION

We have observed that latent neural SDEs underestimate the size of the diffusion. To understand the role of diffusion in the training procedure of latent neural SDEs, recall the loss function of a

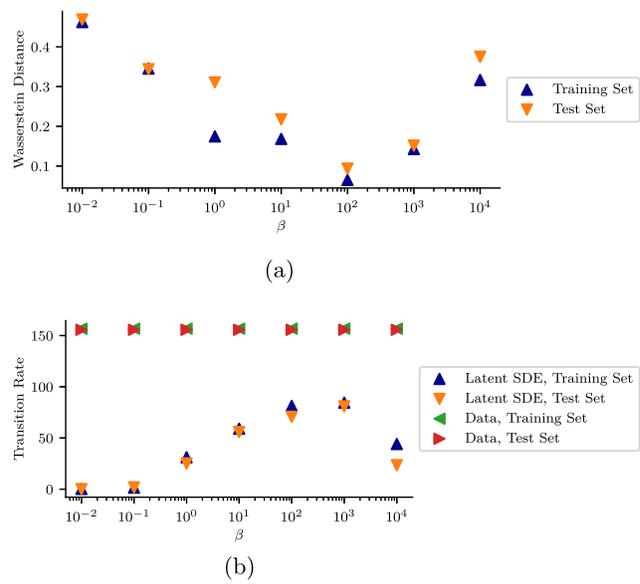


FIG. 3. Comparison of Wasserstein distance and transition rate between trained prior SDE and EBM for different values of hyperparameter  $\beta$  on the training set and test set. (a) Wasserstein distance of marginals, (b) transition rate.

latent neural SDE,

$$\max_{\theta, \phi} \mathbb{E}_{x_{\text{data}}} (\mathcal{L}_E - \beta \mathcal{L}_{KL}), \text{ where}$$

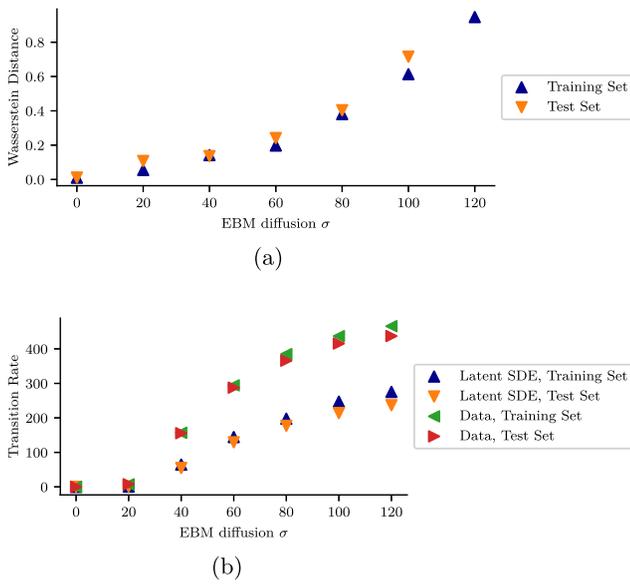
$$\mathcal{L}_{KL} = \mathbb{E}_{B_t} \left[ \int_0^{t_{\text{train}}} \frac{1}{2} \left\| \frac{h_{\theta}(u(t), t, \phi) - f_{\theta}(u(t), t)}{g_{\theta}(u(t), t)} \right\|_2^2 dt \right].$$

The size of the diffusion results from a balance of the two scores: the negative KL divergence  $-\mathcal{L}_{KL}$  correlates to the size of the diffusion, as long as the distance between the prior and posterior SDE’s distributions is not zero, and the likelihoods  $\mathcal{L}_E$  correlate inversely with the size of the diffusion because the highest likelihoods can be achieved with zero diffusion.

The weighing factor  $\beta$  is the key to tuning the balance between the two scores, and thus, we reconsider the transition rate of the trained prior SDE for different values of  $\beta$  in Fig. 3(b) in light of this balance.

For very small values of  $\beta$ , the magnitude of the trained diffusion is close to zero, as only the posterior loss  $\mathcal{L}_E$  is minimized. This is reflected in the transition rate. With zero diffusion, the posterior SDE can trace the data sample with the highest accuracy. In this extreme case, the diffusion size of the latent neural SDE is independent of the amount of noise in the data.

The transition rate significantly decreases with high values of  $\beta$ . This is surprising when considering the formulation of the training loss: as the KL divergence is proportional to the inverse of the diffusion and  $\beta$  increases its importance, we expected the size of the diffusion term to continue to increase with increasing  $\beta$ . A possible explanation is that as  $\beta$  increases so much that the log-likelihoods become unimportant, there is no need to increase the diffusion if



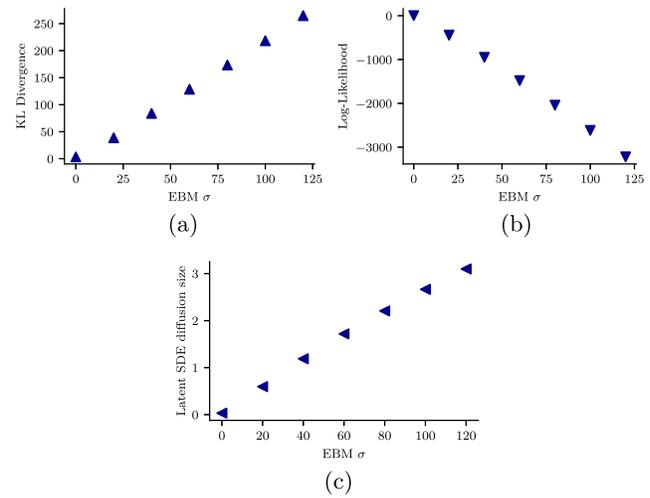
**FIG. 4.** Comparison of Wasserstein distance and transition rate between trained prior SDE and EBM for different EBM diffusion factors  $\sigma$  on the training set and test set. (a) Wasserstein distance (NaN in the last experiment), (b) transition rate.

the prior and posterior SDEs’ distributions are so close that their difference is almost zero. An extreme case of this is that the prior and posterior are equal, making the size of the diffusion arbitrary. We experimentally test this argument in Appendix B 6.

For other values of  $\beta$ , it is unclear how the data noise and the latent neural SDE’s diffusion size relate at first glance. We check if they have any relationship by varying the diffusion size  $\sigma$  of the EBM (cf. Sec. IV) and training a latent neural SDE with  $\beta = 10$  on the resulting data. The results are shown in Figs. 4(a) and 4(b), where “EBM diffusion factor  $\sigma$ ” is the size of the constant in the linear diffusion term of the EBM that is generating the data. We observe a correspondence between  $\sigma$  and the size of the latent neural SDE’s diffusion, although the latent neural SDE’s diffusion is always smaller.

It is more enlightening to plot the size of the KL divergence, log-likelihood, and diffusion of the latent neural SDE, trained on EBM data with different values of the diffusion factor  $\sigma$ , as in 5. First, we observe that the data gets harder to approximate with increasing EBM diffusion factors  $\sigma$ , as both the KL divergence and log-likelihood scores worsen (i.e., the KL divergence increases and the log-likelihoods decrease). The latent neural SDE’s diffusion size grows almost linearly with  $\sigma$ , similar to the KL divergence and log-likelihood.

This can be explained as follows: for latent neural SDEs, there is a “balance point” between the KL divergence and the log-likelihood. The latent neural SDE can freely choose the diffusion size according to these scores while training, and this can be seen in Fig. 5. As such, the diffusion size is a quantification of the “uncertainty” of the latent neural SDE’s drift about the data drift, which is not necessarily the data’s diffusion size. We add that this behavior is perfectly acceptable

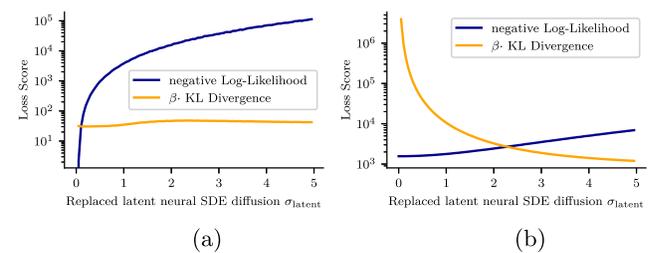


**FIG. 5.** Training information for latent neural SDE for different EBM diffusion factors  $\sigma$ : KL divergence, log-likelihood, and diffusion size (the diffusion size of the latent neural SDE is in the normalized space, while the diffusion size of EBM is in non-normalized space). (a) KL divergence for EBM  $\sigma$ , (b) log-likelihood for EBM  $\sigma$ , (c) diffusion size of latent neural SDE for EBM diffusion factors  $\sigma$ , integrated over training timespan.

in uncertainty quantification, where the diffusion size is a measure of the uncertainty of the model. However, in the context of the EBM, the diffusion size is a measure of the noise in the data, and the latent neural SDE should be able to model this noise.

How is the diffusion size chosen in training? A larger diffusion size makes for a better KL divergence and a worse log-likelihood, and a smaller diffusion size makes for a better log-likelihood and worse KL divergence. To demonstrate this, we now use the latent neural SDE for  $\beta = 10$  and subsequently replace the trained diffusion of the latent neural SDE with constant values from zero to five in Fig. 6. We plot both terms of the loss function for two data diffusion factors  $\sigma$ , zero and 100.

Here, the actual size of the final model’s diffusion turns out to be about the value of  $\sigma$  where these two scores, as weighted in the



**FIG. 6.** Replaced actual trained latent neural SDE’s diffusion with constant value  $\sigma_{\text{latent}}$  and computed both components of the loss function. The latent neural SDE trains to align to the values of  $\sigma_{\text{latent}}$  where the scores, as weighted in this figure, have the smallest sum. (a) Diffusion balance, EBM diffusion  $\sigma = 0$ , (b) diffusion balance, EBM diffusion  $\sigma = 100$ .

figure, have the smallest sum (i.e., the latent neural SDE’s objective is maximized). The harder the problem becomes to predict, the further this point moves to the right, leading to the relationship between the EBM’s diffusion factors  $\sigma$  and the latent neural SDE’s diffusion size we observed.

### VI. INJECTING NOISE INTO LATENT NEURAL SDES

In order to match the transition rate as well as the marginals, we amend the loss function to include a term for the size of diffusion

$$\max_{\theta, \phi} \mathbb{E}_{x_{\text{data}}} (\mathcal{L}_E - \beta \mathcal{L}_{KL} + \gamma \mathcal{L}_G),$$

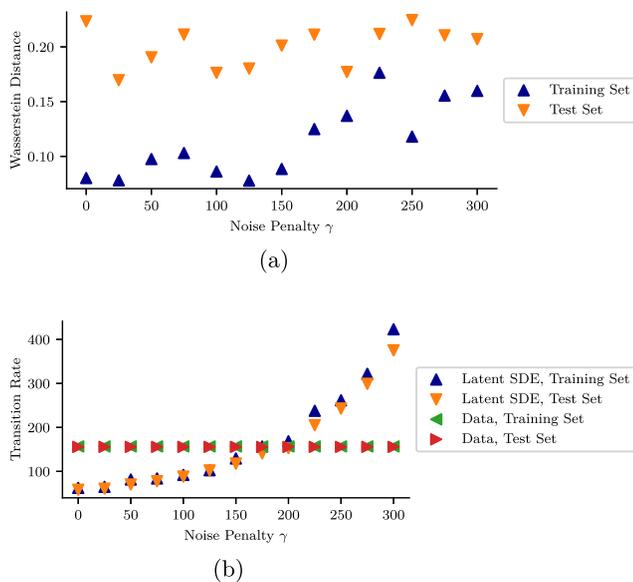
$$\text{where } \mathcal{L}_G = \int_0^T \|g_{\theta}(u(t), t)\| dt.$$

Alternatively, assuming we would know the desired value of  $\mathcal{L}_N$ , we could use a regularization-like penalty

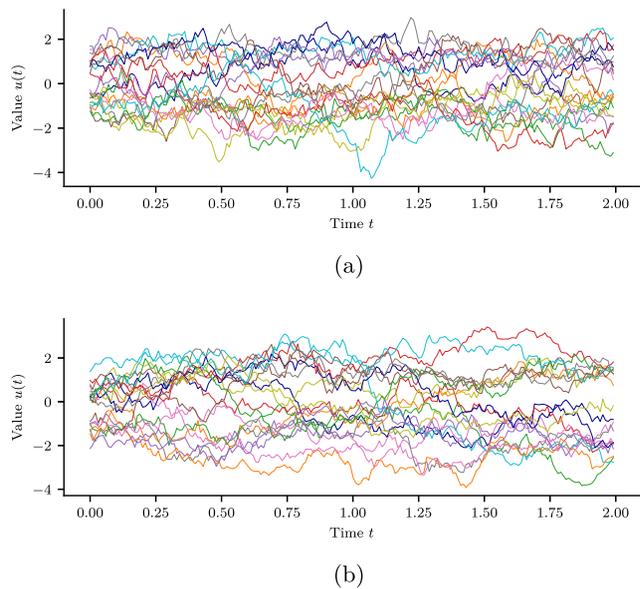
$$\max_{\theta, \phi} \mathbb{E}_{x_{\text{data}}} (\mathcal{L}_E - \beta \mathcal{L}_{KL} + \gamma (\mathcal{L}_G - g_{\text{target}})^2).$$

However, we do not know  $g_{\text{target}}$ , in general, so this adds two hyperparameters instead of one. Thus, we choose the first loss function. We compute  $\mathcal{L}_G$  just like  $\mathcal{L}_{KL}$ , adding another dimension in the integration of the posterior (cf. Ref. 7).

We start by fixing the balance factor  $\beta = 10$  and training a latent neural SDE on the EBM for different values of the new noise penalty  $\gamma$ . In Fig. 7(b), we observe that increasing  $\gamma$  will increase the transition rate. Relating the Wasserstein distance to  $\gamma$  in Fig. 7(a), we cannot observe a clear pattern but see values similar to those that we would obtain by setting  $\gamma = 0$  in the test set. The latent neural SDE



**FIG. 7.** Comparison of Wasserstein distance and transition rate between EBM and trained prior SDE for different values of noise penalty  $\gamma$  on training set and test set. (a) Wasserstein distance, (b) transition rate.

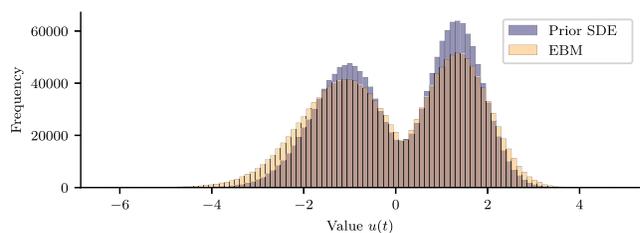


**FIG. 8.** Comparison of best result with the latent neural SDE with noise penalty against noisy EBM. (a) Noisy EBM, (b) latent neural SDE prior,  $\beta = 10$ ,  $\gamma = 200$ .

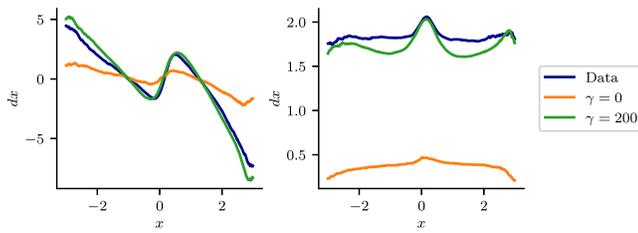
trained with  $\gamma = 200$  matches the transition rate and Wasserstein distances quite well.

We emphasize that this process of finding the correct value of  $\gamma$  can be done without any knowledge of the underlying dynamics, just by analyzing the data. Furthermore, any metric can be chosen to optimize the noise penalty, and this metric will have to be fine-tuned for the use case, such as our choice of the transition rate for the EBM.

Having found a seemingly reasonable value  $\gamma = 200$  for the noise penalty by looking at the transition rates, we now investigate how well the trained latent neural SDE with  $\beta = 10$  and  $\gamma = 200$  fits the EBM. We plot simulations in Fig. 8 and the marginals in Fig. 9. Next, we compare both the KM coefficients of drift and diffusion estimated from the data and the actual values of drift and diffusion between the latent neural SDE’s prior SDE and the EBM. Although we do not explicitly perform a grid search to find the optimal combination of  $\gamma$  and  $\beta$ , it would, in principle, be a good means to find the best possible parameterization.



**FIG. 9.** Marginals on noisy EBM on test set,  $\beta = 10$ ,  $\gamma = 200$ .

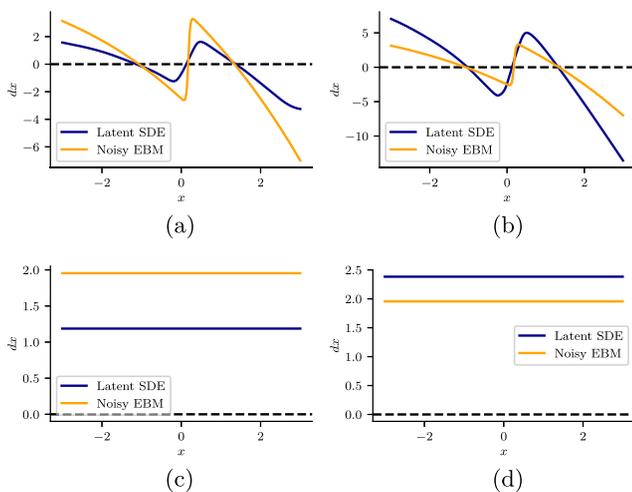


**FIG. 10.** Comparison of the Kramers–Moyal factors (drift left, diffusion right) between EBM and trained prior SDE with noise penalty  $\gamma = 0$  and  $\gamma = 200$  on test set.

In Fig. 10, we use the Epanechnikov kernel found in the `kramersmoyal` Python package<sup>31</sup> to compare Kramers–Moyal coefficients (see also Appendix A). We compute these from all data samples and latent neural SDEs with  $\gamma \in \{0, 200\}$  on the test set. The first and second coefficients represent the drift and diffusion, respectively. Here, the values from the latent neural SDE and the data look very similar.

Not only can we compute the KM coefficients from the data generated by the trained prior SDE and the EBM, but we can directly access their drift and diffusion functions and compare them. We do this for  $\gamma \in \{0, 200\}$  in Fig. 11. Here, we see quite a larger difference than with the KM coefficients. With noise penalty  $\gamma = 200$ , the latent neural SDE learns a drift and diffusion that is significantly different to the EBM’s drift and diffusion, but these values of drift and diffusion yield trajectories that behave very similarly, as can be seen from the KM coefficients.

We conduct some further experiments to test the noise penalty in different scenarios: In Appendix B 1, we demonstrate the



**FIG. 11.** Direct comparison of drift and diffusion functions of prior SDE and EBM for different values of the noise penalty  $\gamma$ . (a) Drift,  $\gamma = 0$ , (b) drift,  $\gamma = 200$ , (c) diffusion,  $\gamma = 0$ , (d) diffusion,  $\gamma = 200$ .

effectiveness of our method on a more complex system, a two-dimensional FitzHugh–Nagumo (FHN) model that models the NGRIP ice core data.<sup>32</sup> In Appendix B 2, we also train on a simple Ornstein–Uhlenbeck process, to show that the noise penalty also fixes the noise underestimation problem for a simple linear example. In Appendix B 4, we also train on a triple well model. In Appendix B 3, we also train on an EBM with a lower level of noise that exhibits tipping more rarely.

We conclude that by setting the hyperparameter  $\gamma$  to the correct value with respect to the modeling goals, we can synthesize an SDE that matches Wasserstein distances and transition rates of the data relatively well in the considered examples with constant diffusion. In this setup,  $\gamma$  cannot be learned automatically, leaving us with two hyperparameters ( $\beta$  and  $\gamma$ ) that must be manually tuned or identified with a grid search. On the other hand, with enough data, correct values of these hyperparameters can be obtained just by analyzing the data without any further knowledge of the underlying dynamics.

The hyperparameter  $\gamma$  only fits diffusion terms globally in its current formulation. The result is thus only accurate for data sets with constant diffusion. We show in Appendix B 5 that our method fails if the noise of the data is linear. To fit more complicated diffusion functions, the loss term  $\mathcal{L}_G$  must be modified in some way to track the diffusion in more detail.

## VII. DISCUSSION

We have shown that with the introduction of an additional noise penalty, latent neural SDEs can fit multistable dynamics, achieving high accuracy in KM coefficients, transition rate, and marginals. The noise penalty is essential to address diffusion underestimation in the stochastic models we considered, which include an energy balance model, the FitzHugh–Nagumo model and an Ornstein–Uhlenbeck process model. We investigated the cause of the underlying noise underestimation effect, which lies in the formulation of the loss function for latent neural SDEs. The unmodified latent neural SDE has to achieve a balance between the KL divergence and the log-likelihood of the data, which in turn results in an underestimated diffusion size.

Our formulation of the noise penalty adds another tunable hyperparameter to the latent neural SDE. More sophisticated methods, such as automatic hyperparameter optimization, might be a successful approach for tuning the noise penalty automatically and, thus, automatically training a latent neural SDE to have correct diffusion.

Using SDE-GANs to estimate these systems might have advantages, although the training procedure will be considerably more difficult. We expect that the noise underestimation problem would not arise for SDE-GANs in the same way, as the discriminator would ideally catch a trajectory with too little noise. The latent neural SDE framework likely has modeling advantages over SDE-GANs, for instance, the possibility of prediction from data samples and out-of-distribution detection through the posterior SDE. One could give the posterior SDE an unseen data sample and make a forecast by switching to the prior after the last posterior point in the state space. Out-of-distribution detection might be possible by exploring

the context returned by the encoder. To the best of our knowledge, this has not yet been attempted in the literature.

By modifying the latent neural SDE framework, we achieve a generative framework that is comparably easy to train and can fit multistable dynamics with an accurate estimation of their diffusion size. It, thus, significantly broadens the applications of the latent neural SDE method.

## ACKNOWLEDGMENTS

N.B. acknowledges funding by the Volkswagen Foundation. This is ClimTip contribution #X; the ClimTip project has received funding from the European Union's Horizon Europe research and innovation programme under Grant Agreement No. 101137601. M.S. acknowledges funding by the Ministry of Culture and Science (MKW) of the German State of North Rhine-Westphalia ("NRW Rückkehrprogramm").

## AUTHOR DECLARATIONS

### Conflict of Interest

The authors have no conflicts to disclose.

### Author Contributions

**L. Heck:** Data curation (equal); Investigation (equal); Methodology (equal); Software (equal); Writing – original draft (equal); Writing – review & editing (equal). **M. Gelbrecht:** Investigation (equal); Methodology (equal); Supervision (equal); Writing – original draft (equal); Writing – review & editing (equal). **M. T. Schaub:** Project administration (equal); Supervision (equal); Writing – review & editing (equal). **N. Boers:** Project administration (lead); Supervision (supporting); Writing – review & editing (equal).

## DATA AVAILABILITY

The data that support the findings of this study are available within the article and its supplementary material.

## APPENDIX A: KRAMERS-MOYAL COEFFICIENTS

In this paper, we use the Kramers–Moyal (KM) coefficients to analyze the behavior of our trained stochastic processes. The idea is that while we know the exact drift and diffusion terms of our ground truth and trained latent neural SDEs, they might behave more similarly as the diffusion terms suggest. Similar KM coefficients indicate this.

Given an  $N$ -dimensional Markovian process  $x(t) \in \mathbb{R}^n$ , the Kramers–Moyal coefficients are given as

$$\mathcal{M}^\sigma(x, t) = \lim_{\Delta t \rightarrow 0} \frac{1}{\Delta t} \int dx' [x'(t) - x(t)]^\sigma P(x', t + \Delta t | x, t),$$

where  $P(x', t + \Delta t | x, t)$  is the transition probability density, square brackets denote dyadic multiplication, and  $\sigma$  is a multi-index describing the orders.<sup>31</sup>

## APPENDIX B: FURTHER EXPERIMENTS

### 1. FitzHugh–Nagumo model

To show the effectiveness of our method on a more complex system, we train a latent neural SDE on data generated by a two-dimensional monostable FitzHugh–Nagumo (FHN) model. The FHN is a generalized van der Pol oscillator designed initially as a simplified model for neuron spikes in nerves. Its state space has two dimensions. The original model has an attractor  $P$ , which FitzHugh calls the resting point. When disturbed far enough from the  $P$ , a trajectory will travel through the model's state space and arrive back at  $P$  at some point.<sup>33</sup>

The governing equation of the FHN in its stochastic form is

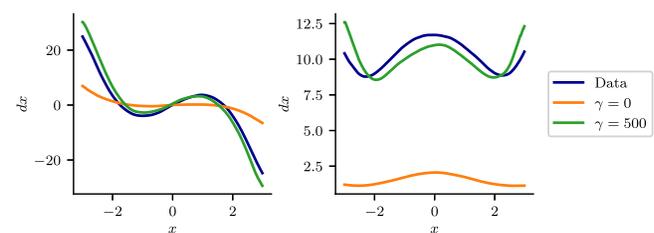
$$\begin{aligned} dx(t) &= \frac{1}{\tau_x} (\alpha_1 x(t) - \alpha_3 x(t)^3 + y(t)) dt + \sigma_x dB_{x,t}, \\ dy(t) &= \frac{1}{\tau_y} (\beta y(t) - x(t) + c) dt + \sigma_y dB_{y,t}. \end{aligned}$$

We choose an automatically synthesized parameterization from Lohmann and Ditlevsen<sup>32</sup> that was selected to match certain behaviors of the NGRIP ice core record, which they call  $\text{FHN}_\gamma$ , where  $\tau_x = \tau_y = 1$ ,  $b = 2.55$ ,  $\alpha_1 = 0.63$ ,  $\alpha_3 = 2.71$ ,  $c = 0.22$ ,  $\sigma_x = 4.80$ ,  $\sigma_y = 11.08$ ,  $\beta = \tan(-0.67)$ . Note that this system is monostable, and while it exhibits tipping-like dynamics, these are not due to the typical relaxation oscillation that the FHN is mostly known for.

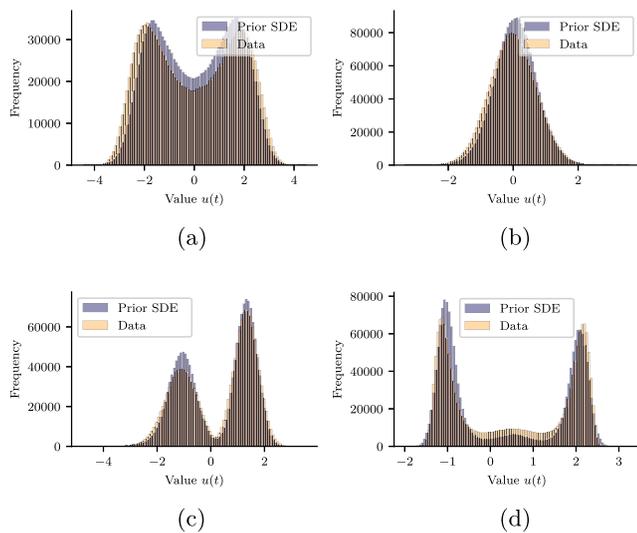
We choose to observe just the first dimension of the FHN model in an analogy to the NGRIP ice core. Thus, we also demonstrate the capabilities of latent neural SDEs to model just partially observed nonlinear chaotic systems.

We train a two-dimensional latent neural SDE on the data from this model, but only let it observe  $x(t)$ , which is the dimension that models the NGRIP ice core record. We choose  $\beta = 10$  and increase the learning rate decay to 0.9997.

By summary statistics, we select  $\gamma = 500$  as the noise penalty. We will restrict our comparative analysis of the fitted model to the first dimension  $x(t)$  of the FHN model, as this dimension models the NGRIP ice core record, and the latent neural SDE only observed this dimension. We plot the KM coefficients for  $\gamma = 0$  and  $\gamma = 500$  in Fig. 12. The latent neural SDE's prior SDE is able to match the behavior of the FHN model well with  $\gamma = 500$ , but not as a standard latent neural SDE with  $\gamma = 0$ . The transition rate of the latent neural SDE with  $\gamma = 500$  is 513.29, which is close to the transition rate



**FIG. 12.** FitzHugh–Nagumo: One-dimensional KM coefficients (drift left, diffusion right) on test set.



**FIG. 13.** Marginals for experiments in Appendixes A and B, all computed on test set. (a) Marginals, FHN  $\beta = 10, \gamma = 500$ , (b) marginals, OU  $\beta = 10, \gamma = 650$ , (c) marginals, EBM with tipping occurring more rarely  $\beta = 10, \gamma = 150$ , (d) marginals, triple well model  $\beta = 10, \gamma = 205$ .

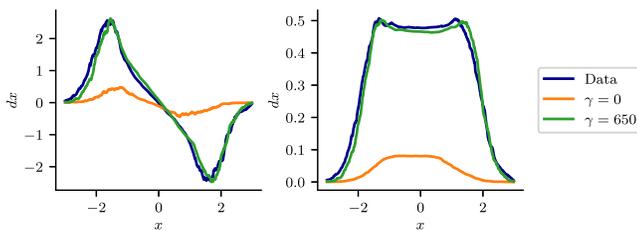
of the FHN model, 458.29, whereas for  $\gamma = 0$ , the transition rate is 232.4. The Wasserstein distances of both latent neural SDEs are similar (0.176 for  $\gamma = 500$  and 0.206 for  $\gamma = 0$ ). We show the marginals of the prior SDE with  $\gamma = 500$  in Fig. 13(a).

**2. Minimal Ornstein-Uhlenbeck process**

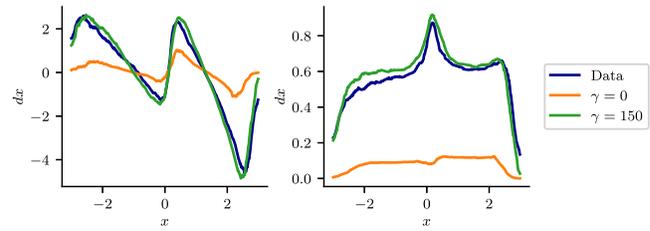
To show that the noise penalty also works with a minimal linear example, we train on an Ornstein-Uhlenbeck process of the form

$$du(t) = -u(t)dt + dB_t.$$

The drift of this process is a simple attractor around zero, and the diffusion is one. We chose  $t_{train} = 5, \Delta t = 0.01$  and trained for 5000 epochs, keeping all other settings the same as for the EBM. We select  $\gamma = 650$  by analyzing summary statistics and plot the KM coefficients for  $\gamma = 0$  and  $\gamma = 650$  in Fig. 14. The latent neural SDE's prior SDE is able to match the behavior of the Ornstein-Uhlenbeck process well with  $\gamma = 650$ , but not as a standard latent neural SDE



**FIG. 14.** Ornstein-Uhlenbeck process: KM coefficients (drift left, diffusion right) on test set.



**FIG. 15.** EBM with rarer tipping: KM coefficients (drift left, diffusion right) on test set.

with  $\gamma = 0$ . We show the marginals of the prior SDE with  $\gamma = 650$  in Fig. 13(b).

**3. EBM with rarer tipping**

We repeat the same experiment as in Sec. VI with an EBM that exhibits tipping more rarely. For this, we keep all parameters the same, but we set  $\sigma = 25$  (in Sec. VI, we set it to 40). We searched for  $\gamma$  as in Sec. VI and found that for  $\gamma = 150$ , the transition rate is similar (47.8 for the trained prior SDE and 50.5 for the noisy EBM). We plot the marginals in Fig. 13(c) and the KM coefficients in Fig. 15. We also show simulations in Fig. 16.

**4. Triple-well model**

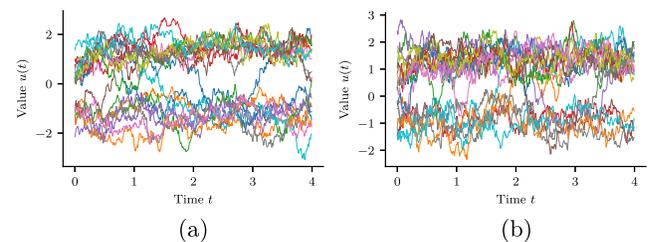
We have considered a double-well model so far. We now consider the following triple-well model:

$$dT(t) = -6T(t)^5 + 20T(t)^3 - 8T(t)dt + 2.5dB_t.$$

We train a one-dimensional latent neural SDE with  $\beta = 10$ . By looking at the KM coefficients, we select  $\gamma = 205$ . We plot the KM coefficients in Fig. 17 and the marginals in Fig. 13(d).

**5. EBM with linear diffusion**

The noise penalty  $\gamma$  globally increases the size of the latent neural SDE's diffusion function, so we expect our noise penalty method to fail to match diffusion functions that are not constant. To test this hypothesis, we build an EBM with a linear diffusion term by slightly



**FIG. 16.** Comparison of best result with the latent neural SDE against noisy EBM with rarer tipping. (a) Noisy EBM with rarer tipping, (b) Latent neural SDE Prior,  $\beta = 10, \gamma = 150$ .

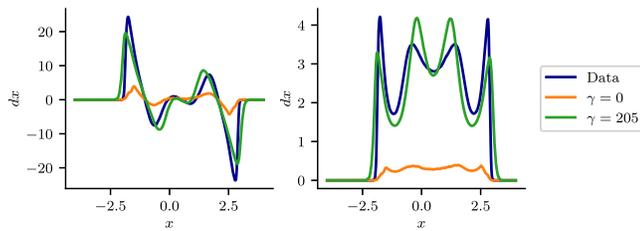


FIG. 17. Triple-well: One-dimensional KM coefficients (drift left, diffusion right) on test set.

modifying the EBM equation in Sec. IV

$$dT(t) = (a_1 + a_2 \tanh(T(t) - T_0) - a_3 T(t)^4)dt + 0.135T(t)dB_t.$$

In this EBM, we have multiplicative noise, meaning that the diffusion depends on the value  $T(t)$ . While our trained latent SDE could exhibit this behavior, our formulation of the noise penalty does not take this into account.

We kept the training settings the same as for the original EBM. We select  $\gamma = 150$  by analyzing summary statistics and plot the KM coefficients for  $\gamma = 0$  and  $\gamma = 150$  in Fig. 18. While the drift's KM coefficients still match between the EBM and the latent neural SDE's prior to noise penalty, the diffusion is clearly linear for the EBM and close to constant for the latent neural SDE. Thus, one can indeed only use the noise penalty method to match constant diffusion functions. A noise penalty for multiplicative noise is an interesting avenue for future research.

### 6. Latent neural SDEs on extreme values of beta

We aim to test our argument from Sec. V experimentally. For this, we train three latent neural SDEs on the EBM from Sec. IV: all without noise penalty, one with  $\beta = 0$ , and two with  $\beta = 10^{10}$ . During the training of the latent neural SDE with  $\beta = 0$ , the noise level consistently moves toward zero, see Fig. 19(a). On the other hand, with  $\beta = 10^{10}$ , the noise level moves in the beginning but then stays almost constant, see Fig. 19(b). The two latent neural SDEs with  $\beta = 10^{10}$  have different noise levels after training: one is around 3.2 while the other is around 3.8. Thus, the noise level indeed seems arbitrary here.

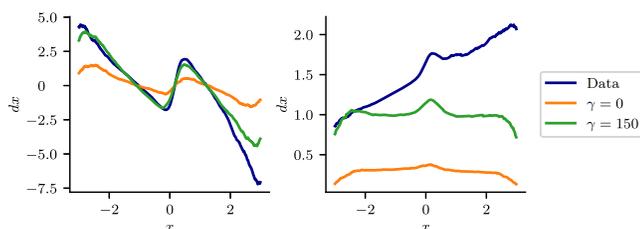


FIG. 18. EBM with linear noise: KM coefficients (drift left, diffusion right) on test set.

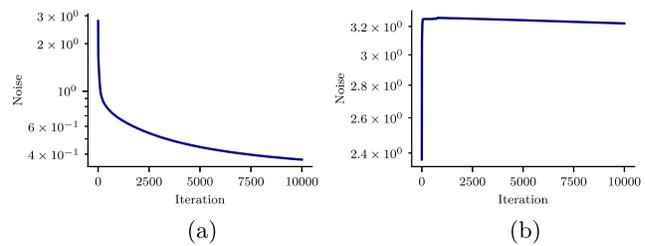


FIG. 19. Noise level of latent neural SDEs. (a) Noise level of latent neural SDE with  $\beta = 0$  during training. (b) Noise level of latent neural SDE with  $\beta = 10^{10}$  during training.

### REFERENCES

- <sup>1</sup>K. Hasselmann, "Stochastic climate models part I. Theory," *Tellus* **28**, 473–485 (1976).
- <sup>2</sup>C. L. E. Franzke, T. J. O’Kane, J. Berner, P. D. Williams, and V. Lucarini, "Stochastic climate theory and modeling," *WIREs Clim. Change* **6**, 63–78 (2015).
- <sup>3</sup>J. K. Møller, M. Zugno, and H. Madsen, "Probabilistic forecasts of wind power generation by stochastic differential equation models," *J. Forecast.* **35**, 189–205 (2016).
- <sup>4</sup>J. Dong, M. M. Olama, T. Kuruganti, A. M. Melin, S. M. Djouadi, Y. Zhang, and Y. Xue, "Novel stochastic methods to predict short-term solar radiation and photovoltaic power," *Renew. Energy* **145**, 333–346 (2020).
- <sup>5</sup>J. Fagin, J. W. Park, H. Best, J. H. H. Chan, K. E. S. Ford, M. J. Graham, V. A. Villar, S. Ho, and M. O’Dowd, "Latent stochastic differential equations for modeling quasar variability and inferring black hole properties," *Astrophys. J.* **965**(2), 104 (2024).
- <sup>6</sup>F. Black and M. Scholes, "The pricing of options and corporate liabilities," *J. Political Econ.* **81**, 637–654 (1973).
- <sup>7</sup>X. Li, T.-K. L. Wong, R. T. Chen, and D. K. Duvenaud, "Scalable gradients and variational inference for stochastic differential equations," in *Symposium on Advances in Approximate Bayesian Inference* (PMLR, 2020), pp. 1–28.
- <sup>8</sup>P. Kidger, "On neural differential equations," Ph.D. thesis (University of Oxford, 2021).
- <sup>9</sup>P. Kidger, J. Foster, X. Li, and T. J. Lyons, "Neural SDEs as infinite-dimensional GANs," in *Proceedings of the 38th International Conference on Machine Learning*, Proceedings of Machine Learning Research (PMLR, 2021), Vol. 139, pp. 5453–5463.
- <sup>10</sup>P. Kidger, J. Foster, X. C. Li, and T. Lyons, "Efficient and accurate gradients for neural SDEs," in *Advances in Neural Information Processing Systems* (Curran Associates, Inc., 2021), Vol. 34, pp. 18747–18761.
- <sup>11</sup>S. Zeng, F. Graf, and R. Kwitt, "Latent SDEs on homogeneous spaces," in *NeurIPS* (Curran Associates, 2023).
- <sup>12</sup>See <https://doi.org/10.5281/zenodo.14534737> for "The artifact is available on Zenodo."
- <sup>13</sup>C. K. Williams and C. E. Rasmussen, *Gaussian Processes for Machine Learning* (MIT Press, Cambridge, MA, 2006).
- <sup>14</sup>W. Xu, R. T. Q. Chen, X. Li, and D. Duvenaud, "Infinitely deep Bayesian neural networks with stochastic differential equations," in *Proceedings of the 25th International Conference on Artificial Intelligence and Statistics*, Proceedings of Machine Learning Research (PMLR, 2022), Vol. 151, pp. 721–738.
- <sup>15</sup>A. Hasan, J. M. Pereira, S. Farsiu, and V. Tarokh, "Identifying latent stochastic differential equations," *IEEE Trans. Signal Process.* **70**, 89–104 (2021).
- <sup>16</sup>J. Jia and A. R. Benson, "Neural jump stochastic differential equations," *Adv. Neural Inf. Process. Syst.* **32**, 9843–9854 (2019).
- <sup>17</sup>C. Herrera, F. Krach, and J. Teichmann, "Neural jump ordinary differential equations: Consistent continuous-time prediction and filtering," in *International Conference on Learning Representations* (2021).
- <sup>18</sup>F. Djeumou, C. Neary, and U. Topcu, "How to learn and generalize from three minutes of data: Physics-constrained and uncertainty-aware neural stochastic

differential equations,” in *CoRL, Proceedings of Machine Learning Research* (PMLR, 2023), Vol. 229, pp. 577–601.

- <sup>19</sup>M. A. Finzi, A. Boral, A. G. Wilson, F. Sha, and L. Zepeda-Núñez, “User-defined event sampling and uncertainty quantification in diffusion models for physical dynamical systems,” in *International Conference on Machine Learning* (PMLR, 2023), pp. 10136–10152.
- <sup>20</sup>D. J. Gagne, H. M. Christensen, A. C. Subramanian, and A. H. Monahan, “Machine learning for stochastic parameterization: Generative adversarial networks in the Lorenz’96 model,” *J. Adv. Model. Earth Syst.* **12**, e2019MS001896 (2020).
- <sup>21</sup>T. Jinno, T. Mitsui, K. Nakai, Y. Saiki, and T. Yoneda, “Long-term prediction of El Niño–Southern Oscillation using reservoir computing with data-driven realtime filter,” [arXiv:2501.17781](https://arxiv.org/abs/2501.17781) (2025).
- <sup>22</sup>B. Tzen and M. Raginsky, “Neural stochastic differential equations: Deep latent Gaussian models in the diffusion limit,” [arXiv:1905.09883](https://arxiv.org/abs/1905.09883) (2019).
- <sup>23</sup>D. P. Kingma and M. Welling, “Auto-encoding variational Bayes,” [arXiv:1312.6114](https://arxiv.org/abs/1312.6114) (2013).
- <sup>24</sup>I. Higgins, L. Matthey, A. Pal, C. Burgess, X. Glorot, M. Botvinick, S. Mohamed, and A. Lerchner, “beta-VAE: Learning basic visual concepts with a constrained variational framework,” in *International Conference on Learning Representations* (2016).
- <sup>25</sup>K. Fraedrich, “Catastrophes and resilience of a zero-dimensional climate system with ice-albedo and greenhouse feedback,” *Q. J. R. Meteorol. Soc.* **105**, 147–167 (1979).
- <sup>26</sup>M. Ghil, “Climate stability for a sellers-type model,” *J. Atmos. Sci.* **33**, 3–20 (1976).
- <sup>27</sup>A. Sutera, “On stochastic perturbation and long-term climate behaviour,” *Q. J. R. Meteorol. Soc.* **107**, 137–151 (1981).
- <sup>28</sup>H. Fu, C. Li, X. Liu, J. Gao, A. Celikyilmaz, and L. Carin, “Cyclical annealing schedule: A simple approach to mitigating KL vanishing,” [arXiv:1903.10145](https://arxiv.org/abs/1903.10145) (2019).
- <sup>29</sup>D. P. Kingma, “Adam: A method for stochastic optimization,” [arXiv:1412.6980](https://arxiv.org/abs/1412.6980) (2014).
- <sup>30</sup>V. M. Panaretos and Y. Zemel, “Statistical aspects of Wasserstein distances,” *Annu. Rev. Stat. Appl.* **6**, 405–431 (2019).
- <sup>31</sup>L. R. Gorjão and F. Meirinhos, “kramersmoyal: Kramers–Moyal coefficients for stochastic processes,” *J. Open Source Softw.* **4**, 1693 (2019).
- <sup>32</sup>J. Lohmann and P. D. Ditlevsen, “A consistent statistical model selection for abrupt glacial climate changes,” *Clim. Dyn.* **52**, 6411–6426 (2019).
- <sup>33</sup>R. FitzHugh, “Impulses and physiological states in theoretical models of nerve membrane,” *Biophys. J.* **1**, 445–466 (1961).