

PHYSICS-AWARE, DEEP PROBABILISTIC MODELING OF MULTISCALE DYNAMICS IN THE SMALL DATA REGIME

Sebastian Kaltenbach¹, Phaedon-Stelios Koutsourelakis¹

¹ Professorship of Continuum Mechanics, Technical University of Munich
Boltzmannstr.15, 85748 Garching
{sebastian.kaltenbach; p.s.koutsourelakis}@tum.de, www.mw.tum.de/contmech

Key words: Bayesian machine learning, virtual observables, multiscale modeling, coarse-graining

Abstract: *The data-based discovery of effective, coarse-grained (CG) models of high-dimensional dynamical systems presents a unique challenge in computational physics and particularly in the context of multiscale problems. The present paper offers a probabilistic perspective that simultaneously identifies predictive, lower-dimensional coarse-grained (CG) variables as well as their dynamics. We make use of the expressive ability of deep neural networks in order to represent the right-hand side of the CG evolution law. Furthermore, we demonstrate how domain knowledge that is very often available in the form of physical constraints (e.g. conservation laws) can be incorporated with the novel concept of virtual observables. Such constraints, apart from leading to physically realistic predictions, can significantly reduce the requisite amount of training data which enables reducing the amount of required, computationally expensive multiscale simulations (Small Data regime). The proposed state-space model is trained using probabilistic inference tools and, in contrast to several other techniques, does not require the prescription of a fine-to-coarse (restriction) projection nor time-derivatives of the state variables. The formulation adopted is capable of quantifying the predictive uncertainty as well as of reconstructing the evolution of the full, fine-scale system which allows to select the quantities of interest a posteriori. We demonstrate the efficacy of the proposed framework in a high-dimensional system of moving particles.*

1 INTRODUCTION

The solution of high-dimensional, multiscale system is challenging as the required computational resources usually grow exponentially with the dimension of the state-space as well as with the smallest time-scale that needs to be resolved. As such systems are ubiquitous in applied physics and engineering, reduced/coarse-grained descriptions and models are necessary that are predictive of various observables or the high-dimensional system, but whose discretization time-scales can be much larger than the inherent ones [1].

We adopt a data-based perspective [2, 3] that relies on data generated by simulations of a fine-grained (FG) system in order to learn a coarse-grained (CG) model. We nevertheless note that such coarse-graining tasks exhibit fundamental differences from large-scale machine learn-

ing tasks [4, 5] as the data involved is usually small due to the expensive data acquisition and as information about the underlying physical structure of the problem is available. When this domain knowledge is incorporated into the CG model it can improve its predictive ability [6, 7].

In contrast to other frameworks for reduced-order modeling (e.g. SINDy [8]) where the dynamics of the CG model is learned based on a large vocabulary of feature functions, we employ a deep neural network for the CG dynamics in order to gain great flexibility and be able to not restrict ourselves to an a priori chosen set of feature functions. This approach is similar to the ideas of Neural ODEs [9] and Neural SDEs [10] which also use neural networks to represent the dynamics. Another possibility would be the use of Gaussian Processes [11] which would allow non-parametric, probabilistic modeling.

In this paper, we combine a generative, probabilistic machine learning framework [12] with virtual observables [6] and deep neural networks for the CG dynamics as well as the mapping from the CG states to the FG states. In doing so, we propose a framework that can make use of the flexibility of neural nets, while still obeying physical laws. We carry out the tasks of model estimation and dimensionality reduction simultaneously and identify the CG states variables, their dynamics as well as a probabilistic coarse-to-fine map based only on small amounts of FG simulation data.

2 METHODOLOGY

In general, the subscript f or lower-case letters are used to denote variables associated with the (high-dimensional) fine-grained(FG) model and the subscript c or upper-case letters are used for quantities of the (lower-dimensional) coarse-grained(CG) description. We also use a circumflex $\hat{}$ to denote observed/known variables.

2.1 The FG and CG model

The fine-grained system considered is a high-dimensional system with state variables \mathbf{x} ($\mathbf{x} \in \mathcal{X}_f \subset \mathbf{R}^{d_f}$), whose dimension d_f is very large ($d_f \gg 1$). We describe the dynamics of such a FG system by a system of deterministic or stochastic ODEs i.e.,

$$\dot{\mathbf{x}}_t = \mathbf{f}(\mathbf{x}_t, t), \quad t > 0 \quad (1)$$

The FG system is moreover considered to have the initial condition \mathbf{x}_0 that might be deterministic or drawn from a specified distribution. In this work, we want to coarse-grain such a system only based on simulated data, i.e. time sequences simulated from Equation (1) with a time-step δt .

Our goal is to simultaneously identify (unknown) CG state variables \mathbf{X} with $\mathbf{X} \in \mathcal{X}_c \subset \mathbf{R}^{d_c}$ as well as the dynamics of those CG variables. The dimension d_c of these CG state variables is intended to be much smaller than d_f . For the CG dynamics a Markovian dynamic is assumed in the form:

$$\dot{\mathbf{X}}_t = \mathbf{F}(\mathbf{X}_t, t), \quad t > 0 \quad (2)$$

2.2 Emission law

In contrast to approaches based on the Mori-Zwanzig formalism [13, 14], which include a mapping from the FG system to the quantities of interest, we employ a probabilistic, generative coarse-to-fine map [15] from the CG state-variables to the FG description. We indicate the associated (conditional) density by:

$$p_{cf}(\mathbf{x}_t | \mathbf{X}_t; \boldsymbol{\theta}_{cf}) \quad (3)$$

where $\boldsymbol{\theta}_{cf}$ denote the (unknown) parameters that we will try to learn from the data. This conditional density p_{cf} can be endowed a priori with domain knowledge by adapting its form to the particulars of the problem or it can be parametrized by deep neural networks to allow for maximum flexibility.

Employing a probabilistic coarse-to-fine map instead of a deterministic, restriction operator has many advantages as e.g. the full FG system's reconstruction and probabilistic predictive estimates.

2.3 Transition law

In the following, we consider discretized time with a fixed time-step Δt and time-related subscripts refer to the number of time-steps.

We model the CG dynamics with the help of a deep neural network in order to gain a great flexibility and be able to express nonlinear functions. Therefore, we assume an explicit discretization of Equation (2) and model the right-hand-side by the deep neural network $NN(\cdot)$ parametrized by $\boldsymbol{\theta}_{NN}$:

$$\mathbf{X}_{t+1} = \mathbf{X}_t + NN(\mathbf{X}_t, \boldsymbol{\theta}_{NN}) + \sigma_r \boldsymbol{\varepsilon}, \quad \boldsymbol{\varepsilon} \sim \mathcal{N}(\mathbf{0}, \mathbf{I}) \quad (4)$$

where the parameter $\sigma_r \geq 0$ is responsible for the stochastic part of the CG dynamics. This leads to the following conditional density:

$$p(\mathbf{X}_{t+1} | \mathbf{X}_t, \boldsymbol{\theta}_{NN}, \sigma_r) = \mathcal{N}(\mathbf{X}_{t+1} | \mathbf{X}_t + NN(\mathbf{X}_t, \boldsymbol{\theta}_{NN}), \sigma_r^2 \mathbf{I}) \quad (5)$$

which effectively represents a discretized version of the neural stochastic ODEs of [10] and is more flexible as compared to approaches in which the right-hand side consists of a restricted amount of first- and second-order interactions of \mathbf{X}_t [6].

2.4 Virtual observables

As the CG state-variables \mathbf{X} employed in multiscale modeling are usually given physical meaning, we employ the concept of *virtual observables* [6] in order to incorporate general physical principles such as conservation of mass, momentum or energy. Let these be expressed as equalities of the form at each time-step l :

$$\mathbf{c}_l(\mathbf{X}_l) = \mathbf{0}, \quad l = 0, 1, \dots \quad (6)$$

where $\mathbf{c}_l : \mathcal{X}_c \subset \mathbb{R}^{d_c} \rightarrow \mathbb{R}^{M_c}$. The only requirement we will impose is that of differentiability of \mathbf{c}_l [6]. We define a new variable $\hat{\mathbf{c}}_l$ which relates to \mathbf{c}_l as follows:

$$\hat{\mathbf{c}}_l = \mathbf{c}_l(\mathbf{X}_l) + \sigma_c \boldsymbol{\varepsilon}_c, \quad \boldsymbol{\varepsilon}_c \sim \mathcal{N}(\mathbf{0}, \mathbf{I}) \quad (7)$$

Now, it is assumed that the $\hat{\mathbf{c}}_l$ have been *virtually* observed and this set of virtual observations $\hat{\mathbf{c}}_l = \mathbf{0}$ leads to an augmented version of the FG data and therefore virtual likelihoods of the type:

$$p(\hat{\mathbf{c}}_l = \mathbf{0} \mid \mathbf{X}_l, \sigma_c) = \mathcal{N}(\mathbf{0} \mid \mathbf{c}_l(\mathbf{X}_l), \sigma_c^2 \mathbf{I}) \quad (8)$$

The “noise” parameter σ_c can be used to account for the intensity of the enforcement of the virtual observations and represents the tolerance parameter with which the constraints would be enforced in a deterministic setting.

We note that the concept of virtual observables is not restricted to physical constraints but could also be applied to residuals of temporal discretization schemes [6] or of PDEs [16]. In both of these cases, it is shown that the incorporation of virtual observables can reduce the amount of training data required and enable training in the Small Data regime.

2.5 Inference and learning

Due to the introduction of virtual observables, we can adopt an enlarged definition of data which we cumulatively denote by $\mathcal{D} = \{\hat{\mathbf{x}}_{0:T}^{(1:n)}, \hat{\mathbf{c}}_{0:T}^{(1:n)}\}$ and which encompasses:

- FG simulation data consisting of n sequences of the FG state-variables. These are denoted by $\hat{\mathbf{x}}_{0:T}^{(1:n)}$ as the likelihood model implied by the p_{cf} in Equation (3) involves only the observables at each coarse time-step.
- Virtual observables $\hat{\mathbf{c}}_l^{(1:n)}$ relating to the CG states \mathbf{X}_l at each time-step l and which relate to the physical constraints as in Equation (7). In the example they pertain to all time-steps from 0 to T and are denoted by $\hat{\mathbf{c}}_{0:T}^{(1:n)}$.

We represent the latent (unobserved) variables of the model by the CG state-variables $\mathbf{X}_{0:T}^{(1:n)}$ and relate them to the FG data through the p_{cf} (in Equation (3)) and to the virtual observables through Equation (8). The parameters of the model are denoted cumulatively by $\boldsymbol{\theta}$ and consist of¹:

- $\boldsymbol{\theta}_{NN}$ which parametrize the neural network for the right-hand-side of the CG evolution law (see section 2.3),
- $\boldsymbol{\theta}_{cf}$ which parametrize the probabilistic coarse-to-fine map (Equation (3)),
- σ_r involved in the stochasticity of the transition law Equation (4) and
- σ_c involved in the enforcement of virtual observables in Equation (7)

¹If any of these parameters are prescribed, then they are omitted from $\boldsymbol{\theta}$.

We follow a fully-Bayesian formulation and express the posterior of the unknowns (i.e. latent variables and parameters) as follows:

$$p(\mathbf{X}_{0:T}^{(1:n)}, \boldsymbol{\theta} \mid \mathcal{D}) = \frac{p(\mathcal{D} \mid \mathbf{X}_{0:T}^{(1:n)}, \boldsymbol{\theta}) p(\mathbf{X}_{0:T}^{(1:n)}, \boldsymbol{\theta})}{p(\mathcal{D})} \quad (9)$$

where $p(\mathbf{X}_{0:T}^{(1:n)}, \boldsymbol{\theta})$ denotes the prior on the latent variables and parameters. The likelihood term $p(\mathcal{D} \mid \mathbf{X}_{0:T}^{(1:n)}, \boldsymbol{\theta})$ involved can be decomposed into the product of two (conditionally) independent terms, one for the FG data and one for the virtual observables, i.e.:

$$p(\mathcal{D} \mid \mathbf{X}_{0:T}^{(1:n)}, \boldsymbol{\theta}) = p(\hat{\mathbf{x}}_{0:T}^{(1:n)} \mid \mathbf{X}_{0:T}^{(1:n)}, \boldsymbol{\theta}) p(\hat{\mathbf{c}}_{0:T}^{(1:n)} \mid \mathbf{X}_{0:T}^{(1:n)}, \boldsymbol{\theta}) \quad (10)$$

We further note that (from Equation (3)):

$$p(\hat{\mathbf{x}}_{0:T}^{(1:n)} \mid \mathbf{X}_{0:T}^{(1:n)}, \boldsymbol{\theta}) = \prod_{i=1}^n \prod_{t=0}^T p_{cf}(\mathbf{x}_t^{(i)} \mid \mathbf{X}_t^{(i)}, \boldsymbol{\theta}_{cf}) \quad (11)$$

and (from Equation (8)):

$$\begin{aligned} p(\hat{\mathbf{c}}_{0:T}^{(1:n)} \mid \mathbf{X}_{0:T}^{(1:n)}, \boldsymbol{\theta}) &= \prod_{i=1}^n \prod_{l=0}^T \mathcal{N}(\mathbf{0} \mid \mathbf{c}_l(\mathbf{X}_l^{(i)}), \sigma_c^2 \mathbf{I}) \\ &\propto \prod_{i=1}^n \prod_{l=0}^T \frac{1}{\sigma_c^{\dim(\mathbf{c})}} \exp \left\{ -\frac{1}{2\sigma_c^2} \left| \mathbf{c}_l(\mathbf{X}_l^{(i)}) \right|^2 \right\} \end{aligned} \quad (12)$$

The prior $p(\mathbf{X}_{0:T}^{(1:n)}, \boldsymbol{\theta})$ can be decomposed into the transition density of Equation (5) and a prior for \mathbf{X}_0 as well as the parameters $\boldsymbol{\theta}$:

$$p(\mathbf{X}_{0:T}^{(1:n)}, \boldsymbol{\theta}) = \prod_{i=1}^n p(\mathbf{X}_0^{(i)}) \prod_{t=0}^{T-1} p(\mathbf{X}_{t+1}^{(i)} \mid \mathbf{X}_t^{(i)}, \boldsymbol{\theta}_{NN}, \sigma_r) p(\boldsymbol{\theta}) \quad (13)$$

We advocate the use of Stochastic Variational Inference [17] for computing an approximate posterior. We select a parameterized family of densities, $q_\phi(\mathbf{X}_{0:T}^{(1:n)}, \boldsymbol{\theta})$ and attempt to find the one that best approximates the posterior by minimizing their Kullback-Leibler divergence. It can be shown [18], that this optimal q_ϕ maximizes the Evidence Lower Bound (ELBO) $\mathcal{F}(q_\phi(\mathbf{X}_{0:T}^{(1:n)}, \boldsymbol{\theta}))$:

$$\begin{aligned} \log p(\mathcal{D}) &= \log \int p(\mathcal{D}, \mathbf{X}_{0:T}^{(1:n)}, \boldsymbol{\theta}) d\mathbf{X}_{0:T}^{(1:n)} d\boldsymbol{\theta} \\ &= \log \int \frac{p(\mathcal{D} \mid \mathbf{X}_{0:T}^{(1:n)}, \boldsymbol{\theta}) p(\mathbf{X}_{0:T}^{(1:n)}, \boldsymbol{\theta})}{q_\phi(\mathbf{X}_{0:T}^{(1:n)}, \boldsymbol{\theta})} q_\phi(\mathbf{X}_{0:T}^{(1:n)}, \boldsymbol{\theta}) d\mathbf{X}_{0:T}^{(1:n)} d\boldsymbol{\theta} \\ &\geq \int \log \frac{p(\mathcal{D} \mid \mathbf{X}_{0:T}^{(1:n)}, \boldsymbol{\theta}) p(\mathbf{X}_{0:T}^{(1:n)}, \boldsymbol{\theta})}{q_\phi(\mathbf{X}_{0:T}^{(1:n)}, \boldsymbol{\theta})} q_\phi(\mathbf{X}_{0:T}^{(1:n)}, \boldsymbol{\theta}) d\mathbf{X}_{0:T}^{(1:n)} d\boldsymbol{\theta} \\ &= \mathcal{F}(q_\phi(\mathbf{X}_{0:T}^{(1:n)}, \boldsymbol{\theta})) \end{aligned} \quad (14)$$

In the following illustrations, we postulate a *mean-field* decomposition:

$$q_{\phi}(\mathbf{X}_{0:T}^{(1:n)}, \boldsymbol{\theta}) = q_{\phi}(\mathbf{X}_{0:T}^{(1:n)}) p_{\phi}(\boldsymbol{\theta}) = \left[\prod_{i=1}^n q_{\phi}(\mathbf{X}_{0:T}^{(i)}) \right] \delta_{\phi}(\boldsymbol{\theta}) \quad (15)$$

where we make use of the (conditional) independence of the time sequences in the likelihood. We further note that we employed Dirac δ_{ϕ} functions for the $q_{\phi}(\boldsymbol{\theta})$ and therefore obtain MAP estimates $\boldsymbol{\theta}_{MAP}$ (i.e. ϕ includes $\boldsymbol{\theta}_{MAP}$) for the unknown parameters.

Gradients of the ELBO with respect to the parameters ϕ involve expectations with respect to q_{ϕ} . These were approximated with Monte Carlo estimates which employ the reparametrization trick [19] and stochastic optimization was carried out with the ADAM algorithm [20].

2.6 Predictions

The proposed framework can produce probabilistic predictive estimates for a sequence which was observed up to time-step T i.e. $\hat{\mathbf{x}}_{0:T}^{(i)}$. This predictive uncertainty reflects not only the information-loss due to the coarse-graining process but also the epistemic uncertainty arising from finite (and small) datasets.

In particular, if $q_{\phi}(\mathbf{X}_T^{(i)})$ is the (marginal) posterior of the last, hidden CG state and $\boldsymbol{\theta}_{MAP}$ the MAP estimate of the model parameters, then we follow the steps described in Algorithm 1. This procedure generates samples of the full FG state evolution but does not necessarily guarantee the enforcement of the constraints for the CG states.

We note that if we would also like to enforce the constraints \mathbf{c}_l for future predictions, then these would need to be included in the posterior density defined in Equation (9). Consequently, future (FG or CG) states would need to be inferred from this augmented posterior and an enlarged inference process is required for predictions.

Algorithm 1: Prediction - Algorithm

- Result:** Sample of $\mathbf{x}_{(T+P)}^{(i)}$
Data: $q_{\phi}(\mathbf{X}_T), \boldsymbol{\theta}_{MAP}$
- 1 Sample from $q_{\phi}(\mathbf{X}_T^{(i)})$;
 - 2 **while** *Time-step* ($T + P$) *not reached* **do**
 - 3 | Sample from the CG evolution law in Equation (4);
 - 4 **end**
 - 5 Sample from $p_{cf}(\mathbf{x}_{(T+P)} | \mathbf{X}_{(T+P)}, \boldsymbol{\theta}_{MAP})$
-

3 NUMERICAL ILLUSTRATIONS

We demonstrate the capabilities of the proposed framework by applying it to a high-dimensional system of stochastically moving particles.

3.1 FG model

For the simulations presented in this section, we used $d_f = 250 \times 10^3$ particles, which, at each microscopic time step $\delta t = 2.5 \times 10^{-3}$ performed random, non-interacting, jumps of size $\delta s = \frac{1}{640}$, either to the left with probability $p_{left} = 0.1875$ or to the right with probability $p_{right} = 0.2125$. The positions were restricted to a domain of $[-1, 1]$ with periodic boundary conditions. It is well-known [21] that in the limit (i.e. $d_f \rightarrow \infty$) the particle density $\rho(s, t)$ can be described with an advection-diffusion PDE with diffusion constant $D = (p_{left} + p_{right}) \frac{\delta s^2}{2\delta t}$ and velocity $v = (p_{right} - p_{left}) \frac{\delta s}{\delta t}$:

$$\frac{\partial \rho}{\partial t} + v \frac{\partial \rho}{\partial s} = D \frac{\partial^2 \rho}{\partial s^2}, \quad s \in (-1, 1).. \quad (16)$$

3.2 CG model specifications

The CG model relates to a discretization of the particle density into $d_c = 25$ equally-sized bins at each coarse time step. The nature of the CG variables \mathbf{X}_t gives rise to a multinomial for the coarse-to-fine density p_{cf} (section 2.2) i.e.:

$$p_{cf}(\mathbf{x}_t | \mathbf{X}_t) = \frac{d_f!}{m_1(\mathbf{x}_t)! m_2(\mathbf{x}_t)! \dots m_{d_c}(\mathbf{x}_t)!} \prod_{j=1}^{d_c} X_{t,j}^{m_j(\mathbf{x}_t)}, \quad (17)$$

where $m_j(\mathbf{x}_t)$ is the number of particles in bin j . We assume that, given the CG state \mathbf{X}_t , the coordinates of the particles \mathbf{x}_t are conditionally independent. This does not imply that they move independently nor that they cannot exhibit coherent behavior [22]. The consequence of Equation (17) is that for this example no parameters need to be learned for p_{cf} .

For the transition law (section 2.3), we assume a coarse time step of $\Delta t = 4$ and employed a two-layered fully connected neural network $NN(\cdot)$ with ReLU activation functions. Each layers consisted of 25 neurons. We enforce conservation of mass, using the following constraint at each time step l :

$$c_l(\mathbf{X}_l) = 1 - \sum_{j=1}^{d_c} X_{l,j} = 0, \quad l = 0, 1, \dots \quad (18)$$

These are complemented by the virtual observables presented earlier and with $\sigma_c^2 = 10^{-9}$ (Equation (7)).

For the family of variational distributions $q_\phi(\mathbf{X}_{(0:T)}^{(i)})$ and since $X_{t,j}^{(i)} > 0, \forall j, t$, we employed multivariate lognormals with a diagonal covariance matrices i.e. we assume $X_{t,j}^{(i)}$ are a posteriori independent. The mean and covariance matrix of the underlying Gaussians for each sequence

i become part of the parameters ϕ with respect to which the ELBO is maximized (see Section 2.5). We note that it would also be possible to use an amortized formulation and explicitly account for the dependence on the data values by employing a neural network for both mean and covariance with the time sequence as an input.

3.3 Results




Figure 1: Particle density: Inferred and predicted posterior mean (bottom) in comparison with the ground truth (top). The red line divides inferred quantities from predicted ones.

We employed $n = 64$ time sequences with $T = 9$ for training and applied our framework in order to infer the unobserved CG states but more importantly the model parameters in right-hand side of the CG dynamics.

In Figure 1 we compare the true particle density with the one predicted by the trained CG model for one illustrative time sequence. We note that the latter is computed by reconstructing the x_t futures. The trained model is able to accurately track first-order statistics well into the future for many more time steps than those contained in the training data.

A more detailed view of the predictive estimates with snapshots of the particle density at selected time instances is presented in Figure 2 and 3 where the predictive posterior mean but also the associated uncertainty is displayed. Inferred as well as predicted particle densities

match accurately the ground-truth and reasonable uncertainty bounds are computed.




Figure 2: Inferred particle density profiles at $t = 0, 5\Delta t, 9\Delta t$ (from left to right).




Figure 3: Predicted particle density profiles at $t = 15\Delta t, 20\Delta t, 25\Delta t, 35\Delta t$ (from left to right and top to bottom).

Finally, in Figure 4, the mass constraint is depicted for inferred as well as predicted particle densities and good agreement with the target value ($= 1$) is observed. This result is particularly important as it demonstrates that the virtual observables were able to find *CG* state variables that agree with an a priori given physical constraint and additionally a transition law has been learned that is able to automatically satisfy the constraint in the future.

4 CONCLUSIONS

We combined a probabilistic generative model with physical constraints and deep neural networks in order to obtain a framework for the automated discovery of coarse-grained variables




Figure 4: Mass based on inferred and predicted particle densities.

and dynamics based on fine-grained simulation data. The FG simulation data are augmented in a fully Bayesian fashion by virtual observables that enable the incorporation of physical constraints at the CG level. These could be for instance conservation laws that are available when CG variables have physical meaning. Deviations from such conservation laws would invalidate predictions. As a result of augmenting the training data with domain knowledge, the model proposed can learn from Small Data (i.e. shorter and fewer FG time-sequences) which is a crucial advantage in multiscale settings where the simulation of the FG dynamics is computationally very expensive.

Our approach learns simultaneously a coarse-to-fine mapping and a transition law for the coarse-grained dynamics by employing probabilistic inference tools for the latent variables and model parameters. Deep neural networks can be used in both of these components in order to endow great expressiveness and flexibility.

The model proposed was successfully tested on a coarse-graining task which involved stochastic particle dynamics. In the example presented, the method was able to accurately predict particle densities at time steps not contained in the training data. Moreover, as it is able to reconstruct the entire FG state vector at any future time instant, it is capable of producing predictions of any FG observable of interest as well as quantify the associated predictive uncertainty.

A shortcoming of presented framework is that the CG dynamics are not fully interpretable and long-term stability is not guaranteed. These limitations have been addressed in [23] where an additional layer of latent variables was employed that ensured the discovery of stable CG dynamics but also promoted the identification of slow-varying processes that are most predictive of the system’s long-term evolution.

REFERENCES

- [1] D.Givon, R.Kupferman and A.Stuart, Extracting Macroscopic Dynamics: Model Problems and Algorithms, Nonlinearity, 2004.
- [2] Z. Ghahramani, Probabilistic machine learning and artificial intelligence, Nature, 2015
- [3] Y. LeCun, Y. Bengio and G. Hinton, Deep Learning, Nature, 2015

- [4] P.-S.Koutsourelakis, N.Zabaras and M. Girolami, Big data and predictive computational modeling, *Journal of Computational Physics*, 2016.
- [5] M.Alber, A.Tepole, W.Cannon, S.De, S.Dura-Bernal, K.Garikipati, G. Karniadakis, W.Lytton, P.Perdikaris, L.Petzold and E.Kuhl, Integrating machine learning and multi-scale modeling - perspectives, challenges, and opportunities in the biological, biomedical, and behavioral sciences, *NPJ digital medicine*, 2019
- [6] S. Kaltenbach and P.-S. Koutsourelakis, Incorporating physical constraints in a deep probabilistic machine learning framework for coarse-graining dynamical systems, *Journal of Computational Physics*, 2020
- [7] P.Stinis, T.Hagge, A.M.Tartakovsky and E.Yeung: Enforcing constraints for interpolation and extrapolation in generative adversarial networks, *Journal of Computational Physics*, 2019
- [8] 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* 113 (15) (2016) 3932–3937.
- [9] R.Chen, Y.Rubanova, J.Bettencourt and D.Duvenaud, Neural ordinary differential equations, *Advances in neural information processing systems*, 2018
- [10] X. Li, T.-K. Wong, R. TQ Chen and D. Duvenaud, Scalable gradients for stochastic differential equations, *arXiv preprint 2001.01328*, 2020
- [11] M. Raissi and G.E. Karniadakis, Hidden physics models: Machine learning of nonlinear partial differential equation, *Journal of Computational Physics* 357 (2018) 125–142.
- [12] P.-S. Koutsourelakis and I. Bilonis, Scalable Bayesian Reduced-Order Models for Simulating High-Dimensional Multiscale Dynamical Systems, *Multiscale Modeling and Simulation*, 2011
- [13] H.Mori, Transport, collective motion, and brownian motion, *Progress of theoretical physics*, 33(3):423–455, 1965.
- [14] R.Zwanzig, Nonlinear generalized langevin equations, *Journal of Statistical Physics*, 9(3):215–220, 1973.
- [15] M. Schöberl, N. Zabaras and P.-S. Koutsourelakis, Predictive coarse-graining, *Journal of Computational Physics* 333 (2017) 49–77.
- [16] M. Rixner and P.-S. Koutsourelakis, A probabilistic generative model for semi-supervised training of coarse-grained surrogates and enforcing physical constraints through virtual observables, *Arxiv Preprint 2006.01789*, 2020
- [17] M.Hoffman, D.Blei, C. Wang and J.Paisley, Stochastic variational inference, *The Journal of Machine Learning Research*, 14(1):1303–1347, 2013.
- [18] C. Bishop, *Pattern Recognition and Machine Learning*. Springer, 2006.
- [19] D.Kingma and M.Welling, Auto-Encoding Variational Bayes, *International Conference on*

- Learning Representations (ICLR), 2014.
- [20] D.Kingma and J.Ba, Adam: A method for stochastic optimization, arXiv preprint 1412.6980, 2014
 - [21] G.Cottet and P.Koumoutsakos, Vortex Methods: Theory and Practice, Cambridge University Press, March 2000. ISBN 978-0-521-62186-1.
 - [22] L. Felsberger and P.-S. Koutsourelakis, Physics-constrained, data-driven discovery of coarse-grained dynamics, Communications in Computational Physics, 2019
 - [23] S. Kaltenbach and P.-S. Koutsourelakis, Physics-aware, probabilistic model order reduction with guaranteed stability, International Conference on Learning Representations (ICLR), 2021