



ELSEVIER

Contents lists available at ScienceDirect

Journal of Computational Physics

www.elsevier.com/locate/jcp



Predictive coarse-graining

Markus Schöberl^a, Nicholas Zabaras^{b,c}, Phaedon-Stelios Koutsourelakis^{a,*}^a Continuum Mechanics Group, Technical University of Munich, Boltzmannstraße 15, 85748 Garching, Germany^b Institute for Advanced Study, Technical University of Munich, Lichtenbergstraße 2a, 85748 Garching, Germany^c Department of Aerospace and Mechanical Engineering, University of Notre Dame, 365 Fitzpatrick Hall, Notre Dame, IN 46556, USA

ARTICLE INFO

Article history:

Received 26 May 2016

Received in revised form 28 September 2016

Accepted 14 October 2016

Available online 21 December 2016

Keywords:

Coarse-graining

Generative models

Bayesian

Uncertainty quantification

SPC/E water

Lattice systems

ABSTRACT

We propose a data-driven, coarse-graining formulation in the context of equilibrium statistical mechanics. In contrast to existing techniques which are based on a fine-to-coarse map, we adopt the opposite strategy by prescribing a *probabilistic coarse-to-fine* map. This corresponds to a directed probabilistic model where the coarse variables play the role of latent generators of the fine scale (all-atom) data. From an information-theoretic perspective, the framework proposed provides an improvement upon the relative entropy method [1] and is capable of quantifying the uncertainty due to the information loss that unavoidably takes place during the coarse-graining process. Furthermore, it can be readily extended to a fully Bayesian model where various sources of uncertainties are reflected in the posterior of the model parameters. The latter can be used to produce not only point estimates of fine-scale reconstructions or macroscopic observables, but more importantly, predictive posterior distributions on these quantities. Predictive posterior distributions reflect the confidence of the model as a function of the amount of data and the level of coarse-graining. The issues of model complexity and model selection are seamlessly addressed by employing a hierarchical prior that favors the discovery of sparse solutions, revealing the most prominent features in the coarse-grained model. A flexible and parallelizable Monte Carlo – Expectation–Maximization (MC-EM) scheme is proposed for carrying out inference and learning tasks. A comparative assessment of the proposed methodology is presented for a lattice spin system and the SPC/E water model.

© 2016 Elsevier Inc. All rights reserved.

1. Introduction

Molecular dynamics simulations [2] are nowadays commonplace in physics, chemistry, and engineering and represent one of the most reliable tools in the analysis of complex processes and the design of new materials [3–5]. Direct simulations are hampered by the gigantic number of degrees of freedom, complex, potentially long-range and high-order interactions, and as a result, are limited to small spatio-temporal scales with current and foreseeable computational resources.

An approach towards making complex simulations practicable over extended time and space scales is coarse-graining (CG) [6]. Coarse-graining methods attempt to summarize the atomistic detail in much fewer degrees of freedom which in turn lead to shorter simulation times, with potentially larger time-steps and enable the analysis of systems that occupy

* Corresponding author.

E-mail addresses: m.schoeberl@tum.de (M. Schöberl), nzabaras@gmail.com (N. Zabaras), p.s.koutsourelakis@tum.de (P.-S. Koutsourelakis).URLs: <http://www.zabarass.com> (N. Zabarass), <http://www.contmech.mw.tum.de> (P.-S. Koutsourelakis).

larger spatial domains. Furthermore, from a reductionist's point of view, they can provide insight into the fundamental components or processes associated with the macroscopic behavior and properties of molecular ensembles.

A systematic strategy towards coarse-graining is offered in the context of free-energy computation methods [7,8]. Nevertheless, their primary goal is to escape deep, free-energy wells and are generally limited to a relatively small number of CG variables. A mathematically rigorous approach to coarse-graining lattice systems and a rich set of multi-level, adaptive algorithms for equilibrium and nonequilibrium settings, has been developed in [9–14]. Inversion-based methods such as the Direct or Iterative Boltzmann Inversion [15,16] and Inverse Monte Carlo [17], represent a popular strategy where the parameters of the CG model are adjusted to reproduce macroscopic observables [18]. Molecular Renormalization Group CG [19] is founded upon the ideas first presented in [20] and is based on matching correlators, obtained from atomistic and coarse-grained simulations, for observables that explicitly enter the coarse-grained Hamiltonian. Data-driven, variational CG methods such as Multiscale CG [21,22], Relative Entropy [1], Ultra GG [23], offer a rigorous way of learning CG models by approximating the Potential of Mean Force (PMF) [24] with respect to the CG variables on the basis of appropriate functionals.

It is obvious that unless there are known redundancies in the all-atom or fine-grained (FG) description, any coarse-graining scheme will result in information loss [25,26]. A manifestation of this can be seen if one attempts to reconstruct the microscopic, FG configurations from the CG states [27,28]. Discrepancies will appear not only because the CG statistics are not captured correctly, but because the CG variables do not encode all the details needed to reproduce the FG picture. Despite this, predictions generated by existing CG schemes are always in the form of *point estimates* that do not reflect any of the predictive uncertainty which the aforementioned information loss induces. It is also reasonable to expect that this information loss increases the larger the difference between the dimension of fine and coarse descriptions becomes. Nevertheless given two competing CG descriptions of the same dimension, it is unlikely that both will capture the FG picture equally well. The discovery of a good set of CG variables (analogous to finding good reaction coordinates or collective variables in free energy computations [29]) is, on one hand, a function of the macroscopic quantities of interest but more importantly of the complex structure of inter-dependencies in the FG model.

The starting point of all CG schemes is the prescription of the coarse variables through a many-to-one, *fine-to-coarse* map. Such maps are dictated by the analysis objectives but also by physical insight on which FG features might be important [30]. For example several atoms/molecules can be lumped into a single, effective, pseudo-molecule with coordinates defined by considering the center of mass. A central component of the present work is the implicit definition of the CG variables through a *coarse-to-fine* map. This is achieved by a *probabilistic generative model* that treats the CG degrees of freedom as latent variables and explicitly quantifies the uncertainty in the reconstruction of the FG states from the CG description. The model is complemented with a distribution for the CG variables. Both densities are parametrized and the optimal values are determined on the basis of an information-theoretic objective (e.g. minimizing a Kullback–Leibler divergence as in [1]) which is shown to be a special case of a more general, Bayesian framework. The latter offers a critical advantage over existing techniques as it enables the prediction of macroscopic observables not only in the form of point estimates, but by providing whole distributions. These reflect the uncertainty due the aforementioned information loss as well as the fact that finite amounts of training data were used.

The emphasis on this amplified predictive ability of the proposed framework is the reason behind the title chosen for the present paper *predictive coarse-graining* (PCG). The Bayesian framework advocated offers a superior setting for model selection. We make use of hierarchical prior models that promote the discovery of a sparse set of features in the aforementioned model components. This enables the search to be carried out over a very large set of feature functions for the CG potential which naturally amplifies the expressivity of the model [30]. We note that a Bayesian framework towards uncertainty quantification for force field parameters in molecular dynamics was introduced in [31,32]. Other Bayesian formulations of coarse-graining problems using macroscopic observables were presented in [33,34] where also the issues of model calibration and validation were discussed.

The structure of the rest of the paper is as follows. Section 2 presents the basic model components, compares them with other CG schemes (primarily the relative entropy method), provides details on the exponential family of distributions employed for which uniqueness of solution can be proven and discusses in detail algorithmic and computational aspects. Numerical evidence of the capabilities of the proposed framework is provided in Section 3 where coarse-graining efforts for an Ising lattice system as well as for the SPC/E water model are documented. In all numerical examples, we report results on the *predictive uncertainty* as a function of the level of coarse graining, and the amount of data available. Finally, Section 4, summarizes the main contributions and discusses natural extensions of the proposed framework.

2. Methodology

This section introduces the notational conventions adopted and presents the proposed modeling and computational frameworks. We frequently draw comparisons with the relative entropy method introduced in [1] and further expanded and studied in [35,36] in order to shed light on the aspects related to information loss and to emphasize the need for quantifying the resulting uncertainty in the predictions.

2.1. Equilibrium statistical mechanics

We consider molecular ensembles in equilibrium described by an n_f -dimensional vector denoted by $\mathbf{x} \in \mathcal{M}_f \subset \mathbb{R}^{n_f}$. This generally consists of the coordinates of the atoms which follow the Boltzmann–Gibbs density¹:

$$p_f(\mathbf{x}|\beta) = \frac{\exp\{-\beta U_f(\mathbf{x})\}}{Z_f(\beta)}, \quad (1)$$

where $U_f(\mathbf{x})$ is the all-atom (fine-grained) potential, $\beta = \frac{1}{k_b T}$ where k_b is the Boltzmann constant and T is the temperature, and $Z_f(\beta)$ is the normalization constant (partition function) given by:

$$Z_f(\beta) = \int_{\mathcal{M}_f} \exp\{-\beta U_f(\mathbf{x})\} d\mathbf{x}. \quad (2)$$

In the following, we assume that the temperature T (or equivalently β) is constant as it is commonly done in coarse-graining literature, even though it is generally of interest to derive coarse-grained descriptions that are suitable for all (or at least a wide range) of temperatures [30]. In this setting and in order to simplify the notation, we drop the temperature dependence.

If $a(\mathbf{x}) : \mathcal{M}_f \rightarrow \mathbb{R}$ denotes an observable (e.g. magnetization in Ising models), then the corresponding macroscopic properties can be computed as an expectation with respect to $p_f(\mathbf{x})$ as follows:

$$\mathbb{E}_{p_f(\mathbf{x})}[a(\mathbf{x})] = \int_{\mathcal{M}_f} a(\mathbf{x}) p_f(\mathbf{x}) d\mathbf{x}. \quad (3)$$

Such expectations are (approximately) computed using long and cumbersome simulations as explained in the introduction e.g. by a long MCMC run [37]. Our goal is two-fold. Firstly, to construct a coarse-grained description of the system that would be easier and faster to simulate, and secondly to use this in order to predict expectations of any observable as in Eq. (3). A distinguishing aspect of the proposed PCG framework is that we also compute quantitative metrics of the predictive uncertainty in those estimates. At a third level, one would also want the coarse-grained description to provide a decomposition of the original, all-atom ensemble into physically interpretable terms and interactions. We defer such a discussion on how the proposed model can achieve this goal for the conclusions.

We denote by \mathbf{X} the coarse-grained variables and assume that they take values in $\mathcal{M}_c \subset \mathbb{R}^{n_c}$. It is obviously desirable that $n_c \ll n_f$. Let also $U_c(\mathbf{X})$ denote the potential associated with \mathbf{X} and $p_c(\mathbf{X})$ the corresponding density:

$$p_c(\mathbf{X}) = \frac{\exp\{-\beta U_c(\mathbf{X})\}}{Z_c}, \quad (4)$$

with the normalization constant,

$$Z_c = \int_{\mathcal{M}_c} \exp\{-\beta U_c(\mathbf{X})\} d\mathbf{X}. \quad (5)$$

In existing coarse-graining formulations, the coarse variables \mathbf{X} are defined using a restriction, fine-to-coarse map $\mathcal{R} : \mathcal{M}_f \rightarrow \mathcal{M}_c$ i.e. $\mathbf{X} = \mathcal{R}(\mathbf{x})$. As this is generally a many-to-one map, it is not invertible [36]. If the observables of interest actually depend on \mathbf{X} i.e. if $a(\mathbf{x}) = A(\mathcal{R}(\mathbf{x})) = A(\mathbf{X})$, then one can readily show that it suffices that $p_c(\mathbf{X})$ is equal to the marginal of \mathbf{X} with respect to $p_f(\mathbf{x})$, or equivalently that $U_c(\mathbf{X}) = U_c^{\text{opt}}(\mathbf{X})$ where:

$$U_c^{\text{opt}}(\mathbf{X}) = -\beta^{-1} \log \int \delta(\mathbf{X} - \mathcal{R}(\mathbf{x})) p_f(\mathbf{x}) d\mathbf{x}. \quad (6)$$

That is the coarse-scale potential $U_c(\mathbf{x})$ coincides with the potential of mean-force of \mathbf{X} . This is a consequence of the following equalities:

$$\begin{aligned} \mathbb{E}_{p_f}[a] &= \int_{\mathcal{M}_f} a(\mathbf{x}) p_f(\mathbf{x}) d\mathbf{x} \\ &= \int_{\mathcal{M}_f} A(\mathcal{R}(\mathbf{x})) p_f(\mathbf{x}) d\mathbf{x} \end{aligned}$$

¹ In the following, we assume all probability measures are absolutely continuous with the Lebesgue measure and therefore work exclusively with the corresponding probability density functions.

$$\begin{aligned}
&= \int_{\mathcal{M}_f} \left(\int_{\mathcal{M}_c} A(\mathbf{X}) \delta(\mathbf{X} - \mathcal{R}(\mathbf{x})) d\mathbf{X} \right) p_f(\mathbf{x}) d\mathbf{x} \\
&= \int_{\mathcal{M}_c} A(\mathbf{X}) \left(\int_{\mathcal{M}_f} \delta(\mathbf{X} - \mathcal{R}(\mathbf{x})) p_f(\mathbf{x}) d\mathbf{x} \right) d\mathbf{X} \\
&= \int_{\mathcal{M}_c} A(\mathbf{X}) p_c(\mathbf{X}) d\mathbf{X}.
\end{aligned}$$

Nevertheless, even if one is able to compute or approximate sufficiently well $U_c^{\text{opt}}(\mathbf{X})$, there is no guarantee that expectations of other observables that do not solely depend on \mathbf{X} can be accurately computed. Consistent reconstructions of the all-atom configurations \mathbf{x} , given \mathbf{X} samples from $p_c(\mathbf{X})$, can be obtained from the conditional:

$$p_{\mathcal{R}}(\mathbf{x}|\mathbf{X}) = \frac{\delta(\mathbf{X} - \mathcal{R}(\mathbf{x}))}{Z_{\mathcal{R}}(\mathbf{X})}, \quad (7)$$

i.e. the uniform density on the manifold in \mathcal{M}_f implied by the map \mathcal{R} ,² where:

$$Z_{\mathcal{R}}(\mathbf{X}) = \int \delta(\mathbf{X} - \mathcal{R}(\mathbf{x})) d\mathbf{x}. \quad (8)$$

Given a coarse-grained potential U_c (not necessarily the optimal as in Eq. (6)) and the density $p_c(\mathbf{X})$ in Eq. (4), the corresponding reconstruction density of the all-atom description consistent with the map $p_{\mathcal{R}}(\mathbf{x}|\mathbf{X})$ (Eq. (7)) is given by:

$$\begin{aligned}
p_{\mathcal{R}}(\mathbf{x}) &= \int p_{\mathcal{R}}(\mathbf{x}|\mathbf{X}) p_c(\mathbf{X}) d\mathbf{X} \\
&= \int \frac{\delta(\mathbf{X} - \mathcal{R}(\mathbf{x}))}{Z_{\mathcal{R}}(\mathbf{X})} p_c(\mathbf{X}) d\mathbf{X} \\
&= \frac{p_c(\mathcal{R}(\mathbf{x}))}{Z_{\mathcal{R}}(\mathcal{R}(\mathbf{x}))}.
\end{aligned} \quad (9)$$

We note that in the context of the relative entropy method [1], which like ours, is data-driven and has an information-theoretic underpinning, the goal is to identify the U_c (within a certain class) that brings $p_{\mathcal{R}}(\mathbf{x})$ (Eq. (9)) as close as possible to the reference, FG density $p_f(\mathbf{x})$ (Eq. (1)). For that purpose the Kullback–Leibler (KL) divergence [39] $\text{KL}(p_f(\mathbf{x})||p_{\mathcal{R}}(\mathbf{x}))$ is employed as the objective which, based on Eq. (9), is given by:

$$\begin{aligned}
0 \leq \text{KL}(p_f(\mathbf{x})||p_{\mathcal{R}}(\mathbf{x})) &= - \int p_f(\mathbf{x}) \log \frac{p_{\mathcal{R}}(\mathbf{x})}{p_f(\mathbf{x})} d\mathbf{x} \\
&= -\mathbb{E}_{p_f(\mathbf{x})}[\log p_c(\mathcal{R}(\mathbf{x}))] + \mathbb{E}_{p_f(\mathbf{x})}[\log Z_{\mathcal{R}}(\mathcal{R}(\mathbf{x}))] - H(p_f),
\end{aligned} \quad (10)$$

where $H(p_f)$ is the entropy of $p_f(\mathbf{x})$, which is independent of U_c and can be ignored in the minimization. As it has been identified in several investigations [35,36,38], while the first term can be reduced by adjusting U_c (it can be shown that the minimum is attained when $U_c(\mathbf{X}) = U_c^{\text{opt}}(\mathbf{X})$), the second term is fixed once the restriction map \mathcal{R} that defines the coarse-grained variables has been selected. It represents a constant penalty reflecting the information loss that takes place due to the coarse-grained (and generally lower-dimensional) description adopted. Our goal is to reduce this component of information loss.

2.2. Probabilistic generative model

We propose a *probabilistic, generative model* [40] in which the coarse description is treated as a latent (hidden) state. In particular, we define a *joint density* $\bar{p}(\mathbf{X}, \mathbf{x})$ for \mathbf{X} and \mathbf{x} as follows:

$$\bar{p}(\mathbf{X}, \mathbf{x}) = p_{\text{cf}}(\mathbf{x}|\mathbf{X}) p_c(\mathbf{X}). \quad (11)$$

This consists of two components i.e.:

- (i) a density $p_c(\mathbf{X})$ describing the statistics of the coarse-grained description \mathbf{X} ,
- (ii) a **probabilistic, coarse-to-fine mapping** implied by the conditional density $p_{\text{cf}}(\mathbf{x}|\mathbf{X})$.

² In [38] this is further generalized by introducing an additional, weighting density.

We discuss the form and parametrization of the aforementioned densities in the sequel. We emphasize at this stage the different definition of the coarse-grained variables as latent generators that give rise to the observables through the probabilistic *lifting* operator implied by p_{cf} [9], in contrast to the restriction operators employed in other schemes explained previously. Such mappings can take various forms (e.g. local or global, linear or nonlinear) and can be extended to many hierarchical levels, as it will be shown. Understanding the meaning of the latent variables can only be done through the prism of this generative mapping. According to this, each FG configuration $\mathbf{x}^{(i)}$ is generated as follows:

- Draw a CG configuration $\mathbf{X}^{(i)}$ from $p_c(\mathbf{X})$.
- Draw $\mathbf{x}^{(i)}$ from $p_{cf}(\mathbf{x}|\mathbf{X}^{(i)})$.

As we will show, an advantage of the proposed framework is that it readily provides a (predictive) probability density for the observables of interest. The marginal density of the FG description \mathbf{x} is given from Eq. (11) by integrating out \mathbf{X} :

$$\bar{p}_f(\mathbf{x}) = \int_{\mathcal{M}_c} p_{cf}(\mathbf{x}|\mathbf{X}) p_c(\mathbf{X}) d\mathbf{X}. \quad (12)$$

Suppose the aforementioned component densities are parametrized by $\theta = (\theta_c, \theta_{cf})$ i.e. $p_c(\mathbf{X}|\theta_c)$ and $p_{cf}(\mathbf{x}|\mathbf{X}, \theta_{cf})$, and we attempt to minimize the KL-divergence between the reference density $p_f(\mathbf{x})$ and the marginal $\bar{p}_f(\mathbf{x}|\theta)$ implied by the generative model proposed :

$$\begin{aligned} \text{KL}(p_f(\mathbf{x})||\bar{p}_f(\mathbf{x}|\theta)) &= - \int_{\mathcal{M}_f} p_f(\mathbf{x}) \log \frac{\bar{p}_f(\mathbf{x}|\theta)}{p_f(\mathbf{x})} d\mathbf{x} \\ &= - \int p_f(\mathbf{x}) \log \bar{p}_f(\mathbf{x}|\theta) d\mathbf{x} + \int p_f(\mathbf{x}) \log p_f(\mathbf{x}) d\mathbf{x}. \end{aligned} \quad (13)$$

This is equivalent to *maximizing* $\int p_f(\mathbf{x}) \log \bar{p}_f(\mathbf{x}|\theta) d\mathbf{x}$ which, given samples $\{\mathbf{x}^{(i)}\}_{i=1}^N$ from $p_f(\mathbf{x})$, is approximated by the *log-likelihood* of $\bar{p}_f(\mathbf{x}|\theta)$ ³:

$$\begin{aligned} \mathcal{L}(\theta) &= \sum_{i=1}^N \log \bar{p}_f(\mathbf{x}^{(i)}|\theta) \\ &= \sum_{i=1}^N \log \left(\int p_{cf}(\mathbf{x}^{(i)}|\mathbf{X}^{(i)}, \theta_{cf}) p_c(\mathbf{X}^{(i)}|\theta_c) d\mathbf{X}^{(i)} \right). \end{aligned} \quad (14)$$

We note in the expression above that we associate a latent, coarse configuration $\mathbf{X}^{(i)}$ to each sample $\mathbf{x}^{(i)}$ which is effectively its pre-image. More importantly, the objective in the aforementioned expression accounts for both the density of the coarse-grained description as well as the reconstruction (lifting) of the all-atom configuration from the (latent) coarse-grained one. Maximizing $\mathcal{L}(\theta)$ naturally leads to the Maximum Likelihood estimate θ_{MLE} .

Furthermore the interpretation of the objective as the log-likelihood makes the progression into Bayesian formulations much more straightforward. If for example we define a prior density $p(\theta)$ then maximizing:

$$\arg \max_{\theta} \{ \mathcal{L}(\theta) + \log p(\theta) \}, \quad (15)$$

is equivalent to obtaining a Maximum a Posteriori (MAP) estimate θ_{MAP} [41]. The next step from point estimates for the model parameters is of course obtaining the full posterior $p(\theta|\mathbf{x}^{(1:N)})$ using Bayes formula as:

$$\begin{aligned} p(\theta|\mathbf{x}^{(1:N)}) &\propto p(\mathbf{x}^{(1:N)}|\theta) p(\theta) \\ &\propto e^{\mathcal{L}(\theta)} p(\theta) \\ &\propto \prod_{i=1}^N \left(\int p_{cf}(\mathbf{x}^{(i)}|\mathbf{X}^{(i)}, \theta_{cf}) p_c(\mathbf{X}^{(i)}|\theta_c) d\mathbf{X}^{(i)} \right) p(\theta). \end{aligned} \quad (16)$$

The aforementioned relationship can be concretely represented in the form of a directed graphical model as depicted in Fig. 1.

We discuss a strategy for approximating this posterior in the next subsections. It is more important to emphasize at this stage that given this posterior, we can produce not just point estimates of the expectation of any observable $a(\mathbf{x})$, but also

³ This result can be obtained (up to $1/N$) by substituting $p_f(\mathbf{x})$ in Eq. (13) by the empirical measure $\frac{1}{N} \sum_{i=1}^N \delta(\mathbf{x} - \mathbf{x}^{(i)})$. The likelihood of N samples drawn from $p_f(\mathbf{x})$ is trivially $\prod_{i=1}^N \bar{p}_f(\mathbf{x}^{(i)}|\theta)$.

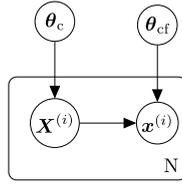


Fig. 1. Probabilistic graphical model representation.

compute its predictive posterior. For that purpose we make use of the predictive posterior $p(\mathbf{x}|\mathbf{x}^{(1:N)})$ of our model which is determined by marginalizing the latent variables \mathbf{X} and the model parameters θ :

$$\begin{aligned} p(\mathbf{x}|\mathbf{x}^{(1:N)}) &= \int p(\mathbf{x}, \mathbf{X}, \theta | \mathbf{x}^{(1:N)}) d\mathbf{X}d\theta \\ &= \int p(\mathbf{x}, \mathbf{X}|\theta, \mathbf{x}^{(1:N)}) p(\theta|\mathbf{x}^{(1:N)}) d\mathbf{X}d\theta. \end{aligned} \quad (17)$$

By replacing the joint density with the proposed generative model in Eq. (11), the predictive posterior $p(\mathbf{x}|\mathbf{x}^{(1:N)})$ becomes:

$$p(\mathbf{x}|\mathbf{x}^{(1:N)}) = \int p_{cf}(\mathbf{x}|\mathbf{X}, \theta_{cf}) p_c(\mathbf{X}|\theta_c) p(\theta|\mathbf{x}^{(1:N)}) d\mathbf{X}d\theta. \quad (18)$$

The latter can be used in place of the FG distribution $p_f(\mathbf{x})$ in Eq. (3), to obtain approximations to the expectation of any observable $a(\mathbf{x})$ as follows:

$$\begin{aligned} \mathbb{E}_{p_f(\mathbf{x})}[a(\mathbf{x})] &\approx \mathbb{E}_{p(\mathbf{x}|\mathbf{x}^{(1:N)})}[a(\mathbf{x})] \\ &= \int a(\mathbf{x}) p(\mathbf{x}|\mathbf{x}^{(1:N)}) d\mathbf{x} \\ &= \int a(\mathbf{x}) \left(\int p_{cf}(\mathbf{x}|\mathbf{X}, \theta_{cf}) p_c(\mathbf{X}|\theta_c) p(\theta|\mathbf{x}^{(1:N)}) d\mathbf{X}d\theta \right) d\mathbf{x} \\ &= \int \underbrace{\left(\int a(\mathbf{x}) p_{cf}(\mathbf{x}|\mathbf{X}, \theta_{cf}) p_c(\mathbf{X}|\theta_c) d\mathbf{X}d\mathbf{x} \right)}_{\hat{a}(\theta)} p(\theta|\mathbf{x}^{(1:N)}) d\theta \\ &= \int \hat{a}(\theta) p(\theta|\mathbf{x}^{(1:N)}) d\theta. \end{aligned} \quad (19)$$

The approximation in the first line reflects the quality of the model as well as the uncertainty arising from the finite data $\mathbf{x}^{(1:N)}$ that were used to calibrate it. This derivation suggests that $\hat{a}(\theta)$ represents the predictive estimate of the expectation of $a(\mathbf{x})$ for a given value θ of the model's parameters. Averaging over the posterior of the latter provides the expected (a posteriori) value of this quantity. More importantly though by propagating the (posterior) uncertainty of θ through $\hat{a}(\theta)$, one can readily obtain the predictive distribution of the observable. In the numerical examples we frequently plot such posterior statistics, usually in the form of credible intervals (see also A.1). Point estimates can be easily recovered if the analyst wishes to do so by employing for example the MAP (or MLE) estimate θ_{MAP} in the aforementioned equation i.e. if $p(\theta|\mathbf{x}^{(1:N)}) \equiv \delta(\theta - \theta_{MAP})$.

2.3. Inference and learning (point estimates)

This section is concerned with the computational aspects of training the proposed model. We pay particular attention to distributions in the exponential family for which the concavity of the maximum-likelihood problem can be analytically shown. Furthermore, we discuss strategies for parallelizing these tasks and improving the computational efficiency. We finally discuss particular prior specifications that are suitable for sparse feature recovery and model selection.

We begin our discussion with a strategy for obtaining point estimates for the model parameters θ by maximizing the log-likelihood (or the log-posterior) as given in Eq. (14) (or Eq. (15)). The difficulty in the optimization problem stems from the intractability of the log-likelihood due to the integration with respect to the latent variables $\mathbf{X}^{(i)}$ (except for trivial cases for p_c, p_{cf}). To address this we employ an *Expectation–Maximization* (EM) scheme [42,43] where MCMC is used to approximate the E-step (MCEM) [44] and stochastic approximations to handle the Monte Carlo noise in the gradient estimates of the M-Step [45,46]. The EM algorithm allows the maximization of the log-likelihood by circumventing the need for repeated evaluations of the aforementioned intractable integrals and normalization constants. To motivate the derivation, we note that for an arbitrary set of densities $q_i(\mathbf{X}^{(i)})$ we can construct lower bounds, denoted by $\mathcal{F}^{(i)}(q_i(\mathbf{X}^{(i)}), \theta)$, for each term in the sum that makes up the log-likelihood as follows:

$$\begin{aligned}
 \mathcal{L}(\boldsymbol{\theta}) &= \sum_{i=1}^N \log \left(\int p_{\text{cf}}(\mathbf{x}^{(i)} | \mathbf{X}^{(i)}, \boldsymbol{\theta}_{\text{cf}}) p_c(\mathbf{X}^{(i)} | \boldsymbol{\theta}_c) d\mathbf{X}^{(i)} \right) \\
 &= \sum_{i=1}^N \log \left(\int \frac{p_{\text{cf}}(\mathbf{x}^{(i)} | \mathbf{X}^{(i)}, \boldsymbol{\theta}_{\text{cf}}) p_c(\mathbf{X}^{(i)} | \boldsymbol{\theta}_c)}{q_i(\mathbf{X}^{(i)})} q_i(\mathbf{X}^{(i)}) d\mathbf{X}^{(i)} \right) \\
 &\geq \sum_{i=1}^N \underbrace{\left(\int q_i(\mathbf{X}^{(i)}) \log \frac{p_{\text{cf}}(\mathbf{x}^{(i)} | \mathbf{X}^{(i)}, \boldsymbol{\theta}_{\text{cf}}) p_c(\mathbf{X}^{(i)} | \boldsymbol{\theta}_c)}{q_i(\mathbf{X}^{(i)})} d\mathbf{X}^{(i)} \right)}_{:= \mathcal{F}^{(i)}(q_i(\mathbf{X}^{(i)}), \boldsymbol{\theta})} \\
 &= \sum_{i=1}^N \mathcal{F}^{(i)}(q_i(\mathbf{X}^{(i)}), \boldsymbol{\theta}) \\
 &= \mathcal{F}(\mathbf{q}(\mathbf{X}), \boldsymbol{\theta}),
 \end{aligned} \tag{20}$$

where $\mathbf{q}(\mathbf{X}) = \prod_{i=1}^N q_i(\mathbf{X}^{(i)})$, and the result in the third step is a consequence of Jensen’s inequality. We note that the optimal $q_i^{\text{opt}}(\mathbf{X}^{(i)})$ for each of the aforementioned terms is:

$$q_i^{\text{opt}}(\mathbf{X}^{(i)}) = q_i(\mathbf{X}^{(i)} | \mathbf{x}^{(i)}, \boldsymbol{\theta}) \propto p_{\text{cf}}(\mathbf{x}^{(i)} | \mathbf{X}^{(i)}, \boldsymbol{\theta}_{\text{cf}}) p_c(\mathbf{X}^{(i)} | \boldsymbol{\theta}_c), \tag{21}$$

i.e. the conditional posterior of the latent variables $\mathbf{X}^{(i)}$ given $\mathbf{x}^{(i)}$ and $\boldsymbol{\theta}$. This is optimal in the sense that the inequality becomes an equality [41] i.e.:

$$\mathcal{F}^{(i)}(q_i^{\text{opt}}(\mathbf{X}^{(i)}), \boldsymbol{\theta}) = \log \left(\int p_{\text{cf}}(\mathbf{x}^{(i)} | \mathbf{X}^{(i)}, \boldsymbol{\theta}_{\text{cf}}) p_c(\mathbf{X}^{(i)} | \boldsymbol{\theta}_c) d\mathbf{X}^{(i)} \right). \tag{22}$$

All other q_i ’s lead to suboptimal schemes that fall under the category of Variational Bayesian Expectation–Maximization (VB-EM, [47]). More importantly, the aforementioned derivation suggests an iterative algorithm where one alternates (until convergence) between the following two steps, i.e. at each iteration t :

E-step: Given the current estimate of $\boldsymbol{\theta} \equiv \boldsymbol{\theta}^{(t)}$, evaluate:

$$\mathcal{F}(\mathbf{q}^{\text{opt}, t}(\mathbf{X}), \boldsymbol{\theta}^{(t)}) = \sum_{i=1}^N \mathcal{F}^{(i)}(q_i^{\text{opt}, t}(\mathbf{X}^{(i)}), \boldsymbol{\theta}^{(t)}), \tag{23}$$

where $q_i^{\text{opt}, t}$ is given in Eq. (21) for $\boldsymbol{\theta} \equiv \boldsymbol{\theta}^{(t)}$.

M-step: Given the current $q_i^{\text{opt}, t}(\mathbf{X}^{(i)})$, find:

$$\begin{aligned}
 \boldsymbol{\theta}^{(t+1)} &= \arg \max_{\boldsymbol{\theta}} \sum_{i=1}^N \mathcal{F}^{(i)}(q_i^{\text{opt}, t}(\mathbf{X}^{(i)}), \boldsymbol{\theta}^{(t)}) \\
 &= \arg \max_{\boldsymbol{\theta}} \sum_{i=1}^N \left(\int q_i^{\text{opt}, t}(\mathbf{X}^{(i)}) \log \left(p_{\text{cf}}(\mathbf{x}^{(i)} | \mathbf{X}^{(i)}, \boldsymbol{\theta}_{\text{cf}}^{(t)}) p_c(\mathbf{X}^{(i)} | \boldsymbol{\theta}_c^{(t)}) \right) d\mathbf{X}^{(i)} \right).
 \end{aligned} \tag{24}$$

We discuss in detail each of the two steps.

- The E-step of the algorithm requires computing expectations with respect to the intractable distributions in Eq. (21). As it can be seen in Eq. (24) only the terms in $\mathcal{F}^{(i)}$ that depends on $\boldsymbol{\theta}$ needs to be computed which we approximate by a Monte Carlo estimator:

$$\begin{aligned}
 \int q_i^{\text{opt}, t}(\mathbf{X}^{(i)}) \log \left(p_{\text{cf}}(\mathbf{x}^{(i)} | \mathbf{X}^{(i)}, \boldsymbol{\theta}_{\text{cf}}^{(t)}) p_c(\mathbf{X}^{(i)} | \boldsymbol{\theta}_c^{(t)}) \right) d\mathbf{X}^{(i)} &\approx \\
 &\approx \frac{1}{m_t} \sum_{j=1}^{m_t} \left(\log p_{\text{cf}}(\mathbf{x}^{(i)} | \mathbf{X}_j^{(i)}, \boldsymbol{\theta}_{\text{cf}}^{(t)}) p_c(\mathbf{X}_j^{(i)} | \boldsymbol{\theta}_c^{(t)}) \right).
 \end{aligned} \tag{25}$$

The m_t samples used at each iteration t are drawn using MCMC from $q_i^{\text{opt}, t}(\mathbf{X}^{(i)})$. Compared to i.i.d. Monte Carlo samples, the use of MCMC introduces theoretical complications with regards to the stability and the error in the approximation [48,49]. A recent treatment of the convergence conditions for such schemes is contained in [50]. The obvious

error source arises from the bias in the MCMC samples which are *approximately* distributed according to the target density. In addition the samples generated are correlated. Such errors can be subdued by increasing the sample size m_t . Heuristically speaking, at the first few iterations t , even a crude estimate of the objective generally suffices to drive the parameter θ -updates toward the region of interest. As the EM iterations proceed, the number of samples should increase in order to zoom-in at the optimum and minimize the oscillatory behavior due to the noise in the estimates. Several strategies have been proposed to optimize m_t or even devise an automatic schedule by making use of error estimates [51–54]. In this work, we used a constant sample size i.e. $m_t = m, \forall t$ that we report in the numerical examples. We found through several cross-validation runs that this had no noticeable effect to the optima identified. We note finally that other Monte Carlo schemes can be utilized. One would expect that Importance Sampling [55], where previously generated samples are re-weighted and re-used, could be quite effective particularly when $\theta^{(t)}$ do not change much and the corresponding $q_i^{\text{opt}, t}$ are quite similar. A more potent alternative is offered by Sequential Monte Carlo schemes (SMC) [8,56] which combine the benefits of MCMC and Importance Sampling.

- The maximization of the lower bound with respect to θ is not analytically tractable even when a Monte Carlo approximation of the objective, as discussed previously, is used. For that purpose, we make use of a gradient ascent scheme that employs the partial derivatives of \mathcal{F} :

$$\begin{aligned} \mathcal{G}(\theta) &= \nabla_{\theta} \mathcal{F} = \sum_{i=1}^N \nabla_{\theta} \mathcal{F}^{(i)} \quad (= \sum_{i=1}^N \mathcal{G}^{(i)}(\theta)) \\ &= \sum_{i=1}^N \nabla_{\theta} \left(\int q_i^{\text{opt}, t}(\mathbf{X}^{(i)}) \log \left(p_{\text{cf}}(\mathbf{x}^{(i)} | \mathbf{X}^{(i)}, \theta_{\text{cf}}^{(t)}) p_c(\mathbf{X}^{(i)} | \theta_c^{(t)}) \right) d\mathbf{X}^{(i)} \right), \end{aligned} \quad (26)$$

where at each iteration t , each term $\mathcal{G}^{(i)}(\theta)$ is approximated by a Monte Carlo estimate (see discussion before) as:

$$\begin{aligned} \mathcal{G}^{(i)}(\theta) &= \nabla_{\theta} \int q_i^{\text{opt}, t}(\mathbf{X}^{(i)}) \log \left(p_{\text{cf}}(\mathbf{x}^{(i)} | \mathbf{X}^{(i)}, \theta_{\text{cf}}^{(t)}) p_c(\mathbf{X}^{(i)} | \theta_c^{(t)}) \right) d\mathbf{X}^{(i)} \\ &\approx \frac{1}{m_t} \sum_{j=1}^{m_t} \nabla_{\theta} \log \left(p_{\text{cf}}(\mathbf{x}^{(i)} | \mathbf{X}_j^{(i)}, \theta_{\text{cf}}^{(t)}) p_c(\mathbf{X}_j^{(i)} | \theta_c^{(t)}) \right) \\ &= \hat{\mathcal{G}}_t^{(i)}. \end{aligned} \quad (27)$$

The latter are used to update θ as follows⁴:

$$\theta^{t+1} = \theta^t + \eta_t \sum_{i=1}^N \hat{\mathcal{G}}_t^{(i)}. \quad (28)$$

The step sizes η_t are defined in the context of the Robbins–Monro scheme [45] which is designed to handle the unavoidable Monte Carlo noise in the gradient estimates. They should satisfy the following conditions [57]:

$$\sum_{t=1}^{\infty} \eta_t = +\infty, \quad \text{and} \quad \sum_{t=1}^{\infty} \eta_t^2 < \infty. \quad (29)$$

In this work, we employ [36]:

$$\eta_t = \frac{\alpha}{(A + t)^{\rho}}, \quad (30)$$

with $\rho \in (0.5, 1]$. The choice for the values α , ρ , and A is problem dependent and is explicitly given in Sections 3.1 and 3.2 for the Ising and water problems, respectively.

- We note finally that the gradient needed for the θ -updates, involves the sum of N independent terms, one for each datum (i.e. FG configuration) available. Apart from the obvious opportunity for parallelization that this offers, it also suggests that fine-scale data can be successively added. Hence the optimization can be initiated with a small number of data points N and the changes in the optimal θ identified can be monitored as more fine-scale data are generated/added to ensure that convergence is achieved with the smallest such effort. Another strategy for reducing the computational effort is to perform the E-step i.e. sample from $q_i^{\text{opt}, t}$ only for a subset of the data $i = 1, \dots, N$ at a time. While this has the potential of reducing the overall number of MCMC steps needed, convergence is still guaranteed [43].

⁴ As discussed in the seminal work of Neal and Hinton [43], more than one updates of θ per EM iteration can be performed.

2.4. Exponential family densities – uniqueness of solution

In order to provide some insight to the log-likelihood maximization, we consider the case of model densities belong to the exponential family [41,58]. As it will be shown in the numerical illustrations, this represents a very large set of flexible densities where by appropriate selection of the feature functions ϕ and ψ in the equations below one can capture interactions of various order (e.g. 2nd, 3rd) [36,38]. Such densities have the form:

$$p_c(\mathbf{X}|\theta_c) = \exp\{\theta_c^T \phi(\mathbf{X}) - A(\theta_c)\}, \tag{31}$$

and:

$$p_{cf}(\mathbf{x}|\mathbf{X}, \theta_{cf}) = \exp\{\theta_{cf}^T \psi(\mathbf{x}, \mathbf{X}) - B(\mathbf{X}, \theta_{cf})\}, \tag{32}$$

where $A(\theta_c)$ and $B(\mathbf{X}, \theta_{cf})$ are the log-partition functions given by:

$$\begin{aligned} A(\theta_c) &= \log \int e^{\theta_c^T \phi(\mathbf{X})} d\mathbf{X}, \\ B(\mathbf{X}, \theta_{cf}) &= \log \int e^{\theta_{cf}^T \psi(\mathbf{x}, \mathbf{X})} d\mathbf{x}. \end{aligned} \tag{33}$$

One can readily show that:

$$\begin{aligned} \frac{\partial A(\theta_c)}{\partial \theta_{c,k}} &= \langle \phi_k(\mathbf{X}) \rangle_{p_c(\mathbf{X}|\theta_c)}, \\ \frac{\partial^2 A(\theta_c)}{\partial \theta_{c,k} \partial \theta_{c,l}} &= \text{Cov}_{p_c(\mathbf{X}|\theta_c)}[\phi_k(\mathbf{X}), \phi_l(\mathbf{X})], \end{aligned} \tag{34}$$

and:

$$\begin{aligned} \frac{\partial B(\mathbf{X}, \theta_{cf})}{\partial \theta_{cf,k}} &= \langle \psi_k(\mathbf{x}, \mathbf{X}) \rangle_{p_{cf}(\mathbf{x}|\mathbf{X}, \theta_{cf})}, \\ \frac{\partial^2 B(\mathbf{X}, \theta_{cf})}{\partial \theta_{cf,k} \partial \theta_{cf,l}} &= \text{Cov}_{p_{cf}(\mathbf{x}|\mathbf{X}, \theta_{cf})}[\psi_k(\mathbf{x}, \mathbf{X}), \psi_l(\mathbf{x}, \mathbf{X})], \end{aligned} \tag{35}$$

where $\langle \cdot \rangle_p$ denotes the expectation with respect to the density p and $\text{Cov}_p[\cdot, \cdot]$ the covariance of the arguments with respect to p . Hence, for p_c and p_{cf} as above, the gradient of the objective \mathcal{F} in Eq. (24) is given by⁵:

$$\frac{\partial \mathcal{F}}{\partial \theta_{c,k}} = \sum_{i=1}^N \left(\langle \phi_k(\mathbf{X}^{(i)}) \rangle_{q_i(\mathbf{X}^{(i)})} - \langle \phi_k(\mathbf{X}) \rangle_{p_c(\mathbf{X}|\theta_c)} \right),$$

and

$$\frac{\partial \mathcal{F}}{\partial \theta_{cf,k}} = \sum_{i=1}^N \left(\langle \psi_k(\mathbf{x}^{(i)}, \mathbf{X}^{(i)}) \rangle_{q_i(\mathbf{X}^{(i)})} - \langle \psi_k(\mathbf{x}, \mathbf{X}^{(i)}) \rangle_{p_{cf}(\mathbf{x}|\mathbf{X}^{(i)}, \theta_{cf})q_i(\mathbf{X}^{(i)})} \right). \tag{36}$$

Furthermore, the Hessian is:

$$\begin{aligned} \frac{\partial^2 \mathcal{F}}{\partial \theta_{c,k} \partial \theta_{c,l}} &= -N \text{Cov}_{p_c(\mathbf{X}|\theta_c)}[\phi_k(\mathbf{X}), \phi_l(\mathbf{X})], \\ \frac{\partial^2 \mathcal{F}}{\partial \theta_{c,k} \partial \theta_{cf,l}} &= 0, \\ \frac{\partial^2 \mathcal{F}}{\partial \theta_{cf,k} \partial \theta_{cf,l}} &= - \sum_{i=1}^N \text{Cov}_{p_{cf}(\mathbf{x}|\mathbf{X}^{(i)}, \theta_{cf})q_i(\mathbf{X}^{(i)})}[\psi_k(\mathbf{x}, \mathbf{X}), \psi_l(\mathbf{x}, \mathbf{X})]. \end{aligned} \tag{37}$$

The block-diagonal Hessian is negative definite (at least when linearly independent feature functions are employed) which ensures that the objective is concave and has a unique maximum (whether arbitrary q_i are employed or q_i^{opt} as in Eq. (21)). We note also that Monte Carlo estimates of the Hessian can also be obtained and used in the θ -updates. These however tend to be more noisy than the gradients and special treatment is needed unless one is willing to generate large numbers of MCMC samples [36]. Finally, there is a wealth of stochastic approximation schemes that have been proposed and exhibit accelerated convergence [59–62].

⁵ We compare gradients of PCG with the relative entropy method in A.2.

2.5. Prior specification

The incorporation of priors for θ does not pose any computational difficulties as their contribution is additive (see Eq. (15)) to the log-likelihood and its partial derivatives. While priors for θ_{cf} , i.e. the parameters in the coarse-to-fine map, are unavoidably problem-dependent due to their special physical meaning, a more general strategy can be adopted for the θ_c , i.e. the parameters associated with the density of the coarse-grained variables \mathbf{X} . For exponential family distributions as in Eq. (31), each $\theta_{c,k}$ is associated with a feature function $\phi_k(\mathbf{X})$. As it will become apparent in the numerical examples, each of these feature functions encapsulates low- or high-order dependencies (or components thereof) between \mathbf{X} . It is obviously impossible to know a priori which of the $\phi(\mathbf{X})$ are relevant for a particular problem and how these depend on the dimension of \mathbf{X} or the coarse-to-fine probabilistic map p_{cf} . This underpins an important *model selection* issue that has been of concern in several coarse-graining studies [30,33,34,38]. One strategy to address this is to initiate the search with a small number of features $\phi(\mathbf{X})$ and progressively add more. These can be selected from a pool of candidates by employing appropriate criteria. In [8,63] for example, the feature function that causes the largest (expected) decrease (or increase) in the KL-divergence (or the log-likelihood) that we seek to minimize (or maximize), is added at each step. In this work, we adopt a different approach whereby *all* available $\phi(\mathbf{X})$ contained in the vocabulary of feature functions, are simultaneously considered. Consequently this leads to a vector of unknowns θ_c of very large dimension which not only impedes computations but can potentially lead to multiple local maxima, if the Hessian in Eq. (37) becomes semi-negative definite i.e. if linear dependencies between the selected $\phi(\mathbf{X})$ are present. More importantly though (at least when the number of data points N is small), it can obstruct the identification of the most salient features of the coarse-grained model which provide valuable physical insight [30].

To address this, we propose the use of sparsity-enforcing priors that are capable of identifying solutions in which only a (small) subset of θ_c are non-zero and therefore only the corresponding $\phi(\mathbf{X})$ are active [64,65]. A lot of the prior models that have been proposed along these lines can be readily cast in the context of *hierarchical Bayesian models* where *hyper-parameters* are introduced in the prior. In this work, we adopt the Automatic Relevance Determination (ARD, [66]) model which consists of the following:

$$p(\theta_c|\tau) \equiv \prod_k \mathcal{N}(\theta_{c,k}|0, \tau_k^{-1}), \quad \tau_k \sim \text{Gamma}(\tau_k|a_0, b_0). \quad (38)$$

This implies that each $\theta_{c,k}$ is modeled (a priori) with an independent, zero-mean, Gaussian, with a precision hyper-parameter τ_k which is in turn modeled (independently) with a (conjugate) Gamma density. We note that when $\tau_k \rightarrow \infty$, then $\theta_{c,k} \rightarrow 0$. The resulting prior for $\theta_{c,k}$ arising by marginalizing the hyper-parameter is a heavy-tailed, Student's t -distribution. For the purposes of learning of θ_c and in order to compute derivatives of the log-prior, we retain the τ_k 's and treat them as latent variables in an inner-loop EM scheme [67] (see derivation in A.3) which consists of:

- E-step: evaluate:

$$\langle \tau_k \rangle_{p(\tau_k|\theta_{c,k})} = \frac{a_0 + \frac{1}{2}}{b_0 + \frac{\theta_{c,k}^2}{2}}. \quad (39)$$

- M-step: evaluate:

$$\frac{\partial \log p(\theta_c)}{\partial \theta_{c,k}} = - \langle \tau_k \rangle_{p(\tau_k|\theta_{c,k})} \theta_{c,k}. \quad (40)$$

We note also that the second derivative of the log-prior with respect to θ_c can be similarly obtained as:

$$\frac{\partial^2 \log p(\theta_c)}{\partial \theta_{c,k} \partial \theta_{c,l}} = \begin{cases} - \langle \tau_k \rangle_{p(\tau_k|\theta_{c,k})}, & \text{if } k = l \\ 0, & \text{otherwise.} \end{cases} \quad (41)$$

2.6. Approximate Bayesian inference – Laplace's approximation

The discussion thus far has been limited to point estimates for θ . A fully Bayesian treatment would pose significant computational challenges. These stem from the intractability of the log-partition function $A(\theta_c)$ of p_c in the exponential family of models (see Eq. (31)). Sampling or approximating the full posterior of θ_c would require repeated evaluations of this and potentially its derivatives, a difficulty which is only amplified when $\dim(\theta_c) \gg 1$. For that reason, we adopt an approximation based on the Laplace's method [68]. According to this, the target posterior $p(\theta|\mathbf{x}^{(1:N)})$ is modeled with a Gaussian (Fig. 2) with mean equal to the MAP estimate θ_{MAP} and a covariance \mathbf{S} equal to the inverse of the negative Hessian of the log-posterior at θ_{MAP} (see Eqs. (37) and (41)). These two quantities are readily obtained at the last iteration (upon convergence) of the MC-EM scheme described previously. Hence:

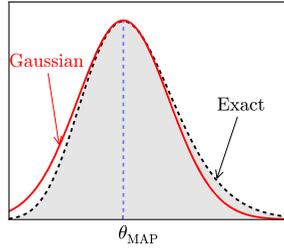


Fig. 2. Schematic illustration of the Laplace's approximation.

$$\mathbf{S}^{-1} = \begin{bmatrix} \mathbf{S}_{cc} & \mathbf{0} \\ \mathbf{0} & \mathbf{S}_{ff} \end{bmatrix}, \quad (42)$$

where the block-matrices above are given by:

$$\begin{aligned} \mathbf{S}_{cc} &= N \text{Cov}_{p_c(\mathbf{X}|\theta_c)}[\boldsymbol{\phi}(\mathbf{X}), \boldsymbol{\phi}_l(\mathbf{X})] + \text{diag}(\langle \tau_k \rangle_{p(\tau_k|\theta_{c,k})}) \\ \mathbf{S}_{ff} &= \sum_{i=1}^N \text{Cov}_{p_{cf}(\mathbf{x}|\mathbf{X}^{(i)}, \theta_{cf})q_i(\mathbf{X}^{(i)})}[\boldsymbol{\psi}(\mathbf{x}, \mathbf{X})]. \end{aligned} \quad (43)$$

Laplace's approximation can also be interpreted as a second-order Taylor series expansion of the log-posterior at θ_{MAP} . Some remarks:

- For $\theta_{c,k}$ that are effectively turned off when using the ARD prior (i.e. $\theta_{c,k,\text{MAP}} = 0$), $\langle \tau_k \rangle_{p(\tau_k|\theta_{c,k})} \rightarrow \infty$ and thus dominate the corresponding terms in \mathbf{S}^{-1} . As a result, the (approximate) posterior covariance of these $\theta_{c,k}$ approaches 0.
- We note that when the number of data points $N \rightarrow \infty$, the corresponding terms in \mathbf{S}^{-1} increase and as a result the (approximate) posterior covariance goes to 0, as one would expect.

Algorithm 1 summarizes the basic steps of the scheme advocated.

Algorithm 1 Proposed MC-EM scheme.

- 1: Initialize $\boldsymbol{\theta}^0 = \{\theta_c^0, \theta_{cf}^0\}$.
 - 2: Select parameters $\{a, \rho, A\}$ for the Robbins–Monro optimization algorithm (Eq. (30)).
 - 3: Step $t = 0$
 - 4: **while** (not converged) **do**
 - 5: MC-E-step:
 - 6: **for** all $i = 1, \dots, N$ **do**
 - 7: Generate MCMC samples from the (conditional) posterior distribution $q_i(\mathbf{X}^{(i)})$ in Eq. (21)
 - 8: **end for**
 - 9: M-step:
 - 10: Construct Monte Carlo gradient estimators $\hat{\mathcal{G}}^{(i)}$ (Eq. (27)) augmented by the prior gradient (Eq. (40)).
 - 11: Update the parameters $\boldsymbol{\theta}$ based on Eq. (28)
 - 12: $t \leftarrow t + 1$
 - 13: **end while**
 - 14: Compute Hessian of the log-posterior Eq. (15) at θ_{MAP} (Eqs. (37), (41)) to construct Laplace's approximation of the posterior $p(\boldsymbol{\theta}|\boldsymbol{\mathbf{x}}^{(1:N)})$ (Eq. (42)).
-

3. Numerical illustrations

We illustrate the proposed PCG framework in two examples. We particularize the definition of coarse-grained variables \mathbf{X} which unavoidably differs from problem to problem. We emphasize through several illustrations the ability of the proposed method to produce predictive estimates of various macroscopic observables as well as quantify the predictive uncertainty as a function of the amount of training data N used and the level of coarse-graining i.e. the ratio of the amount of fine to coarse variables. We also provide comparisons with the results obtained by employing the relative entropy method. Finally, we demonstrate how the ARD prior advocated can lead to the discovery of sparse solutions revealing the most prominent feature functions in the coarse potential and possibly the most significant types of interactions that this should contain. Whenever such a hierarchical prior (ARD) is employed (Eq. (38)) for the parameters θ_c in the coarse potential, the following values were used for the hyperparameters: $a_0 = b_0 = 10^{-5}$.

3.1. Ising model

The Ising model serves as abstraction of various physical problems, e.g. for modeling electromagnetism or lattice gas systems [69,70]. It has been the subject of detailed studies and several strategies for coarse-graining in equilibrium [9, 11–13,35,36] and nonequilibrium [9] settings.

We consider a periodic, one-dimensional lattice consisting of $n_f = 64$ sites. Each site i is associated with a binary variable $x_i, i = 1, \dots, n_f$ which takes values ± 1 . The n_f -dimensional vector $\mathbf{x} = \{x_i\}_{i=1}^{n_f}$ follows $p_f(\mathbf{x}) \propto \exp\{-\beta U_f(\mathbf{x})\}$ with the fine-scale potential given by:

$$U_f(\mathbf{x}) = -\frac{1}{2} \sum_{k=1}^{L_f} J_k \left(\sum_{|i-j|=k} x_i x_j \right) - \mu \sum_{i=1}^{n_f} x_i. \quad (44)$$

The expression $|i - j| = k$ implies a summation over all lattice sites i, j that are k -sites apart (periodic boundary conditions are assumed). The parameter L_f expresses the maximal interaction length. Following [9,28,71], we use a decaying interaction strength J_k with,

$$J_k = \frac{K}{k^a}, \quad (45)$$

and the normalization,

$$K = \frac{J_0}{L_f^{a-1} \sum_{k=1}^{L_f} k^{-a}}. \quad (46)$$

Finally, the parameter μ denotes the external field.

The values $A = 25$, $\alpha = 0.15$, and $\rho = 0.75$ were used for the Robbins–Monro updates (Eq. (28)) based on suggestions given in [36]. We used $m = 170$ samples for the MCMC estimates of the gradients in Eqs. (25) and (27).

3.1.1. Observables

As pointed out previously, the framework proposed readily allows for reconstructions of the whole fine-scale description and therefore probabilistic predictions can be computed for any observable. For comparative purposes, we focus on two such quantities. The first one is the magnetization $m(\mu)$ and its dependence on the external field parameter μ . This is associated with the following observable:

$$a^{(m)}(\mathbf{x}) = \frac{1}{n_f} \sum_i x_i, \quad (47)$$

i.e. $m(\mu) = \mathbb{E}_{p_f(\mathbf{x})}[a^{(m)}(\mathbf{x})]$. The second quantity is the correlation $R(k)$ at various separation distances k which captures second-order statistical information of the fine-scale configurations. The corresponding observable is:

$$a^{(R)}(\mathbf{x}; k) = \frac{1}{n_f} \sum_{|i-j|=k} x_i x_j, \quad (48)$$

i.e. $R(k) = \mathbb{E}_{p_f(\mathbf{x})}[a^{(R)}(\mathbf{x}; k)]$.

3.1.2. Coarse-variables \mathbf{X} and coarse-to-fine map

While the framework proposed offers great flexibility in the definition of the coarse variables \mathbf{X} , in this work we make perhaps the most intuitive choice by assuming that \mathbf{X} are (also) binary and have a *local* dependence on \mathbf{x} . This offers a direct appraisal on the level of coarse-graining as well as a natural, visual interpretation of the coarse variables and their role.

In particular, we assume that each coarse variable $X_I, I = 1, \dots, n_c$ is associated with a one-dimensional lattice that is a coarser version of the fine-scale one, i.e. with $n_c < n_f$ sites (Fig. 3). We can construct such descriptions by regularly coarsening by a factor of 2 such that $n_c = n_f/2^d$, with $d = 1, \dots, D$. We assume that each X_I (parent) is associated with $S = \frac{n_f}{n_c}$ fine-scale variables (children) denoted by $x_{(I-1)S+s} = x_{s,I}$ (where $s = 1, \dots, S$, Fig. 3). We define a coarse-to-fine map of the form:

$$\begin{aligned} p_{\text{cf}}(\mathbf{x}|\mathbf{X}, \boldsymbol{\theta}_{\text{cf}}) &= \prod_{I=1}^{n_c} \prod_{s=1}^S p(x_{s,I}|X_I, \boldsymbol{\theta}_{\text{cf}}) \\ &= \prod_{I=1}^{n_c} \prod_{s=1}^S p_0^{\frac{1+x_{s,I}X_I}{2}} (1-p_0)^{\frac{1-x_{s,I}X_I}{2}} \\ &= p_0^{\sum_{I=1}^{n_c} \sum_{s=1}^S \frac{1+x_{s,I}X_I}{2}} (1-p_0)^{\sum_{I=1}^{n_c} \sum_{s=1}^S \frac{1-x_{s,I}X_I}{2}}. \end{aligned} \quad (49)$$

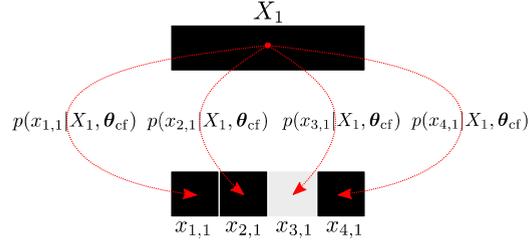


Fig. 3. Probabilistic coarse-to-fine map $p_{cf}(\mathbf{x}|\mathbf{X}, \theta_{cf})$. The coarse-variable X_1 is e.g. associated with $x_{1\dots 4,1}$ fine-scale variables through the probabilistic coarse-to-fine map p_{cf} (Eq. (49)). Each $x_{s,1}$ is *conditionally* independent from the other.

The expression above implies that each $x_{s,l}$ is *conditionally* independent and follows a Bernoulli distribution with probability p_0 of being of the same value as its parent X_l , and probability $(1 - p_0)$ of having the opposite spin. We emphasize that this does not imply that $x_{s,l}$ are also independent. In fact they will be correlated as a result of the dependencies between the coarse variables \mathbf{X} induced by the coarse model p_c which is discussed in the next subsection. The density p_{cf} above belongs to the exponential family (Section 2.4) and is controlled by a single parameter, $p_0 \in [0, 1]$. Given the symmetry of the model, we restrict $p_0 \in [0.5, 1]$. To ensure that it stays within this interval during the MC-EM updates (Algorithm 1), we operate instead on $\theta_{cf} \in \mathbb{R}$ defined as follows:

$$p_0 = \frac{1}{2} \left(1 + \frac{1}{1 + e^{-\theta_{cf}}} \right). \tag{50}$$

The derivatives needed for the updates of the EM-scheme in Eq. (27) and Eq. (37) are:

$$\begin{aligned} \frac{\partial \log p_{cf}}{\partial \theta_{cf}} &= \frac{\partial \log p_{cf}}{\partial p_0} \frac{\partial p_0}{\partial \theta_{cf}}, \\ \frac{\partial^2 \log p_{cf}}{\partial \theta_{cf}^2} &= \frac{\partial^2 \log p_{cf}}{\partial p_0^2} \left(\frac{\partial p_0}{\partial \theta_{cf}} \right)^2 + \frac{\partial \log p_{cf}}{\partial p_0} \frac{\partial^2 p_0}{\partial \theta_{cf}^2}, \end{aligned} \tag{51}$$

where:

$$\begin{aligned} \frac{\partial \log p_{cf}}{\partial p_0} &= \frac{\psi(\mathbf{x}, \mathbf{X})}{p_0} - \frac{1 - \psi(\mathbf{x}, \mathbf{X})}{1 - p_0}, \\ \frac{\partial^2 \log p_{cf}}{\partial p_0^2} &= -\frac{\psi(\mathbf{x}, \mathbf{X})}{p_0^2} - \frac{1 - \psi(\mathbf{x}, \mathbf{X})}{(1 - p_0)^2}, \end{aligned} \tag{52}$$

and $\psi(\mathbf{x}, \mathbf{X}) = \sum_{l=1}^{n_c} \sum_{s=1}^S \frac{1+x_{s,l}X_l}{2}$.

3.1.3. Coarse model

The coarse potential $U_c(\mathbf{X}; \theta_c)$ employed includes first-, second- and third-order interactions with various interaction lengths. In particular, we prescribe:

$$U_c(\mathbf{X}; \theta_c) = -\frac{1}{2} \left\{ \theta_c^{(1)} \sum_i X_i + \sum_i X_i \sum_k \theta_{c,k}^{(2)} X_{i\pm k} + \sum_i X_i \sum_{k=1}^{L_c^{(3)}} \theta_{c,kl}^{(3)} X_{i\pm k} X_{i\pm k\pm l} \right\} - \mu \sum_i X_i. \tag{53}$$

The parameters $L_c^{(2)}$ and $L_c^{(3)}$ denote the maximal second- and third- order interactions, respectively. With superscripts (1), (2), (3) we distinguish between the coarse potential parameters θ_c that are associated with the first, two-body and three-body interactions, respectively. These parameters determine also the number of θ_c which is equal to $1 + L_c^{(2)} + (L_c^{(3)})^2$.

In order to compare the proposed method with the relative entropy method, as briefly summarized in Section 2.1, a deterministic fine-to-coarse mapping $\mathcal{R}(\mathbf{x})$ is needed. We note that in [35,36] such efforts have been made by “coarse-graining” the interactions rather than the degrees of freedom i.e. $\mathbf{x} \equiv \mathbf{X}$. In order to truly assess the performance in cases where the coarse variables are of lower dimension and of the same type as in this study (i.e. binary), we prescribe the following map:

$$X_l = \begin{cases} +1, & \frac{1}{S} \sum_s x_{s,l} \geq 0 \\ -1, & \frac{1}{S} \sum_s x_{s,l} < 0. \end{cases} \tag{54}$$

This implies a “majority rule” where the label of the parent X_l is determined by the majority of the children. The same model as in Eq. (53) was used for the coarse potential. In order to reconstruct the fine configurations \mathbf{x} and estimate the observables of interest from the coarse description \mathbf{X} , a consistent sampling was performed from the conditional in Eq. (7) for the \mathcal{R} above.

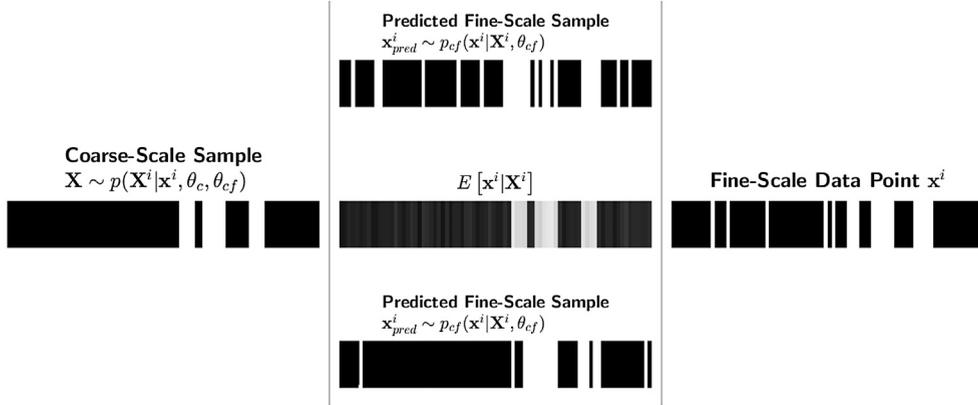


Fig. 4. For the FG datum $\mathbf{x}^{(i)}$ (right), the image on the left shows a sample from the posterior of the CG $\mathbf{X}^{(i)}$ (upon convergence of the Algorithm 1) i.e. one of the possible *pre-images* of $\mathbf{x}^{(i)}$. The three images in the center illustrate the predictions/reconstructions of the fine-scale: the top and bottom are samples drawn from the p_{cf} and the center is the expected FG configuration according to p_{cf} .

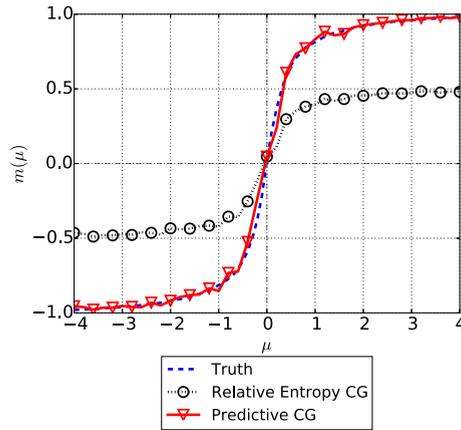


Fig. 5. Comparison of the reference magnetization (computed with the FG configuration) with posterior mean of predictive CG and relative entropy CG. $N = 20$, $\frac{n_f}{n_c} = 2$, $L_c^{(2)} = 15$, $L_c^{(3)} = 3$.

3.1.4. Results

The ensuing results are based on the following values for the fine-scale potential: $J_0 = 1.5$, $a = 0.8$, $L_f = 8$, $\beta = 0.3$, $n_f = 64$. We generated data from the fine scale model for each of 41 values of the external field μ , equidistantly distributed within $[-4, 4]$. A different CG model is trained for every μ value considered. One could also envision introducing a dependence of the CG model's components on μ which would allow a single model to be inferred and to be used for making predictions even for values of μ not contained in the data. Fig. 4 provides some insight on the role of the CG variables, their posterior and their ability to represent/reconstruct the FG configuration. Fig. 5 compares point-estimates of the predicted magnetization as obtained with the proposed method (red) and the relative entropy method (for fine-to-coarse mapping as given in Eq. (54)). While one can claim that better results can be obtained with a different set of CG variables (Eq. (54)), the point in this comparison is to demonstrate the information loss that takes place which can lead to poor predictions when not quantified. Given the same amount of training data N , the information loss in the relative entropy method is driven by the not adjusted map in the consistent density of the fine-scale variables $p_{\mathcal{R}}(\mathbf{x})$ denoted in Eq. (9) compared to PCG. While in PCG the probabilistic map $p_{cf}(\mathbf{x}|\mathbf{X}, \theta_{cf})$ (Eq. (49)) is parametrized and optimized within the parametric family of p_{cf} . We note further that the relative entropy method can lead to good approximations of the Potential of Mean Force, and as a result, accurate estimates (as shown earlier) of expectations of observables that depend solely on \mathbf{X} . We could therefore select \mathbf{X} in such a way that the magnetization is only a function of \mathbf{X} in which case the result of the relative entropy method would probably be good. If however another expectation was sought (that does not depend on the current \mathbf{X}) a new set of \mathbf{X} would need to be defined and a new CG model would need to be retrained.

When $\frac{n_f}{n_c} = 2$, $L_c^{(2)} = 15$, $L_c^{(3)} = 3$, the total number of unknowns parameters θ_c in the potential U_c is $1 + L_c^{(2)} + (L_c^{(3)})^2 = 25$. This is not a particularly large number, but we demonstrate nevertheless the effect of the sparsity enforcing prior in Fig. 6 when $N = 20$ data points are used. In the absence of the ARD prior (Eq. (38)), all θ_c are non-zero and

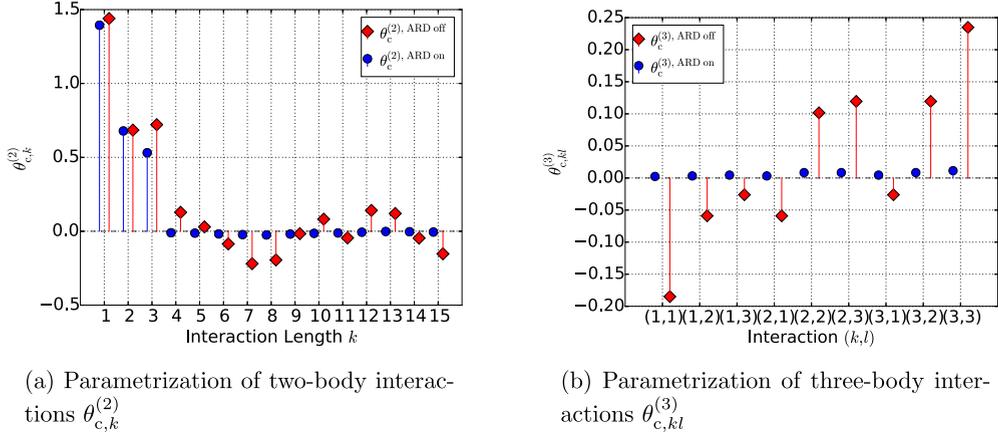


Fig. 6. Parametrization of two- and three-body interactions at $\mu = 0.0$, obtained with and without ARD prior. A sparse solution is obtained with active ARD prior at the same predictive accuracy (see Fig. 7). $N = 20$, $\frac{n_f}{n_c} = 2$, $L_c^{(2)} = 15$, $L_c^{(3)} = 3$.

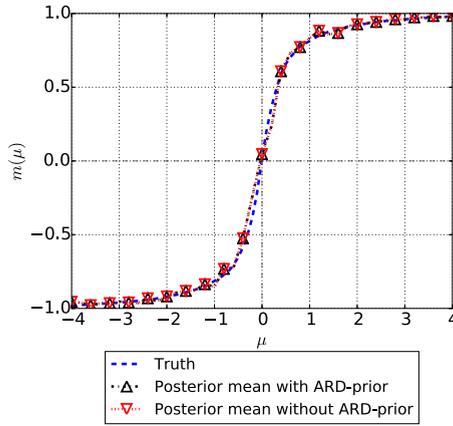


Fig. 7. Comparison of predicted magnetization with and without ARD prior. $N = 20$, $\frac{n_f}{n_c} = 2$, $L_c^{(2)} = 15$, $L_c^{(3)} = 3$.

the corresponding feature functions are all active (Eq. (53)). On the contrary, when the ARD prior is employed, the learning scheme identifies only 3 non-zero θ_c . Interestingly these are associated with two-body interactions up to separation 3 whereas all other terms corresponding to two- and three-body interactions are found to be unnecessary, despite having equal predictive accuracy as shown in Fig. 7 where point estimates of the magnetization are plotted (with and without the ARD prior).

Fig. 8 depicts the effect of adding more training data N in the predictive posterior estimates for the magnetization at various μ values. One observes that as N increases, not only the posterior mean estimates approach the reference solution, but more importantly, the posterior credible intervals shrink around it reflecting the fact that the model becomes more confident. Credible intervals are obtained by sampling the (approximate) posterior distribution $p(\theta | \mathbf{x}^{(1:N)})$ (Eq. (16)) and determining the observable for each sample $\theta^{(i)}$ with the predictive estimator $\hat{a}(\theta^{(i)})$ (Eq. (19)). We use the predictive samples $\hat{a}(\theta^{(i)})$ to determine desired quantiles (see A.1 for more details). The same observations can be made when attempting to predict second-order statistics of the fine-scale i.e. the correlation at various separations k (Fig. 9).

The decreasing variance for increasing N can also be observed in the model parameters e.g. the coarse-to-fine mapping parameter p_0 (Eq. (49)), the (approximate) posterior of which is shown in Fig. 10.

Finally in Figs. 11 and 12, the predictive ability of the model is compared for different levels of coarse-graining. In the formulation adopted, this is quantified by the ratio between the dimension of fine \mathbf{x} and coarse \mathbf{X} descriptions i.e. $\frac{n_f}{n_c}$. We consider two cases i.e. $\frac{n_f}{n_c} = 2, 8$. As one would expect, the posterior mean estimates are superior when $\frac{n_f}{n_c} = 2$ but also the predictive posterior uncertainty increases as the coarse-graining becomes more pronounced. This is easily understood by the fact that the fewer CG variables used, the higher the information loss becomes. It is important to note though that even when $\frac{n_f}{n_c} = 8$, the predictive posterior’s credible intervals always include the reference solution.

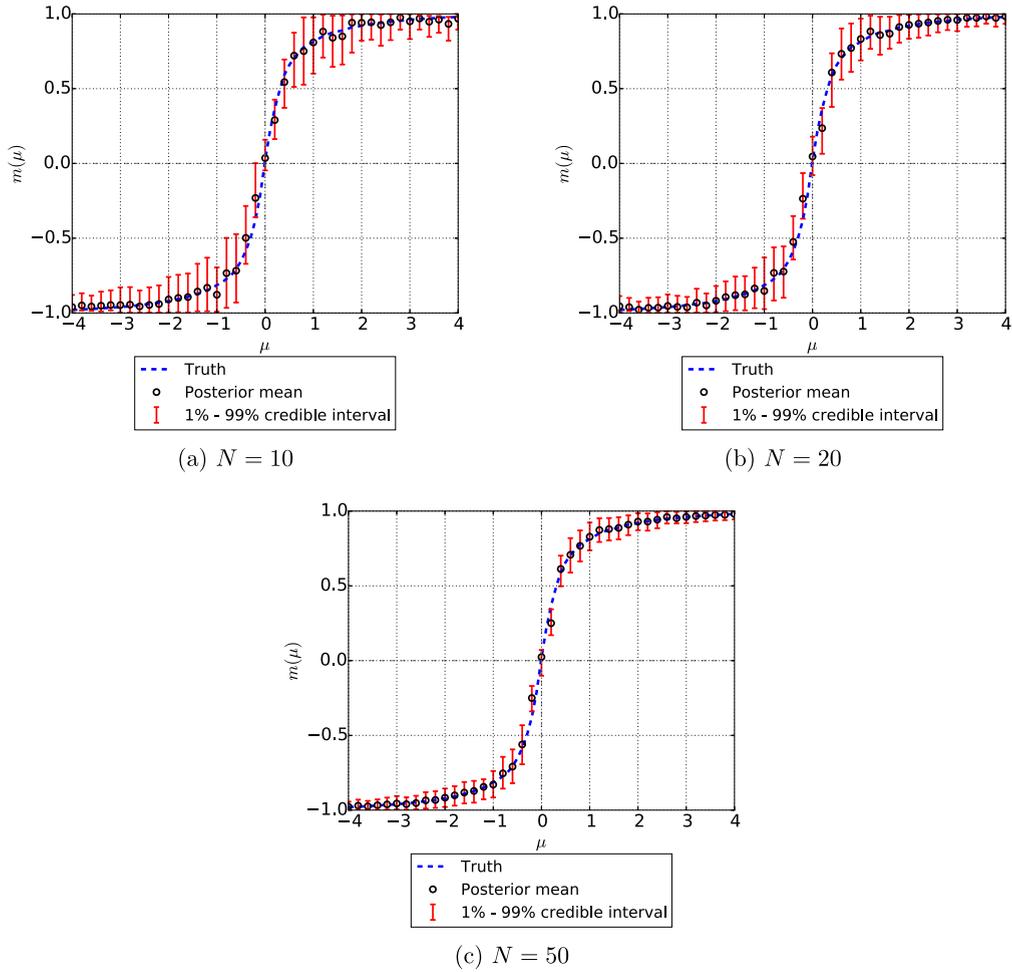


Fig. 8. Comparison of the reference magnetization (computed with the FG configuration) with posterior mean and credible intervals corresponding to 1% and 99% posterior quantiles. $N = 20$, $\frac{n_r}{n_c} = 2$, $L_c^{(2)} = 15$, $L_c^{(3)} = 3$.

3.2. Coarse-graining SPC/E water

The second example addresses the coarse-graining of a water model which is described at the atomistic scale by oxygen and hydrogen atoms. Water has been the focus of several studies in coarse-graining as it plays the role of the solvent in various biological and chemical systems and as a result it can take up to 80% of the total simulation time [30]. Furthermore there exist several well-documented properties which can serve as a measure of comparison. In this study, we employ the Simple Point Charge/Extended (SPC/E) water model introduced in [72,73] for the FG (all-atom) description. In the context of the relative entropy method, coarse-graining of the SPC/E water is addressed in [36,74–76]. In particular, we consider a system of $M = 100$ water molecules at a temperature of $T = 300\text{K}$, and a pressure of $p = 1.0\text{bar}$. The equilibrium box length is $l_{\text{box}} = 14.56\text{\AA}$ and a time step of $\Delta t = 2.0\text{fs}$ is used. Periodic boundary conditions are applied in every dimension while ensuring the NVT ensemble by the Nosé–Hoover thermostat [77,78]. The \mathbf{x} vector contains the coordinates of the 100 oxygen and 200 hydrogen atoms i.e. $\dim(\mathbf{x}) = 900$. The fine-scale potential $U_f(\mathbf{x})$ under the SPC/E model consists of a Lennard–Jones (LJ) potential for non-bonded interactions and a Coulomb potential for long-range interactions. Parameters for the LJ potential,

$$U_f^{\text{LJ}}(\mathbf{x}) = \frac{1}{2} \sum_{j \neq k} 4\epsilon \left(\left(\frac{\sigma}{R_{ij}(\mathbf{x})} \right)^{12} - \left(\frac{\sigma}{R_{ij}(\mathbf{x})} \right)^6 \right), \quad (55)$$

are $\sigma = 3.166\text{\AA}$ and $\epsilon = 0.650 \frac{\text{kJ}}{\text{mol}}$, with the distance between particle i and j denoted as R_{ij} .

The electric load of Hydrogen (H) and Oxygen (O) atoms are given by $q_O = -0.8476e$, $q_H = +0.4238e$ where e represents the elementary charge. The SPC/E model assumes the bonded interaction to be rigid with a bonding angle defined between the two H-atoms and the central O-atom as $\omega_{\text{HOH}} = 109.47^\circ$. The bond-length used in this study is $l_{\text{OH}} = 1.0\text{\AA}$. The

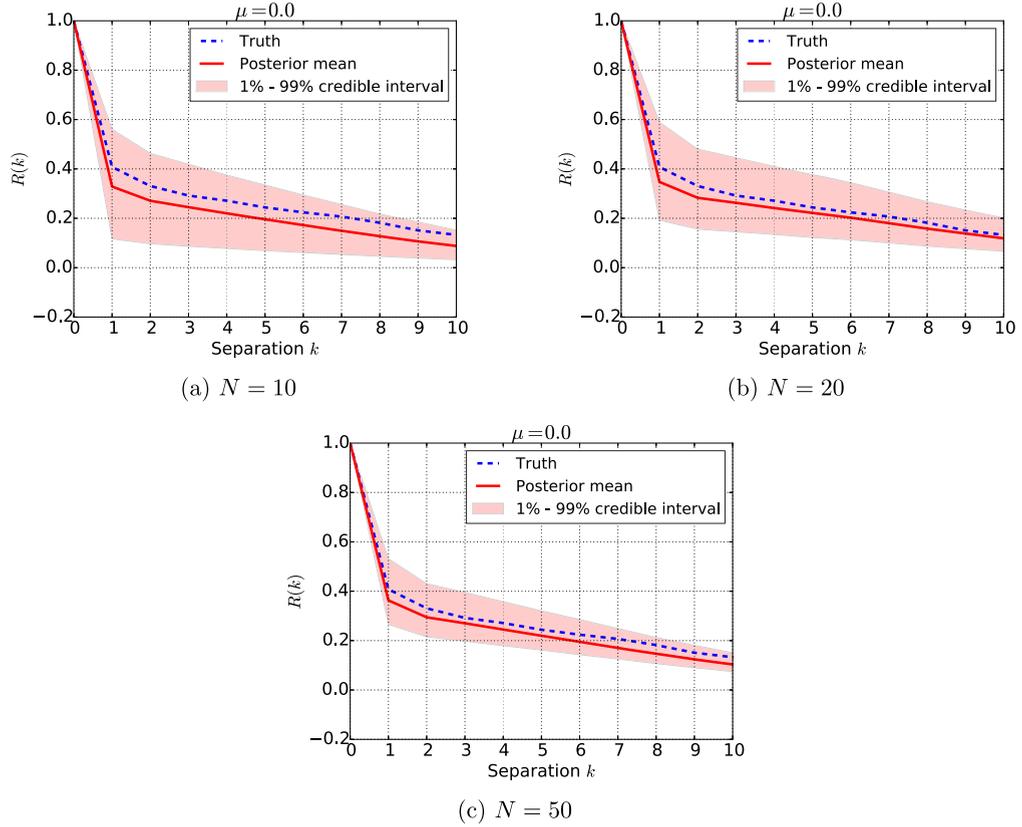


Fig. 9. Comparison of the reference correlation (computed with the FG configuration) with posterior mean and credible intervals corresponding to 1% and 99% posterior quantiles. $N = 20$, $\frac{n_r}{n_c} = 2$, $L_c^{(2)} = 15$, $L_c^{(3)} = 3$.

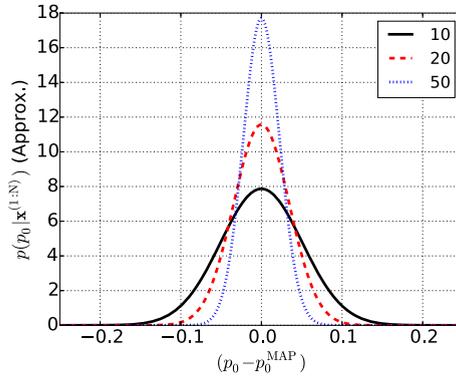


Fig. 10. Posterior $p(p_0 | \mathbf{x}^{1:N})$ at $\mu = 0.0$ for $N = 10, 20, 50$. $\frac{n_r}{n_c} = 2$, $L_c^{(2)} = 15$, $L_c^{(3)} = 3$.

equilibration for the NVT ensemble was performed as in [36,75]. For both fine- and coarse-scale simulations the molecular dynamics software package LAMMPS [79] was used. Further details are contained in B.1.

The values $A = 9$, $\alpha = 0.05$, and $\rho = 0.60$ were used for the Robbins–Monro updates (Eq. (28)) based on suggestions given in [36]. We used $m = 160$ samples for the MCMC estimates of the gradients in Eqs. (25) and (27).

3.2.1. Observables

The first macroscopic observable of interest is the Radial Distribution Function (RDF) $g(r)$ which represents a characteristic and well-studied property in water models. Several computational and experimental results related to the RDF are described in [80]. As a pair correlation function, $g(r)$ depends on the statistics of the distances r_{jk} between each pair of molecules j, k . To compute these distances, we employ the coordinates of the center of mass of each water molecule $\hat{\mathbf{x}}_j$:

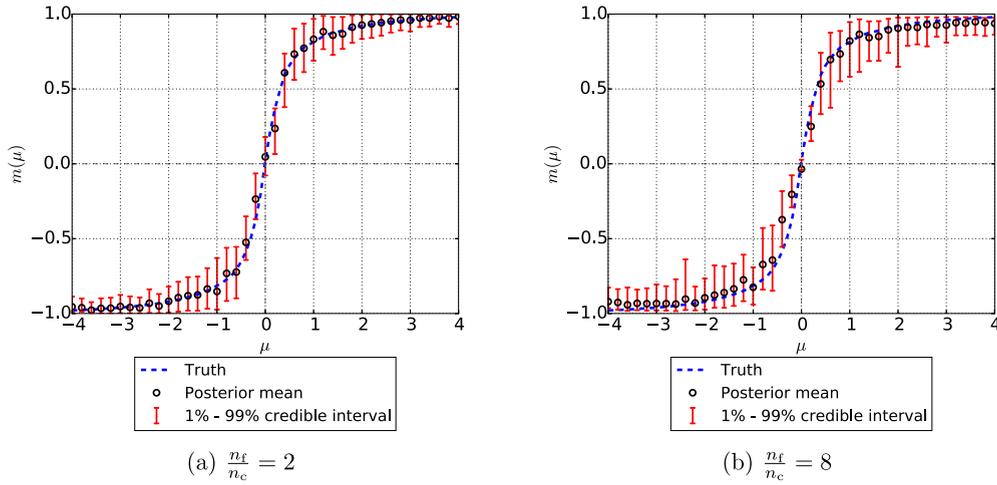


Fig. 11. Magnetization for different level of coarse graining, i.e. ratio of the amount fine/coarse variables $\frac{n_f}{n_c} = 2$ ($L_c^{(2)} = 15$, $L_c^{(3)} = 3$) and $\frac{n_f}{n_c} = 8$ ($L_c^{(2)} = 3$, $L_c^{(3)} = 1$). Both models were trained with the same data $N = 20$.

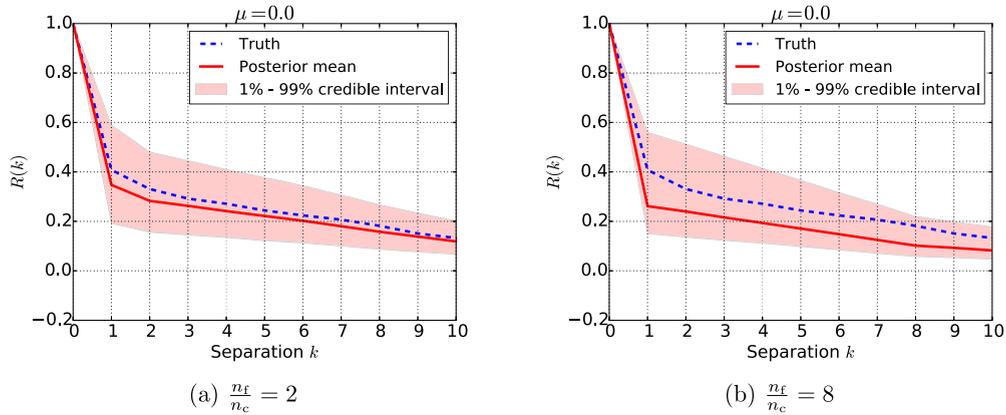


Fig. 12. Correlation for different level of coarse graining, i.e. ratio of the amount fine/coarse variables $\frac{n_f}{n_c} = 2$ ($L_c^{(2)} = 15$, $L_c^{(3)} = 3$) and $\frac{n_f}{n_c} = 8$ ($L_c^{(2)} = 3$, $L_c^{(3)} = 1$). Both models were trained with the same data $N = 20$.

$$\hat{\mathbf{x}}_j = \frac{\mathbf{x}_{O,j}m_O + \mathbf{x}_{H,j_1}m_H + \mathbf{x}_{H,j_2}m_H}{m_O + 2m_H}, \quad (56)$$

where $\mathbf{x}_{O,j}$ are the coordinates of the oxygen atom of molecule j , \mathbf{x}_{H,j_1} , \mathbf{x}_{H,j_2} are the coordinates of the two hydrogen atoms of the same molecule, and m_O, m_H are the masses of oxygen and hydrogen atoms, respectively (see B.1). If $r_{jk} = |\hat{\mathbf{x}}_j - \hat{\mathbf{x}}_k|$, then the corresponding observable of interest is [81]:

$$a^{\text{RDF}}(\mathbf{x}) = \frac{V}{M^2} \sum_j \sum_{j \neq k} \delta(r - r_{jk}), \quad (57)$$

where V denotes the volume of the simulation box (14.56^3 \AA^3) and $M = 100$ the number of molecules in the system. Additional details can be found in B.2.

The second property of interest involves the tetrahedral structure of water. Neighboring water molecules temporarily build such tetrahedral clusters due to the hydrogen bonds. Several measures of tetrahedrality have been proposed which relate to the deviation from the perfect tetrahedral structure $\omega_0 = 109.471^\circ$ [74,82]. In this work, we employ the angular distribution function which considers the eight closest neighbors $n_c = 8$ for a given molecule j . It is defined as follows:

$$a^{\text{tetra}}(\mathbf{x}; \omega) = \frac{1}{Mn_\omega} \sum_{j=1}^M \sum_{k=1}^{n_c} \sum_{l \neq j}^{n_c-1} \delta(\omega - \omega_{jkl}), \quad (58)$$

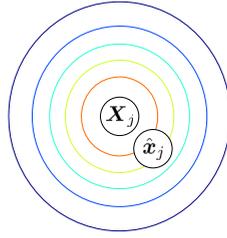


Fig. 13. Probabilistic mapping $p_{\text{cf}}(\hat{\mathbf{x}}_j | \mathbf{X}_j, \boldsymbol{\theta}_{\text{cf}})$, with mean \mathbf{X}_j and predicted fine-scale variable $\hat{\mathbf{x}}_j$. The contours depict the isotropic Gaussian distribution of Eq. (59) with mean \mathbf{X}_j and variance σ^2 .

with ω_{jkl} the angle between molecules j, k, l , with the central molecule j (as computed using the centers of mass $\hat{\mathbf{x}}$ in Eq. (56)) and $n_\omega = \binom{n_c}{3} = 56$. The product (Mn_ω) normalizes a^{tetra} with respect to the considered angular triplets.

We note that since the observables of interest depend only on the centers of mass $\hat{\mathbf{x}} = \hat{\mathbf{x}}(\mathbf{x})$, it suffices to use a coarse-to-fine map that relates the coarse variables \mathbf{X} directly with $\hat{\mathbf{x}}$ (Eq. (19)).

3.2.2. Coarse-variables \mathbf{X} and coarse-to-fine map

Since the observables of interest depend on the centers of mass $\hat{\mathbf{x}}$ (Eq. (56)), the coarse-to-fine probabilistic map assumes the form $p_{\text{cf}}(\hat{\mathbf{x}} | \mathbf{X})$. As frequently done in CG studies of water, each molecule j is represented by a CG variable $\mathbf{X}_j \in \mathbb{R}^3$. We then prescribe a p_{cf} of the following form:

$$p_{\text{cf}}(\hat{\mathbf{x}} | \mathbf{X}, \boldsymbol{\theta}_{\text{cf}}) = \prod_{j=1}^M \mathcal{N}(\hat{\mathbf{x}}_j | \mathbf{X}_j, \sigma^2 \mathbf{I}), \quad (59)$$

where \mathbf{I} is the 3×3 identity matrix. This suggests that each $\mathbf{X}_j, j = 1, \dots, M$ determines the center of mass $\hat{\mathbf{x}}_j$ up to an isotropic Gaussian with mean \mathbf{X}_j and variance σ^2 (see Fig. 13). The latter quantifies the uncertainty in the prediction of the fine-scale (up to centers of mass) from the CG description. Large values of σ^2 imply that \mathbf{X} provides an imprecise reconstruction of $\hat{\mathbf{x}}$ and vice versa. Hence there is only one parameter in the coarse-to-fine map i.e. $\sigma^2 \geq 0$. In order to ensure non-negativity during updates we operate instead on $\boldsymbol{\theta}_{\text{cf}} = -\log \sigma^2$ which leads to the following derivatives needed in Eqs. (27) and (37):

$$\begin{aligned} \frac{\partial \log p_{\text{cf}}}{\partial \boldsymbol{\theta}_{\text{cf}}} &= \frac{3M}{2} - \frac{1}{2\sigma^2} \sum_{j=1}^M |\hat{\mathbf{x}}_j - \mathbf{X}_j|^2, \\ \frac{\partial^2 \log p_{\text{cf}}}{\partial \boldsymbol{\theta}_{\text{cf}}^2} &= -\frac{1}{2\sigma^2} \sum_{j=1}^M |\hat{\mathbf{x}}_j - \mathbf{X}_j|^2. \end{aligned} \quad (60)$$

Naturally, more complex descriptions involving an anisotropic covariance or a mixture of Gaussians could be used.

3.2.3. Coarse model

The coarse potential $U_c(\mathbf{X}; \boldsymbol{\theta}_c)$ employed consists of two- and three-body interactions. It assumes the form:

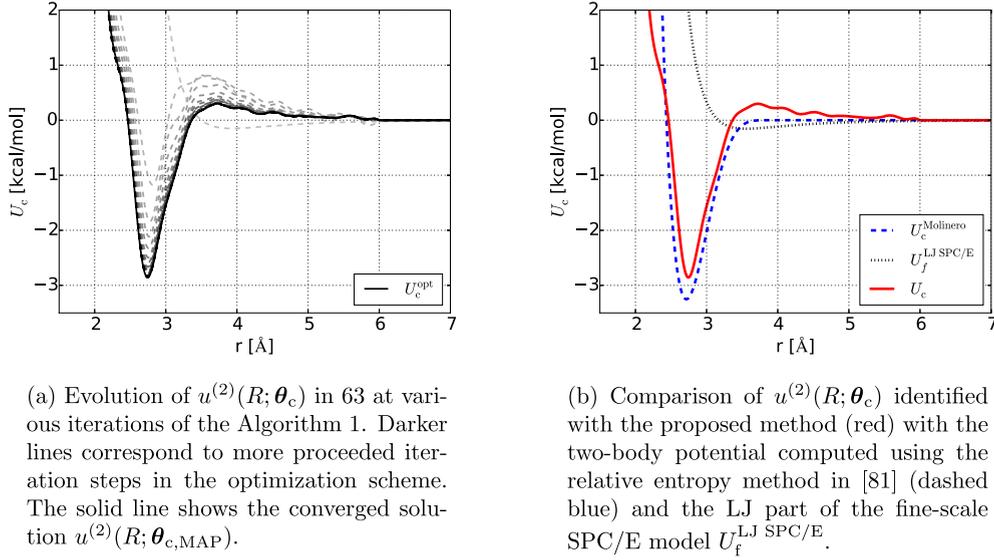
$$U_c(\mathbf{X}; \boldsymbol{\theta}_c) = \underbrace{U^{\text{SW}}(\mathbf{X})}_{\text{fixed}} + \tilde{U}(\mathbf{X}; \boldsymbol{\theta}_c), \quad (61)$$

where $U^{\text{SW}}(\mathbf{X})$ is a fixed term described below and $\tilde{U}(\mathbf{X}; \boldsymbol{\theta}_c)$ represents the “correction” that is learned from the data using the framework advocated. In particular, the fixed term $U^{\text{SW}}(\mathbf{X})$ is given by (a variation of) the Stillinger–Weber (SW) potential proposed in [83] and discussed in B.3. The remaining part $\tilde{U}(\mathbf{X}; \boldsymbol{\theta}_c)$ consists only of two-body interaction terms i.e.

$$\tilde{U}(\mathbf{X}; \boldsymbol{\theta}_c) = \frac{1}{2} \sum_{j \neq k} u^{(2)}(R_{jk}; \boldsymbol{\theta}_c), \quad (62)$$

where $R_{jk} = |\mathbf{X}_j - \mathbf{X}_k|$ and the pairwise potential $u^{(2)}(R; \boldsymbol{\theta}_c)$ is parametrized as follows:

$$u^{(2)}(R; \boldsymbol{\theta}_c) = u^{\text{LJ}}(R; \boldsymbol{\theta}_c^{\text{LJ}}) + \sum_{k=1}^K \theta_{c,k}^{\text{cor}} \phi_k(R), \quad R > 0. \quad (63)$$



(a) Evolution of $u^{(2)}(R; \theta_c)$ in 63 at various iterations of the Algorithm 1. Darker lines correspond to more proceeded iteration steps in the optimization scheme. The solid line shows the converged solution $u^{(2)}(R; \theta_{c, \text{MAP}})$.

(b) Comparison of $u^{(2)}(R; \theta_c)$ identified with the proposed method (red) with the two-body potential computed using the relative entropy method in [81] (dashed blue) and the LJ part of the fine-scale SPC/E model $U_f^{\text{LJ SPC/E}}$.

Fig. 14. Coarse-graining SPC/E water using $N = 20$ training data. Computed two-body, coarse-scale potential $u^{(2)}(R; \theta_c)$ and comparisons.

In the equation above, $u^{\text{LJ}}(R; \theta_c^{\text{LJ}})$ is a Lennard–Jones potential and the feature functions $\phi = \{\phi_k(R)\}_{k=1}^K$ are a combination of sines and cosines truncated in the interval $I_c = [R_{\min} = 2.0 \text{ \AA}, R_{\max} = 6.0 \text{ \AA}]$. The bounds R_{\min}, R_{\max} define an effective window where the LJ potential is corrected to capture the associated CG interactions. In particular:

$$\phi_k(R) = \begin{cases} 1_{I_c}(R) \sin 2\pi \nu_k R, & k = \text{odd}, \\ 1_{I_c}(R) \cos 2\pi \nu_k R, & k = \text{even}, \end{cases} \quad (64)$$

where $1_{I_c}(R)$ is the indicator function of the interval I_c . The wave-numbers ν_k offer a Fourier-like decomposition of the second-order potential and were defined as follows:

$$\nu_{2k'} = \nu_{2k'+1} = 1 + \frac{19}{K/2} k', \quad k' = 0, 2, \dots, K/2 - 1, \quad (65)$$

i.e. at a uniform grid in $[1, 20]$. By increasing the total number K of these terms, one can potentially learn finer fluctuations of this potential. Naturally one would want to use as many feature functions as possible in order to ensure greater flexibility of the model, which gives rise to the need for sparsity-enforcing priors for $\theta_{c,k}^{\text{COR}}$ as discussed previously. In this study, $K = 100$ was used.

The superimposed LJ potential ensures that $\lim_{R \rightarrow 0} u^{(2)}(R; \theta_c) = \infty$ and is of the form:

$$u^{\text{LJ}}(R; \theta_c^{\text{LJ}}) = 4\epsilon \left(\left(\frac{\sigma_{\text{LJ}}}{R} \right)^{12} - \left(\frac{\sigma_{\text{LJ}}}{R} \right)^6 \right), \quad (66)$$

where $\theta_c^{\text{LJ}} = (\sigma_{\text{LJ}}, \epsilon)$. The total number of parameters associated with the two-body term was $K + 2 = 102$ and consists of $\theta_c = (\theta_c^{\text{LJ}}, \theta_c^{\text{COR}})$. The ARD prior is employed only for θ_c^{COR} and an (improper) uniform prior is employed for the rest θ_c^{LJ} . We note that due to the LJ part, the corresponding distribution p_c is not in the exponential family anymore (Section 2.4) and the possibility of multiple local maxima cannot be excluded.

3.2.4. Results

We first run the proposed algorithm for $N = 20$ fine-scale (all-atom) realizations. Fig. 14a depicts the evolution of the inferred coarse-scale potential $u^{(2)}(R; \theta_c)$ (Eq. (63)) at various iterations of the EM-scheme. We initialize with $\theta_c^{\text{COR}} = \mathbf{0}$ and $\theta_c^{\text{LJ}} = (\epsilon = 0.15 \frac{\text{kcal}}{\text{mol}}, \sigma_{\text{LJ}} = 3.5 \text{ \AA})$. After 194 iterations, the converged result $u^{(2)}(R; \theta_{c, \text{MAP}})$ is depicted with a solid black line. In Fig. 14b, we compare this converged result (red) with the two-body potential computed in [81] (dashed blue) using the relative entropy method and the LJ part (black) of the fine-scale SPC/E model. The former two exhibit similarities but also differences which stem from the different structure of these two models. These differences persist even if more training data N are used.

Fig. 15 depicts the effect of the ARD prior on θ_c^{COR} . One observes in Fig. 15a that if no such prior is used (instead a uniform was employed) almost all θ_c^{COR} are non-zero and as a result almost all the corresponding feature functions $\phi_k(R)$ in Eq. (63) are active and the model is unable to distinguish their relative importance (unless N becomes very large). In contrast, the inclusion of the ARD prior in Fig. 15b leads to a sparse solution in which most $\phi_k(R)$ are deactivated (roughly

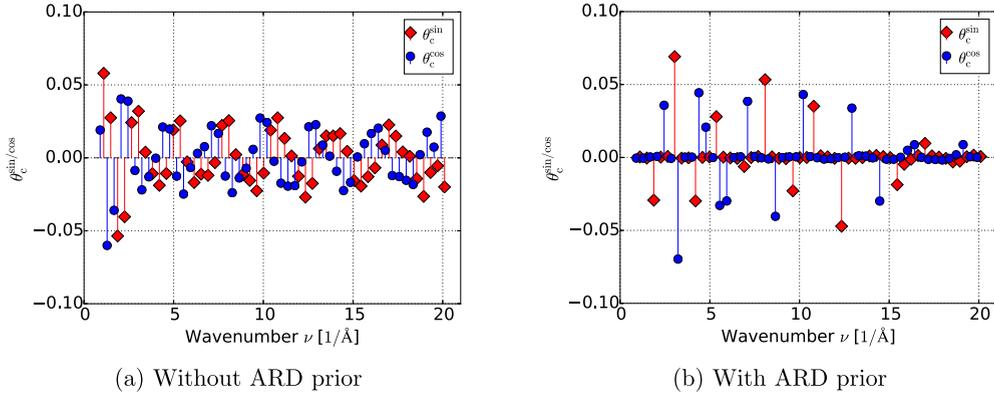


Fig. 15. $\theta_{c,\text{MAP}}^{\text{COR}}$ without and with the ARD prior with respect to the wavenumber ν_k (Eq. (64)). Superscripts sin (red) and cos (blue) indicate whether the corresponding $\theta_{c,k}^{\text{COR}}$ (Eq. (63)) is associated with a sine or cosine feature function respectively.

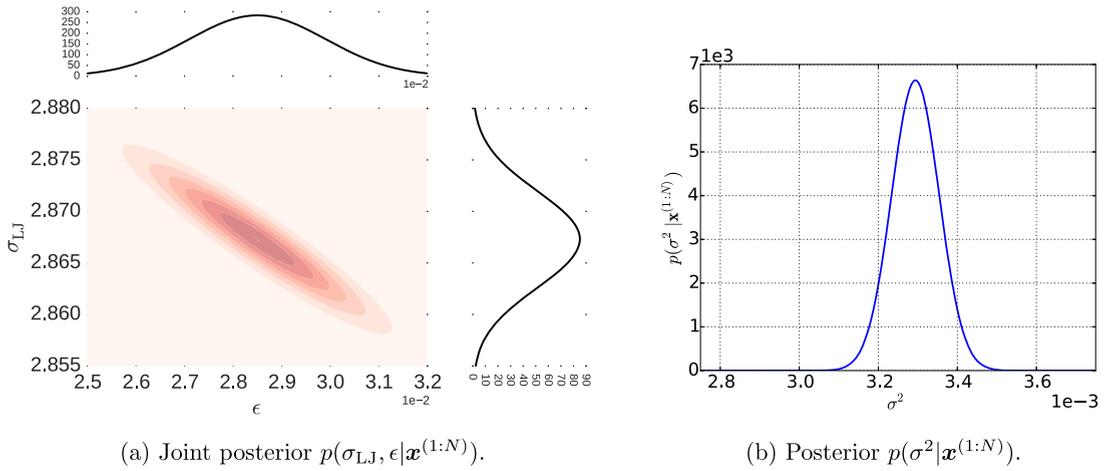
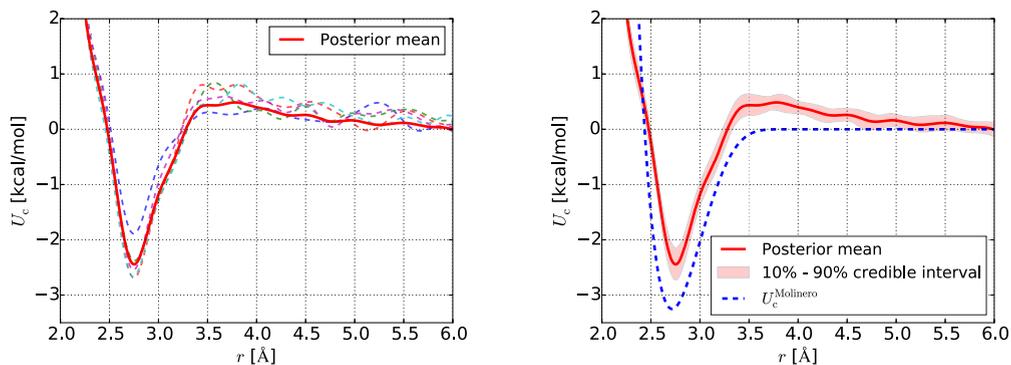


Fig. 16. Posterior of $\theta_c^{\text{LJ}} = (\sigma_{\text{LJ}}, \epsilon)$ in Eq. (66) and σ^2 in Eq. (59) for $N = 20$.

80 out of 100 in this case). It can be clearly seen as well that feature functions (sines/cosines) with high wave-numbers (small wave-lengths) are largely unnecessary for the description of the coarse potential. Although not demonstrated in this run, we envision that this modeling feature will eventually allow us to identify not only the most important terms in each potential term but also the most suitable order of interactions in the coarse potential. Fig. 16 depicts the (approximate) posterior obtained for $\theta_c^{\text{LJ}} = (\sigma_{\text{LJ}}, \epsilon)$ (Eq. (66)) and σ^2 (Eq. (59)) for $N = 20$.

Fig. 17 provides information with regards to the (approximate) posterior of θ_c , computed using the Laplace's approximation proposed, as reflected in the $u^{(2)}(R; \theta_c)$. In particular in Fig. 17a, we plot sample realizations of $u^{(2)}(R; \theta_c)$ corresponding to different samples of θ_c from the (approximate) Gaussian posterior (Section 2.6). We note that all realizations suggest the same location for the minimum of the potential. Variability is observed in the depth of this well as well as in its shape to the right of the minimum. Fig. 17b depicts the posterior mean of $u^{(2)}(R; \theta_c)$ as well as credible intervals at 10% and 90% posterior quantiles which reflect the inferential uncertainties discussed.

We finally report results illustrating the predictive capability of the model in terms of the macroscopic observables of interest i.e. the RDF and the angular distribution function discussed previously. To that end, we consider three data settings with $N = 10, 20$ and 100 fine-scale (all-atom) training data. While the MAP estimates do not exhibit prominent differences, the advantage of the method proposed is the predictive posterior that is furnished (Eq. (19)) and quantifies the uncertainty in the predictions that the coarse-grained model produces. Figs. 18 and 19 depict the posterior means and credible intervals corresponding to 10% and 90% posterior quantiles for the RDF $g(r)$ (i.e. the expected value of the observable in Eq. (57)) and the angular distribution function $p(\omega)$ (i.e. the expected value of the observable in Eq. (58)). In all cases, the posterior means are very close to the reference values obtained by simulating the all-atom SPC/E model. It is interesting to point out that when only $N = 10$ data were used, the posterior mean overestimates the first peak in the RDF (Fig. 18a). Nevertheless the true solution is contained within the credible intervals computed. As one would expect, the breadth of the credible intervals decreases as more training data N is introduced, reflecting the reduction in the predictive uncertainty of the model. Details for the computation of these credible intervals can be found in A.1.



(a) Realizations of $u^{(2)}(R; \theta_c)$ for random samples of θ_c drawn from the approximate posterior.

(b) Posterior mean and credible intervals for $u^{(2)}(R; \theta_c)$. We compare this with the two-body potential computed in [81] (dashed blue) using the relative entropy method.

Fig. 17. Posterior of $u^{(2)}(R; \theta_c)$ for $N = 20$.

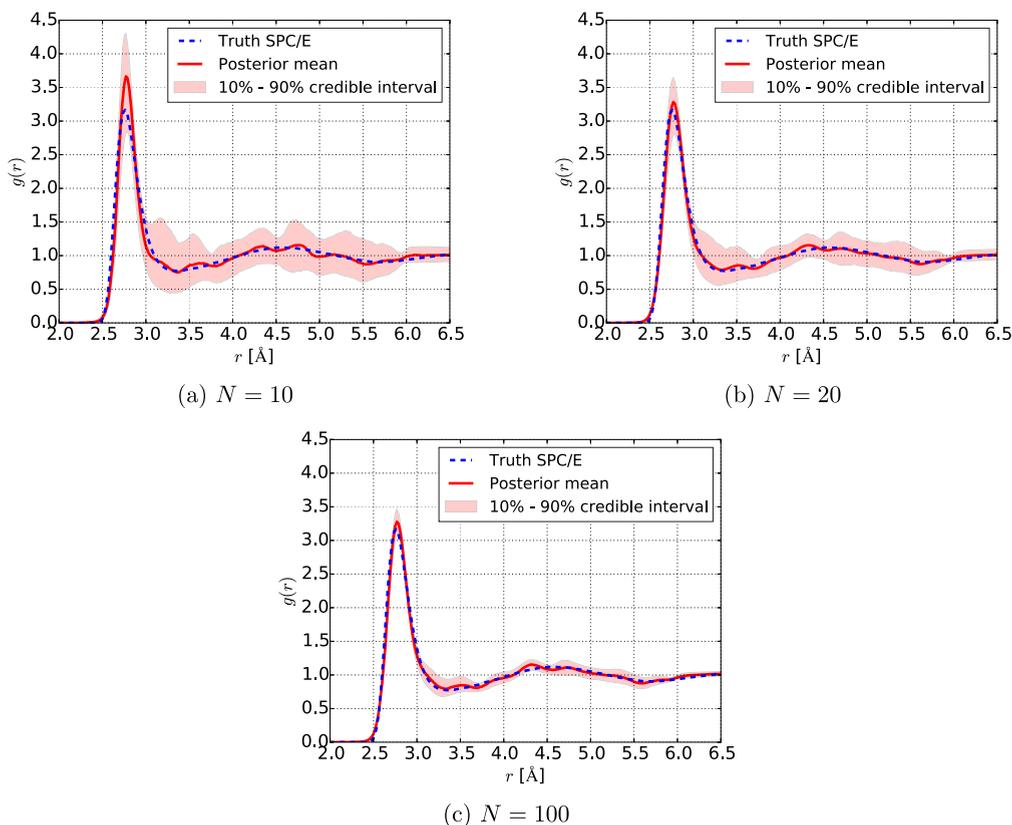


Fig. 18. Comparison of the reference RDF $g(r)$ (computed with all-atom simulations using the SPC/E model) with posterior mean and credible intervals corresponding to 10% and 90% posterior quantiles.

4. Conclusions

We presented a novel, data-driven coarse-graining scheme of atomistic ensembles in equilibrium. In contrast to existing techniques which are based on a restriction, fine-to-coarse map, we adopt the opposite strategy by prescribing a *probabilistic coarse-to-fine* map. This corresponds to a directed probabilistic model where the coarse variables play the role of latent generators of the fine scale (all-atom) data. Such a model can readily quantify the uncertainty due to the information loss

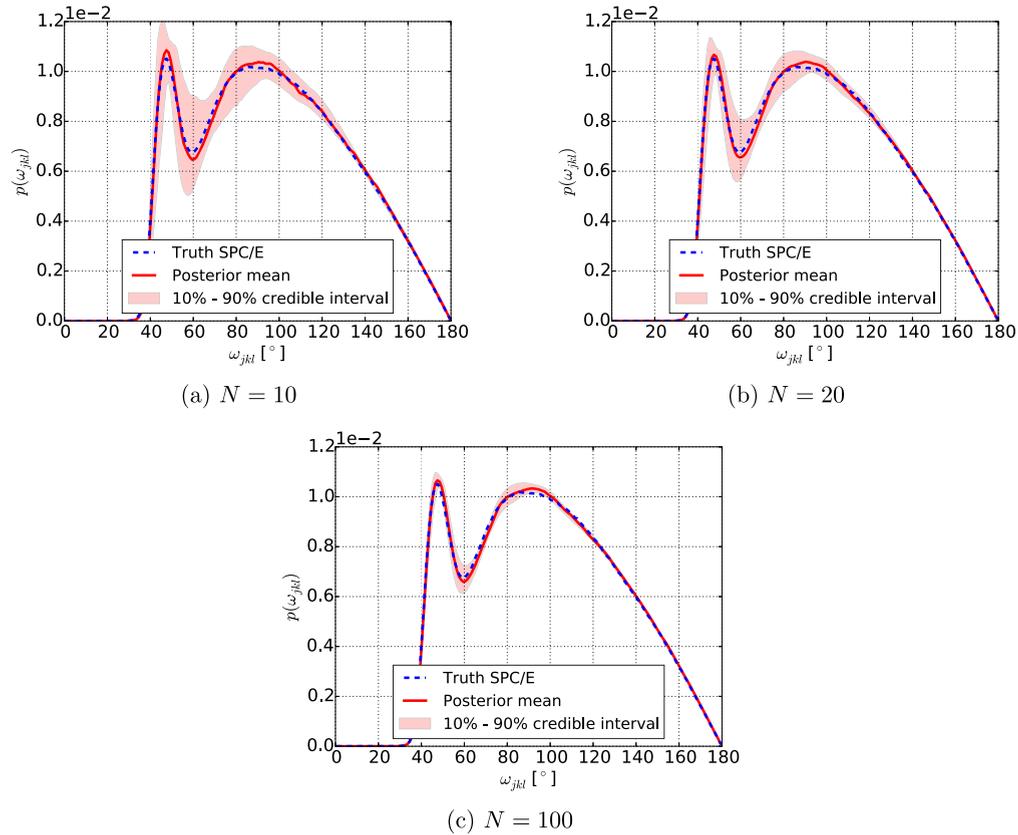


Fig. 19. Comparison of the reference ADF $p(\omega)$ (computed with all-atom simulations using the SPC/E model) with posterior mean and credible intervals corresponding to 10% and 90% posterior quantiles.

that unavoidably occurs during the CG process. We showed that from an information-theoretic perspective, the framework proposed broadens the relative entropy method. Furthermore, it can be readily extended to a fully Bayesian model where various sources of uncertainties are reflected in the posterior of the model parameters. The latter can be used to produce not only point estimates of fine-scale reconstructions or macroscopic observables, but more importantly, predictive posterior distributions on these quantities. We show how these can quantify the confidence of the model as a function of the amount of data and the level of coarse-graining, i.e. the contrast in the dimension between fine and coarse descriptions.

A critical issue in all CG methods pertains to the form of the coarse model or coarse potential. On one hand, it is desirable to introduce not only as many feature functions as possible but also to capture interactions of the highest-order possible. On the other hand, such an intricate representation leads to a large number of unknown parameters, augmented computational cost and an increased possibility of overfitting. Such challenges can be readily addressed within the Bayesian framework adopted by the incorporation of appropriate prior models that promote the discovery of sparse solutions and are capable of revealing the most dominant features in the coarse potential. We demonstrated how such a hierarchical prior model, namely the ARD, is capable of distinguishing the most prominent feature functions.

The computational engine of the proposed framework is based on an MC-EM scheme that alternates between expectations with respect to the posterior of the latent variables and maximization with respect to the model parameters. This leads to MAP estimates of the model parameters which serve as the basis for the Laplace's model that approximates their posterior. We note that this represents a very basic approximation that we intend to extend by exploiting advanced MCMC schemes [84] and/or variational inference schemes [85]. From a practical point of view, we note that the algorithm proposed is embarrassingly parallelizable with regards to the expectation step (which is also the most expensive) and incremental variants can be readily adopted leading to improvements in computational efficiency.

The generative definition of the CG variables through a probabilistic coarse-to-fine map allows for great flexibility in the type and number of CG variables used. For example in [23], the FG configuration space is partitioned and within each of these subdomains a different set of CG variables and CG models is learned. This is a reasonable strategy not only because a *globally-good* set of CG variables is difficult to find, but also because the local CG variables can be lower-dimensional as they need only to work on a limited subdomain. In the context of the directed, probabilistic model advocated, the same effect can be readily achieved by using a mixture model [86]. Consider for example augmenting the set of (latent) CG variables with a discrete-valued variable, S which can take values between 1 and L (which is the number of partitions). The (latent)

variable S characterizes a finite number of discrete states of the system. Depending on the value S takes, the number and type of CG variables \mathbf{X} can change by affecting the two distributions making up the mode, i.e.:

$$p_c(\mathbf{X}, S = s | \theta_c) = p_c(\mathbf{X} | \theta_c^s) p_c(S = s), \quad (67)$$

where each $p_c(\mathbf{X} | \theta_c^s)$ can be of the same or different form (e.g. exponential family) but with different parametrizations $\theta_c^s, s = 1, \dots, L$. Similarly for the coarse-to-fine map, we can define:

$$p_{cf}(\mathbf{x} | \mathbf{X}, S = s, \theta_{cf}) = p_{cf}(\mathbf{x} | \mathbf{X}, \theta_{cf}^s), \quad (68)$$

where again the parametrization can depend or not on $S, \theta_{cf}^s, s = 1, \dots, L$. Infinite mixture models [87–89] based on Dirichlet process priors could provide a rigorous strategy on determining the number L of such hidden states needed to describe the atomistic ensemble. We note finally that, in nonequilibrium settings, by appropriate modeling of the time dependence of S one would recover Hidden Markov Models (HMM, [46]) which have been employed in coarse-graining frameworks [90,91].

Another potentially powerful extension, involves the use of deep, hierarchical models. Deep learning tools have revolutionized various machine learning tasks [92] by stacking multiple layers of simple representations. In the context of coarse-graining, such a scheme could be materialized by augmenting the set of CG variables as $\mathbf{X}_1, \mathbf{X}_2, \dots, \mathbf{X}_L$ and the CG model as:

$$p_c(\mathbf{X}_1, \mathbf{X}_2, \dots, \mathbf{X}_L) = p_{c,1}(\mathbf{X}_1 | \mathbf{X}_2, \theta_c^1) \dots p_{c,L-1}(\mathbf{X}_{L-1} | \mathbf{X}_L, \theta_c^{L-1}) p_{c,L}(\mathbf{X}_L | \theta_c^L). \quad (69)$$

If $\dim(\mathbf{X}_1) > \dim(\mathbf{X}_2) > \dots > \dim(\mathbf{X}_L)$, then such a structure could provide a hierarchical decomposition of the CG picture, starting from a highly coarse description and gradually reaching the more detailed abstraction \mathbf{X}_1 . The coarse-to-fine map could be controlled by \mathbf{X}_1 as $p_{cf}(\mathbf{x} | \mathbf{X}_1, \theta_c)$.

Acknowledgements

We acknowledge the support by the Hans Fisher Senior Fellowship of Nicholas Zabaras of the Technical University of Munich–Institute for Advanced Study, funded by the German Excellence Initiative and the European Union Seventh Framework Programme under grant agreement No. 291763. Nicholas Zabaras also acknowledges support from the Computer Science and Mathematics Division of ORNL under the DARPA EQUIPS program.

Appendix A. Methodology

A.1. Estimating credible intervals

This note summarizes necessary steps for estimating credible intervals. The Bayesian inference algorithms described in Sections 2.3 and 2.6, lead to (Gaussian) approximations of the posterior $p(\theta | \mathbf{x}^{(1:N)})$ (Eq. (16)). The credible intervals shown in Figs. 11, 12, 18, and 19 are constructed from Monte Carlo samples $\hat{a}(\theta^{(i)})$ of the observables of interest. These are generated on the basis of Eq. (19) as follows:

Algorithm 2 Estimating credible intervals.

- 1: **for** all $i = 1, \dots, I$ **do**
- 2: Obtain a posterior sample: $\theta^{(i)} \sim p(\theta | \mathbf{x}^{(1:N)})$ (Eq. (16)).
- 3: Calculate the predictive estimate $\hat{a}(\theta^{(i)})$ shown in Eq. (19):

$$\hat{a}(\theta^{(i)}) = \int a(\mathbf{x}) p_{cf}(\mathbf{x} | \mathbf{X}, \theta_{cf}^{(i)}) p_c(\mathbf{X} | \theta_c^{(i)}) d\mathbf{X} d\mathbf{x}. \quad (A.1)$$

The integrations involved are performed with Monte Carlo sampling. We note that this requires simulating only the CG model as the mapping implied by p_{cf} is straightforward.

- 4: **end for**
 - 5: Compute desired quantiles with the given samples $\hat{a}(\theta^{(1..I)})$.
-

We note that the estimated quantiles of the corresponding predictive posterior are not necessarily symmetric around its MAP estimate $\hat{a}(\theta_{\text{MAP}})$, even in the case of a symmetric posterior of the model's parameters $p(\theta | \mathbf{x}^{(1:N)})$ (Eq. (16)).

A.2. Comparison of gradients between relative entropy method and PCC

This section compares the gradients with respect to the parameters of the coarse potential θ_c , between the proposed scheme and the relative entropy method. These are used for fitting the model parameters θ_c . In our case, the gradient is given by:

$$\frac{\partial \mathcal{F}}{\partial \theta_{c,k}} = \sum_{i=1}^N \left(\langle \phi_k(\mathbf{X}^{(i)}) \rangle_{q_i(\mathbf{X}^{(i)})} - \langle \phi_k(\mathbf{X}) \rangle_{p_c(\mathbf{X}|\theta_c)} \right), \quad (\text{A.2})$$

whereas for the relative entropy method (when the objective \mathcal{F}_{KL} is given as in Eq. (10)):

$$\begin{aligned} \frac{\partial \mathcal{F}_{\text{KL}}}{\partial \theta_{c,k}} &= \langle \phi_k(\mathcal{R}(\mathbf{x})) \rangle_{p_f(\mathbf{x})} - \langle \phi_k(\mathbf{X}) \rangle_{p_c(\mathbf{X}|\theta_c)} \\ &\approx \frac{1}{N} \sum_{i=1}^N \left(\phi_k(\mathcal{R}(\mathbf{x}^{(i)})) - \langle \phi_k(\mathbf{X}) \rangle_{p_c(\mathbf{X}|\theta_c)} \right). \end{aligned} \quad (\text{A.3})$$

In the latter case, the expectations with respect to $p_f(\mathbf{x})$ are estimated using the fine-scale data $\mathbf{x}^{(i)}$ whereas in the former these involve averaging over the *posterior* of the CG variables \mathbf{X} . This emphasizes the role of the CG variables play in our model as latent (hidden) generators of the fine-scale.

A.3. ARD prior

We adopt the Automatic Relevance Determination (ARD, [66]) which is formulated in the context of hierarchical Bayesian models. The prior on the parameters θ_c is modeled as independent Gaussian for each $\theta_{c,k}$ with zero mean and precision hyper-parameter τ_k :

$$p(\theta_c | \boldsymbol{\tau}) \equiv \prod_k \underbrace{\mathcal{N}(\theta_{c,k} | 0, \tau_k^{-1})}_{p(\theta_{c,k} | \tau_k)}. \quad (\text{A.4})$$

The precision (hyper-)parameters τ_k follow a Gamma distribution,

$$\tau_k \sim \text{Gamma}(\tau_k | a_0, b_0). \quad (\text{A.5})$$

Anytime derivatives of the log-prior are needed, an inner-loop Expectation–Maximization scheme can be employed which is based on the same ideas presented previously. In particular, for any set of densities $q_k(\tau_k)$ we can obtain a lower bound on the log-prior as follows:

$$\begin{aligned} \log p(\theta_c) &= \log \left(\prod_k \int p(\theta_{c,k} | \tau_k) p(\tau_k | a_0, b_0) d\tau_k \right) \\ &= \sum_k \log \int q_k(\tau_k) \frac{p(\theta_{c,k} | \tau_k) p(\tau_k | a_0, b_0)}{q_k(\tau_k)} d\tau_k \\ &\geq \sum_k \int q_k(\tau_k) \log \frac{p(\theta_{c,k} | \tau_k) p(\tau_k | a_0, b_0)}{q_k(\tau_k)} d\tau_k \quad (\text{Jensen's inequality}) \end{aligned} \quad (\text{A.6})$$

The optimal q_k i.e. the posteriors $p(\tau_k | \theta_{c,k})$ (for which the lower bound becomes tight) can be analytically computed and are Gamma densities with parameters $a_k = a_0 + \frac{1}{2}$, $b_k = b_0 + \frac{\theta_{c,k}^2}{2}$ [67], where the current values of $\theta_{c,k}$'s are used. This leads to the extremely simple iterations of the following form [67]:

- E-step: evaluate:

$$\langle \tau_k \rangle_{p(\tau_k | \theta_{c,k})} = \frac{a_k}{b_k} = \frac{a_0 + \frac{1}{2}}{b_0 + \frac{\theta_{c,k}^2}{2}}. \quad (\text{A.7})$$

- M-step: evaluate:

$$\begin{aligned} \frac{\partial \log p(\theta_c)}{\partial \theta_{c,k}} &= \frac{\partial}{\partial \theta_{c,k}} \int q_k(\tau_k) \log p(\theta_{c,k} | \tau_k) d\tau_k \\ &= - \int q_k(\tau_k) \tau_k d\tau_k \theta_{c,k} \\ &= - \langle \tau_k \rangle_{p(\tau_k | \theta_{c,k})} \theta_{c,k}. \end{aligned} \quad (\text{A.8})$$

Appendix B. Numerical examples

B.1. SPC/E model, parameters and simulation details

The following SPC/E parameters as given in [36,75] are used for producing the fine-scale data.

- LJ-potential: $\sigma = 3.166 \text{ \AA}$, $\epsilon = 0.650 \frac{\text{kJ}}{\text{mol}}$.
- Electrostatic load: $q_{\text{H}} = +0.4238 e$, $q_{\text{O}} = -0.8476 e$.
- Structural properties of rigid water model: bond-length $l_{\text{OH}} = 1.0 \text{ \AA}$ and bond-angle $\theta_{\text{HOH}} = 109.47^\circ$.
- Masses: $m_{\text{O}} = 15.994 \frac{\text{g}}{\text{mol}}$ and $m_{\text{H}} = 1.00794 \frac{\text{g}}{\text{mol}}$.

B.1.1. Simulation steps

In this work, we consider a system of $N_w = 100$ water molecules at a temperature $T = 300 \text{ K}$. The following steps for obtaining training data are performed:

1. NPT simulation with $p = 1 \text{ bar}$ and a timestep of $\Delta t = 2.0 \text{ fs}$. Simulate the system for $t = 100 \text{ ns}$.
2. Use the last $t = 80 \text{ ns}$ for calculating the equilibrium box size. We found $l_{\text{box}} = 14.5459665 \text{ \AA}$.
3. Fix the box length to the one obtained from the previous step. Simulate the system in NVT ensemble for $t = 45 \text{ ns}$ with a timestep of $\Delta t = 2.0 \text{ fs}$. Use the last $t = 40 \text{ ns}$ and write the trajectory every 200 steps.

B.2. Radial distribution function

The radial distribution function $g(r)$ is defined by,

$$g(r) = \left\langle \frac{V}{N^2} a^{\text{RDF}}(r) \right\rangle.$$

The discrete version follows with the number of bins n_{bin} and a bin size Δr :

$$g(r_1) = \frac{1}{N n_{\text{bin}}} \frac{\langle a^{\text{RDF}}(r_1) \rangle}{\rho_{\text{ideal}}},$$

with,

$$\rho_{\text{ideal}} = N/V,$$

$$a^{\text{RDF}}(r_1) = \frac{n(r_1)}{\Delta V} = \frac{\sum_{ij} \int_{r_1}^{r_1+\Delta r} \delta(r_{ij} - r) dr}{\frac{4}{3}\pi((r_1 + \Delta r)^3 - r_1^3)}.$$

B.3. Stillinger–Weber (SW) potential

The Stillinger–Weber (SW) potential originally proposed in [83] and extended in [81], contains both two- and three-body interactions. In this work, we make use only of the latter three-body contribution:

$$U^{\text{SW}}(\mathbf{X}) = \sum_j \sum_{k \neq j} \sum_{l > k} \phi_3^{\text{SW}}(r_{jk}, r_{jl}, \omega_{jkl}), \quad (\text{B.1})$$

where the three-body term $\phi_3^{\text{SW}}(r_{jk}, r_{jl}, \omega_{jkl})$ is given by:

$$\phi_3^{\text{SW}}(r_{jk}, r_{jl}, \omega_{jkl}) = \lambda \epsilon \left[\cos \omega_{jkl} - \cos \omega_0 \right]^2 \exp\left(\frac{\gamma \sigma}{r_{jk} - a_3 \sigma_{\text{SW}}}\right) \exp\left(\frac{\gamma \sigma}{r_{jl} - a_3 \sigma_{\text{SW}}}\right), \quad (\text{B.2})$$

with r_{jk} being the pairwise distances between molecules j and k and ω_{jkl} is the angle between molecules j, k, l . The following values for the parameters were used [81]: $\lambda = 0.762$, $\epsilon = 83.5737$, $\cos \omega_0 = -0.487217$, $\gamma = 0.291321$, $a_3 = 0.586097$, $\sigma_{\text{SW}} = 6.4144$.

References

- [1] M.S. Shell, The relative entropy is fundamental to multiscale and inverse thermodynamic problems, *J. Chem. Phys.* 129 (14) (2008) 144108.
- [2] B.J. Alder, T.E. Wainwright, Studies in molecular dynamics. I. General method, *J. Chem. Phys.* 31 (2) (1959) 459–466, <http://dx.doi.org/10.1063/1.1730376>, <http://scitation.aip.org/content/aip/journal/jcp/31/2/10.1063/1.1730376>.
- [3] M. Karplus, J.A. McCammon, Molecular dynamics simulations of biomolecules, *Nat. Struct. Biol.* 9 (9) (2002) 646–652, <http://dx.doi.org/10.1038/nsb0902-646>.
- [4] M.J. Buehler (Ed.), *Atomistic Modeling of Materials Failure*, Springer US, Boston, MA, 2008, <http://link.springer.com/10.1007/978-0-387-76426-9>.

- [5] C. Peter, K. Kremer, Multiscale simulation of soft matter systems – from the atomistic to the coarse-grained level and back, *Soft Matter* 5 (2009) 4357–4366, <http://dx.doi.org/10.1039/B912027K>.
- [6] G. Voth.
- [7] T. Lelièvre, M. Rousset, G. Stoltz, *Free Energy Computations: A Mathematical Perspective*, Imperial College Press, London, Hackensack (N.J.), Singapore, 2010, <http://opac.inria.fr/record=b1131369>.
- [8] I. Biliionis, P.S. Koutsourelakis, Free energy computations by minimization of Kullback–Leibler divergence: an efficient adaptive biasing potential method for sparse representations, *J. Comput. Phys.* 231 (9) (2012) 3849–3870, <http://dx.doi.org/10.1016/j.jcp.2012.01.033>, wOS:000302501500020.
- [9] M. Katsoulakis, A. Majda, D. Vlachos, Coarse-grained stochastic processes and Monte Carlo simulations in lattice systems, *J. Comput. Phys.* 186 (1) (2003) 250–278, [http://dx.doi.org/10.1016/S0021-9991\(03\)00051-2](http://dx.doi.org/10.1016/S0021-9991(03)00051-2).
- [10] A. Chatterjee, D.G. Vlachos, M.A. Katsoulakis, Spatially adaptive lattice coarse-grained Monte Carlo simulations for diffusion of interacting molecules, *J. Chem. Phys.* 121 (22) (2004) 11420–11431, <http://scitation.aip.org/content/aip/journal/jcp/121/22/10.1063/1.1811601>.
- [11] M.A. Katsoulakis, P. Plecháč, A. Sopsakis, Error analysis of coarse-graining for stochastic lattice dynamics, *SIAM J. Numer. Anal.* 44 (6) (2006) 2270–2296, <http://epubs.siam.org/doi/abs/10.1137/050637339>.
- [12] M.A. Katsoulakis, P. Plecháč, L. Rey-Bellet, Numerical and statistical methods for the coarse-graining of many-particle stochastic systems, *J. Sci. Comput.* 37 (1) (2008) 43–71, <http://dx.doi.org/10.1007/s10915-008-9216-6>, <http://www.springerlink.com/index/10.1007/s10915-008-9216-6>.
- [13] E. Kalligiannaki, M.A. Katsoulakis, P. Plecháč, D.G. Vlachos, Multilevel coarse graining and nano-pattern discovery in many particle stochastic systems, *J. Comput. Phys.* 231 (6) (2012) 2599–2620, <http://dx.doi.org/10.1007/s10915-008-9216-6>, <http://www.sciencedirect.com/science/article/pii/S0021999111007212>.
- [14] M.A. Katsoulakis, P. Plecháč, Information-theoretic tools for parametrized coarse-graining of non-equilibrium extended systems, *J. Chem. Phys.* 139 (7) (2013) 074115, <http://dx.doi.org/10.1063/1.4818534>, <http://scitation.aip.org/content/aip/journal/jcp/139/7/10.1063/1.4818534>.
- [15] W. Tschöp, K. Kremer, J. Batouli, T. Bürger, O. Hahn, Simulation of polymer melts. I. Coarse-graining procedure for polycarbonates, *Acta Polym.* 49 (2–3) (1998) 61–74.
- [16] D. Reith, M. Pütz, F. Müller-Plathe, Deriving effective mesoscale potentials from atomistic simulations, *J. Comput. Chem.* 24 (13) (2003) 1624–1636, <http://dx.doi.org/10.1002/jcc.10307>.
- [17] A.P. Lyubartsev, A. Laaksonen, Calculation of effective interaction potentials from radial distribution functions: A reverse Monte Carlo approach, *Phys. Rev. E* 52 (1995) 3730–3737.
- [18] J.F. Rudzinski, W.G. Noid, A generalized-Yvon–Born–Green method for coarse-grained modeling, *Eur. Phys. J. Spec. Top.* 224 (12) (2015) 2193–2216, <http://dx.doi.org/10.1140/epjst/e2015-02408-9>, <http://link.springer.com/article/10.1140/epjst/e2015-02408-9>.
- [19] A. Savelyev, G.A. Papoian, Molecular renormalization group coarse-graining of polymer chains: application to double-stranded {DNA}, *Biophys. J.* 96 (10) (2009) 4044–4052, <http://dx.doi.org/10.1016/j.bpj.2009.02.067>, <http://www.sciencedirect.com/science/article/pii/S0006349509006729>.
- [20] R.H. Swendsen, Monte Carlo renormalization group, *Phys. Rev. Lett.* 42 (14) (1979) 859.
- [21] S. Izvekov, G.A. Voth, Multiscale coarse graining of liquid-state systems, *J. Chem. Phys.* 123 (13) (2005) 134105.
- [22] W.G. Noid, J. Chu, G.S. Ayton, G.A. Voth, Multiscale coarse-graining and structural correlations: connections to liquid-state theory, *J. Phys. Chem. B* 111 (16) (2007) 4116–4127.
- [23] J.F. Dama, A.V. Sinitzkiy, M. McCullagh, J. Weare, B. Roux, A.R. Dinner, G.A. Voth, The theory of ultra-coarse-graining. 1. General principles, *J. Chem. Theory Comput.* 9 (5) (2013) 2466–2480, <http://dx.doi.org/10.1021/ct4000444>, PMID: 26583735.
- [24] A.R. Leach, *Molecular Modelling*, Prentice Hall, 2001, http://www.ebook.de/de/product/3246977/andrew_r_leach_molecular_modelling.html.
- [25] M.A. Katsoulakis, J. Trashorras, Information loss in coarse-graining of stochastic particle dynamics, *J. Stat. Phys.* 122 (1) (2006) 115–135, <http://link.springer.com/article/10.1007/s10955-005-8063-1>.
- [26] T.T. Foley, M.S. Shell, W.G. Noid, The impact of resolution upon entropy and information in coarse-grained models, *J. Chem. Phys.* 143 (24) (2015) 243104, <http://dx.doi.org/10.1063/1.4929836>, <http://scitation.aip.org/content/aip/journal/jcp/143/24/10.1063/1.4929836>.
- [27] M.A. Katsoulakis, P. Plecháč, A. Sopsakis, Error analysis of coarse-graining for stochastic lattice dynamics, *SIAM J. Numer. Anal.* 44 (6) (2006) 2270–2296, <http://dx.doi.org/10.1137/050637339>.
- [28] J. Trashorras, D. Tsagkarogiannis, From mesoscale back to microscale: reconstruction schemes for coarse-grained stochastic lattice systems, *SIAM J. Numer. Anal.* 48 (5) (2010) 1647–1677, <http://dx.doi.org/10.1137/080722382>, <https://hal.archives-ouvertes.fr/hal-00275802>.
- [29] M.A. Rohrdanz, W. Zheng, C. Clementi, Discovering mountain passes via torchlight: methods for the definition of reaction coordinates and pathways in complex macromolecular reactions, *Annu. Rev. Phys. Chem.* 64 (2013) 295–316, <http://www.annualreviews.org/doi/abs/10.1146/annurev-physchem-040412-110006>.
- [30] W.G. Noid, Perspective: Coarse-grained models for biomolecular systems, *J. Chem. Phys.* 139 (9) (2013), <http://dx.doi.org/10.1063/1.4818908>, <http://scitation.aip.org/content/aip/journal/jcp/139/9/10.1063/1.4818908>.
- [31] P. Angelikopoulos, C. Papadimitriou, P. Koumoutsakos, Bayesian uncertainty quantification and propagation in molecular dynamics simulations: a high performance computing framework, *J. Chem. Phys.* 137 (14) (2012), <http://dx.doi.org/10.1063/1.4757266>, <http://scitation.aip.org/content/aip/journal/jcp/137/14/10.1063/1.4757266>.
- [32] P. Angelikopoulos, C. Papadimitriou, P. Koumoutsakos, Data driven, predictive molecular dynamics for nanoscale flow simulations under uncertainty, *J. Phys. Chem. B* 117 (47) (2013) 14808–14816.
- [33] K. Farrell, J.T. Oden, Calibration and validation of coarse-grained models of atomic systems: application to semiconductor manufacturing, *Comput. Mech.* 54 (1) (2014) 3–19, <http://dx.doi.org/10.1007/s00466-014-1028-y>.
- [34] K. Farrell, J.T. Oden, D. Faghihi, A Bayesian framework for adaptive selection, calibration, and validation of coarse-grained models of atomistic systems, *J. Comput. Phys.* 295 (2015) 189–208, <http://dx.doi.org/10.1016/j.jcp.2015.03.071>, <http://www.sciencedirect.com/science/article/pii/S0021999115002430>.
- [35] A. Chaimovich, M.S. Shell, Coarse-graining errors and numerical optimization using a relative entropy framework, *J. Chem. Phys.* 134 (9) (2011) 094112, <http://dx.doi.org/10.1063/1.3557038>, <http://www.ncbi.nlm.nih.gov/pubmed/21384955>.
- [36] I. Biliionis, N. Zabarar, A stochastic optimization approach to coarse-graining using a relative-entropy framework, *J. Chem. Phys.* 138 (4) (2013) 044313, <http://dx.doi.org/10.1063/1.4789308>, <http://www.ncbi.nlm.nih.gov/pubmed/23387590>.
- [37] E. Cancès, F. Legoll, G. Stoltz, Theoretical and numerical comparison of some sampling methods for molecular dynamics, *edpsciences.org*, <http://www.edpsciences.org/articles/m2an/abs/2007/02/m2an0588/m2an0588.html>.
- [38] J.F. Rudzinski, W.G. Noid, Coarse-graining entropy, forces, and structures, *J. Chem. Phys.* 135 (2011) 214101, <http://dx.doi.org/10.1063/1.3663709>, <http://scitation.aip.org/content/aip/journal/jcp/135/21/10.1063/1.3663709>.
- [39] T. Cover, J. Thomas, *Elements of Information Theory*, John Wiley & Sons, 1991.
- [40] C. Bishop, Latent variable models, in: M.I. Jordan (Ed.), *Learning in Graphical Models*, MIT Press, 1999, pp. 371–403.
- [41] C.M. Bishop, *Pattern Recognition and Machine Learning (Information Science and Statistics)*, Springer-Verlag New York, Inc., Secaucus, NJ, USA, 2006.
- [42] A. Dempster, N. Laird, D. Rubin, Maximum likelihood from incomplete data via the EM algorithm (with discussion), *J. R. Stat. Soc. B* 39 (1) (1977) 1–38.
- [43] R. Neal, G.E. Hinton, A view of the EM algorithm that justifies incremental, sparse, and other variants, in: *Learning in Graphical Models*, Kluwer Academic Publishers, 1998, pp. 355–368.
- [44] G.C.G. Wei, M.A. Tanner, A Monte Carlo implementation of the EM algorithm and the poor man's data augmentation algorithms, *J. Am. Stat. Assoc.* 85 (411) (1990) 699–704, <http://dx.doi.org/10.1080/01621459.1990.10474930>, <http://www.tandfonline.com/doi/abs/10.1080/01621459.1990.10474930>.

- [45] H. Robbins, S. Monro, A stochastic approximation method, *Ann. Math. Stat.* 22 (3) (1951) 400–407.
- [46] O. Cappé, E. Moulines, T. Ryden, *Inference in Hidden Markov Models*, Springer-Verlag, 2005.
- [47] M.J. Beal, Variational Algorithms for Approximate Bayesian Inference, Ph.D. thesis, Gatsby Computational Neuroscience Unit, University College London, 2003, <http://www.cse.buffalo.edu/faculty/mbeal/thesis/index.html>.
- [48] L. Younes, On the convergence of Markovian stochastic algorithms with rapidly decreasing ergodicity rates, *Stoch. Stoch. Rep.* 65 (3–4) (1999) 177–228, <http://dx.doi.org/10.1080/17442509908834179>, <http://www.tandfonline.com/doi/abs/10.1080/17442509908834179>.
- [49] C. Andrieu, E. Moulines, P. Priouret, Stability of stochastic approximation under verifiable conditions, *SIAM J. Control Optim.* 44 (1) (2005) 283–312, <http://dx.doi.org/10.1137/S0363012902417267>, <http://epubs.siam.org/doi/abs/10.1137/S0363012902417267>.
- [50] G. Fort, E. Moulines, A. Schreck, M. Vihola, Convergence of Markovian stochastic approximation with discontinuous dynamics, *SIAM J. Control Optim.* (2016) 866–893, <http://dx.doi.org/10.1137/140962723>.
- [51] J.G. Booth, J.P. Hobert, Maximizing generalized linear mixed model likelihoods with an automated Monte Carlo EM algorithm, *J. R. Stat. Soc., Ser. B, Stat. Methodol.* 61 (1) (1999) 265–285, <http://www.jstor.org/stable/2680750>.
- [52] R.A. Levine, G. Casella, Implementations of the Monte Carlo EM algorithm, *J. Comput. Graph. Stat.* 10 (3) (2001) 422–439, <http://www.jstor.org/stable/1391097>.
- [53] G. Fort, E. Moulines, Convergence of the Monte Carlo expectation maximization for curved exponential families, *Ann. Stat.* 31 (4) (2003) 1220–1259, <http://dx.doi.org/10.1214/aos/1059655912>, <http://projecteuclid.org/euclid.aos/1059655912>.
- [54] R.A. Levine, J. Fan, An automated (Markov chain) Monte Carlo EM algorithm, *J. Stat. Comput. Simul.* 74 (5) (2004) 349–360, <http://dx.doi.org/10.1080/0094965031000147704>.
- [55] J.S. Liu, *Monte Carlo Strategies in Scientific Computing*, Springer Publishing Company, 2008, Incorporated.
- [56] P. Del Moral, Feynman–Kac Formulae: Genealogical and Interacting Particle Systems with Applications, Springer, New York, 2004.
- [57] J.C. Spall, *Introduction to Stochastic Search and Optimization*, 1st edition, John Wiley & Sons, Inc., New York, NY, USA, 2003.
- [58] S. Mohamed, K.A. Heller, Z. Ghahramani, Bayesian exponential family PCA, in: D. Koller, D. Schuurmans, Y. Bengio, L. Bottou (Eds.), Proceedings of the Twenty-Second Annual Conference on Neural Information Processing Systems, Vancouver, British Columbia, Canada, December 8–11, 2008, in: Advances in Neural Information Processing Systems, vol. 21, Curran Associates, Inc., 2008, pp. 1089–1096, <http://papers.nips.cc/paper/3532-bayesian-exponential-family-pca>.
- [59] P. Moritz, R. Nishihara, M.I. Jordan, A linearly-convergent stochastic L-BFGS algorithm.
- [60] R.H. Byrd, G.M. Chin, W. Neveitt, J. Nocedal, On the use of stochastic Hessian information in optimization methods for machine learning, *SIAM J. Optim.* 21 (3) (2011) 977–995.
- [61] C. Chen, D. Carlson, Z. Gan, C. Li, L. Carin, Bridging the gap between stochastic gradient MCMC and stochastic optimization.
- [62] H. Kushner, G.G. Yin, *Stochastic Approximation and Recursive Algorithms and Applications*, vol. 35, Springer Science & Business Media, 2003.
- [63] S. Della Pietra, V. Della Pietra, J. Lafferty, Inducing features of random fields, *IEEE Trans. Pattern Anal. Mach. Intell.* 19 (4) (1997) 380–393.
- [64] M.A.T. Figueiredo, Adaptive sparseness for supervised learning, *IEEE Trans. Pattern Anal. Mach. Intell.* 25 (9) (2003) 1150–1159.
- [65] M. West, Bayesian factor regression models in the “large p, small n” paradigm, in: J. Bernardo, M. Bayarri, J. Berger, A.P. Dawid, D. Heckerman, A. Smith, M. West (Eds.), *Bayesian Stat.*, vol. 7, 2003.
- [66] D.J.C. MacKay, R.M. Neal, Automatic Relevance Determination for Neural Networks, Tech. rep., University of Cambridge, 1994.
- [67] C.M. Bishop, M.E. Tipping, Variational relevance vector machines, in: UAI, 2000, pp. 46–53.
- [68] D.J.C. MacKay, *Information Theory, Inference & Learning Algorithms*, Cambridge University Press, New York, NY, USA, 2002.
- [69] J.V. Selinger, Introduction to the Theory of Soft Matter, Springer-Verlag GmbH, 2015, http://www.ebook.de/de/product/24265794/jonathan_v_selinger_introduction_to_the_theory_of_soft_matter.html.
- [70] N. Ashcroft, N. Mermin, *Solid State Physics*, Saunders College, Philadelphia, 1976.
- [71] S. Are, M.A. Katsoulakis, P. Plecháč, L.R. Bellet, Multibody interactions in coarse-graining schemes for extended systems, *SIAM J. Sci. Comput.* 31 (2) (2008) 987–1015, <https://dx.doi.org/10.1137/080713276>.
- [72] H.J.C. Berendsen, J.R. Grigera, T.P. Straatsma, The missing term in effective pair potentials, *J. Phys. Chem.* 91 (24) (1987) 6269–6271, <http://dx.doi.org/10.1021/j100308a038>.
- [73] P.G. Kusalik, I.M. Svishchev, The spatial structure in liquid water, *Science* 265 (5176) (1994) 1219–1221, <http://dx.doi.org/10.1126/science.265.5176.1219>, <http://science.sciencemag.org/content/265/5176/1219>.
- [74] A. Chaimovich, M.S. Shell, Tetrahedrality and structural order for hydrophobic interactions in a coarse-grained water model, *Phys. Rev. E* 89 (2014) 022140, <http://dx.doi.org/10.1103/PhysRevE.89.022140>, <http://link.aps.org/doi/10.1103/PhysRevE.89.022140>.
- [75] V. Rühlle, C. Junghans, A. Lukyanov, K. Kremer, D. Andrienko, Versatile object-oriented toolkit for coarse-graining applications, *J. Chem. Theory Comput.* 5 (12) (2009) 3211–3223.
- [76] R. Erban, Coupling all-atom molecular dynamics simulations of ions in water with Brownian dynamics, *Proc. R. Soc. Lond. A, Math. Phys. Eng. Sci.* 472 (2186) (2016), <http://dx.doi.org/10.1098/rspa.2015.0556>, <http://rspa.royalsocietypublishing.org/content/472/2186/20150556>.
- [77] W.G. Hoover, Canonical dynamics: equilibrium phase-space distributions, *Phys. Rev. A* 31 (1985) 1695–1697, <http://dx.doi.org/10.1103/PhysRevA.31.1695>, <http://link.aps.org/doi/10.1103/PhysRevA.31.1695>.
- [78] S. Nosé, A unified formulation of the constant temperature molecular dynamics methods, *J. Chem. Phys.* 81 (1) (1984) 511–519, <http://dx.doi.org/10.1063/1.447334>, <http://scitation.aip.org/content/aip/journal/jcp/81/1/10.1063/1.447334>.
- [79] S. Plimpton, Fast parallel algorithms for short-range molecular dynamics, *J. Comput. Phys.* 117 (1) (1995) 1–19, <http://dx.doi.org/10.1006/jcph.1995.1039>, <http://www.sciencedirect.com/science/article/pii/S002199918571039X>.
- [80] G.N. Clark, C.D. Cappa, J.D. Smith, R.J. Saykally, T. Head-Gordon, The structure of ambient water, *Mol. Phys.* 108 (11) (2010) 1415–1433.
- [81] J. Lu, Y. Qiu, R. Baron, V. Molinero, Coarse-graining of TIP4P/2005, TIP4P-Ew, SPC/E, and TIP3P to monatomic anisotropic water models using relative entropy minimization, *J. Chem. Theory Comput.* 10 (9) (2014) 4104–4120, <http://dx.doi.org/10.1021/ct500487h>, PMID: 26588552.
- [82] H. Wang, C. Junghans, K. Kremer, Comparative atomistic and coarse-grained study of water: what do we lose by coarse-graining?, *Eur. Phys. J. E* 28 (2) (2009) 221–229, <http://dx.doi.org/10.1140/epje/i2008-10413-5>.
- [83] F.H. Stillinger, T.A. Weber, Computer simulation of local order in condensed phases of silicon, *Phys. Rev. B* 31 (1985) 5262–5271, <http://dx.doi.org/10.1103/PhysRevB.31.5262>, <http://link.aps.org/doi/10.1103/PhysRevB.31.5262>.
- [84] F. Liang, A double Metropolis–Hastings sampler for spatial models with intractable normalizing constants, *J. Stat. Comput. Simul.* 80 (9) (2010) 1007–1022, <http://dx.doi.org/10.1080/00949650902882162>.
- [85] M. Wainwright, M. Jordan, Graphical models, exponential families, and variational inference, in: *Foundations and Trends in Machine Learning*, vol. 1, 2008, pp. 1–305.
- [86] C.M. Bishop, M. Svenskn, Bayesian hierarchical mixtures of experts, in: Proceedings of the Nineteenth Conference on Uncertainty in Artificial Intelligence, UAI’03, Morgan Kaufmann Publishers Inc., San Francisco, CA, USA, 2003, pp. 57–64, <http://dl.acm.org/citation.cfm?id=2100584.2100591>.
- [87] C. Antoniak, Mixtures of Dirichlet processes with applications to nonparametric Bayesian problems, *Ann. Stat.* 2 (1974) 1152–1174.
- [88] C.E. Rasmussen, The infinite Gaussian mixture model, in: Advances in Neural Information Processing Systems 12, NIPS Conference, Denver, Colorado, USA, November 29–December 4, 1999, 1999, pp. 554–560, <http://papers.nips.cc/paper/1745-the-infinite-gaussian-mixture-model>.

- [89] P. Chen, N. Zabaras, I. Bilonis, Uncertainty propagation using infinite mixture of Gaussian processes and variational Bayesian inference, *J. Comput. Phys.* 284 (2015) 291–333, <http://dx.doi.org/10.1016/j.jcp.2014.12.028>, <http://www.sciencedirect.com/science/article/pii/S0021999114008456>.
- [90] A. Fischer, S. Waldhausen, I. Horenko, E. Meerbach, C. Schütte, Identification of biomolecular conformations from incomplete torsion angle observations by hidden Markov models, *J. Comput. Chem.* 28 (2007) 2453–2464.
- [91] I. Horenko, F. Noe, C. Hartmann, C. Schütte, Data-based parameter estimation of generalized multidimensional Langevin processes, *Phys. Rev. E* 78 (2007) 016706.
- [92] Y. LeCun, Y. Bengio, G. Hinton, Deep learning, *Nature* 521 (7553) (2015) 436–444, <http://dx.doi.org/10.1038/nature14539>, <http://www.nature.com/nature/journal/v521/n7553/full/nature14539.html>.