Automated determination of hybrid particle-field parameters by machine learning

Morten Ledum , Sigbjørn Løland Bore & Michele Cascella

Published online: 02 Jul 2020.

Submit your article to this journal

Article views: 211

View related articles

View Crossmark data

Taylor & Francis
Taylor & Francis Group

OPEN ACCESS Check for updates

Automated determination of hybrid particle-field parameters by machine learning

Morten Ledum ⓘ, Sigbjørn Løland Bore ⓘ and Michele Cascella ⓘ

Department of Chemistry and Hylleraas Centre for Quantum Molecular Sciences, University of Oslo, Oslo, Norway

ABSTRACT
The hybrid particle-field molecular dynamics method is an efficient alternative to standard particle-based coarse grained approaches. In this work, we propose an automated protocol for optimisation of the effective parameters that define the interaction energy density functional, based on Bayesian optimisation. The machine-learning protocol makes use of an arbitrary fitness function defined upon a set of observables of relevance, which are optimally matched by an iterative process. Employing phospholipid bilayers as test systems, we demonstrate that the parameters obtained through our protocol are able to reproduce reference data better than currently employed sets derived by Flory-Huggins models. The optimisation procedure is robust and yields physically sound values. Moreover, we show that the parameters are satisfactorily transferable among chemically analogous species. Our protocol is general, and does not require heuristic a posteriori rebalancing. Therefore it is particularly suited for optimisation of reliable hybrid particle-field potentials of complex chemical mixtures, and extends the applicability corresponding simulations to all those systems for which calibration of the density functionals may not be done via simple theoretical models.
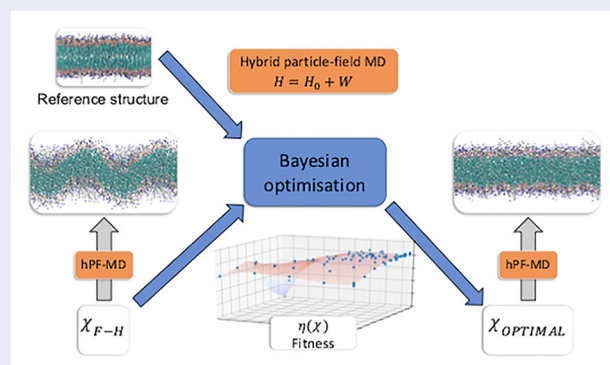
1. Introduction

Hybrid particle-field (hPF) simulations are a class of efficient methods that are well adapted for studying very large soft matter systems with molecular resolution [1–4]. The essence of the hPF methodology is contained in the two terms of the hPF Hamiltonian:

$$H(\{\mathbf{r}\}) = \sum_m H_0(\{\mathbf{r}\}_m) + W[\{\phi(\mathbf{r})\}]. \quad (1)$$

Here H_0, the Hamiltonian of a single molecule m, contains the *kinetic energy* and the *intramolecular potential* as defined in standard particle-based potentials, and W,

the interaction energy functional [5–7] dependent on the density-fields $\phi(\mathbf{r})$ of the different particle species, models all *intermolecular interactions*.

Intramolecular forces, by their very nature, only act on a single molecule, while the density-field interactions manifest as a quasi-instantaneous external potential, coupling the motion of the different molecules. The possibility of computing the external potential using particle-mesh routines allows for a very efficient and highly parallel implementation requiring very little communication among processors, resulting in algorithms formally exhibiting strong-scaling [8,9]. Very recently, a

CONTACT M. Cascella ✉ michele.cascella@kjemi.uio.no ⬛ Department of Chemistry and Hylleraas Centre for Quantum Molecular Sciences, University of Oslo, P.O. 1033 Blindern, 0315 Oslo, Norway

GPU-based implementation of the Monte Carlo-based hPF (*single chain in mean field*) set a new milestone with simulations of polymer melts composed of 10 billion particles [9].

The coupling of hPF to molecular dynamics in efficient parallelised software [8,10] has allowed for the application of hPF simulations on both conventional soft polymer mixtures and biological systems [11–15]. Prominent examples range from nanostructured multiphase materials [16–19] to organised and disorganised lipid/water mixtures [20,21]. Recently, hPF was extended to simulations of polypeptides [15], and to include explicit treatment of electrostatic interactions [22–24], the latter opening up to the formulation of density functional-based computational predictive models of the complex phase behavior of lipopolysaccharides [25].

Despite the growing level of maturity reached by hPF simulations, so far relatively little attention has been put into developing systematic protocols for the parameterisation of the interaction energy functional W. In particular, the quality of hPF models depends on both the physical model chosen for $W[\phi]$, and on the appropriate calibration of all the numerical parameters it may depend upon. The most commonly employed model for W typically takes the form of:

$$W[\phi(\mathbf{r})] = \frac{1}{2\phi_0} \int d\mathbf{r} \left(\sum_{ij} \tilde{\chi}_{ij}\phi_i(\mathbf{r})\phi_j(\mathbf{r}) \right.$$
$$\left. + \frac{1}{\kappa} \left(\sum_j \phi_j(\mathbf{r}) - \phi_0 \right)^2 \right), \qquad (2)$$

where the average number density of the system is denoted ϕ_0, κ is a compressibility term which controls the level of fluctuations of the overall density, and the $\tilde{\chi}_{ij}$ matrix is an energetic parameter that models local mixing energy between species i, j present in the system.

Parameters for the local mixing energy may be derived by different experimental approaches. For example, for simple polymers in a solvent, the $\tilde{\chi}$-parameter can be obtained from thermometric data [26]. This is however not as easily available when considering heteropolymeric systems. Another approach is to estimate $\tilde{\chi}$ by its relationship with the *Hildebrand solubility parameter* [27]. However this can be problematic as solubility parameters are often inaccurate [28]. Most importantly, for the molecular resolution of hPF models, which often adopt coarse grained (CG) representations in the range of four–ten atoms per bead, factorisation of global experimental data into the individual molecular components may not be trivial.

A more effective determination of $\tilde{\chi}$ parameters may be obtained using simple Flory-Huggins (F-H) lattice models:

$$\tilde{\chi}_{ij} = -z \left(\epsilon_{ij} - \tfrac{1}{2} \left(\epsilon_{ii} + \epsilon_{jj} \right) \right), \qquad (3)$$

where ϵ_{ij} is the mixing energy between species i and j, and z is the coordination number, which takes the value of 6 for three-dimensional Cartesian lattices. The mixing energy between two species can be approximated by the two-body interaction energy defined in the potential of the underlying molecular model employed. While this approach has been quite successful so far, there are a few limitations that hamper its general use. Prominently, the F-H model considers contact energies only, sometimes even disregarding entropy contributions to the binding, not taking into account many-body effects, or long-range interactions. The latter are particularly important, for example, in very polar or charged moieties. In practice, F-H parameters provide very good qualitative guesses for the values of $\tilde{\chi}$. Nonetheless, satisfactory quantitative agreement with reference data, especially in chemically complex systems, usually requires an *a posteriori* heuristic fine tuning of at least some of the values of the $\tilde{\chi}$ matrix [21].

Importantly, even though the first term of the interaction energy in (2) accounts in principle for the total energy of mixing, in recent times the addition of other terms to the W functional, for example explicitly describing electrostatics [22–24] or surface interactions [29,30], poses the problem of appropriately factorising such contributions out the mixing $\tilde{\chi}$ term to avoid non-physical double-counting. In these cases, $\tilde{\chi}$ loses a direct physical meaning, and for this reason it is problematic to define plausible values for $\tilde{\chi}$ directly from theoretical models.

The hPF interaction energy is globally dependent on a large set of parameters comprising both the $\tilde{\chi}$ matrix, and any other parameter present in other energy terms eventually employed. Therefore, the determination of an accurate functional W should be addressed as a global optimisation problem. Systematic approaches to parameterisation of ordinary particle-particle potentials in CG force fields, such as force matching [31], Iterative Boltzmann Inversion [32] and Inverse Monte Carlo [33], effectively consider parameterisation as optimisation problems where parameters are chosen to satisfy a given fitness function. For example, Iterative Boltzmann inversion and Inverse Monte Carlo consider a high resolution reference potential of mean force and optimise interaction potentials to reproduce this reference using the CG degrees of freedom. A key observation in such attempts is that the potential of mean force and interactions potentials, due to loss of entropy in the process

of CG, most often are significantly different. Similarly, the $\tilde{\chi}$ parameter of continuum density-field for polymers has not the direct meaning of a potential of mean force, but rather that of a phenomenological energetic term [7].

The determination of hPF $\tilde{\chi}$ force fields parameters poses a particularly challenging optimisation problem. First, these interactions cannot be framed as in a reaction coordinate form; therefore, $\tilde{\chi}$ parameters cannot be optimised through standard state of the art methods, such as Iterative Boltzmann inversion or Inverse Monte Carlo. Second, the gathering of hPF data does not yield derivatives of the model fitness with respect to the parameters, thereby restricting us to gradient-free optimisation. Finally, the $\tilde{\chi}$-matrix may involve a large parameter space for complex chemical mixtures, thus a general optimisation method needs to be capable of dealing with large dimensional parameter spaces.

Given such constraints, the large family of surrogate (or response surface methodology [RSM]) model based approaches, in which a response surface meta-model is introduced and updated through sequential noisy sampling, provides several possible optimisation techniques. Methods in the literature, of particular relevance, are classical sequential RSM [34,35], Lipschitz optimisation [36,37], Trust region methods [34,38], and Bayesian optimisation [39–41] (BO). In addition, various random search methods, such as genetic algorithms [42,43], simulated annealing [44–46], Latin hypercube sampling [47], or straight uniform random sampling, are applicable.

Among the cited methodologies, BO is a versatile scheme for the global optimisation of expensive nonlinear black-box functions for which derivatives with respect to the input parameters are hard or impossible to compute [39–41]. The BO algorithm, developed in the 70s, has in the last decade emerged as a strong solution to derivative-free optimisation of computationally expensive and noisy black-box functions, with powerful performance in many practical applications, especially within the field of machine learning hyper-parameter optimisation [48–52].

In this work we present a protocol for the optimisation of hPF parameters based on BO. The choice of this methodology is based on its strong theoretical convergence properties when paired with an upper-confidence bound (UCB) acquisition function [53], its simple implementation, and its highly data efficient sampling [54]. The effectiveness and robustness of our optimisation protocol is tested against uniform random sampling, the simplest possible optimisation strategy, and previous literature data based on F-H models.

2. Materials and methods

2.1. The hybrid particle-field method

The phase space of a molecular system with total energy (1) may be sampled either by Monte Carlo [1], or by molecular dynamics (hPF-MD) [3]. In this work we employed hPF-MD.

In hPF-MD, the equations of motion for the independent particles are determined by the presence of an external potential obtained as the functional derivative of W. Specifically, the potential acting on each particle species i located at position $\mathbf{r}$ takes the form [3]:

$$V_i^{\text{ext}}(\mathbf{r}) = \frac{\delta W\left[\phi(\mathbf{r})\right]}{\delta \phi_i(\mathbf{r})}$$

$$= \frac{1}{\phi_0}\left(\sum_j \tilde{\chi}_{ij}\phi_j(\mathbf{r}) + \frac{1}{\kappa}\left(\sum_j \phi_j(\mathbf{r}) - \phi_0\right)\right).$$

$$(4)$$

In the OCCAM hPF-MD software [8], which we employ in this paper, the related forces are evaluated via a numerical particle-mesh approach from spatial derivatives of the external potential:

$$\mathbf{F}_i^{\text{ext}}(\mathbf{r}) = -\nabla V_i(\mathbf{r}) = -\frac{1}{\phi_0}\sum_j\left(\tilde{\chi}_{ij} + \frac{1}{\kappa}\right)\nabla\phi_j(\mathbf{r}).$$

$$(5)$$

For more details on the computation of the forces, see ref. [3].

2.2. hPF force field parameterisation protocol

To determine hPF force field parameters, we employ a general iterative automated optimisation framework as depicted in Figure 1. Starting from a force field parameter set $\mathbf{x}$, a hPF trajectory is gathered and analysed giving output data of relevance $\mathbf{y}_{\text{sim.}}$. The output data is then compared to reference data $\mathbf{y}_{\text{ref.}}$, which can be provided by any accurate source, including high(er) resolution simulations or experiment. An objective (or fitness) function $\eta = \eta(\mathbf{y}_{\text{sim.}}, \mathbf{y}_{\text{ref.}}; \mathbf{x})$ assesses the quality of the parameterisation, and from the fitness value, the optimiser proposes a new hPF parameter set $\mathbf{x}$. The full cycle is automated and is repeated until satisfactory convergence of the fitness is reached, yielding the optimal hPF force field.

In principle, any optimiser that is not dependent on gradient values of the fitness values, can be employed. However, given the potentially large dimension of the parameter space $\mathbf{x}$ and the computationally expensive simulations needed to gather $\mathbf{y}_{\text{sim.}}$, it is essential that the optimiser should converge with the fewest possible
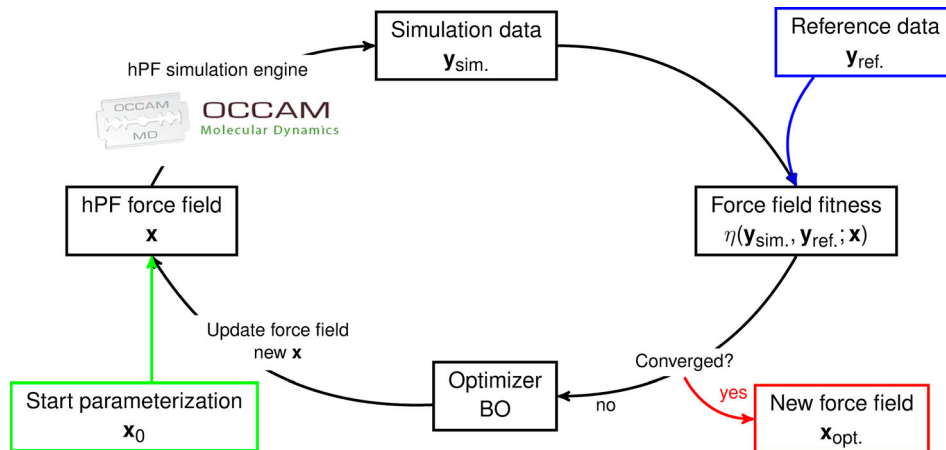
Figure 1. Protocol for optimising hPF force fields.

amount of iterations. Next, simulation data has an element of stochasticity due to only sampling a finite set of system configurations in the ensemble average, therefore the optimiser needs to be robust against noise in the fitness values. Finally, we note that computational expensiveness of gathering $\mathbf{y}_{\text{sim.}}$, makes almost any computational cost of the optimiser itself negligible.

The protocol we propose makes use of BO, a surrogate based model for solving constrained optimisation problems:

$$\mathbf{x}_{\text{opt}} = \arg\max_{\mathbf{x} \in \mathcal{X}} \eta(\mathbf{x}). \tag{6}$$

The space of possible parameter configurations $\mathcal{X}$ is usually a compact subset of $\mathbb{R}$ and the objective function η is in general unknown, non-convex, multimodal, and only accessible through (computationally expensive) point-wise noisy sampling. In the BO algorithm, a Gaussian process (GP) function prior is placed on the underlying true objective and updated via Bayesian posterior updating (Bayes' rule) by sequential probing of η [54]. In this way, a probabilistic response surface is built which represents, at each iteration, the model's beliefs about the objective (μ) and how confident the model is at each point in $\mathcal{X}$ (σ). BO achieves high efficiency in the sampling of the parameter space by leveraging both μ and σ in an *acquisition function* (AF), $a(\mathbf{x}) = a(\mu(\mathbf{x}), \sigma(\mathbf{x}))$. Often, the AF contains a parameter β which governs the trade-off between exploration (sampling areas in $\mathcal{X}$ where the uncertainty is high) and exploitation (sampling areas where *good* $\mathbf{x}$ are known to be located). The AF guides the sampling by picking points $\mathbf{x} \in \mathcal{X}$ to explore according to a strategy for improving upon the currently best found $\mathbf{x}$.

The GP prior is a multivariate Gaussian distribution over functions, uniquely defined by a covariance kernel Σ_0 and a mean function μ_0. The kernel function induces

a metric on $\mathcal{X}$ which defines a measure of the distance (similarity) between points $\mathbf{x}$ and $\mathbf{x}'$. The choice of a specific such Σ_0 represents a priori assumptions about the structure of the underlying true objective.

Often, one or more hyper-parameters in the covariance kernel have to be specified. It is customary to fix the values of these parameters by maximising the *marginal likelihood* of the model, given the observed data. Marginalising out the true noise-free objective function gives the likelihood of the model hyper-parameters. For GPs, the log marginal likelihood integral in question is analytically tractable, and may be easily maximised to determine the optimal kernel hyper-parameters.

2.3. Test case: phospholipid model for bilayers

As test case we consider a hPF-MD model for fully-saturated phospholipid bilayers, using in particular four variants characterised by different lengths of the fatty tail dipalmitoylphosphatidylcholine (DPPC), dimyristoylphosphatidylcholine (DMPC), distearoylphosphatidylcholine (DSPC), and mono-unsaturated dioleoylphosphatidylcholine (DOPC). For direct comparison, we use the same mapping of the model developed by De Nicola *et al.* [21] (Figure 2), which employs a MARTINI CG representation of the phospholipids [55] and explicit solvent.

In this work, we limit our analysis to the optimisation of the $\tilde{\chi}_{ij}$ matrix, while the bonded terms and the compressibility κ are kept the same in the model by De Nicola *et al.* [21].

A $(13 \times 13 \times 14\,\text{nm}^3)$ simulation box containing 528 DPPC lipids and 14,000 water beads is employed. Each simulation in the optimisation look lasts 20 ns. The hPF simulations were run using OCCAM.[1] The simulations are performed under the *NVT* ensemble, using the Andersen thermostat with a coupling time of 0.1 ps

CG mapping

Intramolecular Hamiltonian

$$H_0 = \sum \frac{m_i \dot{\mathbf{r}}_i^2}{2} + \sum \frac{k_r (r_{ij} - r_0)^2}{2} +$$
$$\sum \frac{k_\theta (\cos(\theta_{ijk}) - \cos(\theta_0))^2}{2}$$

Intermolecular interactions

$$W = \frac{1}{2\phi_0} \int d\mathbf{r} \left(\sum_{k\ell} \tilde{\chi}_{k\ell} \phi_k(\mathbf{r}) \phi_\ell(\mathbf{r}) \right.$$
$$\left. + \frac{1}{\kappa} \left(\sum_\ell \phi_\ell(\mathbf{r}) - \phi_0 \right)^2 \right).$$
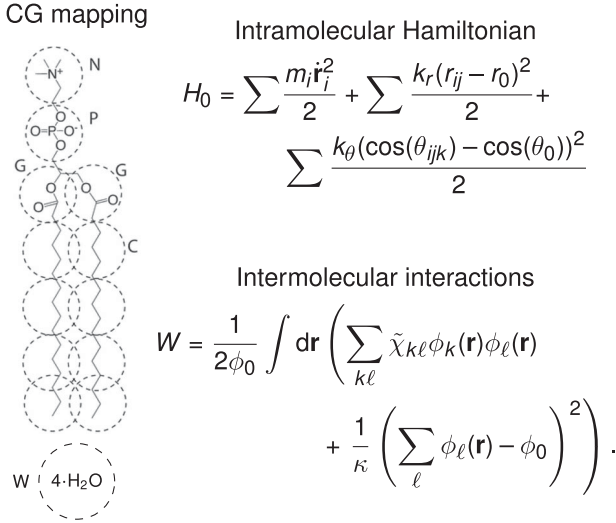
Figure 2. Summary of the hPF phospholipid model. *Left:* CG representation of the DPPC phospholipid and solvent. Right: Outline of the two terms in the hPF Hamiltonian.

and a collision frequency of $7.0 \, \text{ps}^{-1}$. A time step of 0.03 ps was used. The particle-mesh routines for particle-field forces in OCCAM employed a grid size of 0.58 nm (1.25 times the bond length used) and an update period of 0.3 ps (10 time steps). hPF-MD simulations are performed at a temperatures of 335, 325, and 303 K for DSPC, DPPC, and DMPC/DOPC, respectively. These temperatures were chosen to allow direct comparison with the results in [21].

To evaluate the fitness of the model we consider electron density profiles (φ) of the different species, compared to those obtained from reference CG simulations using the MARTINI force field. This choice is made to have the best assessment of the quality of the BO procedure as compared to the F-H parameters used in the literature currently. For optimal determination of hPF parameters for phospholipids, more accurate all-atom models or experimental data may be eventually employed. Note however that using an experimental target for the objective function adds a layer of complexity due to the incompleteness of the data as only the *total* electron density across the membrane normal is available, not the density of individual moieties.

The fitness $\eta(\varphi; \tilde{\chi})$ is defined as the average mean squared error over the electron densities of the different species k:

$$\eta(\varphi; \tilde{\chi}) = \frac{1}{nn_k} \sum_{k=1}^{n_k} \sum_{i=1}^{n} |\varphi_i^k - \hat{\varphi}_i^k|^2, \qquad (7)$$

with φ_i^k being the electron density of species k at a position $z_i = 2i\ell/n - \ell$ along the bilayer normal. The density

profiles are computed relative to the center of mass of all carbon type beads in the simulation, which is taken to be the center of the bilayer. $\hat{\varphi}_i^k$ indicates the reference density to be matched (in our case the MARTINI simulation results). The total number of different particle species is denoted n_k, while n is the number of bins in the chosen density histogram. For a better of comparison with F-H data [21], the absolute deviations S_k are also reported:

$$S_k(\varphi; \tilde{\chi}) = \frac{1}{n} \sum_{i=1}^{n} |\varphi_i^k - \hat{\varphi}_i^k|. \qquad (8)$$

In addition, S_p, the mean percentage error relative to the average electron density φ_0 over the full histogram across all species, is reported:

$$S_p(\varphi; \tilde{\chi}) = \frac{1}{\varphi_0 n n_k} \sum_{k=1}^{n_k} \sum_{i=1}^{n} |\varphi_i^k - \hat{\varphi}_i^k|, \qquad (9)$$

To avoid potential *cold-start* problems, each optimisation run is started with $2d$ (d being the dimension of the parameter space) randomly sampled points. After the initial random sampling period, new points to be probed are selected according to the maximum of the UCB acquisition function [53] (Figure 1). The exploration/exploitation trade-off parameter in the acquisition function is set to $\beta = 2$, favouring exploration of the large parameter space. The Gaussian process underlying the BO uses a Matérn covariance kernel [56,57] with smoothing parameter $\nu = 5/2$, and a constant zero mean function. In addition, a diagonal white noise kernel is added to account for the noisy sampling.

All $\tilde{\chi}$ parameters are constrained to the values used by De Nicola $\pm 10 \, \text{kJ mol}^{-1}$ [21]. Minimising the volume of the $\mathcal{X}$ hypercube is important for reducing the computational cost of the optimisation scheme, however the range used must be large enough to include the true optimum. As our fitness (MSE) is locally convex close to the true minimum, an insufficiently large parameter subset search would result in an optima on the edge of $\mathcal{X}$ being reported by the BO method, as opposed to an interior point. As no such results were found, the value used was deemed to sufficiently balance both concerns.

3. Results and discussion

3.1. Optimisation of hPF parameters for DPPC

DPPC was used as prototypic test systems to assess the effectiveness of BO for the determination of hPF $\tilde{\chi}$ parameters. The choice of such a system was determined both by the presence of a relatively complex chemical

structure, and by the existence of a vast reference literature, including experimental [58–60] and computer simulations [61,62], as well as hPF models [20,21].

Figure 3 reports the density profiles for hPF simulations of DPPC after BO of the $\tilde{\chi}$ parameters with respect to the mean-squared-error objective function, computed between the hPF and reference MARTINI density profiles (7). The density profiles for all the bead types match well those of the reference, with S_k values smaller than 7.3 el./nm^3 for all lipid beads, and with a S_p value less than 2%.

Previously published hPF-MD models for phospholipids are based on the MARTINI CG mapping [55,63,64] and employ a $\tilde{\chi}$ matrix based on the F-H model (3). F-H parameters are extracted from the corresponding Lennard-Jones binding energies of the MARTINI force field ($\tilde{\chi}_{F-H}$ hereafter [21]). As noted in the original work, using the lateral density profile as the benchmark property, heuristic adjustment of the $\tilde{\chi}$ parameter between C and W beads was required to improve the stability and overall structure of the bilayer. Overall, the F-H parameter set produces a satisfactory organisation of the lipid bilayer (Figure 3), evidenced by a very good qualitative agreement of the lateral density profiles for the different moieties compared to reference CG simulations using the MARTINI force field. Nonetheless, the hPF/F-H density profiles are characterised by a S_p of about 4 to 5% for the different lipids [21], and larger values of S_k, reaching a maximum of 20.51 el./nm^3 for the G bead. Comparison

of S_k and S_p values indicates that BO provides a substantial improvement compared to F-H.

Given the use of theoretical models for the derivation of $\tilde{\chi}$, it has been hard so far to discern the origin of any discrepancies from reference data between intrinsic approximations of the hPF method, or the use of non optimal parameter sets. In particular, broader density profiles in lipids were usually understood as a consequence of the intrinsic softness of the field interactions [22]. In fact, using BO parameters, there is an appreciable sharpening of the distributions for all the beads, even though the peaks remain broader than CG simulations based on pair-interactions (Figure 3). This is of particular interest, as it demonstrates that indeed, in phospholipids, the F-H parameterisation is accurate enough to capture the physics of the hPF model; nonetheless, there is still space for significant quantitative improvement by a global optimisation approach.

3.2. Feature importance

Data in Figure 3 show how the performance of the F-H parameters is not equal for all the moieties present in the system. In particular, F-H is better at reproducing the distributions of the lipid head and tails, while the density profile of the glycerol groups (G beads) appears too broad. The physical reason for such discrepancy may be attributed to the fact that glycerol floats at the interface between the phase-separated water and lipid fatty tails.
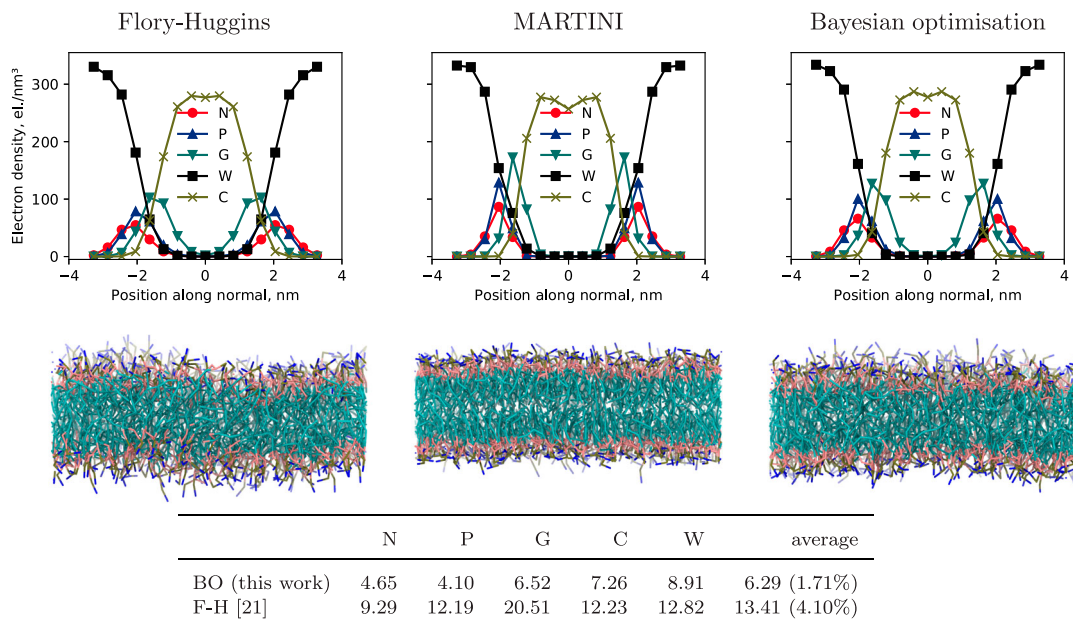

	N	P	G	C	W	average
BO (this work)	4.65	4.10	6.52	7.26	8.91	6.29 (1.71%)
F-H [21]	9.29	12.19	20.51	12.23	12.82	13.41 (4.10%)

Figure 3. Density profiles and representative membrane snapshots from hPF-MD simulations of a DPPC bilayer using $\tilde{\chi}_{F-H}$ parameters [21] (*left*), particle-based simulations using the MARTINI CG force field (*centre*), and hPF-MD $\tilde{\chi}_{BO}$ parameters (*right*). The table presents absolute deviations S_k in the density profiles between the F-H and BO parameter simulations, and the reference MARTINI profile. Percentage deviations S_p are given in parenthesis. S_k values are given in el./nm^3.
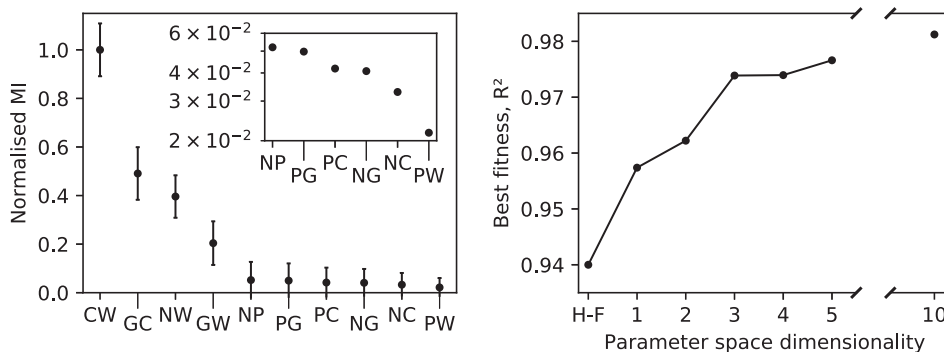
Figure 4. *Left:* Feature importance as ranked by the mutual information measure between the fitness and the individual $\tilde{\chi}_{ij}$ parameters, for hPF-MD simulations of a DPPC bilayer with randomly sampled $\tilde{\chi}$ matrices. Presented values are normalised relative to the most important parameter ($\tilde{\chi}_{CW}$) (arbitrary units). Inset details the low relative MI values found for the last six matrix elements (error bars omitted). $\tilde{\chi}$ parameters are included in order of decreasing feature importance (left). *Right:* Best fitness achieved (here using the average coefficient of determination, R^2, across all bead species) for each dimension of the parameter space subspace used in hPF-MD BO protocol runs on the DPPC bilayer system.

Therefore, its distribution depends more than the others on a delicate balance among all the terms in W. This effect may be difficult to reproduce adopting an independent parameterisation of the individual elements of the $\tilde{\chi}$ mixing energy matrix. On the contrary, the BO approach appears better suited to take into account all competing interactions, producing more balanced $\tilde{\chi}$ values.

The uneven error in the F-H distributions suggests that the hPF model is not equally robust with respect to variations of the different $\tilde{\chi}$ terms. To verify this hypothesis, we calculated the correlation (mutual information, MI) between input $\tilde{\chi}$ parameters and resulting fitness. The MI between two continuous random variables X and Y with probability density functions (PDF) f_X and f_Y (and joint PDF $f_{X,Y}$) is defined as [65,66]

$$I(X;Y) = -\int_{\mathcal{X}}\int_{\mathcal{Y}} dx\, dy f_{X,Y}(x,y) \log \frac{f_{X,Y}(x,y)}{f_X(x)f_Y(x)}, \tag{10}$$

and can be understood as the reduction in uncertainty about the values of Y, once X is revealed. The MI between any input parameter $\tilde{\chi}_{kj}$ and the resulting fitness $\eta(\phi; \tilde{\chi})$, thus yields a measure of the feature importance for the full parameter space.

Figure 4 shows the relative feature importance of the different $\tilde{\chi}_{ij}$ parameters, as well as optimisation results from BO runs which only include the most important ones. The fitness is here represented by the average coefficient of determination, R^2, over the density profiles of all the different beads,

$$\eta^{R^2}(\varphi; \tilde{\chi}) = \frac{1}{n_k}\sum_{k=1}^{n_k} R^2(\varphi^k, \hat{\varphi}^k). \tag{11}$$

Evidently, a subset of just four parameters carry the majority of the feature importance, meaning optimising

only these four, keeping the others at their F-H model value, yields results comparable to the ones obtained after an optimisation over the full 10-dimensional parameter space (Figures 4 and 5). The four relevant parameters have a clear physical meaning, as they are the main determinants for the hydrophilic/hydrophobic character of the polar heads and the fatty tails, respectively ($\tilde{\chi}_{NW}, \tilde{\chi}_{CW}$), and for the amphipathic behaviour of glycerol ($\tilde{\chi}_{GW}, \tilde{\chi}_{GC}$).

3.3. Transferability of BO-hPF parameters

Table 1 reports the parameter sets obtained by BO for DPPC compared to those obtained for two other saturated phospholipids differing in the length of fatty acid chains (DSPC, DMPC), and one unsaturated lipid (DOPC). Overall, the most relevant four $\tilde{\chi}$ matrix elements do not differ significantly from DPPC to DSPC. The less hydrophobic character of the C bead in DOPC may be attributed to the presence of the unsaturated moiety. We remark that for sake of simplicity, the C=C bond was represented by a different bead type, consistent with the MARTINI mapping, and all $\tilde{\chi}$ parameters involving that were kept at the reference F-H values [21].

The transferability of the obtained data sets is tested by performing hPF simulations for a lipid using parameters optimised on other structures. The absolute error on the density profiles obtained exchanging $\tilde{\chi}$ values are presented in Table 2, and show how the global structure of the bilayers remain mostly unaffected, with relatively small changes in S_k and S_p values, which remain systematically lower than those of the F-H parameterisation. This fact indicates that the BO protocol is able to find robust data-sets for chemically similar moieties, also ensuring very good transferability.

4 parameters   **10 parameters**
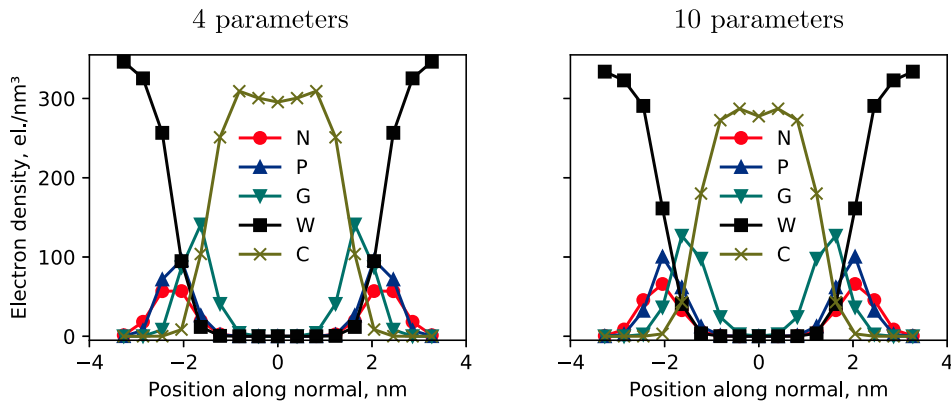
	parameter space dimensionality				
	10	4	3	2	1
C–W	42.24	43.68	43.63	42.09	38.16
G–C	10.47	14.00	15.33	14.69	6.30
N–W	-3.77	1.55	1.82	-8.10	-8.10
G–W	4.53	3.02	4.50	4.50	4.50
N–P	-9.34	-1.50	-1.50	-1.50	-1.50
P–G	8.04	4.50	4.50	4.50	4.50
N–G	1.97	6.30	6.30	6.30	6.30
P–C	14.72	13.50	13.50	13.50	13.50
P–W	-1.51	-3.60	-3.60	-3.60	-3.60
N–C	13.56	9.00	9.00	9.00	9.00
S_{p}	1.71%	1.96%	2.25%	2.29%	2.32%

Figure 5. *Top:* Density profiles for hPF-MD DPPC bilayer simulations ran with Bayesian optimised parameter sets with four (left) and ten (right) included $\tilde{\chi}$ parameters. The four-parameter simulation uses $\tilde{\chi}_{\mathrm{F-H}}$ values for all but the $\tilde{\chi}$ matrix elements with the highest feature importance, namely $\tilde{\chi}_{\mathrm{NW}}$, $\tilde{\chi}_{\mathrm{CW}}$, $\tilde{\chi}_{\mathrm{GW}}$, and $\tilde{\chi}_{\mathrm{GC}}$, c.f. column three of the table (bottom). *Bottom:* Resulting $\tilde{\chi}$ matrices from the BO protocol applied to hPF-MD simulations of a DPPC bilayer. Results reported for selected subspaces of the full 10-dimensional parameter space, with the $\tilde{\chi}_{ij}$s shown in red being fixed and not part of the optimisation run. All $\tilde{\chi}$ values given in kJmol^{-1}. Mean percentage errors, S_{p}, associated with each set of optimised parameters is given in the last row.

Table 1. Optimised $\tilde{\chi}$-matrix parameters found by the BO scheme for hPF-MD simulations of DPPC, DMPC, DSPC, and DOPC bilayer systems. The $\tilde{\chi}$ matrix elements are given in order of decreasing feature importance. Reference parameters are the Flory-Huggins ($\tilde{\chi}_{\mathrm{F-H}}$) parameters used in [21]. All values given in units of kJ mol^{-1}.

		Optimised with BO			
	Ref [21]	DPPC	DMPC	DSPC	DOPC
C–W	33.75	42.24	41.20	40.15	35.00
G–C	6.30	10.47	13.78	14.65	14.61
N–W	−8.10	−3.77	−2.58	−3.02	−2.46
G–W	4.50	4.53	5.91	4.71	9.07
N–P	−1.50	−9.34	−4.34	−5.91	−3.40
P–G	4.50	8.04	5.26	7.25	8.45
N–G	6.30	1.97	3.37	2.99	4.92
P–C	13.50	14.72	19.72	16.16	12.52
P–W	−3.60	−1.51	−1.26	−2.17	−1.27
N–C	9.00	13.56	12.71	10.56	14.39

3.4. Robustness of BO-hPF procedure

Large multidimensional parameter spaces often exhibit multiple locally optimal parameter sets or flat fitness surfaces that can hinder convergence towards the globally optimal parameter set. Figure 6 shows one such example

for the hPF parameters, with a projection of the fitness (estimated by the surrogate fitting function) in terms of $\tilde{\chi}_{\mathrm{CW}}$ and $\tilde{\chi}_{\mathrm{GC}}$. The plot exhibits a narrow region of unacceptable values, and a relatively large flat plateau of high score, were the determination of the position of the maximum is numerically non trivial, and may lead to multiple solutions.

However, our tests on transferability across lipid species do not indicate such problems for the BO-hPF procedure, finding instead systematically consistent parameter sets for the different lipid species. Moreover, as shown in Figure 6, BO converges steadily, and outperforms random sampling protocols in finding the optimal solution, even as both schemes improve upon $\tilde{\chi}_{\mathrm{F-H}}$ after only a handful of iterations. In particular, after a few efficient initial steps, random sampling is not able to converge toward the best solution, and remains confined in the large basin comprising of very different, not fully optimised, combinations of parameters. This is in agreement with results reported in the literature for BO applied to toy model functions [48], and such diverse fields as e.g. chemical design [67], active learning [68],

Table 2. Mean absolute deviations in electron density, S_k, with respect to the MARTINI reference density for the different lipids simulated with Bayesian optimised parameter sets on the different phospholipids (relative percentage deviations S_p in parenthesis). Comparison with data from De Nicola, using the baseline $\tilde{\chi}_{F-H}$ parameter set [21]. All S_k values given el./nm^3.

|	N	P	G	C	W	Average
DPPC						
DPPC-optimised	4.65	4.10	6.52	7.26	8.91	6.29 (1.71%)
DMPC-optimised	4.45	4.05	6.44	8.37	8.36	6.34 (1.97%)
DOPC-optimised	5.39	8.15	12.05	9.81	6.51	8.38 (2.61%)
DSPC-optimised	5.02	6.08	8.71	9.40	9.63	7.76 (2.40%)
Reference [21]	9.29	12.19	20.51	12.23	12.82	13.41 (4.10%)
DMPC						
DPPC-optimised	3.84	4.61	8.59	4.94	6.81	5.76 (1.85%)
DMPC-optimised	4.28	4.15	7.62	5.49	6.51	5.61 (1.81%)
DOPC-optimised	5.89	8.81	13.44	8.29	7.74	8.83 (2.87%)
DSPC-optimised	5.60	7.90	11.63	6.51	6.29	7.58 (2.45%)
Reference [21]	8.53	10.54	13.32	10.00	14.64	11.41 (3.63%)
DOPC						
DPPC-optimised	3.28	4.55	6.27	7.59	8.55	6.05 (2.03%)
DMPC-optimised	3.77	3.61	5.44	8.87	7.22	6.78 (1.96%)
DOPC-optimised	3.21	3.37	5.11	8.41	8.63	5.74 (1.95%)
DSPC-optimised	3.21	2.98	5.25	8.80	10.58	6.16 (2.08%)
Reference [21]	10.33	6.21	13.38	13.98	24.26	13.63 (4.79%)
DSPC						
DPPC-optimised	4.24	3.98	5.13	6.56	11.04	6.19 (1.86%)
DMPC-optimised	4.40	3.52	4.55	7.25	11.59	6.26 (1.88%)
DOPC-optimised	4.90	4.03	5.06	7.36	10.82	6.43 (1.94%)
DSPC-optimised	4.45	3.38	4.17	6.75	11.22	5.99 (1.80%)
Reference [21]	8.60	10.30	11.52	10.85	22.62	12.78 (3.80%)
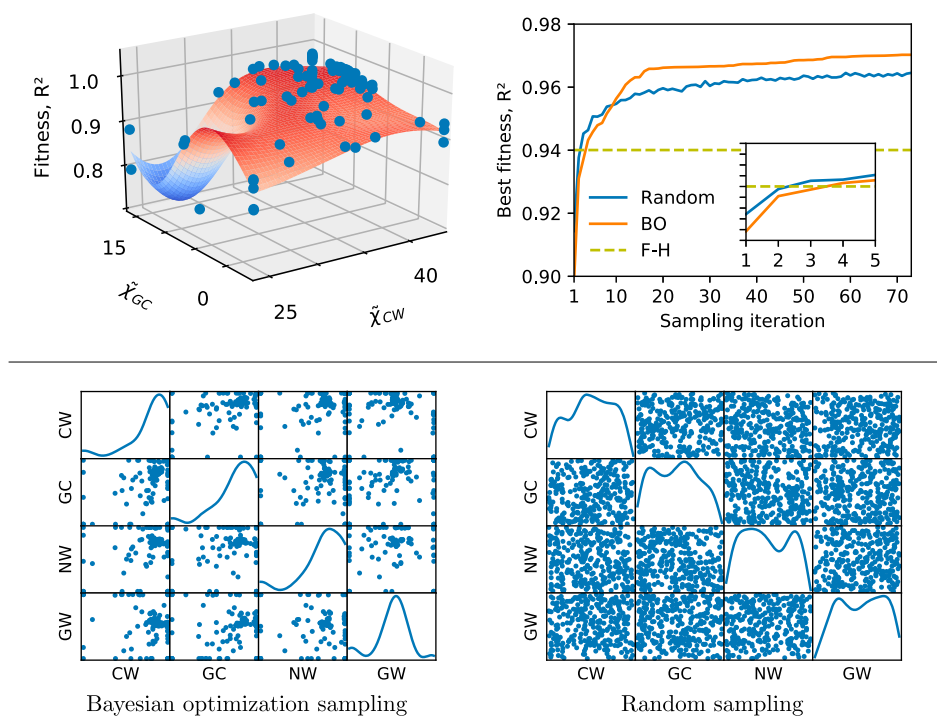
Figure 6. *Top left:* The surrogate objective fitness surface (here using the average coefficient of determination, R^2, across all bead species) in an example DPPC BO run with only the four parameters exhibiting the highest feature importance scores included ($\tilde{\chi}_{CW}$, $\tilde{\chi}_{GC}$, $\tilde{\chi}_{NW}$, and $\tilde{\chi}_{GW}$). Individual samplings with their associated fitnesses are represented as blue dots. A projection onto the subspace spanned by $\tilde{\chi}_{CW}$ and $\tilde{\chi}_{GC}$ shown. All $\tilde{\chi}$ matrix elements are given in kJ mol^{-1}. *Top right:* Best DPPC simulation membrane fitness (average R^2) for BO and random sampling with only the four parameters exhibiting the highest feature importance scores included. Comparison with the fitness achieved by the reference $\tilde{\chi}_{F-H}$ parameter set. Inset details when BO and random sampling surpass the $\tilde{\chi}_{F-H}$ parameter set in terms of R^2 fitness. *Bottom:* Scatter matrices showing correlations between all pairs of $\tilde{\chi}$ parameters in a BO run on a DPPC bilayer (*left*) compared with random sampling (*right*). Only the four parameters exhibiting the highest feature importance scores are included in the sampling. The matrix diagonal shows the density of sampled points for each individual $\tilde{\chi}_{ij}$ parameter. All $\tilde{\chi}$ matrix elements are given in kJ mol^{-1}.

robotics [69,70], and machine learning [71]. The faster and more robust convergence of BO is determined by the intrinsic ability of the algorithm to learn what region of the space is more relevant to sample, disregarding other less relevant regions (Figure 6).

4. Concluding remarks

In this work, we proposed a protocol to determine accurate potentials for hPF simulations, using BO as the main driver for the optimisation of the free parameters. Our scheme requires the definition of an arbitrary fitting function based on any set of relevant observables to be learned. The quantities of relevance may come from experimental data or from more accurate higher-resolution benchmark simulations (for example all-atom or CG), the only requirement being that the pertinent quantities can be straightforwardly estimated with a hPF model.

Using DPPC, DMPC, DSPC, and DOPC phospholipid bilayers as test systems, we showed how this procedure determines sets of parameters for the interaction energy that significantly improve the models present in the literature based on F-H theory. The new Bayesian-optimised potentials also show excellent transferability among chemically similar moieties.

Despite being more complex than F-H, the BO procedure here introduced offers various advantages. First, the procedure does not require the estimate of two-body interaction energies, which may be difficult to determine with good accuracy, for example, in the absence of CG models compatible with the mapping employed in the hPF simulations. Second, the protocol is very general, and can thus be used to concomitantly optimise the mixing terms of the interaction energy ($\tilde{\chi}$) and any other parameter of relevance present in other parts of the energy functional. This is particularly interesting in the view of recent advances for hPF model potentials, which include, for example, specific potentials for peptides, for electrostatics [22–24], or for surface energy terms [29,30]. Finally, being an automatic procedure, BO does not require user-based fine tuning of the parameters, ensuring a more systematic and reproducible determination of the potentials, especially for chemically complex systems.

BO is robust in determining physically meaningful parameters despite the relatively large variable space. This is due to the ability of BO to restrain the search only in a sub-region of the space where the physical solution is contained. Nonetheless, this evidence cannot be assumed as general, and it cannot be excluded that BO of hPF parameters over even higher-dimensional variable spaces would lead to numerical ambiguities. In this respect, we may suggest that the best strategy for the optimisation of hPF parameters implies the formulation of an adequate Ansatz, for example using the F-H method, that would be used as a starting point for the optimisation. In this work, we showed how feature importance can be applied to the BO procedure to identify on-the-fly those parameters that are not relevant for the convergence to the best solution, and which can be thus dropped out of the optimisation protocol. In this way, full BO optimisation can be performed only on a subset of relevant parameters, keeping all the others at (or in the neighbourhood of) their initial F-H values. In case the F-H parameters cannot be determined, or the parameter space is intrinsically too large, we foresee the possibility of introducing penalty terms to the fitting function, similarly to those used in other optimisation procedures like RESP [72], even though this has not been explored in this work.

In conclusion, the establishment of an automated machine-learning procedure for the optimisation of hPF parameters promises to further expand the applicability of this powerful simulation method toward increasingly chemically complex systems.

Note

1. http://www.occammd.org/

Data availability statement

The data that support the findings of this study are available from the corresponding author upon reasonable request.

Disclosure statement

No potential conflict of interest was reported by the authors.

ORCID

Morten Ledum 🔟 http://orcid.org/0000-0003-4244-4876
Sigbjørn Løland Bore 🔟 http://orcid.org/0000-0002-8620-4885
Michele Cascella 🔟 http://orcid.org/0000-0003-2266-5399

References

[1] K.C. Daoulas and M. Müller, J. Chem. Phys. **125** (18), 184904 (2006).

[2] M. Müller, J. Stat. Phys. **145** (4), 967–1016 (2011).

[3] G. Milano and T. Kawakatsu, J. Chem. Phys. **130** (21), 214106 (2009).

[4] G.G. Vogiatzis, G. Megariotis and D.N. Theodorou, Macromolecules **50** (7), 3004–3029 (2017).

[5] F. Schmid, J. Phys.: Condens. Matter **10** (37), 8105 (1998).

[6] T. Kawakatsu, *Statistical Physics of Polymers: An Introduction* (Springer Science & Business Media, Berlin, 2013).

[7] G. Fredrickson, *The Equilibrium Theory of Inhomogeneous Polymers*, Vol. 134 (Oxford University Press, Oxford, 2006).

[8] Y. Zhao, A. De Nicola, T. Kawakatsu and G. Milano, J. Comput. Chem. **33** (8), 868–880 (2012).

[9] L. Schneider and M. Müller, Comput. Phys. Commun. **235**, 463–476 (2019).

[10] Y.L. Zhu, H. Liu, Z.W. Li, H.J. Qian, G. Milano and Z.Y. Lu, J. Comput. Chem. **34** (25), 2197–2211 (2013).

[11] G. Milano, T. Kawakatsu and A. De Nicola, Phys. Biol. **10** (4), 045007 (2013).

[12] T.A. Soares, S. Vanni, G. Milano and M. Cascella, J. Phys. Chem. Lett. **8** (15), 3586–3594 (2017).

[13] M. Cascella and S. Vanni, in *Chemical Modelling: Applications and Theory*, edited by Michael Springborg and Jan-Ole Joswig (Royal Society of Chemistry, London, 2015), Vol. 12, Chap. 1, pp. 1–52.

[14] S.J. Marrink, V. Corradi, P.C.T. Souza, H.I. Ingólfsson, D.P. Tieleman and M.S.P. Sansom, Chem. Rev. **119** (9), 6184–6226 (2019).

[15] S.L. Bore, G. Milano and M. Cascella, J. Chem. Theory Comput. **14** (2), 1120–1130 (2018).

[16] A. De Nicola, T. Kawakatsu, F. Müller-Plathe and G. Milano, Eur. Phys. J. Spec. Top. **225** (8-9), 1817–1841 (2016).

[17] Y. Zhao, M. Byshkin, Y. Cong, T. Kawakatsu, L. Guadagno, A. De Nicola, N. Yu, G. Milano and B. Dong, Nanoscale **8** (34), 15538–15552 (2016).

[18] G. Munaò, A. Pizzirusso, A. Kalogirou, A. De Nicola, T. Kawakatsu, F. Müller-Plathe and G. Milano, Nanoscale **10** (46), 21656–21670 (2018).

[19] G. Munaò, A. De Nicola, F. Müller-Plathe, T. Kawakatsu, A. Kalogirou and G. Milano, Macromolecules **52** (22), 8826–8839 (2019).

[20] A. De Nicola, Y. Zhao, T. Kawakatsu, D. Roccatano and G. Milano, Theor. Chem. Acc. **131** (3), 1167 (2012).

[21] A. De Nicola, Y. Zhao, T. Kawakatsu, D. Roccatano and G. Milano, J. Chem. Theory Comput. **7** (9), 2947–2962 (2011).

[22] Y.L. Zhu, Z.Y. Lu, G. Milano, A.C. Shi and Z.Y. Sun, Phys. Chem. Chem. Phys. **18** (14), 9799–9808 (2016).

[23] H.B. Kolli, A. De Nicola, S.L. Bore, K. Schäfer, G. Diezemann, J. Gauss, T. Kawakatsu, Z.Y. Lu, Y.L. Zhu, G. Milano and M. Cascella, J. Chem. Theory Comput. **14** (9), 4928–4937 (2018).

[24] S.L. Bore, H.B. Kolli, T. Kawakatsu, G. Milano and M. Cascella, J. Chem. Theory Comput. **15** (3), 2033–2041 (2019).

[25] A. De Nicola, T.A. Soares, D.E.S. Santos, S.L. Bore, G.J.A. Sevink, M. Cascella and G. Milano, Biochim. Biophys. Acta **0**, 129570 (2020).

[26] J. Brandrup, E.H. Immergut, E.A. Grulke, A. Abe and D.R. Bloch, *Polymer Handbook*, Vol. 7 (John Wiley & Sons, New York, 1989).

[27] T. Lindvig, M.L. Michelsen and G.M. Kontogeorgis, Fluid Ph. Equilibria **203** (1-2), 247–260 (2002).

[28] S. Venkatram, C. Kim, A. Chandrasekaran and R. Ramprasad, J. Chem. Inf. Model. **59** (10), 4188–4194 (2019).

[29] A.P. Sgouros, A.T. Lakkas, G. Megariotis and D.N. Theodorou, Macromolecules **51** (23), 9798–9815 (2018).

[30] S.L. Bore, H.B. Kolli, A. De Nicola, M. Byshkin, T. Kawakatsu, G. Milano and M. Cascella, J. Chem. Phys. **152** (18), 184908 (2020).

[31] S. Izvekov and G.A. Voth, J. Phys. Chem. B **109** (7), 2469–2473 (2005).

[32] D. Reith, M. Pütz and F. Müller-Plathe, J. Comput. Chem. **24** (13), 1624–1636 (2003).

[33] A.P. Lyubartsev, in *Coarse-Grained Modeling of Biomolecules*, edited by Garegin A. Papoian (CRC Press, Boca Raton, 2017), Chap. 1, pp. 29–54.

[34] S. Amaran, N.V. Sahinidis, B. Sharda and S.J. Bury, Ann. Oper. Res. **240** (1), 351–380 (2016).

[35] R.H. Myers, D.C. Montgomery and C.M. Anderson-Cook, *Response Surface Methodology: Process and Product Optimization Using Designed Experiments* (John Wiley & Sons, New York, 2016).

[36] S.A. Piyavskii, USSR Comput. Math. & Math. Phys. **12** (4), 57–67 (1972).

[37] B.O. Shubert, SIAM J. Numer. Anal. **9** (3), 379–388 (1972).

[38] K.H. Chang, L.J. Hong and H. Wan, Informs. J. Comput. **25** (2), 230–243 (2013).

[39] J. Mockus, V. Tiesis and A. Zilinskas, *The Application of Bayesian Methods for Seeking the Extremum*, Vol. 2 (North-Holland, Amsterdam, 1978), pp. 117–129.

[40] H.J. Kushner, J. Math. Anal. Appl **5** (1), 150–167 (1962).

[41] M. Pelikan, D.E. Goldberg and E. Cantú-Paz, *Proceedings of the 1st Annual Conference on Genetic and Evolutionary Computation*, Vol. 1 (Morgan Kaufmann Publishers Inc., Burlington, 1999), pp. 525–532.

[42] D. Whitley, Stat. Comput. **4** (2), 65–85 (1994).

[43] C.R. Reeves, INFORMS J. Comput. **9** (3), 231–250 (1997).

[44] S. Kirkpatrick, C.D. Gelatt and M.P. Vecchi, Science **220** (4598), 671–680 (1983).

[45] D. Bertsimas and J. Tsitsiklis, Stat. Sci. **8** (1), 10–15 (1993).

[46] T.M. Alkhamis, M.A. Ahmed and V.K. Tuan, Eur. J. Oper. Res. **116** (3), 530–544 (1999).

[47] M.D. McKay, R.J. Beckman and W.J. Conover, Technometrics **21** (2), 239–245 (1979).

[48] J. Snoek, H. Larochelle and R.P. Adams, in *Advances in Neural Information Processing Systems 25*, edited by F. Pereira, C.J.C. Burges, L. Bottou and K.Q. Weinberger (Curran Associates, Inc., Red Hook, 2012), pp. 2951–2959.

[49] D.R. Jones, J. Global Optim. **21** (4), 345–383 (2001).

[50] J. Azimi, A. Jalali and X. Fern, preprint, arXiv:1202.5597 (2012)..

[51] J.S. Bergstra, R. Bardenet, Y. Bengio and B. Kégl, in *Advances in Neural Information Processing Systems 24*, edited by J. Shawe-Taylor, R.S. Zemel, P.L. Bartlett, F. Pereira and K.Q. Weinberger (Curran Associates, Inc., Red Hook, 2011), pp. 2546–2554.

[52] K. Swersky, J. Snoek and R.P. Adams, in *Advances in Neural Information Processing Systems 26*, edited by C.J.C.

Burges, L. Bottou, M. Welling, Z. Ghahramani and K.Q. Weinberger (Curran Associates, Inc., Red Hook, 2013), pp. 2004–2012.

[53] N. Srinivas, A. Krause, S.M. Kakade and M. Seeger, preprint, arXiv:0912.3995 (2009).

[54] B. Shahriari, K. Swersky, Z. Wang, R.P. Adams and N. de Freitas, Proc. IEEE **104** (1), 148–175 (2015).

[55] T.A. Wassenaar, H.I. Ingólfsson, R.A. Böckmann, D.P. Tieleman and S.J. Marrink, J. Chem. Theory Comput. **11** (5), 2144–2155 (2015).

[56] B. Matérn, *Spatial Variation*, Vol. 36 (Springer Science & Business Media, Berlin, 2013).

[57] M.L. Stein, *Interpolation of Spatial Data: Some Theory for Kriging* (Springer Science & Business Media, Berlin, 2012).

[58] H.I. Petrache, S.W. Dodd and M.F. Brown, Biophys. J. **79** (6), 3172–3192 (2000).

[59] Q. Waheed and O. Edholm, Biophys. J. **97** (10), 2754–2760 (2009).

[60] J.F. Nagle and S. Tristram-Nagle, Biochim. Biophys. Acta, Biomembr. **1469** (3), 159–195 (2000).

[61] E. Lindahl and O. Edholm, J. Chem. Phys. **113** (9), 3882–3893 (2000).

[62] Z.A. Levine, R.M. Venable, M.C. Watson, M.G. Lerner, J.E. Shea, R.W. Pastor and F.L. Brown, J. Am. Chem. Soc. **136** (39), 13582–13585 (2014).

[63] S.J. Marrink, A.H. de Vries and A.E. Mark, J. Phys. Chem. B **108** (2), 750–760 (2004).

[64] S.J. Marrink, H.J. Risselada, S. Yefimov, D.P. Tieleman and A.H. de Vries, J. Phys. Chem. B **111** (27), 7812–7824 (2007).

[65] C.E. Shannon, Bell Syst. Tech. J. **27** (3), 379–423 (1948).

[66] B. Frénay, G. Doquire and M. Verleysen, Neural Netw. **48**, 1–7 (2013).

[67] R.R. Griffiths and J.M. Hernández-Lobato, preprint, arXiv:1709.05501 (2017).

[68] E. Brochu, V.M. Cora and N. De Freitas, preprint, arXiv:1012.2599 (2010).

[69] R. Calandra, A. Seyfarth, J. Peters and M.P. Deisenroth, Ann. Math. Artif. Intel. **76** (1-2), 5–23 (2016).

[70] D.J. Lizotte, T. Wang, M.H. Bowling and D. Schuurmans, in *IJCAI*, Vol. 7 (Institute of Electrical and Electronics Engineers, Piscataway, 2007), pp. 944–949.

[71] A. Klein, S. Falkner, S. Bartels, P. Hennig and F. Hutter, Electron. J. Stat. **11** (2), 4945–4968 (2017).

[72] C.I. Bayly, P. Cieplak, W. Cornell and P.A. Kollman, J. Phys. Chem. **97** (40), 10269–10280 (1993).