

Diss. ETH No. 27314

On-the-fly Probability Enhanced Sampling

A dissertation submitted to attain the degree of
DOCTOR OF SCIENCES of ETH ZURICH
(Dr. sc. ETH Zurich)

presented by

MICHELE INVERNIZZI

MSc. in Physics Università degli Studi di Milano

born on 3 April 1990
citizen of Italy

accepted on the recommendation of
Prof. Dr. Michele Parrinello, examiner
Prof. Dr. Giovanni Bussi, co-examiner

2021

Hoc etiam magis haec animum te advertere par est
corpora quae in solis radiis turbare videntur,
quod tales turbae motus quoque materiai
significant clandestinos caecosque subesse.
multa videbis enim plagis ibi percita caecis
commutare viam retroque repulsa reverti
nunc huc nunc illuc in cunctas undique partis.
scilicet hic a principiis est omnibus error.
prima moventur enim per se primordia rerum,
inde ea quae parvo sunt corpora conciliatu
et quasi proxima sunt ad viris principiorum,
ictibus illorum caecis impulsa cientur,
ipsaque proporro paulo maiora lacesunt.
sic a principiis ascendit motus et exit
paulatim nostros ad sensus, ut moveantur
illa quoque, in solis quae lumine cernere quimus
nec quibus id faciant plagis apparet aperte.

Lucretius, *De Rerum Natura*, II, 125-141

Abstract

Molecular simulations have become an important scientific tool with several applications in physics, chemistry, material science and biology. Therefore, the development of more efficient and robust computational methods is of great importance, as any improvement in this respect can potentially have a positive impact on various open research fronts. One of the major challenges for molecular simulations is the sampling of so-called rare events, i.e. microscopic phenomena occurring on macroscopic time scales. Typical examples of rare events are phase transitions, chemical reactions, or protein folding. Here we propose a novel computational technique, called on-the-fly probability enhanced sampling (OPES) method, which can significantly improve molecular simulations of such phenomena. The OPES method can be described as an evolution of another popular enhanced sampling method, metadynamics, with respect to which it brings both conceptual and practical improvements. It goes a long way toward making enhanced sampling not only more efficient, but also more robust and easier to use. Furthermore, OPES unifies in the same approach two traditionally distinct sampling strategies, namely collective-variable-based methods and expanded-ensembles methods. We believe that this perspective will open interesting new possibilities in the field of enhanced sampling.

Riassunto

Le simulazioni molecolari sono diventate un importante strumento scientifico con diverse applicazioni in fisica, chimica, scienza dei materiali e biologia. Pertanto, lo sviluppo di metodi di calcolo più efficienti e robusti è di grande importanza, in quanto qualsiasi miglioramento in questo senso può potenzialmente avere un impatto positivo su vari fronti di ricerca aperti. Una delle maggiori sfide per le simulazioni molecolari è il campionamento dei cosiddetti eventi rari, cioè fenomeni microscopici che si verificano su scale di tempo macroscopiche. Esempi tipici di eventi rari sono le transizioni di fase, le reazioni chimiche o il ripiegamento delle proteine. Qui proponiamo una nuova tecnica di calcolo, chiamata metodo OPES (*on-the-fly probability enhanced sampling*), che può migliorare significativamente le simulazioni molecolari di tali fenomeni. OPES si può considerare come l'evoluzione di un altro popolare metodo di campionamento potenziato, la metadinamica, rispetto a cui apporta miglioramenti sia concettuali che pratici. OPES ha come obiettivo di rendere il campionamento potenziato non solo più efficiente, ma anche più robusto e di più facile impiego. Inoltre, OPES unifica nello stesso approccio due strategie di campionamento tradizionalmente distinte, vale a dire i metodi basati su variabili collettive e quelli basati su insiemi statistici espansi. Riteniamo che questa prospettiva aprirà nuove interessanti possibilità nel campo del campionamento potenziato.

Contents

Contents	v
1 Introduction	1
2 Articles	5
2.1 Coarse graining from variationally enhanced sampling applied to the Ginzburg–Landau model	5
2.2 Making the best of a bad situation: a multiscale approach to free energy calculation	19
2.3 Rethinking metadynamics: from bias potentials to probability distributions	43
2.4 Unified approach to enhanced sampling	65
3 Conclusion and future works	99
3.1 Conclusion	99
3.2 Future works	100
Other articles	103
Bibliography	105
Acknowledgements	117
Curriculum vitae	119

Chapter 1

Introduction

Molecular simulations are playing an ever-increasing role in modern science and are applied to fields as varied as physics, chemistry, biology, and material science. They are used for example to gain insight about the microscopic mechanisms of complex phenomena and better understand experiments or guide them, or to accelerate the discovery of new drugs and materials by pre-screening large databases of possible candidates. Since its early years, molecular simulation had to face two main challenges: the modelling accuracy and the sampling efficiency. Despite the tremendous progress that computer simulations have witnessed, both in hardware and algorithms, these two key challenges are still at the very center of current research and method development. During my doctoral studies, I focused on the sampling problem in molecular simulations and in this thesis I present a novel enhanced sampling method, called *on-the-fly probability enhanced sampling* (OPES), which contributes to tackling this issue.

The sampling problem has its roots in the multiscale nature of many phenomena of interest such as phase transitions, protein folding or chemical reactions, just to name a few. These phenomena, also known as rare events, present long-lived metastable states separated by kinetic bottlenecks. The probability of observing one of such rare transitions can be so small that, when simulating these systems via molecular dynamics or Markov-chain Monte Carlo, one might need an unfeasible amount of computational time to sample all the relevant metastable states. In many practical cases most if not all the simulation is spent in the same free energy basin, and no information about the actual phenomena is obtained. Over the years, several different techniques have been developed to try and solve this central issue in molecular simulations. Among them, we are interested in particular in the so-called collective variables (CVs) based methods, the first of which was proposed in the 70s by Torrie and Valleau [1]. Various popular enhanced sampling methods belong to this category, such as umbrella sampling [2]

and metadynamics [3]. These methods typically add to the system a bias potential that is a function of few CVs and that artificially increases the probability of observing the rare transitions between metastable states.

These methods can efficiently solve the sampling problem only if a set of few good CVs is known for the given system. A good collective variable, or reaction coordinate, should encode all the slow modes of the system and express them as a smooth function of the atomic positions. As a consequence, such CVs should be able to distinguish all the relevant metastable states as well as the transition states in between them. Typically, CVs are carefully devised by trial and error, guided by human insight and expertise, but this process can be difficult and time-consuming and quickly becomes impractical for complex systems. For this reason, in recent years much effort has been devoted to try to apply machine learning techniques to automatically find such good CVs. The interested reader can find in Ref. [4] a good (even though inevitably incomplete) review of this still very active field of research.

With my research I have taken instead a different direction and studied the way the bias potential can be built, rather than focusing on the identification of CVs. This choice was motivated by the experience gained by me and other colleagues in the research group while making use of metadynamics and the variationally enhanced sampling (VES) method [5]. Although these methods can be extremely powerful, allowing the study of elusive phenomena, it was clear to us practitioners that they could be further improved. A first aspect that I wanted to tackle was the slow convergence of the methods when suboptimal CVs are used, thus when some slow mode is not accelerated by biasing. Fewer transitions and hysteresis are clear symptoms of this suboptimality, and in this scenario the choice of the parameters of the methods becomes delicate, often requiring a trial-and-error approach. The outcome of my research has been a novel enhanced sampling method, OPES, that brings not only an improvement over some of these practical aspects, but also a different perspective on enhanced sampling. The method inherits some key characteristics from metadynamics and VES, such as the on-the-fly optimization of the bias, or the concept of target probability distribution, which is pushed even further, bringing new insight into what enhanced sampling can be.

This thesis is made as a compilation of four articles that are connected to the development of the OPES method. The first two articles, Sec. 2.1 and 2.2, do not directly deal with OPES, but are instead based on the VES method. They are included in the thesis because they present some of the key ideas that later have contributed to the formulation of the OPES method, and thus can guide the reader to better understand the underlying issues and the tools that are at the heart of the new method. In the last two articles, Sec. 2.3 and 2.4, the OPES method is formulated and applied to some pro-

totypical systems that are affected by the sampling problem. The OPES method is designed to sample different types of target probability distributions and, while its core elements remain the same, the implementation can change considerably depending on the chosen target. In the first of these two articles, similarly to metadynamics, OPES is used to sample uniform or well-tempered distributions in the CV space and makes use of a kernel density estimation algorithm. On the contrary, in the last paper we explore the exciting possibility of using the OPES formalism to perform the same type of sampling as the popular replica-exchange method[6], thus targeting expanded ensembles. This last article proves that the perspective introduced with OPES has been able to bring new insights into the field of enhanced sampling that go beyond the practical improvements that I was aiming for at the beginning of my research. Finally, in the concluding chapter of this thesis I discuss possible future research routes and applications of the OPES method.

Chapter 2

Articles

2.1 Coarse graining from variationally enhanced sampling applied to the Ginzburg–Landau model

This first article is not directly related to the OPES method, however, thanks to its unusual perspective on enhanced sampling, it touches on some ideas that were later crucial to develop the method. The goal of this paper is to show how the variationally enhanced sampling method can be used to efficiently optimize the parameters of a coarse-graining model, while simultaneously improving the sampling of the underlying full-atoms system. We also explore the use of a nontrivial target distribution and we show that by using a simple functional to model the free energy (only three free parameters), it becomes possible to greatly increase the number of employed CVs (up to 364 in the paper). This idea of taking advantage of a specific functional form to express the bias, was later used for developing the expanded ensemble version of OPES, Sec. 2.4.

My contribution to this article has been implementing the algorithms, performing the simulations, and writing the paper jointly with Valsson and Prof. Parrinello.

Reference: M. Invernizzi, O. Valsson, and M. Parrinello. “Coarse graining from variationally enhanced sampling applied to the Ginzburg-Landau model.” *Proceedings of the National Academy of Sciences* 114.13 (2017): 3370-3374. URL <https://www.pnas.org/content/114/13/3370>

Copyright © 2017 National Academy of Sciences.

Coarse graining from variationally enhanced sampling applied to the Ginzburg–Landau model

Michele Invernizzi^{1,2}, Omar Valsson^{3,2}, and Michele Parrinello^{*3,2}

¹*Department of Physics, ETH Zurich c/o Università della Svizzera italiana, 6900 Lugano, Switzerland*

²*Facoltà di Informatica, Institute of Computational Science, and National Center for Computational Design and Discovery of Novel Materials (MARVEL), Università della Svizzera italiana, 6900 Lugano, Switzerland*

³*Department of Chemistry and Applied Biosciences, ETH Zurich c/o Università della Svizzera italiana, 6900 Lugano, Switzerland*

Abstract

A powerful way to deal with a complex system is to build a coarse-grained model capable of catching its main physical features, while being computationally affordable. Inevitably, such coarse-grained models introduce a set of phenomenological parameters, which are often not easily deducible from the underlying atomistic system. We present a unique approach to the calculation of these parameters, based on the recently introduced variationally enhanced sampling method. It allows us to obtain the parameters from atomistic simulations, providing thus a direct connection between the microscopic and the mesoscopic scale. The coarse-grained model we consider is that of Ginzburg–Landau, valid around a second order critical point. In particular we use it to describe a Lennard-Jones fluid in the region close to the liquid-vapor critical point. The procedure is general and can be adapted to other coarse-grained models.

2.1.1 Introduction

Computer simulations of condensed systems based on an atomistic description of matter are playing an ever-increasing role in many fields of science. Yet, as the complexity of the systems studied increases, so does the awareness that a less detailed, but nevertheless accurate, description of the system is necessary.

This has been long since recognized, and branches of physics like elasticity or hydrodynamics can be classified in modern terms as coarse-grained (CG) models of matter. In more recent times, a field theoretical model suitable to describe second order phase transitions has been introduced by Landau [7], and later perfected by Ginzburg [8]. In recent decades a number of coarse-grained models that aim at describing polymers or biomolecules have also been proposed [9, 10]. In all these approaches some degrees of freedom, deemed not essential to study the phenomenon at hand, are integrated out and the resulting reduced description contains a number of parameters which are not easy to be determined.

Here we shall use the recently developed variationally enhanced sampling (VES) method [5] to suggest a procedure that allows the determination of such parameters, starting from the microscopic Hamiltonian. This will illuminate a somewhat unexpected application of VES, which has been introduced as an enhanced sampling method. We shall show that it also provides a powerful framework for the optimization of CG models, thanks to the combination of its enhanced sampling capabilities and its variational flexibility. Moreover VES takes advantage of its deep connection with relative entropy, a quantity that has been shown to play a key role in multiscale problems [11, 12].

As a first test case for our procedure we shall consider the Ginzburg-Landau (GL) model for continuous phase transitions. An advantage of using this model is that its strengths and limits are well known and that other researchers have already attempted to perform such a calculation [13–19]. A system which undergoes a second order phase transition is described in the GL model by the following free energy, valid in a rotationally invariant one-component real order parameter scenario:

$$F[\psi] = g \int |\nabla \psi(\vec{r})|^2 d^3r + a \int \psi^2(\vec{r}) d^3r + b \int \psi^4(\vec{r}) d^3r, \quad (2.1)$$

where the field $\psi(\vec{r})$ describes the order parameter fluctuations, and g , a and b are phenomenological quantities that we shall call Landau’s parameters. To be defined, we shall apply our method close to the liquid-vapor critical point of a Lennard-Jones fluid.

2.1.2 Variationally Enhanced Sampling

Before discussing our method, we briefly recall some of the ideas at the basis of VES. The VES method shares with metadynamics [3, 20] the idea of enhancing the sampling by adding an external bias potential, which is a function of a few coarse-grained order parameters, so-called collective variables (CVs). Contrary to metadynamics and other similar methods, VES does not generate the bias potential by periodically adding on the fly small contributions, but as the result of the minimization of a convex functional. Various extensions have been proposed, where VES has been employed in different ways [21–26].

Let’s suppose that the physical behavior of interest is well described by a set of CVs $\mathbf{s} = \mathbf{s}(\mathbf{R})$, that are function of the coordinates \mathbf{R} of the N particles that compose the system. Then one can write the associated free energy function as $F(\mathbf{s}) = -(1/\beta) \log \int d\mathbf{R} \delta[\mathbf{s} - \mathbf{s}(\mathbf{R})] e^{-\beta U(\mathbf{R})}$, where $\beta = 1/(k_B T)$ is the inverse temperature and $U(\mathbf{R})$ is the interparticle potential energy.

In VES a functional that depends on an externally applied bias $V(\mathbf{s})$ is introduced:

$$\Omega[V] = \frac{1}{\beta} \log \frac{\int d\mathbf{s} e^{-\beta[F(\mathbf{s})+V(\mathbf{s})]}}{\int d\mathbf{s} e^{-\beta F(\mathbf{s})}} + \int d\mathbf{s} p(\mathbf{s})V(\mathbf{s}), \quad (2.2)$$

where $p(\mathbf{s})$ is a chosen probability distribution, that we will call target distribution. The functional is convex and for the value of $V(\mathbf{s})$ that minimizes $\Omega[V]$, one has

$$P_V(\mathbf{s}) = p(\mathbf{s}), \quad (2.3)$$

where

$$P_V(\mathbf{s}) = \frac{e^{-\beta[F(\mathbf{s})+V(\mathbf{s})]}}{\int d\mathbf{s} e^{-\beta[F(\mathbf{s})+V(\mathbf{s})]}} \quad (2.4)$$

is the distribution of \mathbf{s} in the ensemble biased by $V(\mathbf{s})$. Neglecting immaterial constants, Eq. (2.3) can also be written as

$$F(\mathbf{s}) = -V(\mathbf{s}) - \frac{1}{\beta} \log p(\mathbf{s}). \quad (2.5)$$

Thus, finding the minimizing bias potential $V(\mathbf{s})$ is equivalent to finding the free energy $F(\mathbf{s})$.

We notice here that the functional $\Omega[V]$ has close connections to the relative entropy or Kullback-Leibler divergence D_{KL} [27–29], in particular $\beta\Omega[V] = D_{KL}(p\|P_V) - D_{KL}(p\|P_0)$. Furthermore, at the minimum we have $\Omega[V] \leq 0$, which is a rewrite of the Gibbs-Bogoliubov inequality [30].

In order to minimize $\Omega[V]$ we shall express $V(\mathbf{s})$ as a linear function of a finite set of variational parameters $\alpha = \{\alpha_i\}$, thus $V(\mathbf{s}; \alpha)$. By doing so, our functional becomes a function of α , that can be minimized using the gradient and the Hessian

$$\frac{\partial\Omega(\alpha)}{\partial\alpha_i} = - \left\langle \frac{\partial V(\mathbf{s}; \alpha)}{\partial\alpha_i} \right\rangle_{V(\alpha)} + \left\langle \frac{\partial V(\mathbf{s}; \alpha)}{\partial\alpha_i} \right\rangle_p \quad (2.6)$$

$$\begin{aligned} \frac{\partial^2\Omega(\alpha)}{\partial\alpha_i\partial\alpha_j} = \beta \left[\left\langle \frac{\partial V(\mathbf{s}; \alpha)}{\partial\alpha_i} \frac{\partial V(\mathbf{s}; \alpha)}{\partial\alpha_j} \right\rangle_{V(\alpha)} \right. \\ \left. - \left\langle \frac{\partial V(\mathbf{s}; \alpha)}{\partial\alpha_i} \right\rangle_{V(\alpha)} \left\langle \frac{\partial V(\mathbf{s}; \alpha)}{\partial\alpha_j} \right\rangle_{V(\alpha)} \right], \quad (2.7) \end{aligned}$$

where the averages in the right hand side of the equations are calculated either in the biased ensemble $\langle \cdot \rangle_{V(\alpha)}$, or in the $p(\mathbf{s})$ ensemble $\langle \cdot \rangle_p$, and second derivatives in α have been omitted due to the linearity assumption.

Since the forces in Eq. (2.6,2.7) are calculated as statistical averages, a stochastic optimization method [31] is needed. We describe in detail the minimization procedure in the *Supporting Information* (Sec. 2.1.9).

2.1.3 Free Energy Model

We consider a Lennard-Jones fluid of N particles confined in a periodically repeated cubic box of volume $V = L^3$. The order parameter of such a system is linked to the density $\rho(\vec{r})$:

$$\psi(\vec{r}) = \frac{\rho(\vec{r}) - \rho_c}{\rho_0}, \quad (2.8)$$

where $\rho_0 = N/V$ and ρ_c is the critical density. The order parameter can be expanded in a discrete Fourier series as:

$$\psi(\vec{r}) = \sum_{\vec{k}} e^{i\vec{k}\cdot\vec{r}} \psi_{\vec{k}}, \quad (2.9)$$

where $\vec{k} = \frac{2\pi}{L} \vec{n}$, with $\vec{n} \in \mathbb{Z}^3$. Ginzburg-Landau model implicitly relies on a characteristic CG length Λ . This is defined by including in the series expansion, Eq. (2.9), only those wave vectors \vec{k} whose modulus is less than $k_M = 2\pi/\Lambda$. We will refer to this as “wave vector cutoff of order n_M ”, where n_M is an integer such that $|\vec{k}|^2 \leq k_M^2 = n_M(2\pi/L)^2$ for all the included \vec{k} . Close to T_c the system is dominated by long wavelength fluctuations and the presence of a cutoff should become physically irrelevant.

We shall consider as collective variables these Fourier components of the order parameter, $\mathbf{s} = \{\psi_{\vec{k}}\}_{|\vec{k}| \leq k_M}$. In terms of \mathbf{s} the GL free energy can be rewritten as

$$F(\mathbf{s}) = g I_G(\mathbf{s}) + a I_2(\mathbf{s}) + b I_4(\mathbf{s}), \quad (2.10)$$

where the integrals in Eq. (2.1) are rewritten in Fourier space as:

$$I_G(\mathbf{s}) = V \sum_{|\vec{k}| \leq k_M} k^2 |\psi_{\vec{k}}|^2 \quad (2.11)$$

$$I_2(\mathbf{s}) = V \sum_{|\vec{k}| \leq k_M} |\psi_{\vec{k}}|^2 \quad (2.12)$$

$$I_4(\mathbf{s}) = V \sum_{|\vec{q}| \leq 2k_M} \left| \sum_{|\vec{k}| \leq k_M} \psi_{\vec{k}} \psi_{\vec{q}-\vec{k}} \right|^2 \quad (2.13)$$

and in the last convolution it is implied that if $|\vec{q} - \vec{k}| > k_M$ then $\psi_{\vec{q}-\vec{k}} = 0$. We notice that since the order parameter is a real quantity, its Fourier transform has the property $\psi_{\vec{k}} = \psi_{-\vec{k}}^*$. This symmetry property will be used to reduce the number of independent collective variables. With a cutoff of order $n_M = 3$ one has 26 CVs, but in our simulations we have dealt with up to 364 CVs, in the $n_M = 19$ case (see Table 2.1).




Figure 2.1: A simplified one-dimensional representation of the chosen quadratic target free energy $F_t = -(1/\beta) \log p(\mathbf{s})$ (green dashed line) and the expected Landau's free energy (orange solid line).

2.1.4 Bias Potential and Target Distribution

In order to proceed, we have to choose a variational form for $V(\mathbf{s})$ and introduce a suitable $p(\mathbf{s})$. We shall use for $V(\mathbf{s})$ and $p(\mathbf{s})$ expressions that, when combined as in Eq. (2.5), lead for $F(\mathbf{s})$ to an expression with the same analytical structure as the GL free energy, Eq. (2.10). Guided by this principle we take

$$V(\mathbf{s}) = g_V I_G(\mathbf{s}) + a_V I_2(\mathbf{s}) + b_V I_4(\mathbf{s}) \quad (2.14)$$

and

$$p(\mathbf{s}) = \frac{e^{-\beta[g_t I_G(\mathbf{s}) + a_t I_2(\mathbf{s})]}}{\int d\mathbf{s} e^{-\beta[g_t I_G(\mathbf{s}) + a_t I_2(\mathbf{s})]}}. \quad (2.15)$$

These two expressions can be employed in $\Omega[V]$, which then can be minimized relative to the variational parameters g_V , a_V and b_V . At the minimum the estimated $F(\mathbf{s})$ has the desired GL structure, with parameters $g = g_t - g_V$, $a = a_t - a_V$ and $b = -b_V$.

Since $V(\mathbf{s})$ in Eq. (2.14) has only a limited variational flexibility, our estimate of the free energy $F(\mathbf{s})$ is only approximated. However, given that our $V(\mathbf{s})$ is based on sound physical considerations, we still expect the $F(\mathbf{s})$ obtained to be a good approximation.

Using the $p(\mathbf{s})$ in Eq. (2.15) has some advantages. It is a product of univariate Gaussian distributions, so the averages over $p(\mathbf{s})$ that are needed to minimize $\Omega[V]$ (see Eq. 2.6,2.7) can easily be evaluated. And, more importantly, it can be used to guide sampling and accelerate the optimization convergence.

In order to understand this point we shall refer to Fig. 2.1 that gives an artist's impression of a Landau free energy as a function of a one dimensional order parameter as the system crosses the critical temperature T_c . For $T > T_c$ the order parameter fluctuations are predominantly Gaussian and the effect of the non-Gaussian quartic term is hidden in the tail of the distribution, as pointed out also by other authors [15, 18]. By using a $p(\mathbf{s})$ that is

broader than the natural Gaussian fluctuations, we can enhance the probability with which the tails are sampled, thus improving the computation of the parameter b proportional to the quartic term $I_4(\mathbf{s})$.

For $T < T_c$ one enters into the coexistence region, where different problems arise. The symmetry is broken, and the system spontaneously separates into two slab-shaped regions of finite but different value of the order parameter. These two phases, vapor and liquid, correspond to the two minima in the rightmost curve in Fig. 2.1. By using a $p(\mathbf{s})$ that focuses on the transition region we can keep the system homogeneous and avoid to a certain extent some finite-size drawbacks (see Ref. [32]). However, as the temperature is lowered, the interface between the two phases sharpens up with respect to the CG length Λ . This situation cannot be described quantitatively by a GL model which takes into account only long wavelengths fluctuations of the order parameter. In our simulations there are some clear evidences that the underlying model is not any longer suited for describing the system. In fact the convergence in the optimization of $\Omega[V]$ slows dramatically, and the obtained parameters no more follow the expected linear behavior.

The use of a higher wave vector cutoff, and thus a smaller CG length Λ , can mitigate this issue, but only marginally. In fact there is another important feature that is not taken into account by this simple model, that is the so called *field-mixing* due to the lack of particle-hole symmetry [33], which induces an asymmetry between the two phases at low T . Nevertheless the GL model we adopted can still give us some relevant information, as shown at the end of Sec. 2.1.6.

2.1.5 Computational Setup

We choose to simulate an Argon fluid, described by the classical two body Lennard-Jones potential $\varphi_{LJ}(r) = 4\epsilon[(\sigma/r)^{12} - (\sigma/r)^6]$ truncated and shifted at $R_c = 2.5\sigma$:

$$\Phi_{LJ}(r) = \begin{cases} \varphi_{LJ}(r) - \varphi_{LJ}(R_c) & r < R_c \\ 0 & r \geq R_c \end{cases} \quad (2.16)$$

This simple system has been intensively studied [32–35] and exhibits in its phase diagram the needed second order phase transition, at the liquid-vapor critical point. All the physical quantities in this paper are expressed in terms of Lennard-Jones reduced units.

We perform canonical (NVT) molecular dynamics simulations at a fixed average density equal to the critical one, $\rho_c = 0.317$ (see Ref. [34]). We are aware that in a finite size system the critical density is not exactly ρ_c , but for our purposes we just need to be close to that value. The temperature T is kept constant by the velocity rescaling stochastic thermostat [36]. The time step used is 0.002, and typically we used 10^7 steps for each run.




Figure 2.2: Landau's parameters optimized for systems with: (a) $N = 128$, $n_M = 3$ thus $\Lambda = 4.267$; (b) $N = 2048$, $n_M = 3$ thus $\Lambda = 10.75$; (c) $N = 2048$, $n_M = 19$ thus $\Lambda = 4.272$; (d) $N = 8192$, $n_M = 3$ thus $\Lambda = 17.07$. Data at temperatures where convergence was not clearly reachable are omitted. The linear fits are obtained using only data at $T > 0.95 T_c$, and error bars are smaller than the data points.

All the simulations are performed with the molecular dynamics package GROMACS 5.0.4 [37], patched with a development version of the PLUMED 2 [38] enhance sampling plugin, in which we implemented the VES method for our specific GL free energy model.

The choice of g_t and a_t , that define $p(\mathbf{s})$ in Eq. (2.15), was based on the amplitude of the fluctuations observed in the unbiased runs, and on some trial and error simulations performed on small systems ($N = 128, 256$). The value thus obtained were then used as reference for the larger systems. The adopted values were between 0.05 and 0.2, and generally $g_t > a_t$.

2.1.6 Results

Ginzburg-Landau model expects the parameter a to have a linear temperature dependence close to the critical point, thus

$$a(T) = a_c(T - T_c), \quad (2.17)$$

while the other two parameters should be constant, $g(T) = g_c$ and $b(T) = b_c$. We run multiple simulations varying both, the size of the system and the cutoff order, and consequently also the CG length Λ . As noted at the end of Sec. 2.1.4, depending on the system size and Λ , convergence can become problematic at low temperatures. Nevertheless we can, somewhat conservatively, refer only to the region $T > 0.95 T_c$, where convergence is safely reached for all the systems studied here. In this region our estimates are always in qualitative agreement with Landau's ansatz, as can be seen e.g.

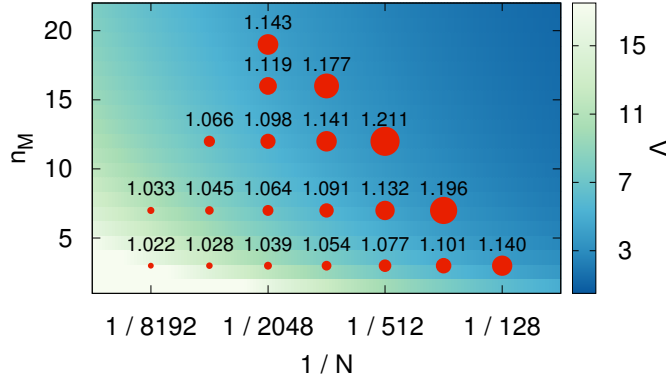


Figure 2.3: The effective critical temperature T_c^*/T_c for different system sizes N and wave vector cutoff orders n_M . The dimension of the points is proportional to $(T_c^*/T_c - 1)$. The background colors are related to the coarse-graining length $\Lambda \propto N^{1/3}/\sqrt{n_M}$, according to the scale on the right.

in Fig. 2.2. The temperature dependence of the three parameters can be well fitted linearly, and all the results obtained are provided in the *Supporting Information* (Sec. 2.1.9). Here we only notice that the slope of $g(T)$ and $b(T)$ is roughly 10 times smaller than the one of $a(T)$. This non-constant behavior has been reported also in previous literature, for other Landau free energy models (see e.g. Ref. [15]).

It is possible to define an effective critical temperature T_c^* , as the temperature at which the parameter $a(T)$ becomes zero. This is different from the critical temperature in the thermodynamic limit, $T_c = 1.085 \pm 0.005$ [34], due to finite size effects. Contrary to the effective critical temperature that one could define by looking at the specific heat, T_c^* is always greater than T_c , as noted already in Ref. [13]. In Fig. 2.3 we show our estimate of T_c^*/T_c as a function of both the system size N , and the wave vector cutoff order n_M . As expected, increasing the size at fixed cutoff we have that T_c^* approaches T_c . Maybe less intuitively, when we increase the number of wave vectors, T_c^* moves away from T_c . What we observe is that T_c^* is mainly linked to the coarse-graining length Λ , in fact we find that roughly $(T_c^* - T_c) \propto 1/\Lambda$. In spite of the great variations in the value of T_c^* , all our data on $a(T)$ can be collapsed into a single line when plotted as a function of $(T - T_c^*)/T_c$, see Fig. 2.4. These are all empirical observations and do not necessarily imply rigorous scaling laws.

2.1.7 Simulations of the coarse-grained model

Having obtained a coarse-grained description, we can ask to what an extent the GL model can represent the microscopic properties of the system. To try and answer this question, we run Monte Carlo simulations, using as




Figure 2.4: Parameter a , the one proportional to the quadratic term in Eq. (2.10), as a function of the rescaled temperature $(T - T_c^*)/T_c$. The effective critical temperature T_c^* is different for each system, as shown in Fig. 2.3. Only points with $T > 0.95 T_c$ are plotted.

variables the Fourier components $\psi_{\vec{k}}$ and as Hamiltonian the GL free energy. We define the model at all temperatures by assuming that the parameters depend linearly on T . Despite the limitations of GL model (see end of section 2.1.4), we find that some of the properties of our system are well described for a wide range of temperatures.

The most remarkable one is maybe the description of the Binder cumulant [39, 40], a quantity often used to obtain a good estimate of the critical temperature T_c , from finite-size simulations. Using our notation, we can handily express the Binder cumulant as follows:

$$U(T) = 1 - \frac{\langle I_4 \rangle}{3\langle I_2 \rangle^2}, \quad (2.18)$$

where $\langle \cdot \rangle$ is the ensemble average at temperature T , and I_2 and I_4 are defined in Eq. (2.12,2.13) respectively. The important property of Binder cumulant is that $U(T_c)$ does not depend on the system size, and thus it is possible to evaluate T_c by looking at the crossing point of the curves obtained with different N .

The behavior of $U(T)$ for different system sizes is reported in Fig. 2.5, as obtained with unbiased molecular dynamics and Monte Carlo coarse-grained simulations. Remarkably the CG simulations well reproduce the $U(T)$ curve, even at $T < T_c$, and gives an estimate of T_c in good agreement with the literature [34]. To calculate I_2 and I_4 for this figure we have used a wave vector cutoff of order $n_M = 3$, but it is important to note that the inclusion of higher order wave vectors does not change the crossing point, even though the shape of the curves $U(T)$ may change. This reflects the fact that the critical behavior is dominated by long wavelength fluctuations.




Figure 2.5: The Binder cumulant $U(T)$ defined in Eq. (2.18) obtained with full molecular dynamics simulations and with Monte Carlo simulations using the optimized GL free energy as Hamiltonian. In both cases a wave vector cutoff of order $n_M = 3$ was used, thus for the CG simulations only 26 independent variables were employed for each system size. The interception of the curves gives an estimate of the critical temperature. We take as reference $T_c = 1.085$ [34]. Error bars are smaller than the data points, and lines are only drawn to guide the eyes.

2.1.8 Conclusions

In this paper we have shown that VES is not only a method for enhanced sampling, but it provides a formidable tool to approach coarse-graining in a systematic and rigorous way. It combines the strength of relative entropy optimization methods with the intriguing ability to orient the sampling towards a desired statistical distribution.

In the present application we treated the case of a second order phase transition, for which the well established GL theory provided a good CG model. In other fields of application careful consideration will have to be paid, and the structure of the CG model might not be easily determinable. However once this is done, VES provides a way to link rigorously and effectively the parameters of the coarse-graining model to the microscopic Hamiltonian.

Acknowledgements

This research was supported by the NCCR MARVEL, funded by the Swiss National Science Foundation, and the European Union Grant No. ERC-2014-AdG-670227/VARMET. The computational time for this work was provided by the Swiss National Supercomputing Center (CSCS). The authors thank Haiyang Niu for carefully reading the manuscript.

2.1.9 Supporting Information

Minimization Procedure

The minimization of the VES functional $\Omega[V]$, Eq. (2.2), was carried out following the iterative procedure of Bach and Moulines [31], as described also in Ref. [5, 21]. In this scheme one considers two distinct sequences of parameters: the instantaneous $\alpha^{(n)}$ and the averaged $\bar{\alpha}^{(n)}$. The instantaneous sequence is updated using the gradient and the full Hessian, see Eq. (2.6,2.7), evaluated at the last available average parameter values $\bar{\alpha}^{(n)}$:

$$\alpha_i^{(n+1)} = \alpha_i^{(n)} - \mu \left[\frac{\partial \Omega(\bar{\alpha}^{(n)})}{\partial \alpha_i} + \sum_j \frac{\partial^2 \Omega(\bar{\alpha}^{(n)})}{\partial \alpha_i \partial \alpha_j} (\alpha_j^{(n)} - \bar{\alpha}_j^{(n)}) \right], \quad (2.19)$$

where μ is a fixed step size, that we keep on the order of 0.1 in all our simulations. An important feature of this algorithm is that it does not require long sampling time for each iteration, in fact we typically estimate the gradient and the Hessian with a stride $S = 2000$ simulation time steps. We also implemented a multiple walkers framework [41] to further improve the sampling, using two independent walkers that share the same bias.

In order to accelerate convergence, we found useful to calculate $\bar{\alpha}$ as an exponentially decaying average, namely

$$\bar{\alpha}_i^{(n+1)} = \bar{\alpha}_i^{(n)} + \frac{\alpha_i^{(n+1)} - \bar{\alpha}_i^{(n)}}{W}, \quad (2.20)$$

that is the discretized equivalent of

$$\bar{\alpha}_i(t) = \int_0^t dt' \alpha_i(t') e^{-(t-t')/\tau}. \quad (2.21)$$

The weight W is linked to the decaying time τ by the relation $1/W = 1 - e^{-S\Delta t/\tau} \approx S\Delta t/\tau$, where Δt is the simulation time step and S is the stride between the parameters update. In the runs here reported we aimed at a τ of about 1/10 of the total simulation length, which amounts at taking $W = 500$.

Within reasonable bounds, the precise choice of the parameters μ , S , and τ do not affect much the rate of convergence of the optimization algorithm. In fact, we found that the choice of the target distribution parameters, g_t and a_t , play a more important role in this.

Collective Variables

We report here explicitly the collective variables (CVs) used for our Lennard-Jones system. We can write the density of the system as a sum of delta

2.1. Coarse graining from variationally enhanced sampling

Table 2.1: Number of independent CVs for a given wave vector cutoff order n_M .

n_M	CVs
3	26
7	80
12	178
16	256
19	364

functions centered in the atom's positions \vec{R}_j :

$$\rho(\vec{r}) = \sum_j^N \delta(\vec{r} - \vec{R}_j). \quad (2.22)$$

This is a safe approximation, because the wave vector cutoff will then actually limit our study to distances greater than Λ . Being the order parameter $\psi(\vec{r}) = (\rho(\vec{r}) - \rho_c)/\rho_0$, its Fourier components $\psi_{\vec{k}}$ are

$$\text{Re}(\psi_{\vec{k}}) = \frac{1}{N} \sum_j^N \cos(\vec{k} \cdot \vec{R}_j), \quad (2.23)$$

$$\text{Im}(\psi_{\vec{k}}) = -\frac{1}{N} \sum_j^N \sin(\vec{k} \cdot \vec{R}_j), \quad (2.24)$$

it is now easy to notice that $\psi_{\vec{k}} = \psi_{-\vec{k}}^*$, which halves the number of independent Fourier components.

The number of independent CVs is directly related to the cutoff order n_M . For each independent Fourier component $\psi_{\vec{k}}$, we have two independent CVs, its real and imaginary part. The number of CVs as a function of n_M is reported in Tab. 2.1.

Landau Parameters

We report here the Landau parameters obtained with the variational enhanced sampling method, as described in the main text. The parameters in Tab. 2.2 are obtained with a linear fit of the optimized parameters at temperatures

$$T = \{1.05, 1.10, 1.15, 1.20, 1.25, 1.30, 1.35, 1.40, 1.45, 1.50\}. \quad (2.25)$$

The lowest temperature for the fit is chosen in such a way that a good convergence is ensured at any T even for the bigger systems. Given Ginzburg-Landau free energy, Eq. (2.10), the fitting parameters are defined as follows:

$$g(T) = g_1 T + g_0 \quad (2.26)$$

$$a(T) = a_c(T - T_c^*) \quad (2.27)$$

$$b(T) = b_1 T + b_0. \quad (2.28)$$

Table 2.2: Landau parameters for different systems sizes N and wave vector cutoff orders n_M . We recall that the coarse-grained length is $\Lambda = N^{1/3}/(\rho_0^{1/3}\sqrt{n_M})$, and that in the thermodynamic limit the critical temperature is $T_c = 1.08(5)$.

N	n_M	Λ	g_1	g_0	a_c	T_c^*	b_1	b_0
128	3	4.27	-0.03(3)	0.13(3)	0.46(2)	1.23(6)	0.06(3)	-0.00(2)
256	3	5.38	-0.03(3)	0.15(1)	0.46(1)	1.19(4)	0.07(0)	-0.02(0)
256	7	3.52	-0.03(2)	0.13(0)	0.46(4)	1.29(7)	0.05(8)	0.00(3)
512	3	6.77	-0.04(4)	0.17(5)	0.45(8)	1.16(8)	0.07(3)	-0.02(4)
512	7	4.43	-0.03(4)	0.14(9)	0.46(9)	1.22(8)	0.06(2)	-0.00(9)
512	12	3.39	-0.03(3)	0.13(0)	0.46(0)	1.31(4)	0.05(9)	-0.00(0)
1024	3	8.53	-0.03(5)	0.17(4)	0.44(4)	1.14(3)	0.05(9)	-0.00(9)
1024	7	5.59	-0.04(1)	0.16(9)	0.46(6)	1.18(4)	0.06(8)	-0.02(2)
1024	12	4.27	-0.03(4)	0.14(7)	0.46(8)	1.23(8)	0.06(1)	-0.00(8)
1024	16	3.70	-0.03(5)	0.14(2)	0.46(8)	1.27(6)	0.06(1)	-0.00(3)
2048	3	10.75	-0.04(8)	0.19(4)	0.43(7)	1.12(7)	0.08(2)	-0.04(0)
2048	7	7.04	-0.03(3)	0.16(8)	0.45(2)	1.15(4)	0.07(1)	-0.03(1)
2048	12	5.38	-0.04(0)	0.16(7)	0.46(7)	1.19(1)	0.06(3)	-0.01(6)
2048	16	4.66	-0.03(3)	0.15(3)	0.46(3)	1.21(4)	0.06(7)	-0.01(8)
2048	19	4.27	-0.03(2)	0.14(5)	0.46(4)	1.24(0)	0.06(3)	-0.01(1)
4096	3	13.55	-0.01(8)	0.15(5)	0.42(9)	1.11(5)	0.06(8)	-0.01(8)
4096	7	8.87	-0.03(4)	0.17(6)	0.44(3)	1.13(3)	0.06(9)	-0.03(1)
4096	12	6.77	-0.03(3)	0.16(8)	0.45(4)	1.15(6)	0.07(2)	-0.03(4)
8192	3	17.07	-0.03(0)	0.16(9)	0.42(3)	1.10(9)	0.07(1)	-0.02(1)
8192	7	11.17	-0.02(4)	0.17(0)	0.43(1)	1.12(0)	0.07(1)	-0.03(3)

The errors in our estimate of the GL parameters arise from two sources: the statistical error coming from the stochastic optimization, and the error provided by the linear fit. The former is significantly smaller than the fit error.


2.2 Making the best of a bad situation: a multiscale approach to free energy calculation

This article, as the previous one, is based on VES and it is not directly connected to the OPES method. However, the paper discusses many of the key concepts that have been at the center of OPES development. Here we suggest the idea that the collective variables typically used in applications are almost always suboptimal and any enhanced sampling method should take into account the resulting multiscale behavior. For example, we argue that when dealing with suboptimal CVs, using a quasi-static bias is more important than having it to precisely reach convergence and fill in all the details of the free energy surface. Thus the proposed method aims to achieve as quickly as possible an adiabatic regimen for the bias potential, which should then be modified only gently, allowing a reweighing procedure not affected by systematic errors to accurately retrieve the free energy and other observables. The OPES method adopts this same biasing strategy. Finally here we write the bias in such a way that it only depends on the free-energy difference between the basins, and thus the parameter to be optimize is exactly the physical quantity we are mostly interested into, as it will be the case in OPES when targeting expanded ensembles, Sec. 2.4.

My contribution to this article has been implementing the algorithms, performing the simulations, and writing the paper jointly with Prof. Parrinello.

Reference: M. Invernizzi, and M. Parrinello. "Making the best of a bad situation: a multiscale approach to free energy calculation." *Journal of chemical theory and computation* 15.4 (2019): 2187-2194. URL <https://pubs.acs.org/doi/10.1021/acs.jctc.9b00032>

Copyright © 2019 American Chemical Society.



Making the Best of a Bad Situation: A Multiscale Approach to Free Energy Calculation

Michele Invernizzi*^{1,2}, and Michele Parrinello^{3,2}

¹*Department of Physics, ETH Zurich c/o Università della Svizzera italiana, 6900 Lugano, Switzerland*

²*Facoltà di Informatica, Institute of Computational Science, and National Center for Computational Design and Discovery of Novel Materials (MARVEL), Università della Svizzera italiana, 6900 Lugano, Switzerland*

³*Department of Chemistry and Applied Biosciences, ETH Zurich c/o Università della Svizzera italiana, 6900 Lugano, Switzerland*

Abstract

Many enhanced sampling techniques rely on the identification of a number of collective variables that describe all the slow modes of the system. By constructing a bias potential in this reduced space, one is then able to sample efficiently and reconstruct the free energy landscape. In methods such as metadynamics, the quality of these collective variables plays a key role in convergence efficiency. Unfortunately in many systems of interest it is not possible to identify an optimal collective variable, and one must deal with the nonideal situation of a system in which some slow modes are not accelerated. We propose a two-step approach in which, by taking into account the residual multiscale nature of the problem, one is able to significantly speed up convergence. To do so, we combine an exploratory metadynamics run with an optimization of the free energy difference between metastable states, based on the recently proposed variationally enhanced sampling method. This new method is well parallelizable and is especially suited for complex systems, because of its simplicity and clear underlying physical picture.

2.2.1 Introduction

Many systems are characterized by metastable states separated by kinetic bottlenecks. Examples of this class of phenomena are chemical reactions, first order phase transitions, and protein folding. Calculating the difference in free energy between these states is of great importance, and various methods have been developed to reach this goal [42, 43]. In this work, we focus on the case in which only two states are relevant. The generalization to a multistate scenario is not too difficult and will be discussed in a future publication.

In the two-state case one can distinguish two time scales, a shorter one in which the system undergoes fluctuations while remaining in one of the minima, and a longer one in which the system moves from one minimum to another. This is the so-called rare event scenario, in which the separation

between the two time scales can be so large as to be inaccessible to direct molecular dynamics simulation, preventing a direct calculation of the free energy surface (FES).

In order to address this issue, a widely used class of free energy methods, including, among others, umbrella sampling [2], metadynamics [3] (MetaD), and variationally enhanced sampling [5, 44] (VES), aim at closing this time scale gap by speeding up the sampling of the slow modes of the system. To achieve this result, one first identifies a set of order parameters or collective variables (CVs) that are functions of the atomic coordinates $\mathbf{s} = \mathbf{s}(\mathbf{R})$. The CVs describe the slow modes whose sampling is accelerated by means of an applied bias potential $V(\mathbf{s})$. An optimal CV is such that if employed for an enhanced simulation run, it closes the time scale gap, so that intrastate and interstates exploration take place on the same time scale. A more common scenario is that the CV does not encode some slow modes, and although much reduced, the gap remains also in the enhanced sampling run. We shall refer to such CV as suboptimal. The limiting case would be that of a CV so poorly chosen that the gap falls outside the computational reach. In this case the CV would be not just suboptimal but also a bad CV. Given the crucial role of CVs, much effort has been devoted to their design and their systematic improvement [45–50].

In metadynamics and variationally enhanced sampling the effect of suboptimal CVs manifests itself in a hysteretic behavior, easily detectable in the CV evolution. In these cases MetaD and VES efficiency suffers [20, 51–53]. The exploration of the system is still enhanced, and the calculation does eventually converge [54], but this can require a substantial computational effort because, even after depositing the bias, some slow modes are present and the system retains a multiscale nature, with two separate time scales. More in detail, in this case the shape of the free energy surface in the basins is easily reconstructed, while their relative height is much harder to converge.

On the basis of this observation, we propose a multiscale approach that separates the calculation in two steps and aims at improving efficiency in a suboptimal CV scenario. First we reconstruct separately the free energy profile of each basin, using MetaD. In a second step we use the properties of VES to optimize their free energy difference and finally reconstruct the global FES. The purpose of our approach is to take into consideration the multiscale nature of a suboptimally enhanced simulation to speed up convergence.

First we shortly present to the unfamiliar reader MetaD and VES, and we introduce a simple illustrative model that is helpful in visualizing the problem; then we explain in detail the proposed method. Finally we successfully apply the new method to some representative systems.

2.2.2 Metadynamics and Variationally Enhanced Sampling

In this section we recall briefly the main features of the MetaD and VES methods, both used in our approach. For a more in-depth review see e.g. Ref. [20]. Both methods enhance sampling by building on the fly a bias potential that is a function of a few collective variables. This bias discourages the system from remaining trapped in a metastable basin and forces it to explore other regions. It also provides an estimate of the free energy, which is related to the probability distribution:

$$\frac{e^{-\beta F(\mathbf{s})}}{Z} = P(\mathbf{s}) \propto \int d\mathbf{R} e^{-\beta U(\mathbf{R})} \delta[\mathbf{s} - \mathbf{s}(\mathbf{R})], \quad (2.29)$$

$Z = \int d\mathbf{s} e^{-\beta F(\mathbf{s})}$ being the normalization partition function, $U(\mathbf{R})$ the potential energy, and $\beta = 1/(k_B T)$.

Metadynamics builds its bias as a sum of Gaussian contributions deposited periodically. In particular we will use its well-tempered variant (WTMetaD) [55], in which the height of the deposited Gaussian decreases exponentially as the bias is deposited:

$$V_n(\mathbf{s}) = \sum_{k=1}^n G(\mathbf{s}, \mathbf{s}_k) e^{-\beta/(\gamma-1) V_{k-1}(\mathbf{s}_k)}, \quad (2.30)$$

where $G(\mathbf{s}, \mathbf{s}_k) = W e^{-\|\mathbf{s} - \mathbf{s}_k\|^2}$ is a Gaussian centered in \mathbf{s}_k , the CVs value at time t_k . Independently from the choice of the Gaussian kernel parameters, this bias converges in the asymptotic limit to

$$V(\mathbf{s}) = -(1 - 1/\gamma)F(\mathbf{s}) + c, \quad (2.31)$$

where $F(\mathbf{s})$ is the FES and $c = c(t)$ does not depend on \mathbf{s} . Once at convergence, the biased ensemble distribution $P_V(\mathbf{s}) \propto e^{-\beta[F(\mathbf{s})+V(\mathbf{s})]}$ is a smoother version of the unbiased one, i.e.: $P_V(\mathbf{s}) = [P(\mathbf{s})]^{1/\gamma}$. This broadening is controlled by the bias factor γ , which can go from 1 (unbiased case) up to infinity (non-well-tempered MetaD).

In the VES method, on the other hand, the bias is obtained as the result of the minimization of a convex functional:

$$\Omega[V] = \frac{1}{\beta} \log \frac{\int d\mathbf{s} e^{-\beta[F(\mathbf{s})+V(\mathbf{s})]}}{\int d\mathbf{s} e^{-\beta F(\mathbf{s})}} + \int d\mathbf{s} p(\mathbf{s})V(\mathbf{s}), \quad (2.32)$$

where $p(\mathbf{s})$ is an arbitrary probability distribution called target distribution. The VES functional is related to the relative entropy, or Kullback-Leibler divergence [56]. The minimum condition is reached when $P_V(\mathbf{s}) = p(\mathbf{s})$, a relation that can also be written as

$$F(\mathbf{s}) = -V(\mathbf{s}) - \frac{1}{\beta} \log p(\mathbf{s}). \quad (2.33)$$




Figure 2.6: 2D free energy surface of the considered model, with the two basins A and B .

It is possible to choose as target distribution the well-tempered one, $p(\mathbf{s}) = [P(\mathbf{s})]^{1/\gamma}$. In this case the target has to be self-consistently adjusted as shown in Ref. [21]. At convergence one obtains the same bias potential as WT-MetaD.

In order to carry out the minimization of $\Omega[V]$, the bias is usually expanded over an orthogonal basis set $V(\mathbf{s}) = \sum_k \alpha_k f_k(\mathbf{s})$, or parametrized according to some physically motivated FES model [23, 26, 56]. In this way the functional becomes a function of the variational parameters and can be minimized using a stochastic gradient descent algorithm as described in Ref. [5].

2.2.3 Illustrative Model

In order to better acquaint the reader with the problem described in the introduction (Sec. 2.2.1), we illustrate the phenomenology associated with the use of suboptimal CVs in a simple model. The model consists of a single particle moving with a Langevin dynamics on a 2D potential energy surface (see Fig. 2.6) with two metastable basins, A and B , such that transitions between them are rare. This system has only two degrees of freedom, namely the x and y positions. If we take x as CV and perform a well-tempered MetaD run (details in the Supporting Information (SI), Sec. 2.2.7), we obtain the result in Fig. 2.7. The fact that we are missing one degree of freedom can be seen from the hysteresis apparent in the CV time evolution. In order to go from one basin to the other, one must wait for a rare y fluctuation, resulting in a delay that cannot be enhanced by adding a bias on x . Nevertheless,

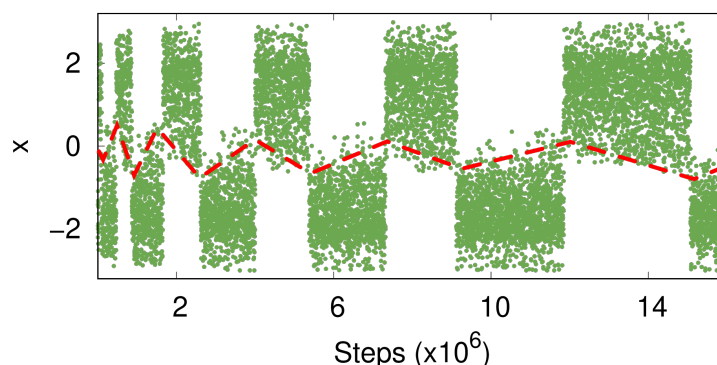


Figure 2.7: Typical trajectory of a suboptimal CV, obtained via a WTMetaD run ($\gamma = 10$). Hysteresis can clearly be seen (highlighted by the red dashed line), and it is also visible how WTMetaD slowly reduces it, by depositing less bias as the simulation proceeds. While it is not possible to draw an horizontal line that separates the two basins, one can easily determine the basin of belonging for each point by inspecting the CV time evolution.

if we put more bias in one basin we open some higher pathways and transitions are observed. However, they do not generally follow the lowest free energy path. This extra bias will also prevent the system from coming back to the same basin until it is properly compensated as the MetaD simulation progresses, and this gives rise to hysteresis.

Well-tempering reduces the amount of bias deposited in such a way that the hysteresis will eventually disappear, and in the asymptotic limit the bias will converge [54].

If we monitor the time dependence of the bias $V(x)$, we can see how the shape of the two basins is learned after just the first few transitions and remains pretty much unchanged throughout the simulation. After this quick initial phase, the vast majority of the computational effort is used to pin down the free energy difference ΔF between the two basins.

2.2.4 Method

The observations made earlier suggest a two-step strategy. First one obtains the shape of the free energy surface in the different metastable basins, then optimizes the free energy difference between them. We will refer to this two-step procedure as $VES\Delta F$.

Free Energy Surface Model

In a rare event scenario, such as the one discussed in Sec. 2.2.3, there are by definition distinct metastable basins and it is thus possible to identify their separate contribution to the global FES. Formally this can be done by

considering the conditional probability $P(\mathbf{s}|A)$, i.e., the probability that the CVs have value \mathbf{s} when the system is in basin A .

$$P(\mathbf{s}|A) \propto \int_A d\mathbf{R} e^{-\beta U(\mathbf{R})} \delta[\mathbf{s} - \mathbf{s}(\mathbf{R})]. \quad (2.34)$$

From $P(\mathbf{s}|A)$ one then can write an associated free energy:

$$\frac{e^{-\beta F_A(\mathbf{s})}}{Z_A} = P(\mathbf{s}|A), \quad (2.35)$$

where Z_A is a normalization constant. With an identical procedure one can define $F_B(\mathbf{s}|B)$.

Notice that in general these probabilities are not mutually exclusive and, especially in the case of suboptimal CVs, there can be an overlapping region between them (see Fig. 2.7).

The global probability distribution is then obtained simply by combining these conditional probabilities:

$$P(\mathbf{s}) = P(\mathbf{s}|A)P_A + P(\mathbf{s}|B)P_B, \quad (2.36)$$

and the global free energy (modulo a constant) can thus be written as

$$F(\mathbf{s}) = -\frac{1}{\beta} \log \left[\frac{e^{-\beta F_A(\mathbf{s})}}{Z_A} + \frac{e^{-\beta F_B(\mathbf{s})}}{Z_B} e^{-\beta \Delta F} \right], \quad (2.37)$$

where the free energy difference between the basins ΔF is defined by

$$e^{-\beta \Delta F} = \frac{P_B}{P_A} = \frac{\int_B d\mathbf{R} e^{-\beta U(\mathbf{R})}}{\int_A d\mathbf{R} e^{-\beta U(\mathbf{R})}}. \quad (2.38)$$

One could also write the global FES as a function of ΔF_h , defined as the difference in height between the two free energy minima. This can be a fairly good and practical approximation to ΔF , especially if the two basins have similar shapes. More precisely, well within the typical statistical uncertainty (see SI), the two quantities differ by a constant that depends only on $F_A(\mathbf{s})$ and $F_B(\mathbf{s})$. If the local free energies are shifted so that $\min[F_A(\mathbf{s})] = \min[F_B(\mathbf{s})] = 0$, one can write

$$\Delta F_h = \Delta F - \frac{1}{\beta} \log \frac{Z_A}{Z_B}, \quad (2.39)$$

$$F(\mathbf{s}) = -\frac{1}{\beta} \log \left[e^{-\beta F_A(\mathbf{s})} + e^{-\beta F_B(\mathbf{s})} e^{-\beta \Delta F_h} \right], \quad (2.40)$$

with $\min[F(\mathbf{s})] = 0$.

Step One: Local Basins

There are different ways of calculating the local free energies, $F_A(\mathbf{s})$ and $F_B(\mathbf{s})$; here we suggest that a good estimate can be obtained using MetaD. We do this by accumulating the bias up to the point in which the system escapes the basin from which the simulation was started. As can be seen, e.g., in Fig. 2.7, the first recrossing usually happens extremely fast, even if the CVs are suboptimal. To increase the accuracy, WTMetaD can be used instead of plain MetaD, and also multiple runs can be combined, starting from different initial conditions.

With this initial estimate it is important to be able to draw the two local FES up to the transition region, because this will ensure that they combine smoothly in Eq. (2.37). Such information is always obtained in the case of suboptimal CVs, thanks to the typical hysteresis.

In our approach we start N independent simulations from basin A and stop each simulation as soon as a transition occurs. The bias $V_i(\mathbf{s})$ deposited by replica i is then averaged with the others to obtain $V_A(\mathbf{s}) = 1/N \sum_i V_i(\mathbf{s})$ which is used to estimate the local free energy as $F_A(\mathbf{s}) = -(1 - 1/\gamma)^{-1} V_A(\mathbf{s})$, following the usual MetaD rules. The same is done for basin B .

A very simple way to automatically stop the simulation at the first recrossing is to stop it when the CV reaches a threshold value set well beyond the transition region (e.g., at the center of the other basin), and then cut a small segment at the end of the trajectory, to make sure that only configurations from the proper basin are kept. Another option for detecting a recrossing is to use a descriptor different from the biased CV. The only requirement for such a descriptor is that it can partition the phase space in two regions, one for each basins; thus, contrary to a CV, it can be a discontinuous function of the atomistic coordinates or even a function of past configurations, such as an exponentially decaying time average, which can greatly improve basins separation.

In case one has already performed an exploratory MetaD simulation with multiple transitions between the basins, another possibility for obtaining $F_A(\mathbf{s})$ and $F_B(\mathbf{s})$ is to separate the CVs trajectory and build a different reweighted free energy for each basin.

It should be noted that at this point an extremely precise determination of the free energy basins is not needed, since a more accurate description can be obtained later with a reweighting procedure of the longer convergence run. In our experience, accuracy of the order of $k_B T$ in $F_A(\mathbf{s})$ and $F_B(\mathbf{s})$ does not slow down the overall convergence.

Step Two: Free Energy Difference

Once the local basins are obtained we still need to know their relative free energy difference, ΔF , in order to build the global FES. We use VES to estimate it.

As target distribution for VES we use the well-tempered one as in Ref. [21], because it remains controllably close to the physical distribution, while at the same time enhancing the transition rate. We can then write the bias potential expansion as a function of a single parameter, ΔF , by combining Eq. (2.37) and (2.33)

$$V(\mathbf{s}) = (1 - 1/\gamma) \frac{1}{\beta} \log \left[\frac{e^{-\beta F_A(\mathbf{s})}}{Z_A} + \frac{e^{-\beta F_B(\mathbf{s})}}{Z_B} e^{-\beta \Delta F} \right]. \quad (2.41)$$

We can then minimize Ω as a function of ΔF , while updating the estimate of the target distribution in a self-consistent way.

Given the suboptimal nature of the CVs, it was somehow natural to consider the introduction of a damping factor in the optimization algorithm commonly used in VES [31]. For this we took inspiration from the AdaGrad [57] stochastic gradient descent algorithm, generally used for training neural networks in a sparse reward scenario. Further details can be found in the Supporting Information (Sec. 2.2.7).

It is relevant to note that during the minimization procedure no extra bias is added; thus, the system remains in the region of CVs space explored during step one and does not spend time unnecessarily in high free energy configurations.

Reweighting Procedure

The reweighting procedure is a key part of the VES ΔF method, because (as we will see) it converges faster than the VES optimization itself and provides a more accurate estimate of ΔF . For reweighting we employ the scheme of Ref. [58] and [59], where the unbiased Boltzmann distribution, $P(\mathbf{s})$, is obtained by sampling the biased one, $P_V(\mathbf{s})$:

$$P(\mathbf{s}) = \frac{e^{-\beta F(\mathbf{s})}}{Z} = \frac{Z_V}{Z} e^{+\beta V(\mathbf{s},t)} \frac{e^{-\beta[F(\mathbf{s})+V(\mathbf{s},t)]}}{Z_V} = e^{-\beta[V(\mathbf{s},t)-c(t)]} P_V(\mathbf{s}), \quad (2.42)$$

where $Z_V = \int d\mathbf{s} e^{-\beta[F(\mathbf{s})+V(\mathbf{s},t)]}$ and the ratio $Z_V/Z = e^{-\beta c(t)}$ is estimated by numerical integration on a grid in the CVs space, using $F(\mathbf{s}) = -(1 - 1/\gamma)^{-1} V(\mathbf{s})$. This reweighting procedure can be used only in the adiabatic limit, when the applied bias is quasi-stationary.

When dealing with an optimal CV, a good biasing strategy is to quickly update the estimate of the underlying FES and correct accordingly the applied

bias. Since the CV incorporates all the relevant slow modes of the system, this strategy ensures a faster transition rate between the metastable basins and a faster convergence. However, if the CV is suboptimal, as soon as we are close enough to a good FES estimate, the bottleneck for making a transition becomes the unbiased slow mode. In such a scenario, the choice of rapidly changing the applied bias is no more convenient, as it will only lead to hysteresis, not to faster convergence. This is one of the reasons why well-tempered metadynamics outperforms plain metadynamics in many applications and why the “first fill, then converge” strategy of transition-tempered metadynamics [60] can be so effective.

In VES ΔF we push this strategy even further, and we design the bias optimization in such a way that it quickly reaches a value close to the optimal one, and from that one we only make gentle adjustments. This allows us to maximize the time spent in the adiabatic limit, thus improving the reweighting efficiency. As can be seen, e.g., in Fig. 2.9, the result is that the estimate coming from the reweighting can reach convergence faster than the direct ΔF optimization.

The other aspect that makes reweighting a convenient tool is accuracy. We do not spend much time in step one; thus the shape of the local basins may not be very accurate. In principle this can lead to a systematic error in the estimate of ΔF , since it depends not only on the relative height of the basins, ΔF_h , but also on their shapes, Eq. (2.39). Thus, a safer way of estimating ΔF is to obtain a reweighted $F(\mathbf{s})$ and to explicitly integrate in the CVs space:

$$\Delta F = -\frac{1}{\beta} \log \left[\frac{\int_B d\mathbf{s} e^{-\beta F(\mathbf{s})}}{\int_A d\mathbf{s} e^{-\beta F(\mathbf{s})}} \right]. \quad (2.43)$$

Such an estimate is based on the much longer step-two run and is a very good approximation provided that the CVs are able to distinguish the two basins, which is usually the case even for suboptimal CVs. This is why in the code implementation we prefer to use the FES model of Eq. (2.40), and to optimize Ω with respect to ΔF_h , and then to use reweighting to estimate ΔF .

2.2.5 Results

The VES ΔF method has been implemented in PLUMED [38], an enhanced sampling plug-in which we used in combination with GROMACS [37] and LAMMPS [61] molecular dynamics packages. Our code is openly available online, in the development version of PLUMED (master branch on GitHub), under the VES module, with the name VES_DELTA.F. The provided implementation can deal with multidimensional CVs and more than two FES basins.

2.2. A multiscale approach to free energy calculation




Figure 2.8: (Left) Local free energies $F_A(x)$ and $F_B(x)$ in comparison with the model one, $F_{model}(x)$, obtained by combining them following Eq. (2.40). (Right) Same model free energy in comparison with the reference one, $F_{ref}(x)$, that can be obtained via the reweighting procedure. Parameter ΔF_h is unknown and is calculated during the VES optimization.

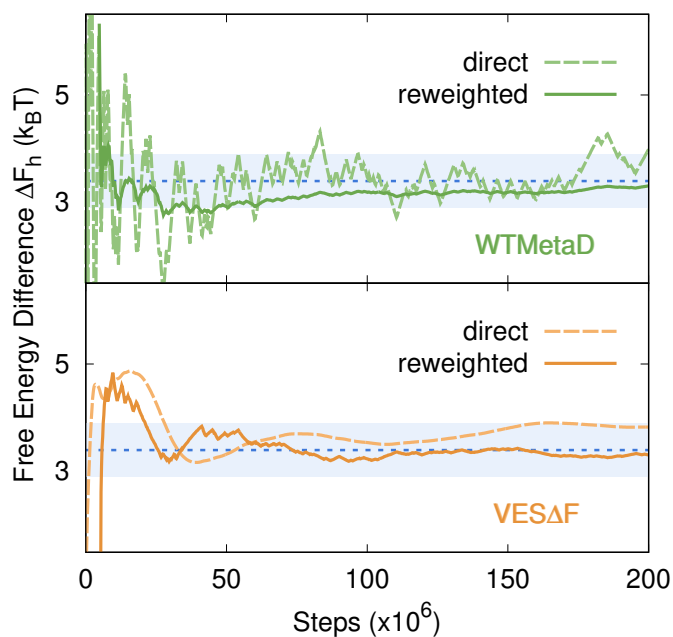


Figure 2.9: Estimate for the illustrative model of Sec. 2.2.3 of ΔF_h along x , obtained directly from the acting bias and from the reweighting procedure (Sec. 2.2.4), in the case of WTMetaD and VES ΔF . In both methods a bias factor $\gamma = 10$ is used. The reference blue stripe is $1 k_B T$ thick.




Figure 2.10: Free energy surface of alanine dipeptide in vacuum, as a function of the φ and ψ angles. On the sides is the free energy projected onto each single angle.

First we apply our approach to the case of the simple model of Sec. 2.2.3. In order to build the local free energy basins we follow step one (Sec. 2.2.4) and run five independent replicas using WTMetaD with $\gamma = 5$ enhancing x fluctuations. We obtain the $F_A(x)$ and $F_B(x)$ shown in Fig. 2.8, with a computational effort that is negligible in comparison to the one needed to converge the global FES.

We use these local basins to perform step two and optimize the free energy difference through VES ($\gamma = 10$, $\mu = 0.05$). Fig. 2.9 shows an example of such optimization, compared to a WTMetaD run. We also show the estimate of ΔF_h obtained with the reweighting procedure.

Both WTMetaD and VES ΔF reach relatively quickly a rough estimate of the ΔF_h between the basins, but while the bias used by the former keeps oscillating, the latter behaves more smoothly. We believe this is ultimately the reason why the reweighting converges faster in the case of VES ΔF .

Alanine Dipeptide

The alanine dipeptide molecule in vacuum is often used as benchmark for enhanced sampling methods. At 300 K it presents two main metastable basins, separated by a kinetic bottleneck. The most stable one, *A*, includes two different conformations, C5 and C7eq, that are separated only by a small barrier. Basin *B* instead hosts only one metastable conformation, known as C7ax.




Figure 2.11: Local free energies obtained for alanine dipeptide. As reference a fully converged $F(\psi)$ is used. The quantity ΔF_h is unknown at this stage but is used here for displaying purposes.

The typical CVs used for this system are the two Ramachandran angles φ and ψ , Fig. 2.10. The first one, φ , is an almost optimal CV, which allows for very fast convergence (~ 10 ns) when used in MetaD. Instead ψ is not only suboptimal, but also a typical example of a bad CV. When biasing ψ , transitions between the basins are enhanced, but the system remains multiscale, with a significant time scale gap between intrabasin and interbasins fluctuations. In this scenario MetaD presents a strong hysteresis and its convergence is rather slow. For the purpose of showing the strength of our method, we will pretend here that we do not know of the existence of the φ angle, and we make the unfortunate choice of biasing only ψ .

It is important to recall that while the free energy surface as a function of ψ is very different from the one relative to φ (and thus the minima relative height ΔF_h is different), the free energy difference, ΔF , between the two basins does not depend on the CV used, and is the same independently of the CV representation of the FES, along ψ , φ , or both (see SI). This holds true because all of these CVs can correctly distinguish the relevant metastable basins.

In order to reconstruct the local free energies, we run 10 independent WT-MetaD simulations for each basin, terminating them as soon as they make the first transition. We use a bias factor $\gamma = 10$, and other simulation details can be found in the Supporting Information (Sec. 2.2.7). On average it takes 0.55 ns to exit the most stable basin A and 0.16 ns to escape basin B . The combined total simulation time employed for determining $F_A(\psi)$ and $F_B(\psi)$ is about 7 ns. This provides a very good estimate, as can be seen in Fig. 2.11. For the ΔF_h optimization we use a target distribution with a well-tempered bias factor $\gamma = 10$ and a minimization step $\mu = 0.05$. As a reference, we perform runs of WTMetaD and transition tempered metady-

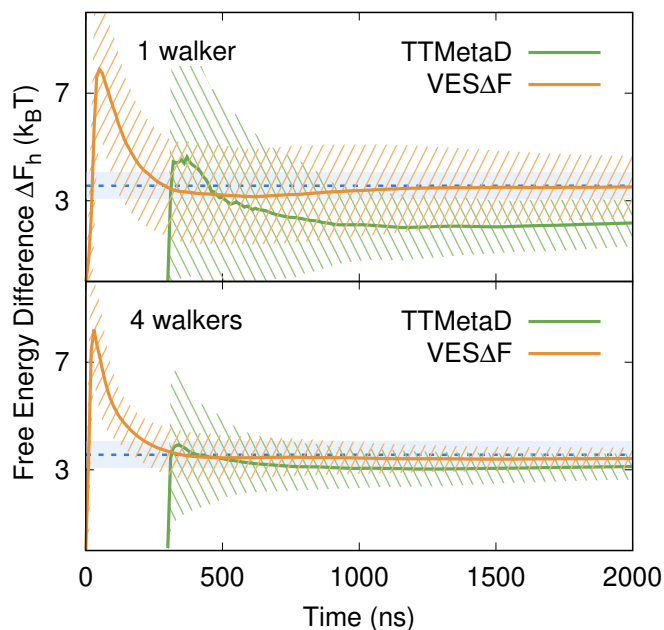


Figure 2.12: Comparison between alanine ΔF_h convergence using TTMetaD and VES ΔF , with one and four multiple walkers ($\gamma = 10$). ΔF_h is obtained through the reweighting procedure described in Sec. 2.2.4. In the case of MetaD we must exclude from the reweighting an initial transient, in which the system is out of equilibrium and the estimate of $c(t)$ is unreliable. We run 100 independent replicas starting from basin A with different initial conditions. Only the averages and the standard deviations are shown. The reference blue stripe is $1 k_B T$ thick.

namics [60] (TTMetaD), with $\gamma = 10$ (see SI). We also perform simulations with different numbers of multiple walkers [41].

In order to have an estimate of the uncertainty on the calculated ΔF_h , we run for each of the three methods multiple completely independent runs, always starting from basin A , but with different initial conditions. We then look at the average and the standard deviation of these replicas, as a function of time. Some of the results are shown in Fig. 2.12, and more can be found in the Supporting Information (Sec. 2.2.7).

On average TTMetaD performs better than plain WTMetaD, while VES ΔF is a great improvement over either method. In particular the VES ΔF reweighting is more accurate even in the case of a single walker, where error bars are large and the optimization is not yet at convergence. MetaD methods show a systematic shift which is not compatible with the estimates coming from the bias and disappears when more statistics are collected (the same happens in Fig. 2.13).

Increasing the number of walkers makes sampling of the neglected slow degrees of freedom more efficient and improves convergence. Our method, VES ΔF , scales particularly well with the number of multiple walkers.

Finally we stress the fact that we are not suggesting to purposely pick bad CVs, because even if our method brings better results when compared to MetaD, a simulation based on φ instead of ψ would still converge with a computational effort smaller by orders of magnitude.

Sodium

As a last example we study the liquid-solid transition of sodium at 350 K. It is a first order phase transition, and the stable solid structure is a body centered cubic (bcc) lattice.

Contrary to the previous case, for this example we use a CV that is one of the best available to describe the system. We consider a recently developed CV [62] based on the peaks of the Debye structure factor (see SI, Sec. 2.2.7). Such a CV is able to clearly identify the liquid and the solid phase, and can drive transitions between the two. Nevertheless it is not an optimal CV and has some hysteresis, as can be seen, e.g., in Fig. 3 of Ref. [62].

Although suboptimal, this CV is able in our system to drive the transitions between liquid and solid avoiding defected states, which might introduce secondary metastable minima in the FES and affect the estimate of the liquid-solid ΔF .

Similarly to what was done in the previous case, we first build the local free energy basins by running five short WMetaD runs ($\gamma = 20$), then use them to optimize the free energy difference. The target distribution is a well-tempered one with $\gamma = 20$, and the minimization step is $\mu = 1$. In Fig. 2.13 we show the convergence of $VES\Delta F$ in comparison with WMetaD. Both simulations run with four parallel multiple walkers; see SI (Sec. 2.2.7) for more details.

Also in this less extreme case of a suboptimal CV, $VES\Delta F$ outperforms the standard approaches.

2.2.6 Conclusions

In this work we propose a new method, based on a combination of MetaD and VES, to calculate the free energy difference between two metastable states. In order to perform such a calculation, MetaD and VES are in principle very appropriate, especially when used with an optimal CV that is able to accelerate all the slow modes of the system. In such a scenario the computational effort needed for convergence is close to that needed for exploring the FES landscape.

However, in real life applications suboptimal CVs are used, and this makes it useful to separate the free energy reconstruction into two steps. First we use MetaD to explore the basins, and then a bespoke version of VES to

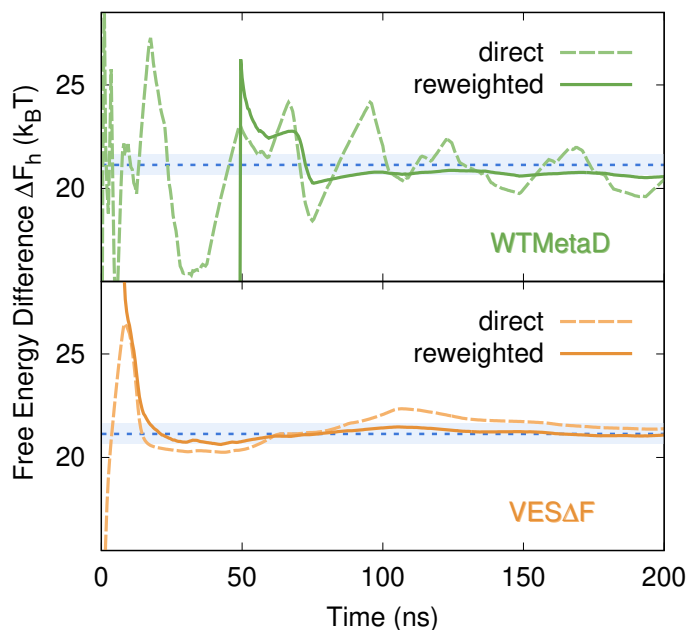


Figure 2.13: Convergence of ΔF_h between liquid and solid for sodium at 350 K, obtained directly from the acting bias and from the reweighting procedure (Sec. 2.2.4), using WTMetaD and VES ΔF . In both methods bias factor $\gamma = 20$ and four multiple walkers were used. The reference blue stripe is $1 k_B T$ thick.

converge the free energy difference. In doing so, our method focuses on approaching a quasi-stationary bias, which improves the efficiency of the reweighting procedure and thus the convergence speed.

In some cases one does not know in advance if the available CV is suboptimal or not, and which are the relevant basins of the system under study. A typical usage scenario would then be to run a first rough non well-tempered MetaD simulation, aimed at discovering the relevant free energy basins and the presence (or absence) of hysteresis. If this preliminary run reveals the presence of a big number of secondary metastable minima, as it can happen e.g. when dealing with defects in a phase transitions, then VES ΔF would not be the best choice.

We notice that in a suboptimal scenario the VES optimization is crucial, and the proper ΔF cannot be retrieved via some reweighting technique, such as weighted histogram analysis (WHAM [63]), applied to the MetaD simulations in the basins. This is because the overlap observed in the CV space is not actually an overlap in the phase space. This issue would not be solved by simply adding new biased windows (in an umbrella sampling spirit), since the hysteresis would still be present [64].

The present method, VES ΔF , is to some extent similar to that proposed in

Ref. [23], but its motivations are different, as are the FES model used and the strategy to obtain the local basins.

Previous authors have proposed variants to the MetaD algorithm to improve its convergence rate [55, 60, 65], but none of them connects directly to the source of the problem, namely, the residual multiscale behavior of suboptimal CVs. We think this is a simple but important consideration, and our approach is meant to fully take into account the multiscale nature of suboptimal CVs enhanced sampling.

Another important feature is that it is very transparent, having only ΔF as a parameter to optimize, that is the very physical quantity one is interested in. It also focuses only on the relevant part of the CV space, avoiding exploring new regions during the convergence phase. This simplicity can be very helpful, especially when dealing with complex systems.

Furthermore, VES ΔF scales very well in the case of multiple walkers, making good use of parallel simulations. For the sake of simplicity we presented here only the case of one-dimensional CVs and two-minima systems, but the implementation of VES ΔF provided in the publicly available PLUMED code can already deal with multidimensional CVs and more than two basins. Our work points to a useful strategy that can be applied also in other circumstances and with other methods.

Acknowledgement

This research was supported by the NCCR MARVEL, funded by the Swiss National Science Foundation, and European Union Grant No. ERC-2014-AdG-670227/VARMET. Calculations were carried out on Euler cluster at ETH Zurich and on the Mönch cluster at the Swiss National Supercomputing Center. We thank Emanuele Grifoni for useful discussions, Giovanni-Maria Piccini for precious help with the artworks, and Omar Valsson for carefully reading the manuscript.

2.2.7 Supporting Information

Sampling with different CVs

We would like to give a simple example of the different sampling that can be accomplished by using different kind of collective variables for biasing. We use alanine dipeptide, and we run three different WMetaD simulations ($\gamma = 10$) biasing respectively only the “bad” angle ψ , only the “good” angle φ , and both angles. The three simulation require very different time to converge, with the bad angle one requiring more than one order of magnitude more time, but we show in Fig. 2.14 the same number of points for each simulation. We used WMetaD in this example, but the point we want to




Figure 2.14: Different sampling obtained by employing different CVs for building the bias in WTMetaD. All simulations are fully at convergence, and only 100000 point are shown for each one. For a reference free energy see Fig. 5 of main text.

make is not limited to this method, it is a characteristic shared with all CV-based enhanced sampling techniques. We would obtain the same figure if we were to use a static bias instead: respectively $V(\psi) = -(1 - 1/\gamma)F(\psi)$, $V(\phi) = -(1 - 1/\gamma)F(\phi)$, and $V(\phi, \psi) = -(1 - 1/\gamma)F(\phi, \psi)$.

It can be seen how, when a slow degree of freedom is missing from the bias, a lot of simulation time is spent in the minima, instead of exploring new configurations, and some high free energy region are almost never visited. Despite this, the three simulations all provide the same estimate for ΔF , when using Eq. (2.43) to calculate it from the reweighting. This is due to the fact that the contribution from these high free energy regions is extremely small and can be safely ignored. If instead one is interested not only in the free energy difference, but on the full FES along a specific CV, a good strategy would be to directly bias also such CV.

Notes on Eqs. (2.40) and (2.39)

We first define \mathbf{s}_A and \mathbf{s}_B as the points where F_A and F_B respectively are minimum, thus $F_A(\mathbf{s}_A) = 0$ and $F_B(\mathbf{s}_B) = 0$. If A is the most stable basin, we can then explicitly write $\Delta F_h = F(\mathbf{s}_B) - F(\mathbf{s}_A)$, where $F(\mathbf{s})$ is the global FES.

Eqs. (2.40) and (2.39) of the main text are strictly true only in the limit of

$$F_A(\mathbf{s}_B) \rightarrow \infty \text{ and } F_B(\mathbf{s}_A) \rightarrow \infty. \quad (2.44)$$

Outside of this limit, Eq. (2.40) holds true if we substitute ΔF_h with another quantity, that we call $\Delta F'_h$. The exact relation between these quantities is the following:

$$\Delta F_h = \Delta F'_h - \frac{1}{\beta} \log \left[\frac{1 + e^{-\beta(F_B(\mathbf{s}_A) - \Delta F'_h)}}}{1 + e^{-\beta(F_A(\mathbf{s}_B) + \Delta F'_h)}} \right]. \quad (2.45)$$

Also, if condition (2.44) does not hold, we have

$$\min[F(\mathbf{s})] = F(\mathbf{s}_A) = -\frac{1}{\beta} \log \left[1 + e^{-\beta(F_B(\mathbf{s}_A) - \Delta F'_h)} \right], \quad (2.46)$$

which in general is non-zero. A similar correction can be found also for Eq. (2.39).

In VES ΔF we use approximate estimates for the local basins and condition (2.44) cannot formally be fulfilled, but from a practical point of view this does not constitute a problem. In fact, the difference between $\Delta F'_h$ and ΔF_h is typically some order of magnitude smaller than the uncertainty of the free energy calculation. For the sodium and the model system examples, this discrepancy stays well below $10^{-5} k_B T$, while for alanine dipeptide, where the employed CV is quite bad at distinguishing the two basins, it stays below $0.1 k_B T$.

It is important to remember that this approximation does not affect in any way the ΔF or ΔF_h calculated through the reweighting procedure (as in Fig. (2.12) of main text).

Optimization algorithm

The optimization algorithm used in the second step of VES ΔF , is a novel combination of the optimization algorithm commonly used in VES [5, 31] and AdaGrad [57] stochastic gradient descent algorithm.

In the present paper we only consider the case of two basins, thus we deal with one optimization parameter ΔF (or ΔF_h , depending on the chosen normalization). Here however we present a more general version of the optimization algorithm, using a multidimensional vector of parameters $\boldsymbol{\alpha}$.

Following Ref. [5], we write gradient and Hessian of the VES functional:

$$\frac{\partial \Omega(\boldsymbol{\alpha})}{\partial \alpha_i} = - \left\langle \frac{\partial V(\mathbf{s}; \boldsymbol{\alpha})}{\partial \alpha_i} \right\rangle_{V(\boldsymbol{\alpha})} + \left\langle \frac{\partial V(\mathbf{s}; \boldsymbol{\alpha})}{\partial \alpha_i} \right\rangle_p, \quad (2.47)$$

$$\begin{aligned} \frac{\partial^2 \Omega(\boldsymbol{\alpha})}{\partial \alpha_i \partial \alpha_j} = & \beta \left[\left\langle \frac{\partial V(\mathbf{s}; \boldsymbol{\alpha})}{\partial \alpha_i} \frac{\partial V(\mathbf{s}; \boldsymbol{\alpha})}{\partial \alpha_j} \right\rangle_{V(\boldsymbol{\alpha})} - \left\langle \frac{\partial V(\mathbf{s}; \boldsymbol{\alpha})}{\partial \alpha_i} \right\rangle_{V(\boldsymbol{\alpha})} \left\langle \frac{\partial V(\mathbf{s}; \boldsymbol{\alpha})}{\partial \alpha_j} \right\rangle_{V(\boldsymbol{\alpha})} \right] + \\ & - \left\langle \frac{\partial^2 V(\mathbf{s}; \boldsymbol{\alpha})}{\partial \alpha_i \partial \alpha_j} \right\rangle_{V(\boldsymbol{\alpha})} + \left\langle \frac{\partial^2 V(\mathbf{s}; \boldsymbol{\alpha})}{\partial \alpha_i \partial \alpha_j} \right\rangle_p, \end{aligned} \quad (2.48)$$

where the averages are calculated either in the biased ensemble $\langle \cdot \rangle_{V(\boldsymbol{\alpha})}$ (by running the molecular dynamics), or in the target $p(\mathbf{s})$ ensemble $\langle \cdot \rangle_p$ (by explicit integration on a grid).

We make use of a second set of auxiliary parameters for updating the main one. At each iteration n , we have some auxiliary instantaneous iterate $\alpha^{(n)}$, while the actual parameters are obtained as their averages

$$\bar{\alpha}^{(n)} = \frac{1}{n+1} \sum_{k=0}^n \alpha^{(k)}. \quad (2.49)$$

Gradient and Hessian are always evaluated using this set of averaged parameters.

We use the above gradient and Hessian, Eq. (2.47) and (2.48), to define an effective gradient $\mathbf{g}^{(n)}$:

$$g_i^{(n)} = \frac{\partial \Omega(\bar{\alpha}^{(n)})}{\partial \alpha_i} + \sum_j \frac{\partial^2 \Omega(\bar{\alpha}^{(n)})}{\partial \alpha_i \partial \alpha_j} (\alpha_j^{(n)} - \bar{\alpha}_j^{(n)}). \quad (2.50)$$

In the spirit of the AdaGrad algorithm, we then introduce an history dependent damping factor $\mathbf{d}^{(n)}$:

$$d_i^{(n)} = \sqrt{[d_i^{(n-1)}]^2 + [g_i^{(n)}]^2}. \quad (2.51)$$

We can now write the update rule for the auxiliary parameters:

$$\alpha_i^{(n+1)} = \alpha_i^{(n)} - \frac{\mu}{d_i^{(n)}} g_i^{(n)}, \quad (2.52)$$

where μ is a fixed optimization hyperparameter. The standard VES optimization algorithm is retrieved if we set $d_i^{(n)} = 1$ for each i and n .

Hyperparameters choice The hyperparameters the user needs to set in VES ΔF are essentially three:

- the bias factor for the target distribution, γ
- the update stride for the parameters, during which the ensemble averages are estimated
- the optimization step μ

For the choice of the bias factor we use criteria similar to those employed for well-tempered VES [21] and WMetaD [55], the main difference being that our method is less sensitive to the choice of γ . In particular, contrary to what happens in WMetaD, the value of γ does not have a direct impact on the speed at which bias is added, allowing for more flexibility. A similar effect is obtained also in globally tempered MetaD [54].

We chose an update stride of 1 ps that is a typical value for both MetaD and standard VES.

The choice of the optimization step μ is in our experience the most crucial, but it was not a hard one. We did not try systematically to optimize this choice, since there is a reasonable range of “good” values. We did notice, though, that an extreme choice of this hyperparameter can make convergence terribly slow. In the systems studied here, we noticed that the choice of μ had a direct impact on the speed at which the parameter α (thus ΔF or ΔF_n) grows, independently of other factors like γ and the number of multiple walkers. As an empirical rule at the beginning of the optimization, before any transition has taken place, we have $\alpha(t) \approx 1.2\mu\sqrt{t}$. Further investigation is needed in order to provide a simple rule of thumb for the choice of μ .

Illustrative model

To run the illustrative model considered in Sec. 3 of the main text we use a simple molecular dynamics code implemented in PLUMED [38] (version 2.4 or higher), called `ves_md_linearexpansion`. We used the default parameters for the simulation (`tstep=0.005`, `temperature=1`, `friction=10`) and the Wolfe-Quapp potential,

$$U(x, y) = x^4 + y^4 - 2x^2 - 4y^2 + xy + 0.3x + 0.1y, \quad (2.53)$$

which we rotated of an angle $\theta = -0.15\pi$ by rotating the coordinates, thus:

$$\begin{aligned} x &\rightarrow x \cos \theta - y \sin \theta \\ y &\rightarrow x \sin \theta + y \cos \theta \end{aligned} \quad (2.54)$$

The resulting potential is:

$$\begin{aligned} U(x, y) = & 1.34549 x^4 + 1.90211 x^3 y + 3.92705 x^2 y^2 \\ & - 6.44246 x^2 - 1.90211 x y^3 + 5.58721 x y + 1.33481 x \\ & + 1.34549 y^4 - 5.55754 y^2 + 0.904586 y + 18.5598, \end{aligned} \quad (2.55)$$

where the last shift term is added just to put the minimum at zero.

For the MetaD simulations we always use the following parameters: `PACE=500`, `HEIGHT=1.2`, `SIGMA=0.35`, `BIASFACTOR=10`, and store the bias on a grid of `GRID_BIN=300`, `GRID_MIN=-3`, `GRID_MAX=3`. The local free energy basins obtained for $\text{VES}\Delta F$ are stored on the same kind of grid, and required a total combined number of 2.11×10^6 simulation steps. The parameters used for the VES optimization are: `AV_STRIDE=500`, `M_STEP=0.05`, `BIASFACTOR=10`.

Alanine dipeptide

For the Alanine dipeptide simulations we use GROMACS [37] patched with PLUMED. The setup is the same of Ref. [5], namely: *NVT* simulation in a vacuum, Amber99-SB [66] force field, time step 2 fs, temperature 300 K, velocity rescaling thermostat [36].

The MetaD simulations use the following parameters: `PACE=500`, `HEIGHT=1.2`, `SIGMA=0.35`, `BIASFACTOR=10`, and store the bias on a grid of `GRID_BIN=100`, `GRID_MIN=-pi`, `GRID_MAX=pi`. For the TTMetaD the same parameters are used, plus `TTBIASFACTOR=10`, `TRANSITIONWELLO=0.8`, `TRANSITIONWELL1=2.7`. The local free energy basins obtained for $VES\Delta F$ are stored on the same kind of grid, and required a total combined simulation time of 7.10 ns. The parameters used for the VES optimization are: `AV_STRIDE=500`, `M_STEP=0.05`, `BIASFACTOR=10`.

Results We show here more in detail some of the results from our alanine dipeptide simulations. We performed a set of calculations with different number of walkers, comparing WTMetaD, TTMetaD and $VES\Delta F$ methods. Each calculation is repeated 10 times to gauge the error. All replicas and walkers start from basin *A* with different initial conditions and random seed. The same initial configurations are used for the three methods.

We notice how in this case TTMetaD converges better than WTMetaD. It can be seen how $VES\Delta F$ convergence is smoother than MetaD, and gains more from the reweighting procedure.

Finally, as noticed in the main text, increasing the number of walkers allows for a better sampling of the unbiased slow degree of freedom, and thus ameliorates convergence. This is typical of suboptimal CVs, whereas if an optimal CV is used one soon reaches a plateau in sampling efficiency when increasing the number of multiple walkers.

Sodium

For the sodium simulations we use LAMMPS [61] patched with PLUMED. An embedded atom model (EAM) is used as interatomic potential [67], molecular dynamics time step is 2 fs, and temperature 350 K. We use the stochastic velocity rescaling thermostat [36] (0.1 ps relaxation time) and an isotropic Parrinello-Rahman barostat [68] (1 ps relaxation time). The system size is 250 atoms.

The Debye structure factor CV we used is implemented in a development version of PLUMED, but it is openly available upon request and we plan to make it public in the near future. This collective variable is extensively described in Ref. [62]. We consider the first structure factor peak of the

2.2. A multiscale approach to free energy calculation




Figure 2.15: Comparison between alanine ΔF_h convergence for 10 replicas using WTMetaD, TTMetaD and VES ΔF , with 1, 4 and 16 multiple walkers ($\gamma = 10$). The ΔF_h is estimated directly from the bias applied, via the relation $F(s) = -(1 - 1/\gamma)^{-1}V(s)$. The reference blue stripe is $1 k_B T$ thick.

bcc solid, at a scattering frequency $Q = 2.070595 \text{ \AA}^{-1}$ and with a cutoff $R_c = 10.5 \text{ \AA}$.

For the MetaD simulations we use the following parameters: PACE=500, HEIGHT=10, SIGMA=0.2, BIASFACTOR=20, and store the bias on a grid of GRID.BIN=200, GRID_MIN=1, GRID_MAX=3. The local free energy basins obtained for VES ΔF are stored on the same kind of grid, and required a total combined simulation time of 6.75 ns. The parameters used for the VES optimization are: AV_STRIDE=500, M_STEP=1, BIASFACTOR=20. In both cases we used 4 walkers, all initialized in different random configurations inside basin A , obtained from an unbiased run. The reference ΔF_h value was obtained from a longer WTMetaD 10 walkers run.




Figure 2.16: Comparison between alanine ΔF_h convergence for 10 replicas using WTMetaD, TTMetaD and VES ΔF , with 1, 4 and 16 multiple walkers ($\gamma = 10$). The ΔF_h is estimated through the reweighting procedure described in the main text. In the case of MetaD we must exclude from the reweighting an initial transient, in which the system is out of equilibrium and the estimate of $c(t)$ is unreliable. The reference blue stripe is $1 k_B T$ thick.


2.3 Rethinking metadynamics: from bias potentials to probability distributions

This is the article that first presents the OPES method. Here it is used to sample the well-tempered target distribution, the same typically sampled by metadynamics. When targeting this kind of distributions, the OPES method needs to build an explicit estimate of the marginal probability along the chosen CVs. To do so, we use a weighted kernel density estimator together with a kernel compression algorithm that allows us to quickly build a coarse estimate and then only adiabatically adjust the finer details of the probability distribution. Another peculiar feature we introduce is a normalization factor with respect to the explored CV space that significantly speeds up the discovery of new metastable states, especially in case of multidimensional CVs. Following a common practice, the new enhanced sampling method is tested on the alanine dipeptide molecule using as CVs the torsional angles. However, the advantages of the OPES method become more evident in the more realistic case of suboptimal collective variables. The interested reader can find in Refs. [69–71] applications of OPES to other systems.

My contribution to this article has been implementing the algorithms, performing the simulations, and writing the paper jointly with Prof. Parrinello.

Reference: M. Invernizzi, and M. Parrinello. “Rethinking Metadynamics: From Bias Potentials to Probability Distributions.” *The Journal of Physical Chemistry Letters* 11.7 (2020): 2731-2736. URL <https://pubs.acs.org/doi/10.1021/acs.jpcllett.0c00497>

Copyright © 2020 American Chemical Society.



Rethinking Metadynamics: From Bias Potentials to Probability Distributions

Michele Invernizzi^{1,2}, and Michele Parrinello^{*3,2}

¹*Department of Physics, ETH Zurich c/o Università della Svizzera italiana, 6900 Lugano, Switzerland*

²*Facoltà di Informatica, Institute of Computational Science, and National Center for Computational Design and Discovery of Novel Materials (MARVEL), Università della Svizzera italiana, 6900 Lugano, Switzerland*

³*Department of Chemistry and Applied Biosciences, ETH Zurich c/o Università della Svizzera italiana, 6900 Lugano, Switzerland
and Italian Institute of Technology, 16163 Genova, Italy*

Abstract

Metadynamics is an enhanced sampling method of great popularity, based on the on-the-fly construction of a bias potential that is a function of a selected number of collective variables. We propose here a change in perspective that shifts the focus from the bias to the probability distribution reconstruction while retaining some of the key characteristics of metadynamics, such as flexible on-the-fly adjustments to the free-energy estimate. The result is an enhanced sampling method that presents a drastic improvement in convergence speed, especially when dealing with suboptimal and/or multidimensional sets of collective variables. The method is especially robust and easy to use and in fact requires only a few simple parameters to be set, and it has a straightforward reweighting scheme to recover the statistics of the unbiased ensemble. Furthermore, it gives more control of the desired exploration of the phase space since the deposited bias is not allowed to grow indefinitely and it does not push the simulation to uninteresting high free-energy regions. We demonstrate the performance of the method in a number of representative examples.

Enhanced sampling plays a crucial role in modern simulation techniques and is a very active area of research [72]. Of particular historical importance has been the work of Torrie and Valleau [2]. They consider a system with an interaction potential $U(\mathbf{R})$, where \mathbf{R} denotes the atomic coordinates. Sampling is accelerated by adding a bias potential $V(\mathbf{s})$ that depends on \mathbf{R} via a set of collective variables (CVs), $\mathbf{s} = \mathbf{s}(\mathbf{R})$. The CVs are chosen so as to describe the modes of the system that are more difficult to sample. The choice of a proper set of CVs is critical because it determines the efficiency of the method. The properties of the unbiased system are then calculated by using a reweighting procedure. In fact, the unbiased probability density $P(\mathbf{s}) = \langle \delta[\mathbf{s} - \mathbf{s}(\mathbf{R})] \rangle \propto \int d\mathbf{R} e^{-\beta U(\mathbf{R})} \delta[\mathbf{s} - \mathbf{s}(\mathbf{R})]$ can be written as an average over the biased ensemble

$$P(\mathbf{s}) = \frac{\langle \delta[\mathbf{s} - \mathbf{s}(\mathbf{R})] e^{\beta V(\mathbf{s})} \rangle_V}{\langle e^{\beta V(\mathbf{s})} \rangle_V}, \quad (2.56)$$

where β is the inverse temperature. In this way, it is also possible to reconstruct the free-energy surface (FES), defined as $F(\mathbf{s}) = -\frac{1}{\beta} \log P(\mathbf{s})$.

Since the work of Torrie and Valleau, a large number of CV-based enhanced sampling methods have been proposed. Among them is metadynamics [3, 55] (MetaD), which builds the bias $V(\mathbf{s})$ by adding at fixed intervals repulsive Gaussians centered at the instantaneous point sampled. At the n th iteration, the bias is given by

$$V_n(\mathbf{s}) = \sum_k^n e^{-\beta V_{k-1}(\mathbf{s}_k)/(\gamma-1)} G(\mathbf{s}, \mathbf{s}_k), \quad (2.57)$$

where the parameter $\gamma > 1$ is called the bias factor, and the Gaussian function is defined as $G(\mathbf{s}, \mathbf{s}') = h \exp[-\frac{1}{2}(\mathbf{s} - \mathbf{s}')^T \boldsymbol{\Sigma}^{-1}(\mathbf{s} - \mathbf{s}')]$, with height h and variance $\boldsymbol{\Sigma}$ set by the user. Typically only diagonal variances $\Sigma_{ij} = \sigma_i^2 \delta_{ij}$ are employed, but more general choices have also been suggested [65]. It has been proven [54] that at convergence there is a simple relationship between the bias and the free energy, $V(\mathbf{s}) = -(1 - 1/\gamma)F(\mathbf{s})$, and the sampled ensemble is a smoothed version of the unbiased one, with FES barriers lowered by a factor of γ .

Arguably, a major development of MetaD has been its well-tempered variant [55]. With only a simple change to the original MetaD equations, it brought about many improvements, especially regarding the following points. (1) By damping the bias oscillations, it allows for better handling of suboptimal CVs, that is, CVs that do not include all of the slow modes of the system. This is a crucial issue since finding a good CV for a complex system is nontrivial, and even a good CV is usually suboptimal [52]. (2) It opens up the possibility of performing reweighting, which is a fundamental aspect of any enhanced sampling method, since it allows the retrieval of unbiased statistics of any quantity of interest. (3) It gives more control over the regions explored since the bias does not push the system to extremely high free-energy regions. (4) Thanks to this property, it also improves the handling of multiple CVs by reducing the volume of CV space that is sampled at convergence.

Despite the success of MetaD, there is certainly room for further improvement. In fact, over the years many new MetaD variants have been proposed, which put particular emphasis on one of the above-mentioned issues [41, 50, 60, 65, 73, 74]. Particular attention has been paid to reweighting, and many different solutions have also been proposed in recent years [58, 65, 75–79]. With this letter, we want to take a step back and propose a new perspective on MetaD in order to provide a general improvement to all of these issues, as has been the case for well-tempered MetaD.

We start from the observation that in the case of suboptimal CVs, the FES estimate obtained via reweighting can converge faster than the bias itself

[80]. In particular, it is more robust and does not present the strong oscillations typical of MetaD. Furthermore, a more static bias can help with the reweighting procedure itself, giving rise to a positive feedback loop. Thus, we develop a method that is based on the reconstruction of the probability distribution via reweighting and uses this estimate to define the bias potential rather than directly building it as in Eq. (2.57).

Enhanced sampling based on the probability reconstruction is not a new idea. It was first proposed in the adaptive umbrella sampling method [81] and later by many others [82–84]. Typically, in such methods the bias at n th iteration is defined as

$$V_n(\mathbf{s}) = \frac{1}{\beta} \log \hat{P}_n(\mathbf{s}), \quad (2.58)$$

where $\hat{P}_n(\mathbf{s})$ is an estimate of the probability obtained via a weighted histogram or some more elaborate method [82] and is updated iteratively or on the fly [83]. In building our method, we will introduce few key differences that come from the long experience with MetaD, which allow us to overcome some of the limitations of previous probability-based methods.

First, we explicitly introduce a target distribution $p^{\text{tg}}(\mathbf{s})$ that will be sampled once the method reaches convergence. This can be obtained with the following bias:

$$V(\mathbf{s}) = \frac{1}{\beta} \log \frac{P(\mathbf{s})}{p^{\text{tg}}(\mathbf{s})}. \quad (2.59)$$

In adaptive umbrella sampling, the target distribution is uniform, $p^{\text{tg}}(\mathbf{s}) \propto 1$, while in MetaD, it is the well-tempered distribution, $p^{\text{tg}}(\mathbf{s}) \propto [P(\mathbf{s})]^{1/\gamma}$. It is possible to modify MetaD in order to reach any arbitrary target [74], and in general the concept of a target distribution has proven to be very useful, especially in the context of variationally enhanced sampling [5, 21, 24, 56, 85, 86]. In the present work, we will limit ourselves to a well-tempered target (or a flat target in the $\gamma \rightarrow \infty$ limit), leaving other interesting possibilities for future work. We notice here that a well-tempered target leads to more efficient importance sampling compared to the common choice of a flat target, and despite lowering the FES barriers by γ instead of flattening them, it generally does not give rise to a slower transition rate between the metastable states. In fact, in most applications suboptimal CVs are employed and the transition rate is limited by the slow modes not accelerated by $V(\mathbf{s})$ rather than by the small FES barriers left along \mathbf{s} (Supporting Information - SI, Sec. 2.3.1) [80].

Since we can express the target distribution as a function of the unbiased one, $p^{\text{tg}}(\mathbf{s}) \propto [P(\mathbf{s})]^{1/\gamma}$, we only need to estimate $P(\mathbf{s})$ via reweighting in order to calculate the bias. We build our probability distribution estimate on the fly by periodically depositing Gaussians, similarly to how MetaD builds the bias potential. This is indeed a common way of reconstructing a

probability, known as kernel density estimation (KDE), and we shall draw from the vast literature on the subject [87]. Each new Gaussian is weighted according to the previously deposited bias potential

$$\tilde{P}_n(\mathbf{s}) = \frac{\sum_k^n w_k G(\mathbf{s}, \mathbf{s}_k)}{\sum_k^n w_k}, \quad (2.60)$$

where the weights w_k are given by $w_k = e^{\beta V_{k-1}(\mathbf{s}_k)}$.

We write the estimator in Eq. (2.60) with a tilde, $\tilde{P}_n(\mathbf{s})$, to indicate that it is not properly normalized, and we will take care of the normalization separately. $G(\mathbf{s}, \mathbf{s}_k)$ represents Gaussians such as those defined previously for MetaD, with diagonal variance $\Sigma_{ij} = \sigma_i^2 \delta_{ij}$ and fixed height $h = \prod_i (\sigma_i \sqrt{2\pi})^{-1}$. Contrary to MetaD, here the height of the Gaussians is not a free parameter, and changing it simply corresponds to changing the overall normalization.

It has been shown [87] that in KDE the most relevant parameter is the bandwidth (i.e. the width of the Gaussians). A good choice of the bandwidth should depend on the amount of available data: the larger the sampling, the smaller the bandwidth. Thus, we choose to shrink the bandwidth as the simulation proceeds according to the popular Silverman's rule of thumb [87]. At the n th iteration

$$\sigma_i^{(n)} = \sigma_i^{(0)} [N_{\text{eff}}^{(n)} (d+2)/4]^{-1/(d+4)}, \quad (2.61)$$

where $\sigma_i^{(0)}$ is the initial standard deviation estimated from a short unbiased simulation, d is the dimensionality of the CV space, and $N_{\text{eff}}^{(n)} = (\sum_k^n w_k)^2 / \sum_k^n w_k^2$ is the effective sample size [88]. The KDE literature presents many other promising alternatives for the bandwidth selection, but we leave their study to future investigation.

The number of kernels accumulated during the simulation quickly becomes very large, and summing all of them at each time step is prohibitive. To avoid this problem, we adapt to our needs a simple on-the-fly kernel compression algorithm [89] that allows the insertion of new kernels only in newly explored regions and otherwise merges them with existing ones. In the Supporting Information (Sec. 2.3.1), we discuss this choice in detail, and we show the advantages over the more common approach of storing the bias on a grid [38].

In Figure 2.17 we show how the FES estimate evolves during a typical simulation with our new method. Our choice of the probability estimator aims at quickly obtaining a coarse representation of the FES and then slowly converging the finer details, and it is one of the key novelties of our method.

We can now discuss the normalization problem. Any constant overall normalization of the probability estimate $\tilde{P}_n(\mathbf{s})$ would simply result in a global




Figure 2.17: Time evolution of a typical OPES simulation of alanine dipeptide in vacuum, using as CVs the dihedral angles φ and ψ . In the top row, the compressed kernels forming $\tilde{P}_n(\varphi, \psi)$ are shown, with the point size indicating the bandwidth, while in the bottom one is the corresponding free-energy estimate $F_n(\varphi, \psi) = -\frac{1}{\beta} \log \tilde{P}_n(\varphi, \psi)$, shifted to have a zero minimum. See the SI (Sec. 2.3.1) for the computational details and a performance comparison with MetaD.

shift of the bias and thus would not have any influence over the simulation. However, $\tilde{P}_n(\mathbf{s})$ should be normalized not with respect to the full CV space but only over the CV space actually explored up to step n , which we call Ω_n . Thus, we introduce the normalization factor

$$Z_n = \frac{1}{|\Omega_n|} \int_{\Omega_n} \tilde{P}_n(\mathbf{s}) d\mathbf{s}, \quad (2.62)$$

that will change over time as the system explores new regions of the CV space, and it will have an impact on the biasing scheme. This impact becomes particularly relevant in CV spaces of dimension $d \gg 1$ since the volume explored $|\Omega_n|$ grows with a power of d . Adding such a normalization, together with the chosen probability estimator, helps us overcome the limitations in exploration speed that have affected some previously proposed on-the-fly probability-based methods [90]. To estimate Z_n , we take advantage of our compressed kernels representation and substitute the integral in Eq. (2.62) with a sum over the positions of the compressed kernels (SI, Sec. 2.3.1).

Finally, we can explicitly write the bias at the n th step as

$$V_n(\mathbf{s}) = (1 - 1/\gamma) \frac{1}{\beta} \log \left(\frac{\tilde{P}_n(\mathbf{s})}{Z_n} + \varepsilon \right), \quad (2.63)$$

where $\varepsilon \ll 1$ can be seen as a regularization term that ensures that the argument of the logarithm is always greater than zero. We notice that the addition of this term is not merely a technicality to solve a numerical issue but rather it allows one to set a limit on the bias, thus providing better control over the desired exploration. It can be chosen to be $\varepsilon = e^{-\beta\Delta E/(1-1/\gamma)}$, where ΔE is the height of the free-energy barriers one wishes to overcome during the enhanced sampling (SI, Sec. 2.3.1). We could have obtained the same effect of controlling the exploration by properly modifying the target distribution $p^{\text{tg}}(\mathbf{s})$, but we believe that introducing ε as a separate term makes for cleaner equations. By comparing Eqs. (2.63) and (2.58) it should be clear that our method distinguishes itself from previous adaptive umbrella sampling methods not only for the employed probability estimator but also for some other novel key components.

An important feature of our method is that it allows for a simple and straightforward reweighting scheme. In fact, reweighting can be performed in the usual umbrella sampling way [Eq. (2.56)] without the need for further postprocessing analysis. The method has a rapid initial exploration phase, after which a quasi-static regime is reached, but it is by construction robust with respect to the initial nonadiabatic part of the trajectory so that the reweighting can in practice be performed without cropping out the initial transient (SI, Sec. 2.3.1).

We implemented the new method, called on-the-fly probability enhanced sampling (OPES), in the enhanced sampling library PLUMED [38] and tested it on a variety of different systems. Here we provide only a quick overview of these tests, but the full results are presented in detail in the Supporting Information (Sec. 2.3.1). The code and all of the files needed to reproduce the simulations are openly available on the PLUMED-NEST website [91] as `plumID:19.068`.

A full comparison of different enhanced sampling methods is a nontrivial task and is not the goal of the present letter. However, in order to give a better idea of our method, we compare it with standard well-tempered metadynamics whose performances might already be familiar to many readers.

We want to test the methods in an agnostic fashion, using very standard input parameters rather than running multiple different simulations and choosing the best-performing ones. One strength of OPES is that it is very simple to set up and needs just three main parameters: the pace at which the bias is updated, the initial bandwidth of the Gaussian kernels, and the




Figure 2.18: Potential energy of the suboptimal double-well 2D model system and its free energy along the x coordinate.




Figure 2.19: Typical time evolution of the free-energy difference between the two basins of the model in Figure 2.18. We run the same simulation with 100 different initial conditions, and in the SI (Sec. 2.3.1) we show the average and uncertainty obtained from these estimates. The reference blue stripe is $1k_B T$ thick.

approximate height of the barriers one wishes to cross. In our tests, we always keep the deposition pace equal to the one used in MetaD, typically 500 simulation steps. The initial bandwidth is simply chosen to be equal to the smaller standard deviation of the CVs in the minima, which can be measured in a short unbiased run. The choice of the barrier parameter requires a minimal knowledge of the system, but only a vague idea is usually enough. This barrier parameter is used to set a reasonable default of both the regularization factor ε and the bias factor γ (SI, Sec. 2.3.1). It should be noticed that the choice of γ is not as critical as in MetaD since here it does not directly influence the convergence speed but only the shape of the target distribution. In fact, in OPES the limit $\gamma \rightarrow \infty$ is not problematic, as in MetaD [92], and OPES can also converge to the flat target distribution. In our tests, we always used the same value of γ for both OPES and MetaD.

In order to test the convergence speed in the case of suboptimal CVs, we consider Langevin dynamics on a 2D model potential [80], Figure 2.18, and

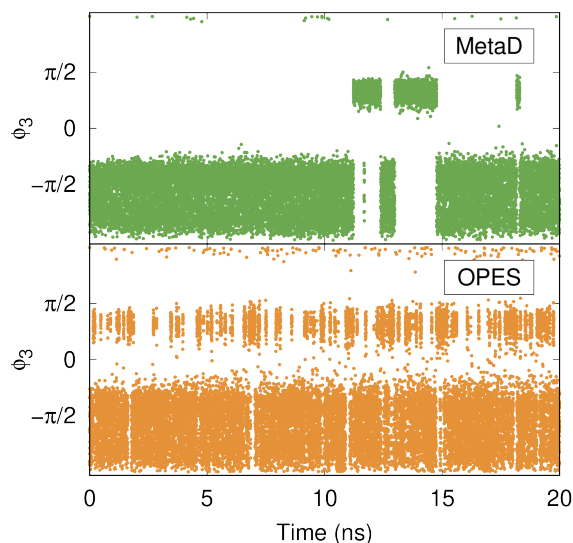


Figure 2.20: Φ_3 trajectory for alanine tetrapeptide in vacuum, obtained by biasing all the six dihedral angles ($\varphi_1, \varphi_2, \varphi_3, \psi_1, \psi_2, \text{ and } \psi_3$) using MetaD and OPES, with the same input parameters as the standard ones used for alanine dipeptide. We show the φ_3 angle because it is the hardest one to sample, but in the SI (Sec. 2.3.1), all of them are presented, together with the reconstructed FES.

bias only the x coordinate. In Figure 2.19, we compare MetaD and OPES by plotting the estimate of the free-energy difference between the two basins as a function of time. Such estimates are obtained directly from the applied bias, as $F_n(x) = V_n(x)/(1/\gamma - 1)$. From Figure 2.19, one can see that as OPES converges it does not present the strong bias oscillations typical of MetaD.

We also run the typical benchmark system for novel enhanced sampling methods, alanine dipeptide, biasing dihedral angles φ and ψ . The results are in Figure 2.17 and the SI (Sec. 2.3.1).

As an example of multidimensional bias, we run simulations of alanine tetrapeptide in vacuum, biasing all six dihedral angles. In Figure 2.20, we show how OPES is able to explore this high-dimensional CV space much more efficiently than MetaD. It is important to notice that we use the same input parameters for alanine dipeptide and alanine tetrapeptide. This should be a reasonable choice since the two systems are very similar from the point of view of the physics involved, the main difference being the increased dimensionality of the CV space. However, many CV-based enhanced sampling methods would require further tuning or a different set of inputs in order to perform well in both systems. This is not the case for OPES, thanks to its robustness to the choice of the input parameters.

In conclusion, within this letter we present a new enhanced sampling meth-

od, OPES, based on an on-the-fly reconstruction of the probability distribution. It performs such reconstruction via a weighted kernel density estimation [Eq. (2.60)] with an on-the-fly compression algorithm that allows it to be most effective, starting from a coarse-grained guess of the free-energy surface and then converging the finer details. Thanks to this strategy and to the introduction of a normalization over the explored CV space [Eq. (2.62)], the method provides an extremely fast exploration, also at relatively high dimensions. Another peculiarity of the method is the presence of an upper limit to the applied bias [Eq. (2.63)], which can be useful for avoiding the sampling of unphysical states. Most importantly, the proposed method requires few simple and robust input parameters, has very good convergence performance, and presents a straightforward reweighting scheme. Finally, we considered here only the case of a well-tempered target, but OPES provides a general framework in which different targets can also be considered.

We believe that this new method can become a handy tool in addressing enhanced sampling problems and has the potential for further interesting developments.

Acknowledgements

The authors thank Sandro Bottaro and Barak Hirshberg for useful feedback on the manuscript. M.I. thanks Dario Azzimonti and Manuel Schuerch for drawing his attention to kernel compression algorithms. This research was supported by the NCCR MARVEL, funded by the Swiss National Science Foundation, and European Union grant no. ERC-2014-AdG-670227/VARMET. Calculations were carried out on the Euler cluster at ETH Zurich and are openly available in the Materials Cloud Archive (www.materialscloud.org) with ID `materialscloud:2019.0063`.

2.3.1 Supporting Information

Algorithmic details

Kernel density compression As discussed in the main text, it is impractical to describe the probabilities as a sums of all the Gaussians deposited. In metadynamics (MetaD) and in adaptive umbrella sampling this difficulty is usually circumvented by mapping the bias on an auxiliary grid. Here we shall reduce the number of Gaussians needed to represent the probability by using a kernel compression algorithm borrowed from signal processing literature [89]. This algorithm has been previously used only for post-processing estimation, and never with a shrinking bandwidth, but it turned out to be especially convenient in our setup, thanks to the recursive nature of our method.

We restrict ourselves to the consideration of Gaussian kernels with diagonal bandwidths, the generalisation to more complex kernels being straightforward. Thus, as in the main text, we write:

$$G(\mathbf{s}, \mathbf{s}') = h e^{-\frac{1}{2} \sum_i \left(\frac{s_i - s'_i}{\sigma_i} \right)^2}. \quad (2.64)$$

To measure of the distance between a point \mathbf{s}' and a kernel G we consider the Mahalanobis distance, that in our case reads:

$$d(\mathbf{s}', G) = \sqrt{\sum_i \left(\frac{s_i - s'_i}{\sigma_i} \right)^2}. \quad (2.65)$$

Let us suppose that we have to deposit the n th kernel $G(\mathbf{s}, \mathbf{s}^{(n)})$ at position $\mathbf{s}^{(n)}$, after the ones previously deposited have been compressed to N kernels G_k , where in general as n grows we expect $N \ll n$. The algorithm then goes as follows:

1. find the minimum distance between the new point and the compressed kernels, $d_{\min} = \min_k d(\mathbf{s}^{(n)}, G_k)$
2. check if this distance is smaller than a given threshold, d_t :
 - a. if $d_{\min} < d_t$ merge the new kernel $G(\mathbf{s}, \mathbf{s}^{(n)})$ with the closest one $G_{k_{\min}}$. Then go back to step 1, using instead of $\mathbf{s}^{(n)}$ the new center of the merged kernel
 - b. if $d_{\min} \geq d_t$ add the new kernel $G(\mathbf{s}, \mathbf{s}^{(n)})$ to the compressed ones

When merging two Gaussians G_1 and G_2 the new Gaussian $G(\mathbf{s}, \mathbf{s}')$ will have the following parameters:

$$h = h_1 + h_2 \quad (2.66)$$

$$\mathbf{s}' = h^{-1}(h_1 \mathbf{s}_1 + h_2 \mathbf{s}_2) \quad (2.67)$$

$$\sigma'^2 = h^{-1}[h_1(\sigma_1^2 + \mathbf{s}_1^2) + h_2(\sigma_2^2 + \mathbf{s}_2^2)] - \mathbf{s}'^2 \quad (2.68)$$

This simple merging rule applies only to Gaussians with diagonal bandwidth, however is possible to extend it also to the general case [89].

We found that $d_t = 1$ is a good default value. However, sometimes it might be useful to adopt a higher threshold, in order to reduce the total number of kernels, at the cost of obtaining a coarser estimate of the probability density.

This compression algorithm differs from that of Ref. [89] mostly for the addition of the recursive rule at the end of 2.a. We found this step very useful especially when the CVs dimensionality is greater than one, because without it the minimum distance between the deposited kernels quickly becomes

smaller than d_t , and the total number of kernels may grow much faster. This extra step does add an overhead to the algorithm, so in our implementation we also provide the option not to perform it.

The total number of compressed kernels will always grow in time, but it soon reaches a sort of plateau, with minimal growth. It is in general not easy to guess a priori how many kernels will be needed, since it depends not only on the dimensionality of the CV space, but also on the features of the free-energy surface and on the range of exploration. However, in our experience the total number of kernels is always reasonable, and generally much smaller than the number of points in a typical grid of the same dimensionality.

The overall computational cost of the biasing scheme scales roughly linearly with the total number of compressed kernels, since at each time step we need to sum over all them in order to estimate the bias and the forces due to the bias. A way of greatly reduce the computational cost would be to implement a neighbor list scheme, but at this point this has been left to future development.

This bias representation has some advantages when compared to the grid representation more commonly used, in particular:

- the user does not need to guess in advance the CV region that will be explored, kernels will be deposited only where necessary
- it can handle a higher dimensional CV space, while grids typically already struggle at three/four dimensions
- it allows changing the scale used for the description of the bias, a feature that we found very useful

Bandwidth rescaling The kernels bandwidth is rescaled according to Eq. (2.61) of the main text:

$$\sigma_i^{(n)} = \sigma_i^{(0)} [N_{\text{eff}}^{(n)} (d + 2) / 4]^{-1/(d+4)}, \quad (2.69)$$

where the effective sample size at n th step is:

$$N_{\text{eff}}^{(n)} = \frac{(\sum_k^n w_k)^2}{\sum_k^n w_k^2}. \quad (2.70)$$

We find useful to define $w_0 = 1$, so that at the beginning of the simulation $N_{\text{eff}}^{(0)} = 1$. This choice makes the estimate of the effective sample size more robust, especially at the very first steps of the simulation.

We notice that the fact that the bandwidth is rescaled, requires that also the height of the deposited Gaussian is adjusted accordingly, since

$$h = \prod_i \frac{1}{\sigma_i \sqrt{2\pi}}. \quad (2.71)$$

Normalization factor According to its definition, in Eq. (2.60) of the main text, the probability estimator $\tilde{P}(\mathbf{s})$ is normalized to 1 over a boundless CV space. By adding the normalization factor Z_n , defined as in Eq. (2.62) of the main text, we take into account for the actually explored CV space, Ω_n . To calculate Z_n we take advantage of our compressed kernel representation, and consider the centers of the kernels as points for a Monte Carlo integration:

$$Z_n = \frac{1}{N} \sum_k \tilde{P}(\mathbf{s}_k) = \frac{1}{NS} \sum_{k,k'} G(\mathbf{s}_k, \mathbf{s}_{k'}), \quad (2.72)$$

where $G(\mathbf{s}_k, \mathbf{s}_{k'})$ are the compressed Gaussians, N is their total number, and $S = \sum_k w_k$ is the global normalization of the KDE. Z_n changes accordingly to the explored space Ω_n , and typically converges as soon as no new CV space region is explored. The above normalization factor also helps correcting the possible changes in the global normalization S due to the compression algorithm.

The estimate in Eq. (2.72) scales quadratically with the number of compressed kernels N . In our implementation however, we avoid this by calculating only the changes with respect to the previous value, $Z_n = Z_{n-1} + \Delta Z_n$, which instead scales linearly in N .

The barrier parameter In the current implementation of the method, there is a parameter called barrier, that should be roughly equal to the free-energy barrier ΔE that the bias should help overcome. We use this value to set the bias factor $\gamma = \beta\Delta E$ and the regularization parameter $\varepsilon = e^{-\beta\Delta E/(1-1/\gamma)}$.

According to Eq. (2.63) of the main text, the minimum value that the bias can assume is reached when $\tilde{P}(\mathbf{s}) = 0$, and with the given choice of ε we have:

$$\min_{\mathbf{s}} V_n(\mathbf{s}) = -\Delta E \quad (2.73)$$

The maximum value of the bias is instead harder to obtain, because it changes according to the normalization factor Z_n . However it is typically only few $k_B T$, so that in practice the maximum deposited bias, ΔV_{\max} , is not too much higher than ΔE .

It is in principle possible to set a hard limit to ΔV_{\max} , e.g. by adaptively adjusting the ε value, but we prefer to have some tolerance with respect

to the provided estimate of the barrier ΔE , since it is generally not easy to guess it correctly before running simulations.

We chose to link the bias factor γ to the value of the barrier ΔE because γ is not a crucial parameter for the convergence speed, as it is the case for MetaD. However it is also possible to set the value of γ independently, in order to fine tune the shape of the target distribution. This can be useful e.g. in case a high dimensional CV space, where one typically wants to restrict as much as possible the explored space to speed up sampling.

Results on model systems

We implemented our method in a development version of the PLUMED [38] plugin, and we plan to add it to the official PLUMED code. All the code and the input files needed to reproduce the following results are openly available in the Materials Cloud Archive (www.materialscloud.org), as `materialscloud:2019.0063`, and on the PLUMED-NEST [91] (www.plumed-nest.org), as `plumID:19.068`.

Suboptimal double well We first test our new method on the simple toy model, first presented in Ref. [80], that consists in a particle moving on a 2D potential $U(x, y)$ with Langevin dynamics, see Fig. 2.18 of main text. The potential is a rotated Wolfe-Quapp ($\theta = -0.6\frac{\pi}{4}$), the temperature of the Langevin dynamics is $T = \beta = 1$ and the friction coefficient is 10.

We want to model a double well system with a suboptimal CV, thus we bias only x . In this simple system we can directly calculate the free energy profile along x by numerical integration on a grid:

$$F(x) = -\frac{1}{\beta} \log \int_{-3}^3 e^{-\beta U(x,y)} dy, \quad (2.74)$$

and also the free-energy difference ΔF , that we define as follows:

$$\Delta F = -\frac{1}{\beta} \log \frac{\int_0^3 e^{-\beta F(x)} dx}{\int_{-3}^0 e^{-\beta F(x)} dx}. \quad (2.75)$$

It is important to notice that the free energy profile $F(x)$ can change significantly with the choice of the CV, while the free-energy difference ΔF is not much influenced by the CV choice, provided that the chosen CV is able to distinguish between the two basins. To illustrate this, in Fig. 2.21 we consider other possible rotations of the potential $U(x, y)$, and we plot the free energy along the new axis x' , together with the barrier height. As expected, for all these different rotations the ΔF estimate obtained via Eq. (2.75) is constant, changing by less than 0.03%.




Figure 2.21: The free energy of the double well model, for different rotations of the $x - y$ plane (see Fig. 2.18 of main text). The profile along x' changes, and so does the highest free-energy barrier between the two states, while the ΔF is almost identical for all the shown rotations.

We run simulations with both OPES and MetaD, using input parameters as similar as possible. In particular the bias factor $\gamma = 10$, the deposition stride is 500 time steps and the initial bandwidth is $\sigma = 0.185815$ (equal to the unbiased standard deviation in the minimum). Using the reference from Eq. (2.74) we also run simulations with a static bias $V(x) = -(1 - 1/\gamma)F(x)$, that is the bias that both OPES and MetaD reach at convergence. The MetaD simulations use Gaussians with initial height $h = 1$ and deposit bias only in the interval $[-2.9, 2.9]$, to avoid pushing the system out of the $[-3, 3]$ range in which the potential is defined. This last precaution is not needed for the OPES and the static bias simulations.

In order to compare the different methods we select 100 initial conditions from an unbiased run in which the system remained in the deeper basin. These configurations are used to start 100 independent runs for each of the methods investigated. The free-energy difference ΔF is calculated from the free energy estimator given by $F_n(x) = -(1 - 1/\gamma)^{-1}V_n(x)$, except for the static bias case where reweighting [Eq. (2.56) of main text] is used. Fig. 2.22 shows the average ΔF obtained, together with the standard deviation, as a function of time.

It is interesting to notice that OPES yields the best estimate, outperforming also the reference run with a perfect static bias. This is due to the fact that all the 100 replicas start from the same basin (thus $\Delta F_n = 0$ at the beginning), and a dynamic bias allows for a faster first transition than a static one, thus more quickly correcting this wrong initial estimate.

The number of compressed kernels that build the probability estimate is always smaller than 80, for all of the OPES simulations.

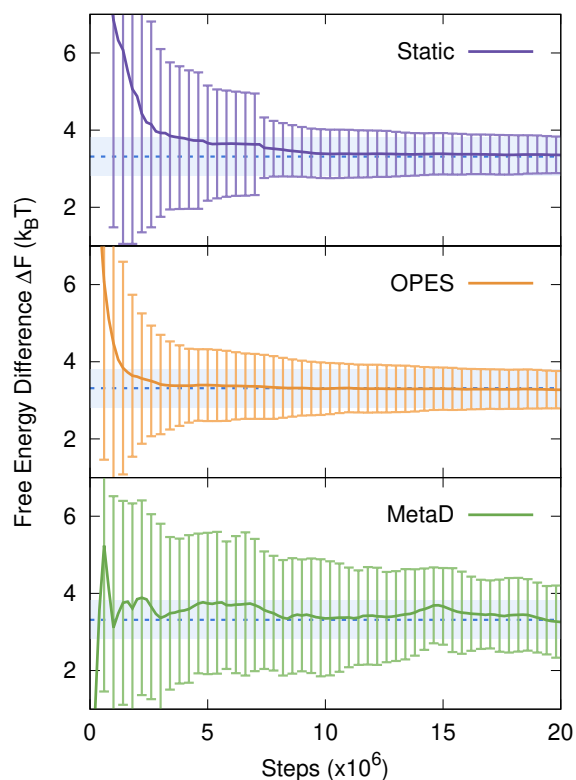


Figure 2.22: The free-energy difference of the double well model, calculated with different methods. For the static bias case the reweight estimate is shown, while for the others the estimate obtained from the applied bias. Average and standard deviation over 100 independent runs are shown.

Reweighting. It is also possible to obtain an estimate of the free energy and ΔF via reweighting, instead of using the one obtained from the bias potential. To perform reweighting we follow Eq. (2.56) of main text, using as probability estimator a weighted KDE. In order to reweight a nonstatic bias one usually must discard the initial nonadiabatic part of the trajectory, where the bias changes too fast. Unfortunately it is not always clear how much one should cut out. For this reason we find it useful to invert the trajectory and plot the reweighted estimate as a function of time, as in Fig. 2.23. The plot shows all the possible choices for initial truncation, with on the abscissa axis the amount of steps truncated. At the far left is the estimate obtained without discarding any point, at half of the plot is the reweighting obtained only from the second half of the trajectory and so on. Thus the final part of the trajectory is always used, which makes sense, since it is the most close to convergence.

There are different ways of reweighting MetaD. Since we do know the real $F(x)$ of our toy model, we can use the real $c(t)$, instead of one of its estimates

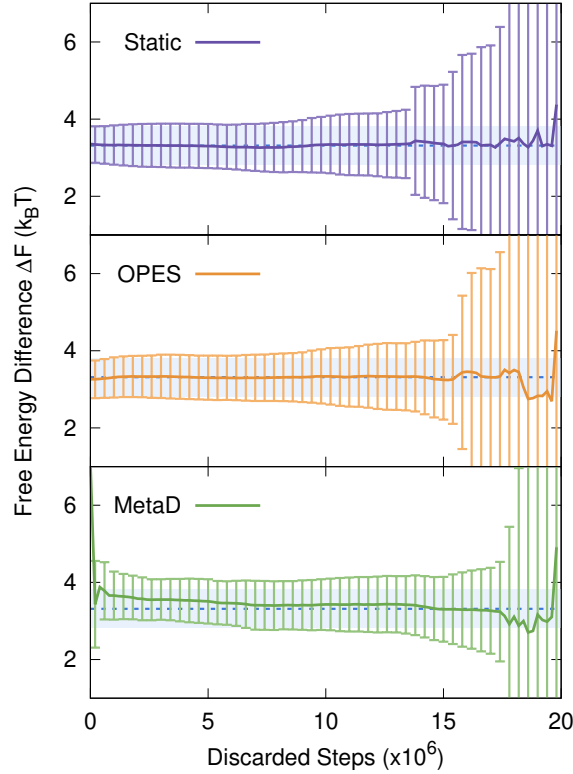


Figure 2.23: The ΔF of the double well model from the same simulations used for Fig. 2.22, but obtained through reweighting. The trajectory has been time-inverted before reweighting, in order to show all the possible initial truncations (see explanation in the text). The MetaD reweighting is obtained using a perfect $c(t)$, obtained using the real $F(x)$.

[58, 79]. At step n the “real” $c(t)$ is obtained via a grid integration:

$$c_n = -\frac{1}{\beta} \log \frac{\int_{-3}^3 e^{-\beta[F(x)+V_n(x)]} dx}{\int_{-3}^3 e^{-\beta F(x)} dx}. \quad (2.76)$$

This provides a better reweighting, because it removes the noise from the $c(t)$ estimation.

By looking at Fig. 2.23 we can see that, as expected, in the static bias case there is no need to cut any initial transient. In OPES the reweighting estimate is not significantly influenced by keeping the initial part of the simulation, even though it is nonadiabatic, and it gives an estimated ΔF very similar to the one actually used during the biasing. In MetaD instead the reweighting estimate has a smaller standard deviation compared to the direct estimate, but one should discard a significant part of the trajectory in order to obtain a correct estimate.




Figure 2.24: The free-energy difference for alanine dipeptide, calculated with different methods. The estimate is obtained from the applied bias, average and standard deviation over 10 independent simulations are shown.

Sampling efficiency: well-tempered vs uniform. Using this toy model we can also quantify the gain in sampling efficiency provided by using a well-tempered target instead of a uniform one. To do so, we look at the effective sample size $N_{\text{eff}} = (\sum_k w_k)^2 / \sum_k w_k^2$ which provides an estimate of the efficiency of the importance sampling we are doing. We run as before with 100 different initial conditions, using a well-tempered static bias $V(x) = -(1 - 1/\gamma)F(x)$ with ($\gamma = 10$) and a uniform one $V(x) = -F(x)$. For the well-tempered case we obtain $N_{\text{eff}}/N = 0.0236 \pm 0.042$ while for the uniform $N_{\text{eff}}/N = 0.0117 \pm 0.022$. This means that given the same number of steps N , using a well-tempered target in this case provides roughly twice the effective sample size compared to a uniform target, and thus any quantity estimated from it will have a smaller statistical uncertainty.

Alanine dipeptide For the alanine dipeptide simulations we use GROMACS [37] patched with PLUMED. The setup is the same of Ref. [80], namely: canonical (NVT) simulation in a vacuum, Amber99-SB [66] force field, time step 0.002 fs, temperature 300 K, and velocity rescaling thermostat [36].

We run OPES and MetaD, using the same bias factor $\gamma = 10$ and deposition stride of 500 time steps. For MetaD we used the typical standard parameters for alanine dipeptide, namely $\sigma_\varphi = \sigma_\psi = 0.35$ rad, and Gaussian height $h = 1.2$ kJ/mol. For the OPES simulations we choose an initial bandwidth $\sigma_\varphi = \sigma_\psi = 0.15$ rad, which is roughly the smaller of the standard deviations one gets by running a short unbiased simulation in the two basins, and a




Figure 2.25: The typical number of compressed kernels used in the alanine dipeptide simulations. The black dotted line shows the total number of deposited kernels. After 100 ns with compression threshold $d_t = 1$ about 68% of the deposited kernels have been merged, while with $d_t = 3$ about 99%. In this $d_t = 3$ case the bias is very coarse, but still allows for good sampling and the reweighting estimate of ΔF falls within $1 k_B T$.

barrier parameter $\Delta E = 50$ kJ/mol. Such ΔE value would lead to a default bias factor $\gamma \approx 20$, but we set it instead to $\gamma = 10$, in order to facilitate the comparison with MetaD. We run 10 independent simulations for each method, starting from initial configurations taken from a long unbiased run in the most stable basin (basin A). Fig. 2.24 shows the average and the standard deviation of the free-energy difference ΔF , for each method. We define the free-energy difference in units of $k_B T$ as:

$$\Delta F = -\log \frac{\int_A e^{-\beta F(\varphi, \psi)} d\varphi d\psi}{\int_B e^{-\beta F(\varphi, \psi)} d\varphi d\psi}, \quad (2.77)$$

where $A = \{\varphi \in [0, 2.3], \psi \in [-\pi, \pi]\}$ and B is the complementary region in the CVs space.

The number of kernels used in the OPES simulations for representing the probability distribution, and thus the bias, is shown in Fig. 2.25. In this case we also test some other values of the compression threshold d_t , but all the results shown in other figures and in the main text are obtained with $d_t = 1$. For the MetaD simulations of alanine dipeptide we use a 100x100 grid with spline interpolation.

Alanine tetrapeptide With the same computational setup of alanine dipeptide, we simulate alanine tetrapeptide, which has 3 φ and 3 ψ angles, see Fig. 2.26. Using the 3 φ angles as CVs would be enough to get a good sampling, but in order to test the performance in higher dimensions we run simulations with 6 CVs, using all the angles.




Figure 2.26: The alanine dipeptide and alanine tetrapeptide molecules, with their Ramachandran angles.

Fig. 2.27 shows the trajectories of the 6 dihedral angles for OPES and MetaD. It can be seen how OPES provides an extremely fast exploration also in such high dimensional case. In principle one could decide to fix the OPES bias after a few nanoseconds, and use it as static bias to perform umbrella sampling, in the spirit of Refs. [56, 93].

We notice that the performance of MetaD can be improved by tweaking the parameters, e.g. using a bigger initial height for the Gaussians, but we could not find any combination of parameters giving an exploration qualitatively similar to OPES. Using parallel bias metadynamics [50] (PBMetaD) brings much better performances in such high CV space. From the point of view of the exploration however it was still less efficient than OPES, requiring a few nanoseconds to find all the basins (but we only tested with the standard parameters, without any tweaking). It would be interesting to implement in OPES the same target distribution that PBMetaD reaches at convergence, because it covers a much smaller CV volume than well-tempered, while still sampling all the relevant basins. Having a smaller CV volume to sample can help a lot in converging, especially when a big number of CVs is employed.

In Figure 2.28 a free-energy estimate obtained from reweighting the previous shown OPES simulation is shown, compared with a reference one obtained from a much longer simulation performed with only the 3 φ angles. We chose to show the free energy projected onto the φ_2 - φ_3 space, because these are the two angles harder to sample. It is remarkable that after only 1 nanosecond (1000 kernels deposited) all the metastable basins have been sampled, and the free energy already shows all the relevant features.




Figure 2.27: The trajectory of the 6 dihedral angles of alanine tetrapeptide, obtained by biasing all of them with OPES and MetaD respectively.




Figure 2.28: Free energy for alanine tetrapeptide, obtained by reweighting over the φ_2, φ_3 plane the OPES simulation shown in Fig. 2.27. The reference free energy comes from a 100ns OPES simulation performed using only the three φ angles as CVs.

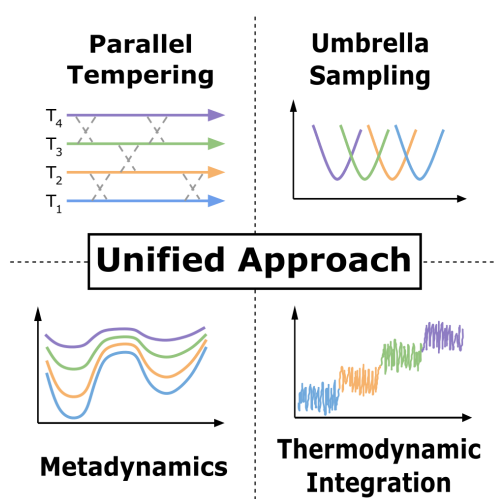
2.4 Unified approach to enhanced sampling

This final article shows how the OPES method can be used for sampling expanded ensembles that are the type of ensembles sampled e.g. by the replica-exchange method. To this end, we introduce the concept of expansion collective variables, that are used to uniquely define a nonweighted expanded target distribution, together with the needed free-energy differences and the target bias. Remarkably, when an expanded ensemble is chosen as target, the kernel density estimation used in the previous paper is no longer needed and thus the number of external parameters OPES needs is further reduced, making the method even more robust. In the proposed scheme, combining tempering-based and CV-based enhanced sampling is straightforward and not limited to the expanded ensembles considered in the paper. With this article we bring a unified perspective on enhanced sampling that allows for a conceptual and practical simplification, and also paves the way for new types of sampling.

My contribution to this article has been implementing the algorithms, performing the simulations, and writing the paper jointly with Piaggi and Prof. Parrinello.

Reference: M. Invernizzi, P. M. Piaggi, and M. Parrinello. “Unified approach to enhanced sampling.” *Physical Review X* 10.4 (2020): 041034. URL <https://journals.aps.org/prx/abstract/10.1103/PhysRevX.10.041034>

Copyright © 2020 American Physical Society.



Unified Approach to Enhanced Sampling

Michele Invernizzi*^{1,2}, Pablo M. Piaggi³, and Michele Parrinello*^{4,2}

¹*Department of Physics, ETH Zurich c/o Università della Svizzera italiana, 6900 Lugano, Switzerland*

²*Facoltà di Informatica, Institute of Computational Science, and National Center for Computational Design and Discovery of Novel Materials (MARVEL), Università della Svizzera italiana, 6900 Lugano, Switzerland*

³*Department of Chemistry, Princeton University, Princeton, New Jersey 08544 USA*

⁴*Department of Chemistry and Applied Biosciences, ETH Zurich c/o Università della Svizzera italiana, 6900 Lugano, Switzerland
and Italian Institute of Technology, 16163 Genova, Italy*

Abstract

The sampling problem lies at the heart of atomistic simulations and over the years many different enhanced sampling methods have been suggested toward its solution. These methods are often grouped into two broad families. On the one hand, are methods such as umbrella sampling and metadynamics that build a bias potential based on few order parameters or collective variables. On the other hand, are tempering methods such as replica exchange that combine different thermodynamic ensembles in one single expanded ensemble. We instead adopt a unifying perspective, focusing on the target probability distribution sampled by the different methods. This allows us to introduce a new class of collective-variables-based bias potentials that can be used to sample any of the expanded ensembles normally sampled via replica exchange. We also provide a practical implementation by properly adapting the iterative scheme of the recently developed on-the-fly probability enhanced sampling method [M. Invernizzi and M. Parrinello, *J. Phys. Chem. Lett.* **11**, 2731 (2020)], which was originally introduced for metadynamicslike sampling. The resulting method is very general and can be used to achieve different types of enhanced sampling. It is also reliable and simple to use, since it presents only few and robust external parameters and has a straightforward reweighting scheme. Furthermore, it can be used with any number of parallel replicas. We show the versatility of our approach with applications to multicanonical and multithermal-multibaric simulations, thermodynamic integration, umbrella sampling, and combinations thereof.

2.4.1 Introduction

Sampling is one of the main challenges in atomistic simulations. In fact, even the most accurate models cannot produce high-quality results if the phase space is not properly sampled. The sampling issue is due to the large gap between the physical macroscopic timescales and the actual time that can be explored in standard atomistic simulations. This results in an ergodicity problem that can be encountered in fields as varied as materials

science, chemistry, and biology. One facet of this problem is the existence of different metastable states separated by kinetic bottlenecks, that make the transition from one state to another a rare event. Enhanced sampling methods are a possible solution to this problem. Instead of extracting configurations from the relevant physical ensemble, these methods create an *ad hoc* modified ensemble in which the probability of sampling rare events is greatly enhanced. One kind of such target ensembles is obtained by combining multiple subensembles that differ only for the temperature or some other quantity, a typical example being parallel tempering [94]. We refer to these ensembles as expanded ensembles [95].

In the present paper we formulate the problem of generating such expanded ensembles in a way that allows us to use collective-variables-based methods. We find that the recently developed on-the-fly probability enhanced sampling (OPES) [96] can be adapted to the scope and provides an efficient implementation. This method was introduced as an evolution of metadynamics [3], since it can provide the same type of enhanced sampling, but presents in most cases a faster convergence and has only few and robust adjustable parameters. These properties of OPES are retained when it is applied to sample expanded ensembles. This provides us with a general and reliable method, that can be easily applied to sample many different ensembles.

We accompany this paper with a number of general considerations (Secs. 2.4.2 and 2.4.7), but the reader mostly interested in the method itself and its practical implementation can go directly to Sec. 2.4.4. Section 2.4.3 briefly recalls OPES in its original formulation for metadynamicslike sampling. We also present a variety of simulations to show the versatility of the new scheme, in particular (Sec. 2.4.5) multicanonical, multithermal-multibaric, thermodynamic integration, and also (Sec. 2.4.6) enhanced sampling based on an order parameter, both alone and in combination with the previous ensembles.

2.4.2 Unified Approach

The most popular approaches to enhanced sampling follow mainly two strategies. A first one was proposed in a pioneering work by Torrie and Valleau and is referred to as umbrella sampling [1, 2]. This method starts by identifying one or few order parameters, or collective variables (CVs), $\mathbf{s} = \mathbf{s}(\mathbf{x})$, that are a function of the microscopic configuration and encode some of the slow modes of the system. Then a bias potential that is function of the CVs is added to the energy of the system, so that the sampling of the slow modes encoded by the CVs is accelerated. Many have followed this approach, and nowadays one of the most popular methods in this class is metadynamics [3].

A different perspective to enhanced sampling is that of parallel or simulated tempering [97, 98]. In this case the idea is to combine in the same generalized ensemble the configurations explored by the system at different temperatures. This can improve the sampling because at higher temperatures the exploration of the phase space is often more efficient, and the system is less likely to remain stuck in metastable states. Over the years this approach has been extended and implemented in a variety of different methods, among which replica exchange [6] is probably the most widely employed.

These two families of enhanced sampling methods often have been seen as distinct and complementary. Although there are some papers in which the two perspectives are combined [73, 99–101], typically they have been perceived as hybrid approaches [43, 102]. Here we want to take a closer look at these two families and show that it is possible to provide a unified perspective to the enhanced sampling problem.

For a start we must specify that we are not interested in looking at the specific computational technique the various enhanced sampling methods use, since according to this criterion there would be many more than two families. There are methods that use a bias potential and others that use specific Monte Carlo moves [6], but also methods that introduce a fictitious dynamics [103], or that focus on directly modifying the atomic forces [104], to name just a few. This kind of classification is of course perfectly legitimate, but we find it of limited relevance for our purposes.

The distinction between the two families cannot be based on the fact that one uses system-specific CVs, while the other makes use of general thermodynamic properties. For instance, it is known that Hamiltonian replica exchange can be used to enhance the fluctuations of any chosen CV [105], and on the other hand, that it is possible with metadynamics to use the potential energy itself as CV and sample a multithermal ensemble [106].

Thus we prefer to focus on the target distribution $p^{\text{tg}}(\mathbf{x})$ that the different methods sample. In fact, each enhanced sampling method explicitly or implicitly aims at sampling a specific probability distribution in the configuration space that is not the physical one, but assigns a higher probability to some rare event. Designing such target distributions so that they are effective is far from trivial, and we can relate the two families of methods to the type of target distribution they imply.

A first class of enhanced sampling methods defines the target distribution by setting a constraint on its marginal distribution along some chosen CVs, $p^{\text{tg}}(\mathbf{s}) = \int \delta(s(\mathbf{x}) - s) p^{\text{tg}}(\mathbf{x}) d\mathbf{x}$. The most common choice is to impose a uniform marginal, $p^{\text{tg}}(\mathbf{s}) = \text{const}$. Among the methods that adopt this strategy are adaptive umbrella sampling [81] and metadynamics in its original formulation [3]. Also, the Wang-Landau algorithm [107] and various multicanonical algorithms [108, 109] chose to sample a flat marginal distribution,

using the potential energy as CV. An interesting case is the one of well-tempered metadynamics [55] that aims at sampling an s distribution that is a smoothed version of the original one. Contrary to the uniform case and in general to the fixed target case [74], the well-tempered target explicitly depends on the unbiased probability, and is thus not completely known beforehand. Other kinds of targets are also used in the $1/k$ ensemble [110] and in nested sampling [111].

Another class of methods will be the main focus of this paper, and it is the one that aims at sampling the so-called expanded ensembles [95]. These targets are not defined explicitly as a function of some CVs, but rather consist in the sum of overlapping probability distributions. A typical enhanced sampling technique that targets expanded distributions is, for example, replica exchange [6]. Expanded ensembles can be obtained by combining the same system at different temperatures, or, more in general, different Hamiltonians [105, 112, 113]. They can be sampled also with single replica approaches, such as simulated tempering [98] and integrated tempering [114, 115]. Broadly speaking, one could consider as part of this expanded ensemble class also methods like multiple-windows umbrella sampling [116], or thermodynamic integration [117], where multiple separated ensembles are simulated and then combined into one via some postprocessing procedure such as the weighted histogram analysis method (WHAM) [63].

It is important to note that by classifying enhanced sampling methods with respect to $p^{\text{tg}}(\mathbf{x})$ we are not implying that methods in the same class are equivalent. Different methods in fact can use very different strategies to reach their target, each having its own strengths and weaknesses. However, this target-distribution perspective suggests that there is not a fundamental difference between the two traditional families, and that a unified approach is possible.

From this point of view, a special place is occupied by variationally enhanced sampling (VES) [5], targeted metadynamics [74], and on-the-fly probability enhanced sampling [96], since in these methods one has to explicitly choose a target distribution. This makes them particularly suited for developing a unified approach, since they are in principle capable of sampling the targets of both of the two families of enhanced sampling, and also combine them in new ways. In VES the usefulness of various target distributions has already been explored [24, 56, 85]. In particular, a target distribution has been proposed for sampling multithermal-multibaric ensembles [86] and also for combining them with a CV that drives a phase transition [118]. It is this paper that inspired us to try a generalized unified approach.

Our goal here is to introduce explicitly the expanded ensemble target distribution and show that it can be sampled by using a CV-based bias potential

method such as VES or OPES. In doing so we will introduce the concept of expansion CVs, that allows us to define both the target expanded distribution and the bias needed to sample it. The method we propose is thus capable of sampling both kinds of target distributions, those typical of replica exchange, but also the uniform and well-tempered distributions typical of metadynamics. In this sense it provides a unified approach to enhanced sampling.

2.4.3 On-the-fly Probability Enhanced Sampling

The recently developed on-the-fly probability enhanced sampling [96] is a collective-variables-based method. Collective variables are a function of the microscopic configuration, $\mathbf{s} = \mathbf{s}(\mathbf{x})$, that provide a low-dimensional description of the system. In OPES we aim at modifying the physical probability distribution of \mathbf{s} , $P(\mathbf{s})$, in order to reach a given target probability distribution, $p^{\text{tg}}(\mathbf{s})$. To achieve this we must add the following bias potential

$$V(\mathbf{s}) = -\frac{1}{\beta} \log \frac{p^{\text{tg}}(\mathbf{s})}{P(\mathbf{s})}, \quad (2.78)$$

where β is the inverse temperature. OPES has been introduced as an evolution of metadynamics and in this spirit we first have used the well-tempered target distribution, defined as $p^{\text{WT}}(\mathbf{s}) \propto [P(\mathbf{s})]^{1/\gamma}$, where $\gamma > 1$ is known as the bias factor. This target distribution aims at increasing the transition rate between metastable states of the system, by lowering of a factor γ the free-energy barriers along the CVs. In the limit of $\gamma = \infty$ it is equivalent to choosing a uniform target.

Since $P(\mathbf{s})$ is not known *a priori*, we resort to an iterative scheme [81, 119]. The core idea in OPES is to estimate the probability distribution at each step n , $P_n(\mathbf{s})$, by reweighting on the fly a simulation that is biased with $V_n(\mathbf{s})$, which is itself constructed from such $P_n(\mathbf{s})$ estimate according to Eq. (2.78). The $P_n(\mathbf{s})$ is obtained via a weighted kernel density estimation, adding a new kernel at a fixed small interval, similarly to metadynamics.

We refer the interested reader to Ref. [96], where the OPES iterative equations for the case of a well-tempered and a uniform target are presented in detail, and to Refs. [69, 70] for some initial applications. In the present paper we introduce a class of target distributions that allows sampling any expanded ensemble. We also present in detail the OPES iterative scheme for this class of targets. While the core ingredients of OPES remain the same, the resulting method will look quite different from the one presented in Ref. [96]. In particular, when targeting expanded ensembles we will not need to use the kernel density estimation that plays instead a crucial role in the well-tempered case.

In applying OPES to sample expanded ensembles, we find a method that is similar in spirit to that of Ref. [95] and to other more recent methods, such as integrated tempering sampling [114], infinite switch simulated tempering [120], and variationally-derived intermediates [121].

2.4.4 Targeting Expanded Ensembles

Let us call $u(\mathbf{x})$ the adimensional reduced potential that contains all the terms depending on the thermodynamic constraints, such as temperature, pressure, or others. With \mathbf{x} we concisely indicate the atomic coordinates and any other configurational variable that the reduced potential might depend on, such as the volume or the box tensor. As an example, in the case of the canonical ensemble one has $u(\mathbf{x}) = \beta U(\mathbf{x})$, where β is the inverse temperature and $U(\mathbf{x})$ is the potential energy of the system. Let us consider a system with a reduced potential $u_\lambda(\mathbf{x})$ that is a function of λ , where λ could be either a single parameter or a set of parameters, and might indicate, e.g., a thermodynamic property such as the temperature. At equilibrium its probability distribution follows Boltzmann statistics:

$$P_\lambda(\mathbf{x}) = \frac{e^{-u_\lambda(\mathbf{x})}}{Z_\lambda}, \quad (2.79)$$

where Z_λ is the partition function, $Z_\lambda = \int e^{-u_\lambda(\mathbf{x})} d\mathbf{x}$.

We are interested in sampling configurations in a range $\Delta\lambda$ of λ values. Instead of running multiple independent simulations at different values of λ , we can sample a generalized ensemble which contains all the relevant microscopic configurations, and then reweight them to retrieve the correct statistics for any $\lambda \in \Delta\lambda$. Sampling such an ensemble over $\Delta\lambda$ instead of the separate single λ ensembles is more efficient when different λ ensembles overlap in the coordinate space, and it can also help in solving ergodicity problems.

We must choose as target a distribution that covers all the microscopic configurations relevant to the chosen $\Delta\lambda$ range. Similarly to what is done in replica exchange, we choose a set $\{\lambda\}$ of $N_{\{\lambda\}}$ values $\lambda \in \Delta\lambda$ such that there is a good overlap between contiguous $P_\lambda(\mathbf{x})$. We then define our target distribution as

$$p_{\{\lambda\}}(\mathbf{x}) = \frac{1}{N_{\{\lambda\}}} \sum_{\lambda} P_\lambda(\mathbf{x}). \quad (2.80)$$

We refer to this class of target probability distributions as expanded ensemble target distributions. In the present paper we limit ourselves to considering nonweighted expanded targets, that assign the same $1/N_{\{\lambda\}}$ weight to all the subensembles, but it is also possible to add some λ -dependent weights and give different importance to different $P_\lambda(\mathbf{x})$.

Without loss of generality, one can consider $\lambda = 0$ to be inside the desired interval $\Delta\lambda$. It is then possible to run a simulation at $\lambda = 0$ and use the OPES scheme to iteratively optimize a bias that allows one to sample $p_{\{\lambda\}}(\mathbf{x})$. Before proceeding to explicitly write the target distribution and the bias potential, we express $P_\lambda(\mathbf{x})$ as

$$P_\lambda(\mathbf{x}) = P_0(\mathbf{x}) e^{-u_\lambda(\mathbf{x})+u_0(\mathbf{x})} \frac{Z_0}{Z_\lambda} = P_0(\mathbf{x}) e^{-\Delta u_\lambda(\mathbf{x})+\Delta F(\lambda)}, \quad (2.81)$$

where $\Delta u_\lambda(\mathbf{x}) = u_\lambda(\mathbf{x}) - u_0(\mathbf{x})$ is the potential energy difference and

$$\Delta F(\lambda) = -\log \frac{Z_\lambda}{Z_0} = -\log \langle e^{-\Delta u_\lambda} \rangle_{u_0}, \quad (2.82)$$

is the dimensionless free-energy difference from the reference system u_0 , that can be expressed also as an ensemble average, indicated with the notation $\langle \cdot \rangle_{u_0}$. Our expanded target thus becomes

$$p_{\{\lambda\}}(\mathbf{x}) = P_0(\mathbf{x}) \frac{1}{N_{\{\lambda\}}} \sum_\lambda e^{-\Delta u_\lambda(\mathbf{x})+\Delta F(\lambda)}. \quad (2.83)$$

In order to define the target bias, we first rewrite Eq. (2.78) as

$$v(\mathbf{x}) = -\log \frac{p^{\text{tg}}(\mathbf{x})}{P_0(\mathbf{x})}. \quad (2.84)$$

Finally the adimensional bias needed to sample the expanded target $p_{\{\lambda\}}(\mathbf{x})$ is:

$$v(\mathbf{x}) = -\log \left(\frac{1}{N_{\{\lambda\}}} \sum_\lambda e^{-\Delta u_\lambda(\mathbf{x})+\Delta F(\lambda)} \right). \quad (2.85)$$

Note that in writing the bias in this way $P_0(\mathbf{x})$ cancels out. It follows that the bias potential $v(\mathbf{x})$ depends on the coordinates \mathbf{x} only through the $N_{\{\lambda\}}$ quantities $\Delta u_\lambda(\mathbf{x})$. We refer to these Δu_λ as expansion collective variables. The expansion CVs completely characterize the expansion, since not only the bias, but also $\Delta F(\lambda)$, Eq. (2.82), and the expanded target distribution $p_{\{\lambda\}}(\mathbf{x})$, Eq. (2.83), are unambiguously defined through these quantities. We will see how, by properly choosing the expansion CVs $\Delta u_\lambda(\mathbf{x})$, it is possible to sample different kinds of expanded ensembles. For each of them we also highlight the connection between these expansion CVs and more traditional CVs that have a straightforward physical interpretation.

Our target bias, Eq. (2.85), depends on the free energy along the λ parameter, $\Delta F(\lambda)$, that is in general unknown. In the OPES spirit we will reach the target bias iteratively, by estimating on the fly $\Delta F(\lambda)$ via a reweighting procedure, and using such an estimate to define the applied bias.

Iterative OPES scheme

The free-energy difference $\Delta F(\lambda)$ defined in Eq. (2.82) can be written using an ensemble average over the reference unbiased system u_0 [122]. However, estimating $\langle e^{-\Delta u_\lambda} \rangle_{u_0}$ from an unbiased trajectory is practically impossible due to the fact that $e^{-\Delta u_\lambda}$ can be significantly large in regions where P_0 is extremely small, and that are thus not properly sampled. For this reason we use reweighting to write it as an average over the biased ensemble

$$e^{-\Delta F(\lambda)} = \langle e^{-\Delta u_\lambda} \rangle_{u_0} = \frac{\langle e^{-\Delta u_\lambda + v} \rangle_{u_0 + v}}{\langle e^v \rangle_{u_0 + v}}, \quad (2.86)$$

where the ensemble average $\langle \cdot \rangle_{u_0 + v}$ is computed as a time average over a biased trajectory. In this way, one can obtain a much more accurate estimate of $\Delta F(\lambda)$.

The problem with this procedure is that the target bias v , Eq. (2.85), is itself a function of $\Delta F(\lambda)$. Therefore, we set up an iterative scheme that consists in running a biased simulation whose bias is based on the estimate of the free-energy difference that we obtain via on-the-fly reweighting. At step n the simulation runs with potential $u_0(\mathbf{x}) + v_n(\mathbf{x})$, where

$$v_n(\mathbf{x}) = -\log \left(\frac{1}{N_{\{\lambda\}}} \sum_{\lambda} e^{-\Delta u_\lambda(\mathbf{x}) + \Delta F_n(\lambda)} \right). \quad (2.87)$$

The reweighted estimate $\Delta F_n(\lambda)$ is updated at every iteration step n . In between the iteration steps there is a fixed short stride where the simulation proceeds and both $\Delta F_n(\lambda)$ and the bias $v_n(\mathbf{x})$ are kept constant. The free-energy estimate at the n th step can be explicitly written as

$$\Delta F_n(\lambda) = -\log \left(\frac{\sum_{k=1}^n e^{-\Delta u_\lambda^{(k)} + v_{k-1}^{(k)}}}{\sum_{k=1}^n e^{v_{k-1}^{(k)}}} \right), \quad (2.88)$$

where we use the notation $\Delta u_\lambda^{(k)} \equiv \Delta u_\lambda(\mathbf{x}_k)$ and $v_n^{(k)} \equiv v_n(\mathbf{x}_k)$, with \mathbf{x}_k the configuration at the k th iteration step.

As the bias approaches convergence, the ensemble sampled approaches the target one, and the $\Delta F_n(\lambda)$ estimates become more and more accurate. Thus not only do we obtain the target bias, but we also have an estimate of the free energy as a function of the λ parameter, i.e., $\Delta F(\lambda)$. Our iterative scheme is similar in spirit to the one used in integrated tempering sampling [114], but the two differ both in their implementation and in their applications.

Equations (2.87) and (2.88) are the explicit OPES iterative equations used for sampling the expanded ensemble defined by the target distribution $p_{\{\lambda\}}(\mathbf{x})$,

Eq. (2.80), and are at the heart of our method. In the following sections we show how these equations can be used to sample different expanded ensembles, simply by specifying different expansion CVs $\Delta u_\lambda(\mathbf{x})$. Once these are chosen, the only free parameter of the method is the stride between the iterations. This should be set so that consecutive steps are not too correlated, as it is the case for the on-the-fly Gaussians' deposition in metadynamics.

It is possible to parallelize the procedure using multiple replicas of the system, as is done in multiple walkers metadynamics [41], where each replica shares the same bias and all contribute to the ensemble average in Eq. (2.88). At variance with replica exchange, the number of parallel simulations does not have to be equal to the number $N_{\{\lambda\}}$ of λ points that define the target.

Finally, we note that one could consider expressing the free energy $\Delta F(\lambda)$ via a cumulant expansion [123, 124]. This generally provides a very good estimate close to the reference $\lambda = 0$, but can be very inaccurate when the range is broad, requiring a great number of terms in the expansion. Furthermore, we found it can introduce artificial barriers that might stop the system from visiting all the relevant configurations, thus making the OPES self-consistent procedure much less efficient.

Reweighting

Until now we have seen how to sample expanded ensembles by applying a bias potential. We now need a reweighting procedure in order to retrieve statistics at any desired value of λ . To this effect one can use standard umbrella sampling reweighting [2]. Given any observable $O = O(\mathbf{x})$ that is a function of the atomic coordinates, we can calculate its average in the ensemble λ via the following reweighting equation:

$$\langle O \rangle_{u_\lambda} = \frac{\langle O e^{-\Delta u_\lambda + v} \rangle_{u_0 + v}}{\langle e^{-\Delta u_\lambda + v} \rangle_{u_0 + v}} \approx \frac{\sum_k^n O_k w_k(\lambda)}{\sum_k^n w_k(\lambda)}, \quad (2.89)$$

where $O_k \equiv O(\mathbf{x}_k)$ and the weight $w_k(\lambda)$ is defined as $w_k(\lambda) \equiv e^{-\Delta u_\lambda^{(k)} + v_{k-1}^{(k)}}$.

This equation assumes that the applied bias is static or quasistatic, meaning that it is updated in an adiabatic fashion. It is thus good practice to discard an initial transient of the simulation, where the bias changes quickly, and not use it for reweighting. Determining the exact length to be discarded might not be intuitive; however, OPES generally assigns a very low weight to this initial transient, and thus the result will not be significantly affected by this choice [96].

A useful diagnostic tool when performing reweighting is the so-called effective sample size, defined as the number of sampled points n times the ratio between the variance of an observable in the unbiased ensemble and

its variance in the reweighted ensemble [125]. We use a popular estimator for the effective sample size, defined as [88, 125]

$$n_{\text{eff}}(\lambda) = \frac{[\sum_k w_k(\lambda)]^2}{\sum_k w_k^2(\lambda)}, \quad (2.90)$$

where $w_k(\lambda)$ are the importance sampling weights. Intuitively, the effective sample size for a given λ will be smaller than the total number of samples, $n_{\text{eff}}(\lambda) < n$. One should expect the efficiency to be roughly $n_{\text{eff}}(\lambda)/n \propto 1/N_{\{\lambda\}}$, given a minimal choice of λ points that properly covers the target range. Plotting n_{eff}/n as a function of λ can be a good diagnostic tool to monitor the consistency of the iterative procedure.

Finally, we note that the estimate of uncertainties requires some extra care in case of weighted samples [126]. In Appendix 2.4.9 we describe in detail the weighted block averaging procedure we adopt, and show how the effective sample size plays a role.

2.4.5 Linearly Expanded Ensembles

An important type of expanded ensemble is the one obtained by linearly perturbing the reduced potential of the system, $u_\lambda(\mathbf{x}) = u_0(\mathbf{x}) + \lambda\Delta u(\mathbf{x})$. It is defined by the following expansion CVs:

$$\Delta u_\lambda(\mathbf{x}) = \lambda\Delta u(\mathbf{x}). \quad (2.91)$$

Various different ensembles can be obtained in this way, such as the multi-canonical ensemble and the multibaric ensemble, and also alchemical transformations, and others. Recently an interesting “multiforce” ensemble that falls in this category has been proposed [127]. We present in detail some of these ensembles in the following sections.

It can be useful to group together these linearly expanded ensembles because they share some interesting properties. In particular, for these ensembles we can propose a simple automatic way to choose the λ points that define the target $p_{\{\lambda\}}(\mathbf{x})$. The idea is to have the λ points uniformly distributed in the $\Delta\lambda$ interval with a spacing $\Delta\lambda/N_{\{\lambda\}}$ estimated from the effective sample size as a function of λ , $n_{\text{eff}}(\lambda)$. In practice what we do is to run a short unbiased simulation of n steps at $\lambda = 0$ and use a root finding algorithm to determine the points $\lambda_+ > 0$ and $\lambda_- < 0$ such that $n_{\text{eff}}(\lambda_\pm)/n \approx 0.5$. Then one can use a total of $N_{\{\lambda\}} = \Delta\lambda/(\lambda_+ + \lambda_-)$ equally spaced λ points to define the target $p_{\{\lambda\}}(\mathbf{x})$.

This heuristic way of choosing the λ points is not optimal, and more elaborate options have been explored in the replica-exchange literature [128]. However, in our case this choice is less critical, since within our scheme one

can increase $N_{\{\lambda\}}$ without the need to simulate additional replicas of the system. Thus this procedure provides an easy and automatic guess for linearly expanded ensembles that can be practically useful in many scenarios (see Sec. 2.4.7 for further considerations on the choice of the λ points).

Multicanonical ensemble

We start by considering as an example of linearly expanded ensemble the case of the multicanonical ensemble, which is probably the one with the longest history. The goal is to sample all the configurations relevant for canonical simulations in a given range of temperatures. In a canonical simulation the reduced potential is $u(\mathbf{x}) = \beta U(\mathbf{x})$, where $U(\mathbf{x})$ is the potential energy, $\beta = 1/(k_B T)$ is the inverse thermodynamic temperature, and k_B is Boltzmann constant. It is possible to define a multicanonical linearly expanded ensemble, by putting $\Delta u(\mathbf{x}) = u_0(\mathbf{x}) = \beta_0 U(\mathbf{x})$ and $\lambda = (\beta - \beta_0)/\beta_0$, where β_0 is the inverse temperature set by the thermostat of the simulation and β spans the target range $\beta_{\min} < \beta < \beta_{\max}$.

The expansion CVs that define such a target are

$$\Delta u_\lambda(\mathbf{x}) = \lambda \beta_0 U(\mathbf{x}) = (\beta - \beta_0) U(\mathbf{x}), \quad (2.92)$$

and by using them in the OPES iterative equations, Eqs. (2.87) and (2.88), we obtain our multicanonical simulation. Given the physical significance of the inverse temperature β , it is more natural to directly consider β as parameter instead of the dimensionless λ . We thus write Δu_β and $\Delta F = \Delta F(\beta)$, where we have set $\Delta F(\beta_0) = 0$. Similarly, it is natural to consider the potential energy $U(\mathbf{x})$ as a collective variable, and thus write the bias as

$$v(U) = -\log \left(\frac{1}{N_{\{\beta\}}} \sum_{\beta} e^{-(\beta - \beta_0)U + \Delta F(\beta)} \right). \quad (2.93)$$

It is important to notice that we did not require the bias to be a function of a single CV, but rather we find it to be the case when we set as target the temperature-expanded ensemble. This is in fact a general property of linearly expanded ensembles. When expanding according to a given λ , the resulting bias will be a function only of the thermodynamic conjugate variable Δu . To define the bias $v = v(\Delta u)$ we then need to estimate the free energy along λ , $\Delta F(\lambda)$.

Other multicanonical methods aim instead at sampling a flat energy distribution [86, 108, 109]. In order to do so, they need to estimate the free energy as a function of U (or equivalently the density of states along U), while in our method, as in other tempering approaches [120], we instead need to estimate the free energy as a function of temperature, $\Delta F(\beta)$.




Figure 2.29: Alanine dipeptide in the multicanonical ensemble ($T_{\min} = 300$ K, $T_{\max} = 1000$ K). (a) Explored configurations as a function of the dihedral angles. Sampled points are colored according to their reweighting weight at $T = T_0 = 300$ K, $w_k(\beta) = e^{-(\beta-\beta_0)U_k + v_{k-1}^{(k)}}$. Notice how all the points in the transition state, close to $\phi = 0$, have extremely low probability of being sampled in an isothermal simulation at 300 K. (b) Free-energy difference between the two basins ΔF_{AB} as a function of temperature. (c) Reweighted free-energy surface at two different temperatures.

Example: Alanine dipeptide As an example of multicanonical sampling, we consider alanine dipeptide in a vacuum, at temperature $T_0 = 300$ K. This is a typical toy model for testing enhanced sampling methods, since at room temperature it presents two metastable states with an extremely low transition probability. A possible way of enhancing the sampling is to bias the ϕ and ψ dihedral angles, using as target a flat uniform distribution or the well-tempered distribution [96]. Here instead we bias the potential energy U , and use as a target the multicanonical ensemble over a temperature range from 300 up to 1000 K.

Simulations are performed with the molecular dynamics software GRO-MACS [37], patched with the enhanced sampling library PLUMED [38] (see Supplemental Material, Sec. 2.4.10, for computational details). The only input needed for OPES, besides the temperature range we are interested in sampling, is the bias update pace that is set to 500 simulation steps (1 ps).

Figure 2.29(a) shows on the ϕ, ψ plane the configurations sampled during the 100 ns multicanonical run. It is interesting to notice that the potential energy U would be considered a bad CV in enhanced sampling methods such as metadynamics, since it cannot distinguish between the two basins that have roughly the same energy. However, when using the multicanonical ensemble as target, by biasing U we can enhance the probability of visiting the transition state (roughly the region where $\phi = 0$), and thus observe multiple transitions between the basins and converge the free-energy difference between them, ΔF_{AB} [Fig. 2.29(b)]. We can use the angle ϕ to define this

free-energy difference between the two basins:

$$\Delta F_{AB} = -\log \left(\frac{\langle \chi_{\varphi \in [0, \pi]} \rangle}{\langle \chi_{\varphi \in [-\pi, 0]} \rangle} \right), \quad (2.94)$$

where χ is a characteristic function, equal to 1 if the variable is in the proper range and 0 otherwise.

In the Supplemental Material (Sec. 2.4.10) we discuss the differences with biasing the energy with metadynamics [100, 106] and show a comparison between this multicanonical run and a well-tempered run biasing the two dihedral angles. As expected, the latter is much more efficient (roughly 10 times) in converging the free-energy difference at 300 K, due to the fact that it focuses on the relevant degrees of freedom and on a single temperature.

Multithermal-multibaric ensemble

Within our scheme, combining different linearly expanded ensembles is straightforward. One simply has a two-dimensional parameter $\lambda = \{\lambda_1, \lambda_2\}$, and considers $u_\lambda(\mathbf{x}) = u_0 + \lambda_1 \Delta u_1(\mathbf{x}) + \lambda_2 \Delta u_2(\mathbf{x})$. This can be useful, for example, to sample multiple temperatures and multiple pressures in a single multithermal-multibaric simulation.

In this case we consider *NPT* simulations with a reference reduced potential $u_0(\mathbf{x}) = \beta_0 U(\mathbf{x}) + \beta_0 p_0 V(\mathbf{x})$, where p is the pressure and $V(\mathbf{x})$ the volume. Similarly to what was done before, it is more natural to use as λ parameters directly the temperature β and the pressure p , and write the expansion CVs $\Delta u_\lambda(\mathbf{x})$ as

$$\Delta u_{\beta,p}(\mathbf{x}) = (\beta - \beta_0)U(\mathbf{x}) + (\beta p - \beta_0 p_0)V(\mathbf{x}). \quad (2.95)$$

The target distribution is defined by a set of $N_{\{\beta\}}$ temperatures $\beta \in [\beta_{\min}, \beta_{\max}]$ and $N_{\{p\}}$ pressures $p \in [p_{\min}, p_{\max}]$, for a total of $N_{\{\beta,p\}} = N_{\{\beta\}} \times N_{\{p\}}$ different $\Delta F(\beta, p)$ to be estimated. We will also express the bias, Eq. (2.85), as a function of the potential energy and the volume $v = v(U, V)$, which come as a natural CVs choice. As already discussed, the intermediate temperatures β and pressures p that define the target can be chosen automatically from a short unbiased simulation. We can do this independently for the two parameters, despite the fact that the pressure term p is multiplied by β in Eq. (2.95).

Finally, we notice how the choice of β_0 and p_0 is completely free. As long as they lie inside the range of temperatures and pressures that we aim at sampling, no matter what thermodynamic conditions we start from at convergence we will sample the same configurational space. However, when the target range is very broad, choosing β_0 and p_0 roughly at the center can help to speed up convergence.

2.4. Unified approach to enhanced sampling




Figure 2.30: Multithermal-multibaric simulation of chignolin. (a) Sampled target distribution in the CV space of potential energy and volume. (b) Relative effective sample size at different temperatures and pressures. The bottom corner of high temperatures and low pressures is excluded from the target to avoid vaporization of the system. (c) Time evolution of the two biased CVs and of the $C\alpha$ RMSD for one of the 40 walkers. A RMSD threshold between folded and unfolded is highlighted with a dashed line.

Example: Chignolin As an example of a multithermal-multibaric simulation we consider the miniprotein chignolin (CLN025) with CHARMM22* force field [129] and the three-site transferable intermolecular potential (TIP3P) water model (about 2800 molecules), over a temperature range from $T_{\min} = 270$ K to $T_{\max} = 800$ K and a pressure range from $p_{\min} = 1$ bar to $p_{\max} = 4000$ bar. The velocity-rescaled thermostat [36] is set at $T_0 = 500$ K and the Parrinello-Rahman barostat [68] at $p_0 = 2000$ bar. The $\Delta F_n(\beta, p)$ estimates and the bias are updated every 500 simulation steps (1 ps). The $N_{\{\beta\}}$ temperature steps and $N_{\{p\}}$ pressure steps are chosen automatically based on a short 100 ps unbiased run. This results in 92 steps in temperature and 26 in pressure, for a total of $N_{\{\beta, p\}} = 2392$ points. In order to avoid the region of low pressure and high temperature where water could evaporate, we discard any (β, p) -point lying below the line from (500 K, 1 bar) and (800 K, 1000 bar). In this way, 91 (β, p) points are discarded.

The simulation is performed in parallel using 40 multiple walkers, and runs for a total of 300 ns per walker, of which roughly 10 ns are needed to converge the bias. Also in this case we use GROMACS patched with PLUMED.

In Fig. 2.30(a) we show the distribution sampled in the energy-volume space, while in Fig. 2.30(b) the corresponding effective sample size n_{eff} is plotted, as a function of temperature and pressure, and rescaled over the total number of samples n . The n_{eff}/n is not perfectly uniform, but it has the same order of magnitude over the whole target region. On the right, Fig. 2.30(c), we show




Figure 2.31: The fraction folded of chignolin estimated from the multithermal-multibaric simulation. The inset shows the same quantity over a smaller range of temperatures at 1 bar (highlighted with a gray box), in order to compare it with the reference results from Lindorff-Larsen *et al.* [131].

for one of the 40 replicas the energy and volume trajectory, together with the trajectory of the $C\alpha$ root-mean-square deviation (RMSD) to the experimental NMR structure [130].

In Fig. 2.31 we show the chignolin folded fraction at the different temperatures and pressures we targeted. The folded fraction is defined as in Ref. [131], using a dual cutoff on the $C\alpha$ RMSD based on the CLN025 experimental NMR structure. A configuration is considered folded when the RMSD goes below 0.1 nm, and unfolded when it goes above 0.4 nm. In the inset we compare our results with those of Ref. [131] that considered a smaller temperature range at standard pressure. The confidence interval of our estimate is calculated with the block analysis described in Appendix 2.4.9.

The stability diagram of chignolin (Fig. 2.31) does not present striking features, in qualitative agreement with Ref. [132]. However, it has been recently shown [133] that other miniproteins can have a nontrivial phase diagram, with unfolding both at low temperature and at high pressure.

Thermodynamic integration

Another interesting application of our method is its use for performing thermodynamic integration [134]. Let us consider a system with reduced potential energy $u_0(\mathbf{x})$ and free energy $F_0 = -\log Z_0$ and another similar system with potential $u_1(\mathbf{x})$ and free energy F_1 . We are interested in calculating the

free-energy difference $\Delta F_{0 \rightarrow 1} = F_1 - F_0$, for instance, because we know the free energy of one of the two systems and in this way we can retrieve the other one. The key idea of thermodynamic integration is to define a ladder of intermediate systems with reduced potentials $u_\lambda(\mathbf{x})$ and $0 < \lambda < 1$, to connect the two systems. The free-energy difference $\Delta F_{0 \rightarrow 1}$ to go from the u_0 system to the u_1 can be calculated via the following integral:

$$\Delta F_{0 \rightarrow 1} = \int_0^1 \left\langle \frac{\partial u_\lambda(\mathbf{x})}{\partial \lambda} \right\rangle_{u_\lambda} d\lambda. \quad (2.96)$$

Typically individual simulations are run using $u_\lambda(\mathbf{x})$ for different values of λ and the ensemble average $\left\langle \frac{\partial u_\lambda(\mathbf{x})}{\partial \lambda} \right\rangle_{u_\lambda}$ is estimated for each of them. Then the integration in Eq. (2.96) can be carried out numerically, e.g., using the trapezoid rule or a Gaussian quadrature.

The most common way to define the intermediate states $u_\lambda(\mathbf{x})$ is via a linear interpolation,

$$u_\lambda(\mathbf{x}) = u_0(\mathbf{x}) + \lambda \Delta u(\mathbf{x}), \quad (2.97)$$

where $\Delta u(\mathbf{x}) \equiv u_1(\mathbf{x}) - u_0(\mathbf{x})$. In this case we have $\partial u_\lambda / \partial \lambda = \Delta u$.

In the spirit of the present paper, we aim at performing a single simulation that samples all values of λ simultaneously, similarly to other methods [95, 135, 136]. It is then possible to reweight for any λ and calculate the integral in Eq. (2.96). Thus we simulate the system at $u_0(\mathbf{x})$ and build a target $p_{\{\lambda\}}(\mathbf{x})$ as in Eq. (2.80) using $N_{\{\lambda\}}$ λ points in the interval $0 < \lambda < 1$. The OPES iterative equations, Eqs. (2.87) and (2.88), can be written using the expansion CVs $\Delta u_\lambda = \lambda \Delta u$ as defined in Eq. (2.97).

Finally, we note that thermodynamic integration can be performed using interpolation schemes different from the linear one, and our method is general and can be applied also in those scenarios, simply by properly defining the expansion CVs $\Delta u_\lambda(\mathbf{x})$.

Example: TIP4P water to Lennard-Jones fluid We now use the thermodynamic integration formalism described above to calculate the free energy of TIP4P water, relative to a reference Lennard-Jones system. The TIP4P potential energy (U_{TIP4P}) is made of an electrostatic energy term and a van der Waals-type interaction between the oxygens described by a Lennard-Jones potential (U_{LJ}). The free energy of a Lennard-Jones fluid with the same U_{LJ} potential has been fit to an equation of state and thus is a good reference [137]. For the simulations we use the LAMMPS [61] molecular dynamics software, patched with PLUMED. We perform an NVT canonical simulation at 443 K using the TIP4P water potential, thus $u_0(\mathbf{x}) = \beta U_{\text{TIP4P}}(\mathbf{x})$, with $N = 384$ molecules. The reference system is characterized by $u_1(\mathbf{x}) = \beta U_{\text{LJ}}(\mathbf{x})$.




Figure 2.32: Calculation of the free energy of liquid TIP4P water using a single-simulation thermodynamic integration. (a) Evolution of the collective variable ΔU as a function of simulation time. (b) Integrand for the thermodynamic integration obtained through reweighting.

With β a constant, we consider as collective variable $\Delta U(\mathbf{x}) \equiv U_{\text{LJ}}(\mathbf{x}) - U_{\text{TIP4P}}(\mathbf{x})$, and write the bias according to Eq. (2.85):

$$v(\Delta U) = -\log \left(\frac{1}{N_{\{\lambda\}}} \sum_{\lambda} e^{-\lambda\beta\Delta U + \Delta F(\lambda)} \right). \quad (2.98)$$

From a short 20 ps unbiased run we obtain with the usual automatic procedure (Sec. 2.4.5) 30 equispaced points in the interval $0 < \lambda < 1$, that define our target distribution. The evolution of ΔU as a function of simulation time is shown in Fig. 2.32(a). There is an initial transient of about 3 ns until the bias potential is optimized and then the system diffuses freely. This has to be compared with a simulation for a given value of λ in which the fluctuations of ΔU would be very small. From this simulation the integrand $\left\langle \frac{\partial u_\lambda(\mathbf{x})}{\partial \lambda} \right\rangle_{u_\lambda} = \beta \langle \Delta U \rangle_\lambda$ can be calculated via reweighting, Eq. (2.89),

$$\langle \Delta U \rangle_\lambda = \frac{\sum_k^n \Delta U_k w_k(\lambda)}{\sum_k^n w_k(\lambda)}, \quad (2.99)$$

where $\Delta U_k = U_{\text{LJ}}(\mathbf{x}_k) - U_{\text{TIP4P}}(\mathbf{x}_k)$ and $w_k(\lambda) = e^{-\lambda\beta\Delta U_k + v_{k-1}^{(k)}}$. The values of $\langle \Delta U \rangle_\lambda$ thus calculated are shown in Fig. 2.32(b). Using these results and Eq. (2.96) we find a free-energy difference $\Delta F_{\text{TIP4P} \rightarrow \text{LJ}} = F_{\text{LJ}} - F_{\text{TIP4P}} = 7.00(1)$ ($Nk_B T$ units) in agreement with the result reported in Ref. [117].

2.4.6 Beyond linearity: Multiumbrella Ensemble

We consider now another important kind of expanded ensemble, namely the one obtained by combining all the different windows of a typical umbrella sampling simulation [105]. We refer to such an ensemble as a multiumbrella ensemble.

Multiple-windows umbrella sampling [116] allows for the free-energy surface (FES) reconstruction along some collective variable $s = s(\mathbf{x})$, that can be the reaction coordinate or some slow mode of the system. Typically one simulates multiple copies of the system, each one with a parabolic bias potential centered at a different s_λ -point, in such a way that the resulting probability distributions have an overlap and cover the whole CV range. Postprocessing via WHAM [63] or other methods is then needed to estimate the relative free-energy differences and combine the data in a single FES estimate. Here instead, we aim at sampling all the umbrella windows in the same simulation, by estimating on the fly these free-energy differences and building a single global potential. For this reason, the FES can then be obtained with the simple reweighting scheme described in Sec 2.4.4. Both approaches have their own strengths and weaknesses, as we discuss while presenting the example below.

Given a system with reduced potential $u_0(\mathbf{x})$ and equilibrium Boltzmann distribution $P_0(\mathbf{x})$, we can write the reduced potential of each umbrella window as $u_\lambda(\mathbf{x}) = u_0(\mathbf{x}) + \Delta u_\lambda(\mathbf{x})$, with expansion CVs:

$$\Delta u_\lambda(\mathbf{x}) = \frac{(s(\mathbf{x}) - s_\lambda)^2}{2\sigma^2}. \quad (2.100)$$

The associated probability distribution is $P_\lambda(\mathbf{x}) \propto P_0(\mathbf{x})G_\sigma(s(\mathbf{x}), s_\lambda)$, where $G_\sigma(s, s_\lambda)$ is a Gaussian of width σ centered in s_λ . The resulting expanded target $p_{\{\lambda\}}(\mathbf{x}) = \frac{1}{N_{\{\lambda\}}} \sum_\lambda P_\lambda(\mathbf{x})$ is clearly not linear in λ , and in fact requires an extra parameter σ to be defined. The width σ can in principle vary with λ , but we consider here only the case of uniform umbrellas.

Since the expansion CVs, Eq. (2.100), depend on \mathbf{x} only through $s = s(\mathbf{x})$, it is natural to write the bias as a function of the s CV

$$v(s) = -\log \left(\frac{1}{N_{\{\lambda\}}} \sum_\lambda e^{-(s-s_\lambda)^2/2\sigma^2 + \Delta F(s_\lambda)} \right). \quad (2.101)$$

Contrary to the linear case, in this multiumbrella case both the bias $v(s)$ and the free-energy differences $\Delta F(s) = -\log \langle G_\sigma(s(\mathbf{x}), s) \rangle_{u_0}$ are expressed as functions of the same CV.

The $N_{\{\lambda\}}$ s_λ points can be chosen to be uniformly distributed in the desired $\Delta s = s_{\max} - s_{\min}$ interval, in such a way to be at most at a distance of σ , ensuring overlap between contiguous P_λ . For a small enough σ , the estimate $\Delta F_n(s)$ converges precisely to the free-energy surface (FES), while if σ is too broad there will be small artifacts, similarly to what happens when a too broad bandwidth is used in kernel density estimation.

It is instructive to consider the marginal of the target probability with respect to the CV, $p_{\{\lambda\}}(s)$. In the limit of infinitely small σ and thus infinitely large

$N_{\{\lambda\}}$, the multiumbrella target $p_{\{\lambda\}}(s)$ is a uniform flat distribution over the Δs interval. In the opposite limit, of a very broad σ , the target distribution will look like the original, hard-to-sample P_0 . As a rule of thumb σ should be as small as the smallest features of the FES we are interested in. We note that this is the same criterion used to choose the σ parameter in metadynamics [20], and it can typically be guessed from a short unbiased run. For this reason we prefer to use as parameter σ instead of the more commonly used strength of the harmonic umbrella potential $K = 1/\sigma^2$ [116].

In some cases it proved useful to introduce two small modifications to make the multiumbrella iterative optimization scheme more robust. We leave the explanation of them to Appendix 2.4.9, since they have not been necessary for the examples presented in the paper.

For simplifying the exposition we presented the procedure in the case of a 1D CV, but it is straightforward to extend it to higher dimension, by using multidimensional Gaussians and placing the s_λ points on an appropriate multidimensional grid. When dealing with higher dimensions it might be interesting to use some more elaborate shapes for the umbrellas, e.g., a Gaussian mixture in a similar but complementary way to Ref. [138], or to follow a specific path, as in Ref. [45].

Example: Double-well model As an example for the multiumbrella ensemble we consider a Langevin dynamics on a 2D model potential [80] using as CV the x coordinate only, Fig. 2.33(a). Such CV is suboptimal, in the sense that it correctly describes the metastable basins, but does not include all the slow modes of the system [80]. One of the positive features of multiple-windows umbrella sampling is that it allows one to estimate the FES by running many short simulations in parallel. However, in this case some of the windows—for instance the one centered at $x = 0$ —cannot be efficiently sampled in a single short simulation due to the barrier along y . This is a known problem that can be diagnosed with dedicated consistency tests and is typically solved by performing more sampling [64]. With our approach, instead, we have a single combined simulation that in this example is long enough for the y mode to diffuse, and thus the CV suboptimality does not constitute a problem, and no extra care is required to handle it.

Figure 2.33(b) shows how the target distribution changes for different σ choices, expressed in units of the unbiased standard deviation in the basins, $\sigma_0 \approx 0.18$. The FES estimate could be directly obtained from the $\Delta F_n(s)$, but in the case of large σ this would lead to an estimate in which features are oversmoothed (see Supplemental Material, Sec. 2.4.10). Also, due to the CV suboptimality this estimate might not be very precise, but simply good enough for having that the main bottleneck for the transition is the unbiased




Figure 2.33: Double-well model in the multumbrella ensemble. (a) Potential energy of the double-well 2D model system, and its free energy along the x coordinate. (b) Multumbrella target distribution, for different values of umbrella width σ . The black dotted line is the unbiased probability distribution $P_0(x)$.

slow mode y [80]. As a general rule, it is better to estimate the FES via the reweighting procedure.

In the Supplemental Material (Sec. 2.4.10) we provide all the simulation details and show the convergence of the free energy, comparing it to well-tempered OPES and metadynamics. While in metadynamics and well-tempered OPES the bias is constructed in such a way to push the system out of the visited areas, with multumbrella OPES we are forcing the system to stay in a chosen CV region. Despite this difference in both cases we have similar target distributions and the resulting sampling allows us to reconstruct the FES.

Combining thermodynamic and order parameter expansions

An important characteristic of the present scheme is that it allows for a straightforward combination of different expanded ensembles. In particular it allows for a rigorous and efficient combination of thermodynamic generalized ensembles with enhanced sampling along a system-specific order parameter.

To understand why this is important one can think about a first-order phase transition, where there is a kinetic bottleneck between the two phases that is responsible for an ergodicity problem. Increasing the temperature typically changes the relative stability of the two phases, but the free-energy barrier separating them might remain high along the whole coexistence line, thus making convergence very slow. A possible solution is to identify a suitable order parameter and biasing it to increase the transition probability.

Combining the two approaches might actually outperform both [139, 140]. This kind of combination can be useful not only for phase transitions; for instance, also in alchemical free energy calculations an open problem is how to properly handle barriers orthogonal to the transformation [141].

We have already cited some hybrid methods that combine a replica-exchange approach with metadynamics, in order to enhance the sampling along both a thermodynamic quantity and an order parameter [73, 99–101, 139]. A nonhybrid approach has been first proposed with multidimensional replica exchange [105], but it has the drawback of requiring a sometimes impractical number of parallel replicas, due to the multidimensionality of the expansion. With OPES we can sample the same target distribution of multidimensional replica exchange, but using a bias potential and without requiring a minimum number of parallel replicas. In developing our method we followed the footsteps of another nonhybrid approach that has been recently proposed by our group, using the VES formalism and a custom target distribution [118, 142]. Compared to the very flexible and customizable VES approach, OPES has the advantage of having much fewer free parameters and thus being simpler to set up and use.

Example: Sodium We consider here as an example the calculation of the liquid-bcc phase diagram of a model of sodium [67], the same studied in Ref. [118]. We sample the liquid and solid phase over a range of temperatures and pressures, using a recently proposed order parameter s , called environment similarity collective variable [118]. Such CV provides a measure of the crystallinity of the system, by comparing the local environment of the atoms to a reference one. For this reason we refer to it as crystallinity CV, but it is actually more general and can be used to describe a variety of phase transitions [142, 143].

Using LAMMPS patched with PLUMED, we perform NPT simulations, $u_0(\mathbf{x}) = \beta_0 U(\mathbf{x}) + \beta_0 p_0 V(\mathbf{x})$. We can write the OPES equations, Eqs. (2.87) and (2.88), via the following expansion CVs

$$\Delta u_{\beta,p,s}(\mathbf{x}) = (\beta - \beta_0)U(\mathbf{x}) + (\beta p - \beta_0 p_0)V + \frac{(s(\mathbf{x}) - s)^2}{2\sigma^2}. \quad (2.102)$$

The free-energy estimates $\Delta F_n(\beta, p, s)$ are expressed as a function of the inverse temperature β , the pressure p and the crystallinity CV s . The bias $v = v(U, V, s)$ is expressed as a function of the potential energy U , the volume V and the crystallinity CV s .

The simulation is performed with 250 atoms at $T_0 = 400$ K and $p_0 = 0.5$ GPa (5 kbar), using 4 multiple walkers that share the same bias and contribute to the same ensemble averages to update the $\Delta F_n(\beta, p, s)$ estimate. The aim is to sample liquid and solid configurations in the temperature range from

2.4. Unified approach to enhanced sampling




Figure 2.34: Configurations of sodium sampled during the multithermal-multibaric-multiumbrella simulation. (a) The points are colored accordingly to the value of the bias $v(U, V, s)$. (b) The points are shown in the energy-volume space and colored accordingly to the value of the crystallinity CV s . As reference we also show in grey the region sampled during an unbiased simulation in the bcc phase.

350 to 450 K and pressures from 0 to 1 GPa (10 kbar). The uniform grid over β and p to define the target distribution is automatically generated from a short 100 ps unbiased run, and consists of 4 temperature steps and 8 pressure steps. We chose as σ for the multiumbrella target a value of about 2.5 times the unbiased standard deviation in the basins, and it determines the presence of 26 umbrellas uniformly placed between $s_{\min} = 0$ (liquid) and $s_{\max} = 1$ (solid). In total the $\Delta F_n(\beta, p, s)$ to be estimated are $4 \times 8 \times 26 = 832$. After less than 3 ns the bias has clearly reached the adiabatic regime, and could be kept constant and used as static bias. However, since there is no drawback in doing so, we keep updating the bias and run 25 ns per walker, for a total combined simulation of 100 ns.

In Fig. 2.34(a) we show the points sampled in the CV space during the simulation, colored according to the value of the bias potential $v(U, V, s)$. We can clearly distinguish the hourglass shape described in Ref. [118]. In Fig. 2.34(b) we see the same trajectory plotted on the energy-volume plane. For comparison we show the configuration sampled in an unbiased simulation at a single temperature and pressure, in which the system remains all the time in the bcc crystal phase.

We can define the free-energy difference between the two phases as in Ref. [118]:

$$\Delta F_{\text{liq} \rightarrow \text{bcc}}(T, p) = -\log \left(\frac{\langle \chi_{s \in [0.5, 1]} \rangle_{T, p}}{\langle \chi_{s \in [0, 0.5]} \rangle_{T, p}} \right), \quad (2.103)$$

where χ is a characteristic function, equal to 1 if the variable is in the proper range and 0 otherwise, and $\langle \cdot \rangle_{T, p}$ is the ensemble average at temperature T and pressure p .




Figure 2.35: Phase equilibrium of liquid and solid (bcc) sodium using a combination of thermodynamic and order parameter expansions. (a) Free-energy difference between the phases, $\Delta F_{\text{liq} \rightarrow \text{bcc}}$, at different thermodynamic conditions. The coexistence line is shown as a grey dashed line. (b) Free-energy surfaces as a function of the crystallinity CV , for three representative thermodynamic conditions. Error bars are smaller than the linewidth.

In Fig. 2.35(a) we show $\Delta F_{\text{liq} \rightarrow \text{bcc}}(T, p)$ obtained by reweighting at different temperatures and pressures. The coexistence line $\Delta F_{\text{liq} \rightarrow \text{bcc}}(T, p) = 0$ is shown with a dotted grey line. On the right-hand side, Fig. 2.35(b), we provide the free-energy surface as a function of the CV, $F(s)$, at different thermodynamic conditions. Error bars are calculated with a weighted block average (Appendix 2.4.9) and all the results are in agreement with Ref. [118], see Supplemental Material (Sec. 2.4.10).

It is important to notice that while the relative stability between liquid and solid changes across the range considered here, the probability of being in the transition state between the two is always extremely small, as can be seen from the high FES values around $s = 0.5$ in Fig. 2.35(b). By actively biasing the CV s we allow the system to efficiently sample the transition region as well, and this makes it possible to quickly converge the multithermal-multibaric simulation despite the presence of a first-order phase transition.

2.4.7 About the Optimal Target Distribution

Before reaching the conclusion of the paper we would like to add a final remark. At the beginning of Sec. 2.4.4 we presented the nonweighted expanded ensemble target distribution, $p_{\{\lambda\}}(\mathbf{x}) = \frac{1}{N_{\{\lambda\}}} \sum_{\lambda} P_{\lambda}(\mathbf{x})$. It is reasonable to wonder if this is an optimal target and in which sense. We argue here that the effective sample size can be used to define an optimality criterion.

Let us say our goal is to sample from a generalized ensemble that contains all the relevant microscopical configurations for a give range of the parameter λ . While the expanded target distribution $p_{\{\lambda\}}(\mathbf{x})$, Eq. (2.80), fulfills such a goal, there are in principle other possible choices.

It is useful to look at the special case of multicanonical ensembles, that has a long history (see also Sec. 2.4.5). In this context various different target distributions have been used, other than the nonweighted expanded ensembles one. One option is to have a uniform sampling in the energy [86, 108, 109], another one is to have a uniform sampling in the entropy [110], and a third one is to define the target by integrating the probability over the temperature, as in Refs. [114, 120]. In this last approach one often approximates the integral with a sum, and effectively uses a target similar to our, Eq. (2.80), which is also the one used in temperature replica exchange. Another interesting perspective is presented in Ref. [144], where a Riemann metric is introduced to define optimality.

We believe that if our goal is to reweight at different temperatures, then the optimal target distribution is the one that yields the highest possible uniform effective sample size over the whole considered $\Delta\lambda$ range. Here we will not further explore such optimal target distribution nor dig deeper in the definition of effective sample size. However, simply by defining this criterion we can notice that one should not see the sum in Eq. (2.80) as an approximation to an integral. As a matter of fact, using as few as possible intermediate λ points brings us closer to this optimal target than having more, at least in the systems we studied (see Supplemental Material, Sec. 2.4.10).

It might also be the case that one is not interested in obtaining statistic for the whole $\Delta\lambda$ range, but only for a subset of λ states. In this case the optimal target would be the one that maximizes the effective sample size for those λ states while ensuring ergodicity. According to this criterion we argued in Ref. [96] that the well-tempered target is better than a uniform one, since it allows for an ergodic sampling while providing a higher n_{eff}/n ratio and avoids unimportant high free energy regions.

2.4.8 Conclusion

In this paper we presented a general framework that provides a unified approach to enhanced sampling. To implement our method we leveraged the iterative scheme of OPES, an enhanced sampling method based on the construction of a bias potential along a set of collective variables, that was originally introduced for metadynamicslike sampling. We showed how this approach can be used to sample the same expanded ensembles typically sampled by a different family of enhanced sampling methods.

We also introduced the concept of expansion CVs, $\Delta u_\lambda(\mathbf{x})$, that can be used to fully characterize a nonweighted expanded target distribution $p_{\{\lambda\}}(\mathbf{x})$, Eq. (2.80), together with the free-energy differences to be iteratively estimated, Eq. (2.82), and the target bias, Eq. (2.85).

We then presented various examples of the application of the method to

Table 2.3: Some of the most common expanded ensembles, together with the expansion collective variables $\Delta u_\lambda(\mathbf{x})$ that define the OPES target bias, Eq. (2.85), and the free-energy differences $\Delta F(\lambda)$, Eq. (2.82). Each of the considered target biases can in turn be expressed as a function of one or two CVs. It is also possible to easily combine these ensembles to form new ones, as shown in Sec. 2.4.6.

Target ensemble	Expansion CVs	Parameters	CVs
Linearly expanded	$\lambda \Delta u(\mathbf{x})$	$\{\lambda\}$	$\Delta u(\mathbf{x})$
Multicanonical	$(\beta - \beta_0)U(\mathbf{x})$	$\{\beta\}$	$U(\mathbf{x})$
Multibaric	$\beta_0(p - p_0)V(\mathbf{x})$	$\{p\}$	$V(\mathbf{x})$
Multithermal-baric	$(\beta - \beta_0)U(\mathbf{x}) + (\beta p - \beta_0 p_0)V(\mathbf{x})$	$\{\beta, p\}$	$U(\mathbf{x}), V(\mathbf{x})$
Multiumbrella	$(s(\mathbf{x}) - s_\lambda)^2 / 2\sigma^2$	$\{s_\lambda\}$	$s(\mathbf{x})$

sample the most common expanded ensembles. These ensembles are summarized in Table 2.3. In particular, we have shown how OPES can be used to enhance at the same time temperature-related fluctuations and a system-specific order parameter, Sec. 2.4.6.

We notice that in defining the target distribution $p^{\text{tg}}(\mathbf{x})$ we consider only the positional degrees of freedom, and not the atomic velocities. Thus the ensembles sampled by our method are not identical to the ones sampled, for instance, by replica exchange, even though the target distribution is the same. In fact, the two methods sample the same configuration space, but a different velocity space. This does not have an effect on any statistical average of observables that are function of the coordinates only, but might be an interesting point for further research.

In the future it would be interesting to combine expanded target distributions with well-tempered-like distributions, which can scale better with higher dimensionality. Also weighted expanded targets might be of interest, where each subensemble λ has a specific different normalized weight. It might also be useful to implement the target distribution implicitly used by bias-exchange and parallel-bias metadynamics [50, 73], that scales efficiently with the number of CVs. More generally, we believe that our perspective of focusing on the target distribution has further potential that should be explored.

Data Availability Statement

The method is implemented in the open source PLUMED enhanced sampling library [38], and will be available as a contributed module called OPES. All the input files and postprocessing scripts used for this paper are openly available on the PLUMED-NEST [91] (www.plumed-nest.org, plumID:20.022), and in the Materials Cloud Archive (www.materialscloud.org, materialscloud:2020.81), where also the trajectories of the simulations are stored.

Acknowledgments

The authors thank Riccardo Capelli and Sandro Bottaro for useful discussions, and Andrea Rizzi for helpful feedback on the manuscript. This research was supported by the NCCR MARVEL, funded by the Swiss National Science Foundation, and European Union Grant No. ERC-2014-AdG-670227/VARMET. P.M.P. acknowledges support from the Swiss National Science Foundation through an Early Postdoc.Mobility fellowship. Calculations were carried out on Euler cluster at ETH Zurich and in Swiss National Supercomputing Center (CSCS) cluster Piz Daint.

2.4.9 Appendix

A. Algorithmic Scheme

We schematically present the OPES algorithm used for sampling expanded ensembles.

- **INITIALIZATION:** each set of expansion CVs can be initialized independently, following one of two strategies
 - **manual** Choose the expansion CVs $\Delta u_\lambda(\mathbf{x})$ and initialize $\Delta F_0(\lambda) = \Delta u_\lambda(\mathbf{x}_0)$
 - **automatic** (linear expansions only) Chose λ_{\min} and λ_{\max} . Run n_{init} unbiased simulation steps, and use $n_{\text{eff}}(\lambda)$ as described in Sec. 2.4.5 to define the $\{\lambda\}$ set. Initialize $\Delta F_0(\lambda)$ via Eq. (2.88), using the n_{init} points and $v = 0$
- **SIMULATION:** run the simulation for τ steps with potential $u(\mathbf{x}) + v_n(\mathbf{x})$. Add the point $\mathbf{x}_{n+1} = \mathbf{x}(n\tau)$ to the estimate $\Delta F_{n+1}(\lambda)$ and thus define $v_{n+1}(\mathbf{x})$, according to Eqs. (2.88) and (2.87). If multiple walkers are present, use each different configuration \mathbf{x}_{n+1} to update the same $\Delta F_{n+1}(\lambda)$ estimate.
- **REWEIGHTING:** to reweight for an observable $O(\mathbf{x})$ store its value during the simulation, together with the value of the potential v and use Eq. (2.89)

B. Weighted Block Average

When estimating the uncertainty related to an ensemble average obtained from molecular dynamics, one must take into account for the time correlation between the samples. Methods such as block averaging [145] are commonly used to properly handle this. The effect of such time correlation is to make the sample size effectively smaller, thus simply taking the square root of the variance divided by the number of samples would underestimate the actual uncertainty. When dealing with weighted samples, as it is the case when a bias potential is used, the effective sample size is further reduced

by the presence of these weights. To account also for this effect we should slightly modify the block averaging procedure. We report here the protocol that we follow to estimate uncertainties, which is the same as the one presented in Ref. [126], but here we highlight the role played by the effective sample size.

We are interested in estimating the ensemble average of an observable $O = O(\mathbf{x})$ from a biased ensemble

$$\langle O \rangle \approx \hat{O} = \frac{\sum_{k=1}^n O_k w_k}{\sum_{k=1}^n w_k}, \quad (2.104)$$

where w_k are the weights due to the bias potential, as in Eq. (2.89). In order to estimate the uncertainty we divide the data into M sub-sets or blocks, each containing an equal number of samples n/M . We then calculate the estimate \hat{O}_i from each block, via a weighted average, and also the weight of the i th block

$$W_i = \sum_{k=(i-1)(n/M)}^{i(n/M)} w_k. \quad (2.105)$$

In this way the total estimate can be obtained as the weighted average of the blocks,

$$\hat{O} = \frac{\sum_{i=1}^M W_i \hat{O}_i}{\sum_{i=1}^M W_i}. \quad (2.106)$$

Then according to the usual block average procedure we estimate the unbiased variance between the blocks, that in this case is a weighted variance:

$$\sigma_O^2 = \frac{M_{\text{eff}}}{M_{\text{eff}} - 1} \frac{\sum_{i=1}^M W_i (\hat{O}_i - \hat{O})^2}{\sum_{i=1}^M W_i}, \quad (2.107)$$

where instead of the total number of blocks M , we use the effective block size $M_{\text{eff}} < M$:

$$M_{\text{eff}} = \frac{\left(\sum_{i=1}^M W_i\right)^2}{\sum_{i=1}^M W_i^2}, \quad (2.108)$$

that is the same as Eq. (2.90). The statistical error on the \hat{O} estimate is then $\sigma_O / \sqrt{M_{\text{eff}}}$. When the number of blocks M is small, or when the weight W_i are unbalanced, using M instead of M_{eff} can introduce a considerable underestimate of the real uncertainty.

The usual block average procedure can then be carried out, repeating the analysis using different number of blocks M and looking for a plateau in the error estimate.

C. Improving Robustness for the Multumbrella Target

In some cases, when targeting the multumbrella ensemble it can be important to introduce two small modifications, to make the iterative optimization scheme more effective. We did not use these modifications in any of the examples presented, but we found them useful for dealing with more challenging systems and thus have been implemented in the code.

When the iterative scheme starts, the first guess of the $\Delta F_n(s_\lambda)$ comes from just one single point, and is thus very inaccurate for CV values far away from the visited one. In particular it tends to become extremely large, which causes the bias to be very strong in pushing the system to the farthest s_λ -value. This initial bias is stronger the smaller the σ , and might even cause the simulation to fail during the very first biased steps. To avoid this, we can limit the initial value of the $\Delta F_n(s_\lambda)$ estimates to be always smaller than a given value, thus $\Delta F_0(s_\lambda) \leq \Delta E$. This ΔE value can be set to be equal to an estimate of the free-energy barrier that has to be overcome. Thus, similarly to Ref. [96], we add an extra optional parameter called “barrier”, that sets the value of ΔE . This barrier guess does not have to be perfect and a very rough estimate typically suffices. Also, we did not observe significant change in the convergence speed, thus we suggest to use this extra parameter only in case of an initial failure of the simulation, due to a too strong initial bias. The barrier parameter can in principle be used also with types of expansion other than the multumbrella one, even though in those cases it might be less useful.

The second modification, comes from the observation that, contrary to the previously considered expanded ensembles, the multumbrella one is not guaranteed to sample the full unbiased distribution $P_0(\mathbf{x})$. This can be a problem for the iterative scheme, because all the $\Delta F_n(s_\lambda)$ use as reference the unbiased free energy $F_0 = -\log Z_0$, whose estimate can vary substantially if P_0 is not properly sampled. Typically it should be easy to choose an interval Δs that contains all the s values relevant for P_0 , but if this is not the case then the problem can be fixed by simply adding P_0 itself to the target distribution. To add P_0 to the target, it is sufficient to add an extra expansion CV $\Delta u_0(\mathbf{x}) \equiv 0$, that always returns zero.

2.4.10 Supplemental Material

Alanine Dipeptide

The simulations for the alanine dipeptide example are performed with GRO-MACS [37] 2018.6 patched with a custom version of PLUMED [38] 2.6, that is available on the PLUMED-NEST website [91]. The setup is the same of Refs. [80, 96], thus canonical ensemble (*NVT*) in a vacuum, Amber99-SB force field [66], time step 2 fs, temperature $T_0 = 300$ K, velocity rescaling




Figure 2.36: Trajectories for the potential energy U and the dihedral angle ϕ of alanine dipeptide for two different OPES runs. The difference between the two is the target distribution: (a) is the multicanonical expanded one, while (b) is the well-tempered one with respect to the ϕ and ψ angles.

thermostat [36]. The only input needed to perform the multicanonical simulation presented in the main text was the temperature range, $T_{\min} = 300$ K and $T_{\max} = 1000$ K, and the pace at which updating the bias potential, which was set to 500 simulation steps, thus 1 ps. For this simple system the bias converges so quickly that we did not have to drop any initial transient before performing reweighting, e.g. in producing Fig. 2.29 of the main text.

We show in Fig. 2.36 the comparison between two alanine dipeptide simulations at 300 K where OPES has been used for sampling two different target distributions. On the right, Fig. 2.36a, we use the multicanonical expanded target presented in Sec. 2.4.5 of the main text, while on the left, Fig. 2.36b, the well-tempered over the dihedral angles is used, with same parameters as in Ref. [96]. While the two sampled ensembles are very different they can be both properly reweighted to obtain for instance the free energy difference between the two metastable ΔF_{AB} at 300 K. While for calculating ΔF_{AB} at a single temperature it is more convenient to directly bias ϕ and ψ , this requires some knowledge of the system, and in general it is far from trivial to identify good CVs, thus a more generic approach such as the multicanonical one can be useful.

We notice that the sampling obtained with our multicanonical simulation is very different from the one that can be achieved by biasing the energy with metadynamics, as proposed in Refs. [100, 106]. This is due to the very different target distributions they aim to sample. Biasing the energy in metadynamics enhances the energy fluctuations without seeking a specific temperature range. For instance, by using the well-tempered ensemble [100] in the above alanine example, one would sample also lower energies (associated to lower temperatures) and it would not be trivial to decide how




Figure 2.37: Difference in using only 3 sub-temperatures, or 1000. (a) shows the energy histogram of the sampled ensemble, while (b) shows the relative effective sample size as a function of the temperature. Once the intermediate steps are enough to ensure good overlap, adding more does not lead to a better target distribution.

much bias to deposit in order to sample all the relevant energies for the given temperature range.

Next in Fig. 2.37 we show how the multicanonical target distribution for alanine dipeptide changes when varying the number $N_{\{\beta\}}$ of intermediate temperatures. By default the number of intermediate temperatures is determined automatically via a short (100 ps in this case) unbiased simulation (see Sec. 2.4.5). In the considered alanine example this suggest the use of just 3 temperatures, equispaced in β : $T_1 = 300$ K, $T_2 = 461.538$ K, $T_3 = 1000$ K. As discussed in Sec. 2.4.7 of the main text adding more intermediate temperatures might not bring better sampling. We can see this by plotting the effective sample size n_{eff} relative to the total sample size n . We see that by using $N_{\{\beta\}} = 3$ we can cover in a more uniform way the temperature interval, compared to the case of $N_{\{\beta\}} = 100$. This suggest that the optimal target distribution is not the one obtained by substituting the sum with an integral.

Finally, we notice that figures identical to Figs. 2.36a and 2.37 can be obtained by running parallel tempering simulations and combining all the trajectories into one, provided that the same intermediate temperatures are used.

Double-well Model

The double-well Langevin model is the same use in Refs. [80, 96], and is implemented in PLUMED. The 2D potential energy is defined by rotating of an angle $\theta = -0.6\frac{\pi}{4}$ the following potential:

$$u(x, y) = x^4 + y^4 - 2x^2 - 4y^2 + xy + 0.3x + 0.1y. \quad (2.109)$$




Figure 2.38: The $\Delta F(x_\lambda)$ of the double-well model, for different choices of σ . The black dotted line is the reference free energy surface. The case $\sigma = 0.1 \sigma_0$, $N_{\{\lambda\}} = 270$, is not shown because it perfectly matches the reference line.

Fig. 2.38 shows how the estimated $\Delta F(x_\lambda)$ changes when varying the width of the umbrellas σ . The corresponding target distribution along x are shown in Fig. 2.33(b) of the main text. The estimate of the free energy surface $F(x)$ obtained from reweighting is instead independent from the choice of σ .

Fig. 2.39 presents a comparison between OPES using different kind of target distributions, namely the multumbrella one and the well-tempered one, and metadynamics. In a one dimensional problem the multumbrellas target can be slightly more efficient than the well-tempered one. However, OPES with a well-tempered target is more versatile, since it can be used also without knowing in advance the CV range, and can be scaled up to higher dimensions more efficiently, as shown in Ref. [96].

Sodium

The sodium simulations were performed in the same setting as Ref. [118]. We show the sampled space in Figs. 2.40 and 2.41.




Figure 2.39: The estimate of the free energy difference ΔF_{AB} between the two basins of the model, obtained with different methods, by averaging 10 independent runs. The run obtained using OPES with the multumbrella target uses $\sigma = \sigma_0 = 0.185815$. The results for OPES with a well-tempered target and metadynamics are taken from Ref. [96].




Figure 2.40: Sampled sodium configurations in the energy-volume space, (a) at fixed temperature and pressure, biasing the crystallinity CV with a multiumbrella target, and (b) multithermal, multibaric, and multiumbrella simulation, the same used in Fig. 2.34 in the main text.




Figure 2.41: Sampled sodium configurations for the multithermal-multibaric-multiumbrella simulation in the energy-volume-crystallinity space. (a) Colored accordingly to the crystallinity CV s . (b) Colored accordingly to the value of the bias $v(U, V, s)$. The hourglass shape described in Ref. [118] is clearly visible.

Conclusion and future works

3.1 Conclusion

In this thesis I have presented OPES, a novel computational method for performing enhanced sampling molecular simulations. The method works by estimating on-the-fly the ratio between the underlying probability distribution and the chosen target distribution, and using its logarithm as bias potential. The estimate is obtained via reweighting and is updated one sample at the time, so that after a brief initial transient the bias changes in an adiabatic fashion. The OPES method aims at making enhanced sampling simpler to use, with few robust input parameters and a straightforward reweighting scheme. Furthermore, it has been developed with suboptimal collective variables in mind, to efficiently deal with them. The inspiration for OPES has come from the experience with metadynamics and VES, and the method inherits many of their traits, while improving on some of their weaknesses. OPES is a collective variable based method, but it has been shown that it is capable also of performing expanded ensemble simulations, thus opening a novel unified perspective on enhanced sampling.

The first two articles included in this thesis show the preliminary development of the ideas that are at the basis of the OPES method, while the last two articles introduce the actual OPES method, applying it to two different families of target distributions. On the one hand, we have distributions defined by their projection into a smaller CV space, similar to those sampled by adaptive umbrella sampling or metadynamics. On the other hand instead, are expanded target distributions obtained by summing distributions that differ only by a small amount, as it is the case for temperature tempering distributions, or more in general distributions sampled by the replica exchange method.

3.2 Future works

3.2.1 Target distributions

In future works I would like to further extend OPES by applying it to sample other types of target distributions. One interesting possibility is to target uniform-entropy ensembles, similar to those sampled by nested sampling [111]. Another type of target distribution that could find practical applications is the one sampled by bias-exchange or parallel-bias metadynamics [50, 73], which allows for a more efficient scaling of the sampled volume with the CVs dimensionality, compared to the well-tempered distribution.

Yet another exciting possibility would be to use machine learning for designing novel efficient target distributions, instead of limiting its usage to the search of collective variables. This is interesting because we have evidence that even without modifying the CVs, one can achieve better performances simply by properly changing the target distribution. As an example, when the potential energy is used as CV, much better sampling can be obtained by targeting the expanded multicanonical ensemble rather than the uniform one, as discussed in Sec. 2.4. Such machine learning algorithms should aim at target distributions that increase the probability of sampling transition states, while maximizing the effective sample size of the reweighting to the original distribution.

3.2.2 Exploration

Finally, there is another OPES variant that I have been developing and will be presented in a future article.¹ To understand the need for this new variant, we must first distinguish between suboptimal CVs and degenerate CVs. A degenerate CV is a suboptimal CV that not only ignores some slow mode of the system, but it also maps different physical states into the same CV-space region. An exemplification of a degenerate CV is presented in Fig. 3 of Ref. [52]. A typical metadynamics simulation when using a degenerate CV will present multiple transitions between the identified metastable states, but the bias will be extremely slow to converge, and any estimate, such as the one of the free-energy difference between the states, will have both a significant statistical error and an uncontrolled systematic error. If instead one uses such degenerate CV with the OPES method as presented in Sec. 2.3, typically the initial coarse estimate of the free energy surface will be wrong, and thus after only one or few transitions, the simulation will likely remain stuck in the metastable state that is mislabeled by the CV. This can be a great advantage when developing a new CV, because it gives a clear indication that the chosen CV is degenerate and should be improved. However, it is

¹A development version of the code is already available in the PLUMED library www.plumed.org/doc-master/user-doc/html/_o_p_e_s__m_e_t_a_d__e_x_p_l_o_r_e.html.

not always possible to improve a CV, and it can be useful to have a method that is robust with respect to degenerate CVs, thus still capable of exploring the free energy surface and give a rough estimate of it even in this scenario.

In order to obtain a method that is less sensitive to degenerate CVs, one can simply base the OPES procedure on the estimate of the sampled distribution, instead of the reweighted one. This is easy to do when the target is the well-tempered distribution $p^{\text{WT}}(\mathbf{s}) \propto [P(\mathbf{s})]^{1/\gamma}$, that can be estimated as in Eq. (2.60) but with uniform weights, $w_k = 1$, thus

$$\tilde{p}_n^{\text{WT}}(\mathbf{s}) = \frac{1}{n} \sum_k^n G(\mathbf{s}, \mathbf{s}_k). \quad (3.1)$$

The bias, Eq. (2.63), can then be written as

$$V_n(\mathbf{s}) = (\gamma - 1) \frac{1}{\beta} \log \left(\frac{\tilde{p}_n^{\text{WT}}(\mathbf{s})}{Z_n} + \varepsilon \right). \quad (3.2)$$

By making just this simple change, we obtain a method that is better suited for exploring the CV space even in case of a degenerate CV. It can be shown that also this OPES variant correctly converges to the well-tempered distribution, and it can be reweighted with the same simple procedure, after discarding the initial non-adiabatic part of the simulation.

This novel variant of OPES has a convergence behaviour more similar to the one of well-tempered metadynamics rather than the standard OPES scheme presented in Sec. 2.3. However, it retains some of the good features of OPES, in particular it requires fewer input parameters, scales better to higher dimensional CVs and has a straightforward reweighing scheme. I believe that this new OPES variant can become a handy tool for quickly exploring the free energy surface of new systems, making enhanced sampling even easier to use also for non-experts.

Other articles

This chapter contains all the other articles that I co-authored during my PhD and were not included in the thesis. Rather than listing them in chronological order, here I group them according to their subject.

First there is a paper based on a brief presentation of the OPES method that I submitted to the 2020 National Congress of the Italian Physical Society and that was selected for publication.

- **M. Invernizzi**. "OPES: On-the-fly Probability Enhanced Sampling Method." submitted to *Il Nuovo Cimento C* (2021) preprint on arXiv. URL <https://arxiv.org/abs/2101.06991>

The successive article uses the variationally enhanced sampling method and explores a novel kind of target probability distribution.

- J. Debnath, **M. Invernizzi**, and M. Parrinello. "Enhanced sampling of transition states." *Journal of chemical theory and computation* 15.4 (2019): 2454-2459. URL <https://pubs.acs.org/doi/abs/10.1021/acs.jctc.8b01283>

The papers that follow next deal with the phenomenon of crystallization. My contribution to them is related to the work on the density Fourier transform as a collective variable, reported in Sec. 2.1.

- H. Niu, P. M. Piaggi, **M. Invernizzi**, and M. Parrinello. "Molecular dynamics simulations of liquid silica crystallization." *Proceedings of the National Academy of Sciences* 115.21 (2018): 5348-5352. URL <https://www.pnas.org/content/115/21/5348>
- P. Ahlawat, H. Lu, A. Ummadisingu, H. Niu, **M. Invernizzi**, S. M. Zakeeruddin, A. Hagfeldt, M. Graetzel, U. Rothlisberger, and M. Parrinello. "Molecular dynamics simulations of two-step process enable room-temperature synthesis of α -FAPbI₃." submitted to *Science Advances*

(2020) preprint on ChemRxiv. URL <https://doi.org/10.26434/chemrxiv.13084265.v1>

- T. Karmakat, **M. Invernizzi**, V. Rizzi, and M. Parrinello. "Collective Variables for the Study of Crystallization." submitted to *Molecular Physics* (2021) preprint on arXiv. URL <https://arxiv.org/abs/2101.03150>

The following papers instead, are about path integral simulations of fermionic particles. My contribution is related to applying reweighting and other techniques typical of enhanced sampling in order to ease the fermion sign problem.

- B. Hirshberg, **M. Invernizzi**, and M. Parrinello. "Path Integral Molecular Dynamics for Fermions: Alleviating the Sign Problem with the Bogoliubov Inequality." *The Journal of Chemical Physics* 152.17 (2020): 171102. URL <https://aip.scitation.org/doi/10.1063/5.0008720>
- T. Dornheim, **M. Invernizzi**, J. Vorberger, and B. Hirshberg. "Attenuating the fermion sign problem in path integral Monte Carlo simulations using the Bogoliubov inequality and thermodynamic integration." *The Journal of Chemical Physics* 153.23 (2020): 234104. URL <https://aip.scitation.org/doi/10.1063/5.0030760>

Finally, I am part of the PLUMED consortium (www.plumed-nest.org) that promotes open science and reproducibility of molecular dynamics protocols. I contributed by sharing all the inputs of the simulations in my papers, and by implementing into the open-source library PLUMED (www.plumed.org) the methods I developed.

- M. Bonomi, G. Bussi, C. Camilloni, G.A. Tribello *et al.* "Promoting transparency and reproducibility in enhanced molecular simulations." *Nature methods* 16.8 (2019): 670-673. URL <https://www.nature.com/articles/s41592-019-0506-8>
- Implementation of the method presented in Sec. 2.2. URL www.plumed.org/doc-master/user-doc/html/_v_e_s__d_e_l_t_a__f.html
- Contributed module, where OPES is implemented. URL www.plumed.org/doc-master/user-doc/html/_o_p_e_s.html

Bibliography

- [1] G. M. Torrie and J. P. Valleau, "Monte Carlo free energy estimates using non-Boltzmann sampling: Application to the sub-critical Lennard-Jones fluid," *Chemical Physics Letters* **28**, 578–581 (1974).
- [2] G. M. Torrie and J. Valleau, "Nonphysical sampling distributions in Monte Carlo free-energy estimation: Umbrella sampling," *Journal of Computational Physics* **23**, 187–199 (1977).
- [3] A. Laio and M. Parrinello, "Escaping free-energy minima," *Proceedings of the National Academy of Sciences* **99**, 12562–12566 (2002).
- [4] H. Sidky, W. Chen, and A. L. Ferguson, "Machine learning for collective variable discovery and enhanced sampling in biomolecular simulation," *Molecular Physics* **118**, e1737742 (2020).
- [5] O. Valsson and M. Parrinello, "Variational Approach to Enhanced Sampling and Free Energy Calculations," *Physical Review Letters* **113**, 090601 (2014).
- [6] Y. Sugita and Y. Okamoto, "Replica-exchange molecular dynamics method for protein folding," *Chemical Physics Letters* **314**, 141–151 (1999).
- [7] L. D. Landau, "On the theory of phase transitions," in *Collected Papers of L.D. Landau*, edited by D. Ter Haar (Pergamon, 1965) pp. 193 – 216.
- [8] L. D. Landau and E. M. Lifshitz, *Statistical Physics*, Vol. 5 (Elsevier Science, 2013).
- [9] W. G. Noid, J. W. Chu, G. S. Ayton, V. Krishna, S. Izvekov, G. A. Voth, A. Das, and H. C. Andersen, "The multiscale coarse-graining method. I. A rigorous bridge between atomistic and coarse-grained models," *Journal of Chemical Physics* **128** (2008), 10.1063/1.2938860.

- [10] M. G. Saunders and G. A. Voth, "Coarse-Graining Methods for Computational Biology," *Annual Review of Biophysics* **42**, 73–93 (2013).
- [11] M. S. Shell, "The relative entropy is fundamental to multiscale and inverse thermodynamic problems," *Journal of Chemical Physics* **129**, 1–7 (2008).
- [12] A. Chaimovich and M. S. Shell, "Coarse-graining errors and numerical optimization using a relative entropy framework," *Journal of Chemical Physics* **134** (2011).
- [13] K. Kaski, K. Binder, and J. D. Gunton, "A study of a coarse-grained free energy functional for the three-dimensional Ising model," *Journal of Physics A: Mathematical and General* **16**, L623 (1983).
- [14] A. P. Giddy, M. T. Dove, and V. Heine, "What do Landau free energies really look like for structural phase transitions?" *Journal of Physics: Condensed Matter* **1**, 8327–8335 (1989).
- [15] S. Radescu, I. Etxebarria, and J. M. Perez-Mato, "The Landau free energy of the three-dimensional Φ^4 model in wide temperature intervals," *Journal of Physics: Condensed Matter* **7**, 585–595 (1995).
- [16] M. E. Gracheva, J. M. Rickman, and J. D. Gunton, "Coarse-grained Ginzburg-Landau free energy for Lennard-Jones systems," *Journal of Chemical Physics* **113**, 3525–3529 (2000).
- [17] J. Iñiguez, S. Ivantchev, J. M. Perez-Mato, and A. García, "Devonshire-Landau free energy of BaTiO_3 from first principles," *Physical Review B* **63**, 144103–10 (2001).
- [18] A. Tröster, C. Dellago, and W. Schranz, "Free energies of the φ^4 model from Wang-Landau simulations," *Physical Review B - Condensed Matter and Materials Physics* **72**, 1–11 (2005).
- [19] A. Tröster and C. Dellago, "Coarse Graining the φ^4 Model: Landau-Ginzburg Potentials from Computer Simulations," *Ferroelectrics* **354**, 225–237 (2007).
- [20] O. Valsson, P. Tiwary, and M. Parrinello, "Enhancing Important Fluctuations: Rare Events and Metadynamics from a Conceptual Viewpoint," *Annual Review of Physical Chemistry* **67**, 159–184 (2016).
- [21] O. Valsson and M. Parrinello, "Well-tempered variational approach to enhanced sampling," *Journal of Chemical Theory and Computation* **11**, 1996–2002 (2015).
- [22] J. McCarty, O. Valsson, P. Tiwary, and M. Parrinello, "Variationally Optimized Free-Energy Flooding for Rate Calculation," *Physical Review Letters* **115**, 1–5 (2015).

-
- [23] J. McCarty, O. Valsson, and M. Parrinello, "Bespoke Bias for Obtaining Free Energy Differences within Variationally Enhanced Sampling," *Journal of Chemical Theory and Computation* **12**, 2162–2169 (2016).
- [24] P. Shaffer, O. Valsson, and M. Parrinello, "Enhanced, targeted sampling of high-dimensional free-energy landscapes using variationally enhanced sampling, with an application to chignolin," *Proceedings of the National Academy of Sciences* **113**, 1150–1155 (2016).
- [25] P. Shaffer, O. Valsson, and M. Parrinello, "Hierarchical Protein Free Energy Landscapes from Variationally Enhanced Sampling," *Journal of Chemical Theory and Computation* **12**, 5751–5757 (2016).
- [26] P. M. Piaggi, O. Valsson, and M. Parrinello, "A variational approach to nucleation simulation," *Faraday Discussions* **195**, 557–568 (2016).
- [27] R. Rubinstein, "The Cross-Entropy Method for Combinatorial and Continuous Optimization," *Methodology And Computing In Applied Probability* **1**, 127–190 (1999).
- [28] I. Billionis and P. Koutsourelakis, "Free energy computations by minimization of Kullback–Leibler divergence: An efficient adaptive biasing potential method for sparse representations," *Journal of Computational Physics* **231**, 3849–3870 (2012).
- [29] W. Zhang, H. Wang, C. Hartmann, M. Weber, and C. Schütte, "Applications of the cross-entropy method to importance sampling and optimal control of diffusions," *SIAM Journal on Scientific Computing* **36**, A2654–A2672 (2014).
- [30] a. Isihara, "The Gibbs-Bogoliubov inequality dagger," *Journal of Physics A: General Physics* **1**, 305 (1968).
- [31] F. Bach and E. Moulines, "Non-strongly-convex smooth stochastic approximation with convergence rate $O(1/n)$," in *Advances in Neural Information Processing Systems 26*, edited by C. J. C. Burges, L. Bottou, Z. G. Welling, and K. Q. Weinberger (Curran Associates, Inc., 2013) pp. 773–781.
- [32] K. Binder, B. J. Block, P. Virnau, and A. Tröster, "Beyond the Van Der Waals loop: What can be learned from simulating Lennard-Jones fluids inside the region of phase coexistence," *American Journal of Physics* **80**, 1099–1109 (2012).
- [33] N. B. Wilding, "Critical-point and coexistence-curve properties of the Lennard-Jones fluid: A finite-size scaling study," *Physical Review E* **52**, 602–611 (1995).
- [34] B. Smit, "Phase diagrams of Lennard-Jones fluids," *The Journal of Chemical Physics* **96**, 8639–8640 (1992).

- [35] J. Pérez-Pellitero, P. Ungerer, G. Orkoulas, and A. D. Mackie, "Critical point estimation of the Lennard-Jones pure fluid and binary mixtures," *Journal of Chemical Physics* **125** (2006).
- [36] G. Bussi, D. Donadio, and M. Parrinello, "Canonical sampling through velocity rescaling," *The Journal of Chemical Physics* **126**, 014101 (2007).
- [37] M. J. Abraham, T. Murtola, R. Schulz, S. Páll, J. C. Smith, B. Hess, and E. Lindahl, "GROMACS: High performance molecular simulations through multi-level parallelism from laptops to supercomputers," *SoftwareX* **1-2**, 19–25 (2015).
- [38] G. A. Tribello, M. Bonomi, D. Branduardi, C. Camilloni, and G. Bussi, "PLUMED 2: New feathers for an old bird," *Computer Physics Communications* **185**, 604–613 (2014).
- [39] K. Binder, "Finite size scaling analysis of ising model block distribution functions," *Zeitschrift für Physik B Condensed Matter* **43**, 119–140 (1981).
- [40] D. P. Landau and K. Binder, *A Guide to Monte Carlo Simulations in Statistical Physics* (Cambridge University Press, 2014).
- [41] P. Raiteri, A. Laio, F. L. Gervasio, C. Micheletti, and M. Parrinello, "Efficient Reconstruction of Complex Free Energy Landscapes by Multiple Walkers Metadynamics," *The Journal of Physical Chemistry B* **110**, 3533–3539 (2006).
- [42] C. Chipot and A. Pohorille, *Free Energy Calculations*, Springer Series in Chemical Physics, Vol. 86 (Springer Berlin Heidelberg, Berlin, Heidelberg, 2007).
- [43] C. Abrams and G. Bussi, "Enhanced sampling in molecular dynamics using metadynamics, replica-exchange, and temperature-acceleration," *Entropy* **16**, 163–199 (2014).
- [44] O. Valsson and M. Parrinello, "Variationally Enhanced Sampling," in *Handbook of Materials Modeling* (Springer International Publishing, Cham, 2018) pp. 1–14.
- [45] D. Branduardi, F. L. Gervasio, and M. Parrinello, "From A to B in free energy space," *The Journal of Chemical Physics* **126**, 054103 (2007).
- [46] F. Pietrucci and W. Andreoni, "Graph Theory Meets Ab Initio Molecular Dynamics: Atomic Structures and Transformations at the Nanoscale," *Physical Review Letters* **107**, 085504 (2011).
- [47] P. Tiwary and B. J. Berne, "Spectral gap optimization of order parameters for sampling complex molecular systems," *Proceedings of the National Academy of Sciences* **113**, 2839–2844 (2016).

-
- [48] J. McCarty and M. Parrinello, "A variational conformational dynamics approach to the selection of collective variables in metadynamics," *The Journal of Chemical Physics* **147**, 204109 (2017).
- [49] D. Mendels, G. Piccini, and M. Parrinello, "Collective Variables from Local Fluctuations," *Journal of Physical Chemistry Letters* **9**, 2776–2781 (2018).
- [50] J. Pfandtner and M. Bonomi, "Efficient Sampling of High-Dimensional Free-Energy Landscapes with Parallel Bias Metadynamics," *Journal of Chemical Theory and Computation* **11**, 5062–5067 (2015).
- [51] G. Bussi and D. Branduardi, "Free-Energy Calculations with Metadynamics: Theory and Practice," in *Reviews in Computational Chemistry Volume 28* (John Wiley & Sons, Inc, 2015) pp. 1–49.
- [52] F. Pietrucci, "Strategies for the exploration of free energy landscapes: Unity in diversity and challenges ahead," *Reviews in Physics* **2**, 32–45 (2017).
- [53] G. Bussi, A. Laio, and P. Tiwary, "Metadynamics: A Unified Framework for Accelerating Rare Events and Sampling Thermodynamics and Kinetics," in *Handbook of Materials Modeling* (Springer International Publishing, Cham, 2018) pp. 1–31.
- [54] J. F. Dama, M. Parrinello, and G. A. Voth, "Well-tempered metadynamics converges asymptotically." *Physical Review Letters* **112**, 240602 (2014).
- [55] A. Barducci, G. Bussi, and M. Parrinello, "Well-Tempered Metadynamics: A Smoothly Converging and Tunable Free-Energy Method," *Physical Review Letters* **100**, 020603 (2008).
- [56] M. Invernizzi, O. Valsson, and M. Parrinello, "Coarse graining from variationally enhanced sampling applied to the Ginzburg–Landau model," *Proceedings of the National Academy of Sciences* **114**, 3370–3374 (2017).
- [57] J. Duchi, E. Hazan, and Y. Singer, "Adaptive subgradient methods for online learning and stochastic optimization," in *The Journal of Machine Learning Research*, Vol. 12 (2011) pp. 2121–2159.
- [58] P. Tiwary and M. Parrinello, "A time-independent free energy estimator for metadynamics," *Journal of Physical Chemistry B* **119**, 736–742 (2015).
- [59] Y. I. Yang and M. Parrinello, "Refining Collective Coordinates and Improving Free Energy Representation in Variational Enhanced Sam-

- pling," *Journal of Chemical Theory and Computation* **14**, 2889–2894 (2018).
- [60] J. F. Dama, G. Rotskoff, M. Parrinello, and G. A. Voth, "Transition-tempered metadynamics: Robust, convergent metadynamics via on-the-fly transition barrier estimation," *Journal of Chemical Theory and Computation* **10**, 3626–3633 (2014).
- [61] S. Plimpton, "Fast Parallel Algorithms for Short-Range Molecular Dynamics," *Journal of Computational Physics* **117**, 1–19 (1995).
- [62] H. Niu, P. M. Piaggi, M. Invernizzi, and M. Parrinello, "Molecular dynamics simulations of liquid silica crystallization," *Proceedings of the National Academy of Sciences* **115**, 5348–5352 (2018).
- [63] S. Kumar, J. M. Rosenberg, D. Bouzida, R. H. Swendsen, and P. A. Kollman, "The weighted histogram analysis method for free-energy calculations on biomolecules. I. The method," *Journal of Computational Chemistry* **13**, 1011–1021 (1992).
- [64] F. Zhu and G. Hummer, "Convergence and error estimation in free energy calculations using the weighted histogram analysis method," *Journal of Computational Chemistry* **33**, 453–465 (2012).
- [65] D. Branduardi, G. Bussi, and M. Parrinello, "Metadynamics with Adaptive Gaussians," *Journal of Chemical Theory and Computation* **8**, 2247–2254 (2012).
- [66] V. Hornak, R. Abel, A. Okur, B. Strockbine, A. Roitberg, and C. Simmerling, "Comparison of multiple Amber force fields and development of improved protein backbone parameters," *Proteins: Structure, Function, and Bioinformatics* **65**, 712–725 (2006).
- [67] S. R. Wilson, K. G. S. H. Gunawardana, and M. I. Mendeleev, "Solid-liquid interface free energies of pure bcc metals and B2 phases," *The Journal of Chemical Physics* **142**, 134705 (2015).
- [68] M. Parrinello and A. Rahman, "Polymorphic transitions in single crystals: A new molecular dynamics method," *Journal of Applied Physics* **52**, 7182–7190 (1981).
- [69] D. Mandelli, B. Hirshberg, and M. Parrinello, "Metadynamics of Paths," *Physical Review Letters* **125**, 026001 (2020).
- [70] L. Bonati, V. Rizzi, and M. Parrinello, "Data-Driven Collective Variables for Enhanced Sampling," *The Journal of Physical Chemistry Letters* **11**, 2998–3004 (2020).
- [71] V. Rizzi, L. Bonati, N. Ansari, and M. Parrinello, "The role of water in host-guest interaction," *Nature Communications* **12**, 93 (2021).

-
- [72] B. Peters, *Reaction rate theory and rare events* (Elsevier, 2017) p. 619.
- [73] S. Piana and A. Laio, "A Bias-Exchange Approach to Protein Folding," *The Journal of Physical Chemistry B* **111**, 4553–4559 (2007).
- [74] A. D. White, J. F. Dama, and G. A. Voth, "Designing Free Energy Surfaces That Match Experimental Data with Metadynamics," *Journal of Chemical Theory and Computation* **11**, 2451–2460 (2015).
- [75] M. Bonomi, A. Barducci, and M. Parrinello, "Reconstructing the equilibrium Boltzmann distribution from well-tempered metadynamics," *Journal of Computational Chemistry* **30**, 1615–1621 (2009).
- [76] L. Mones, N. Bernstein, and G. Csányi, "Exploration, Sampling, And Reconstruction of Free Energy Surfaces with Gaussian Process Regression," *Journal of Chemical Theory and Computation* **12**, 5100–5110 (2016).
- [77] L. Donati and B. G. Keller, "Girsanov reweighting for metadynamics simulations," *The Journal of Chemical Physics* **149**, 072335 (2018).
- [78] V. Marinova and M. Salvalaglio, "Time-independent free energies from metadynamics via mean force integration," *The Journal of Chemical Physics* **151**, 164115 (2019).
- [79] F. Giberti, B. Cheng, G. A. Tribello, and M. Ceriotti, "Iterative Unbiasing of Quasi-Equilibrium Sampling," *Journal of Chemical Theory and Computation* **16**, 100–107 (2020).
- [80] M. Invernizzi and M. Parrinello, "Making the Best of a Bad Situation: A Multiscale Approach to Free Energy Calculation," *Journal of Chemical Theory and Computation* **15**, 2187–2194 (2019).
- [81] M. Mezei, "Adaptive umbrella sampling: Self-consistent determination of the non-Boltzmann bias," *Journal of Computational Physics* **68**, 237–248 (1987).
- [82] P. Maragakis, A. van der Vaart, and M. Karplus, "Gaussian-Mixture Umbrella Sampling," *The Journal of Physical Chemistry B* **113**, 4664–4673 (2009).
- [83] S. Marsili, A. Barducci, R. Chelli, P. Procacci, and V. Schettino, "Self-healing Umbrella Sampling: A Non-equilibrium Approach for Quantitative Free Energy Calculations," *The Journal of Physical Chemistry B* **110**, 14011–14013 (2006).
- [84] B. M. Dickson, F. Legoll, T. Lelièvre, G. Stoltz, and P. Fleurat-Lessard, "Free Energy Calculations: An Efficient Adaptive Biasing Potential Method," *The Journal of Physical Chemistry B* **114**, 5823–5830 (2010).

- [85] J. Debnath, M. Invernizzi, and M. Parrinello, "Enhanced Sampling of Transition States," *Journal of Chemical Theory and Computation* **15**, 2454–2459 (2019).
- [86] P. M. Piaggi and M. Parrinello, "Multithermal-Multibarc Molecular Simulations from a Variational Principle," *Physical Review Letters* **122**, 050601 (2019).
- [87] B. Silverman, *Density Estimation for Statistics and Data Analysis* (Routledge, New York, 1998).
- [88] L. Kish, "Sampling Organizations and Groups of Unequal Sizes," *American Sociological Review* **30**, 564 (1965).
- [89] D. Sodkomkham, D. Ciliberti, M. A. Wilson, K.-I. Fukui, K. Moriyama, M. Numao, and F. Kloosterman, "Kernel density compression for real-time Bayesian encoding/decoding of unsorted hippocampal spikes," *Knowledge-Based Systems* **94**, 1–12 (2016).
- [90] G. Fort, B. Jourdain, T. Lelièvre, and G. Stoltz, "Self-healing umbrella sampling: convergence and efficiency," *Statistics and Computing* **27**, 147–168 (2017).
- [91] The PLUMED consortium, "Promoting transparency and reproducibility in enhanced molecular simulations," *Nature Methods* **16**, 670–673 (2019).
- [92] T. Lelièvre, M. Rousset, and G. Stoltz, *Free Energy Computations* (Imperial College Press, 2010).
- [93] L. Bonati, Y.-Y. Zhang, and M. Parrinello, "Neural networks-based variationally enhanced sampling," *Proceedings of the National Academy of Sciences*, 201907975 (2019).
- [94] D. J. Earl and M. W. Deem, "Parallel tempering: Theory, applications, and new perspectives," *Physical Chemistry Chemical Physics* **7**, 3910 (2005).
- [95] A. P. Lyubartsev, A. A. Martsinovski, S. V. Shevkunov, and P. N. Vorontsov-Velyaminov, "New approach to Monte Carlo calculation of the free energy: Method of expanded ensembles," *The Journal of Chemical Physics* **96**, 1776–1783 (1992).
- [96] M. Invernizzi and M. Parrinello, "Rethinking Metadynamics: From Bias Potentials to Probability Distributions," *The Journal of Physical Chemistry Letters* **11**, 2731–2736 (2020).
- [97] R. H. Swendsen and J.-S. Wang, "Replica Monte Carlo Simulation of Spin-Glasses," *Physical Review Letters* **57**, 2607–2609 (1986).

-
- [98] E. Marinari and G. Parisi, "Simulated Tempering: A New Monte Carlo Scheme," *Europhysics Letters (EPL)* **19**, 451–458 (1992).
- [99] G. Bussi, F. L. Gervasio, A. Laio, and M. Parrinello, "Free-Energy Landscape for β Hairpin Folding from Combined Parallel Tempering and Metadynamics," *Journal of the American Chemical Society* **128**, 13435–13441 (2006).
- [100] M. Bonomi and M. Parrinello, "Enhanced Sampling in the Well-Tempered Ensemble," *Physical Review Letters* **104**, 190601 (2010).
- [101] M. Nava, R. Quhe, F. Palazzesi, P. Tiwary, and M. Parrinello, "de Broglie Swapping Metadynamics for Quantum and Classical Sampling," *Journal of Chemical Theory and Computation* **11**, 5114–5119 (2015).
- [102] Y. I. Yang, Q. Shao, J. Zhang, L. Yang, and Y. Q. Gao, "Enhanced sampling in molecular dynamics," *The Journal of Chemical Physics* **151**, 070902 (2019).
- [103] L. Rosso, P. Mináry, Z. Zhu, and M. E. Tuckerman, "On the use of the adiabatic molecular dynamics technique in the calculation of free energy profiles," *The Journal of Chemical Physics* **116**, 4389–4402 (2002).
- [104] E. Darve and A. Pohorille, "Calculating free energies using average force," *The Journal of Chemical Physics* **115**, 9169–9183 (2001).
- [105] Y. Sugita, A. Kitao, and Y. Okamoto, "Multidimensional replica-exchange method for free-energy calculations," *The Journal of Chemical Physics* **113**, 6042–6051 (2000).
- [106] C. Micheletti, A. Laio, and M. Parrinello, "Reconstructing the Density of States by History-Dependent Metadynamics," *Physical Review Letters* **92**, 170601 (2004).
- [107] F. Wang and D. P. Landau, "Efficient, Multiple-Range Random Walk Algorithm to Calculate the Density of States," *Physical Review Letters* **86**, 2050–2053 (2001).
- [108] B. A. Berg and T. Neuhaus, "Multicanonical algorithms for first order phase transitions," *Physics Letters B* **267**, 249–253 (1991).
- [109] J. Lee, "New Monte Carlo algorithm: Entropic sampling," *Physical Review Letters* **71**, 211–214 (1993).
- [110] B. Hesselbo and R. B. Stinchcombe, "Monte Carlo Simulation and Global Optimization without Parameters," *Physical Review Letters* **74**, 2151–2155 (1995).

- [111] L. B. Pártay, A. P. Bartók, and G. Csányi, "Efficient Sampling of Atomic Configurational Spaces," *The Journal of Physical Chemistry B* **114**, 10502–10512 (2010).
- [112] P. Liu, B. Kim, R. A. Friesner, and B. J. Berne, "Replica exchange with solute tempering: A method for sampling biological systems in explicit water," *Proceedings of the National Academy of Sciences* **102**, 13749–13754 (2005).
- [113] G. Bussi, "Hamiltonian replica exchange in GROMACS: a flexible implementation," *Molecular Physics* **112**, 379–384 (2014).
- [114] Y. Q. Gao, "An integrate-over-temperature approach for enhanced sampling," *The Journal of Chemical Physics* **128**, 064105 (2008).
- [115] H. Nymeyer, "Serial tempering without exchange," *The Journal of Chemical Physics* **133**, 114113 (2010).
- [116] J. Kästner, "Umbrella sampling," *Wiley Interdisciplinary Reviews: Computational Molecular Science* **1**, 932–942 (2011).
- [117] C. Vega, E. Sanz, J. L. F. Abascal, and E. G. Noya, "Determination of phase diagrams via computer simulation: methodology and applications to water, electrolytes and proteins," *Journal of Physics: Condensed Matter* **20**, 153101 (2008).
- [118] P. M. Piaggi and M. Parrinello, "Calculation of phase diagrams in the multithermal-multibaric ensemble," *The Journal of Chemical Physics* **150**, 244119 (2019).
- [119] D. Reith, M. Pütz, and F. Müller-Plathe, "Deriving effective mesoscale potentials from atomistic simulations," *Journal of Computational Chemistry* **24**, 1624–1636 (2003).
- [120] A. Martinsson, J. Lu, B. Leimkuhler, and E. Vanden-Eijnden, "The simulated tempering method in the infinite switch limit with adaptive weight learning," *Journal of Statistical Mechanics: Theory and Experiment* **2019**, 013207 (2019).
- [121] M. Reinhardt and H. Grubmüller, "Determining Free Energy Differences Through Variationally-Derived Intermediates," *Journal of Chemical Theory and Computation* **116**, acs.jctc.0c00106 (2020).
- [122] R. W. Zwanzig, "High-Temperature Equation of State by a Perturbation Method. I. Nonpolar Gases," *The Journal of Chemical Physics* **22**, 1420–1426 (1954).
- [123] S. Park, "Comparison of the serial and parallel algorithms of generalized ensemble simulations: An analytical approach," *Physical Review E* **77**, 016709 (2008).

- [124] B. K. Radak, D. Suh, and B. Roux, "A generalized linear response framework for expanded ensemble and replica exchange simulations," *The Journal of Chemical Physics* **149**, 072315 (2018).
- [125] A. Kong, J. S. Liu, and W. H. Wong, "Sequential Imputations and Bayesian Missing Data Problems," *Journal of the American Statistical Association* **89**, 278–288 (1994).
- [126] G. Bussi and G. A. Tribello, "Analyzing and Biasing Simulations with PLUMED," in *Biomolecular Simulations: Methods and Protocols*, edited by M. Bonomi and C. Camilloni (Springer New York, New York, NY, 2019) pp. 529–578.
- [127] M. J. Hartmann, Y. Singh, E. Vanden-Eijnden, and G. M. Hocky, "Infinite switch simulated tempering in force (FISST)," *The Journal of Chemical Physics* **152**, 244120 (2020).
- [128] M. K. Prakash, A. Barducci, and M. Parrinello, "Replica Temperatures for Uniform Exchange and Efficient Roundtrip Times in Explicit Solvent Parallel Tempering Simulations," *Journal of Chemical Theory and Computation* **7**, 2025–2027 (2011).
- [129] S. Piana, K. Lindorff-Larsen, and D. E. Shaw, "How Robust Are Protein Folding Simulations with Respect to Force Field Parameterization?" *Biophysical Journal* **100**, L47–L49 (2011).
- [130] S. Honda, T. Akiba, Y. S. Kato, Y. Sawada, M. Sekijima, M. Ishimura, A. Ooishi, H. Watanabe, T. Odahara, and K. Harata, "Crystal Structure of a Ten-Amino Acid Protein," *Journal of the American Chemical Society* **130**, 15327–15331 (2008).
- [131] K. Lindorff-Larsen, P. Maragakis, S. Piana, M. P. Eastwood, R. O. Dror, and D. E. Shaw, "Systematic Validation of Protein Force Fields against Experimental Data," *PLoS ONE* **7**, e32131 (2012).
- [132] H. Okumura, "Temperature and pressure denaturation of chignolin: Folding and unfolding simulation by multibaric-multithermal molecular dynamics method," *Proteins: Structure, Function, and Bioinformatics* **80**, 2397–2416 (2012).
- [133] B. Uralcan and P. G. Debenedetti, "Computational Investigation of the Effect of Pressure on Protein Stability," *The Journal of Physical Chemistry Letters* **10**, 1894–1899 (2019).
- [134] D. Frenkel and B. Smit, *Understanding Molecular Simulation: from algorithms to applications* (Academic Press, New York, 2001) p. 664.
- [135] X. Kong and C. L. Brooks, " λ -dynamics: A new approach to free energy calculations," *The Journal of Chemical Physics* **105**, 2414–2423 (1996).

- [136] R. Bitetti-Putzer, W. Yang, and M. Karplus, "Generalized ensembles serve to improve the convergence of free energy simulations," *Chemical Physics Letters* **377**, 633–641 (2003).
- [137] J. K. Johnson, J. A. Zollweg, and K. E. Gubbins, "The Lennard-Jones equation of state revisited," *Molecular Physics* **78**, 591–618 (1993).
- [138] J. Debnath and M. Parrinello, "Gaussian Mixture-Based Enhanced Sampling for Statics and Dynamics," *The Journal of Physical Chemistry Letters* **11**, 5076–5080 (2020).
- [139] Y. I. Yang, H. Niu, and M. Parrinello, "Combining Metadynamics and Integrated Tempering Sampling," *The Journal of Physical Chemistry Letters* **9**, 6426–6430 (2018).
- [140] H. Niu, Y. I. Yang, and M. Parrinello, "Temperature Dependence of Homogeneous Nucleation in Ice," *Physical Review Letters* **122**, 245501 (2019).
- [141] D. F. Hahn, G. König, and P. H. Hünenberger, "Overcoming Orthogonal Barriers in Alchemical Free Energy Calculations: On the Relative Merits of λ -Variations, λ -Extrapolations, and Biasing," *Journal of Chemical Theory and Computation* **16**, 1630–1645 (2020).
- [142] H. Niu, L. Bonati, P. M. Piaggi, and M. Parrinello, "Ab initio phase diagram and nucleation of gallium," *Nature Communications* **11**, 2654 (2020).
- [143] P. M. Piaggi and R. Car, "Phase equilibrium of liquid water and hexagonal ice from enhanced sampling molecular dynamics simulations," *The Journal of Chemical Physics* **152**, 204116 (2020).
- [144] V. Lindahl, J. Lidmar, and B. Hess, "Riemann metric approach to optimal sampling of multidimensional free-energy landscapes," *Physical Review E* **98**, 023312 (2018).
- [145] H. Flyvbjerg and H. G. Petersen, "Error estimates on averages of correlated data," *The Journal of Chemical Physics* **91**, 461–466 (1989).
- [146] B. Hirshberg, M. Invernizzi, and M. Parrinello, "Path integral molecular dynamics for fermions: Alleviating the sign problem with the Bogoliubov inequality," *The Journal of Chemical Physics* **152**, 171102 (2020).
- [147] T. Dornheim, M. Invernizzi, J. Vorberger, and B. Hirshberg, "Attenuating the fermion sign problem in path integral Monte Carlo simulations using the Bogoliubov inequality and thermodynamic integration," *The Journal of Chemical Physics* **153**, 234104 (2020).
- [148] M. Invernizzi, P. M. Piaggi, and M. Parrinello, "Unified Approach to Enhanced Sampling," *Physical Review X* **10**, 41034 (2020).

Acknowledgements

Many are the people that have been close to me during these wonderful years in Lugano, and I regret not being able to properly address each one of them in these few lines.

I am grateful to professor Parrinello for the opportunities he has given me and for all the things he taught me about science and research, through the many hours we spent together discussing. I thank the various colleagues that during these years made my PhD extremely lively and intellectually stimulating. I would also like to acknowledge the organizers of MARVEL national center of competence in research, founded by the Swiss National Science Foundation, that sponsored my PhD and gave me several opportunities to meet great computational scientists from other Swiss universities.

Curriculum vitae

Personal data

BIRTH DATE April 3rd, 1990
BIRTH PLACE Cernusco sul Naviglio (Milano), Italy
NATIONALITY Italian
ORCID 0000-0002-3328-6557
GOOGLE SCHOLAR Michele Invernizzi

Education

May '16 - Dec '20 **Doctorate in Physics**
ETH Zurich c/o Università della Svizzera Italiana,
Lugano, Switzerland
THESIS: On-the-fly Probability Enhanced Sampling
SUPERVISOR: Prof. M. Parrinello

Mar '13 - Apr '16 **Master Degree in Physics**
Università degli Studi di Milano, Milano, Italy
FINAL GRADE: 110/110 cum laude
THESIS: Microscopic calculation of Landau's
free-energy parameters
SUPERVISORS: Prof. M. Parrinello, Prof. D. Galli

Sep '09 - Feb '13 **Bachelor Degree in Physics**
Università degli Studi di Milano, Milano, Italy
FINAL GRADE: 110/110 cum laude
THESIS: Depinning and structural transitions in a
2D model of colloids
SUPERVISORS: Prof. N. Manini, Prof. A. Vanossi

- Aug '07 - Jun '08 **Exchange student** (1 year)
Christopher Columbus High School, Miami, USA
- Sep '04 - Jun '09 **Humanistic High School Diploma**
Liceo Classico Fondazione Sacro Cuore, Milano, Italy
FINAL GRADE: 100/100
DISSERTATION: Scientific knowledge – retracing
Husserl's analysis

Teaching

- Fall 2016 Teaching assistant, *Programming Techniques for Scientific Simulations*, Prof. Matthias Troyer,
Institute for Theoretical Physics, ETH Zurich

Awards

- 2020 Early Postdoc Mobility fellowship, by the Swiss
National Science Foundation
- 2018 Giovanni Polvani award for young researchers,
by the Italian Physical Society
- 2016 Poster prize, by Paris International School on
Advanced Computational Materials Science
- 2016 Award for Master Degree by BCC Cernusco
- 2013 Award for Bachelor Degree by BCC Cernusco
- 2010/'11/'12 Prizes for physics undergraduate students
by Università degli Studi di Milano
- 2009 Award for High School Degree by BCC Cernusco