JCTC Journal of Chemical Theory and Computation

pubs.acs.org/JCTC Article

Exploration vs Convergence Speed in Adaptive-Bias Enhanced Sampling

Michele Invernizzi* and Michele Parrinello

Cite This: J. Chem. Theory Comput. 2022, 18, 3988−3996 Read Online

ACCESS | Metrics & More | Article Recommendations | Supporting Information

ABSTRACT: In adaptive-bias enhanced sampling methods, a bias potential is added to the system to drive transitions between metastable states. The bias potential is a function of a few collective variables and is gradually modified according to the underlying free energy surface. We show that when the collective variables are suboptimal, there is an exploration−convergence tradeoff, and one must choose between a quickly converging bias that will lead to fewer transitions or a slower to converge bias that can explore the phase space more efficiently but might require a much longer time to produce an accurate free energy estimate. The recently proposed on-the-fly probability enhanced sampling (OPES) method focuses on fast convergence, but there are cases where fast exploration is preferred instead. For this reason, we introduce a new variant of the OPES method that focuses on quickly escaping metastable states at the expense of convergence speed. We illustrate the benefits of this approach in prototypical systems and show that it outperforms the popular metadynamics method.

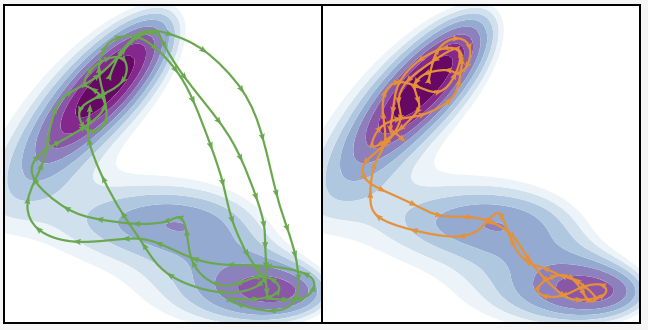

1. INTRODUCTION

Molecular dynamics has become a valuable tool in the study of a variety of phenomena in physics, chemistry, biology, and materials science. One of the long-standing challenges in this important field is the sampling of rare events, such as chemical reactions or conformational changes in biomolecules. To simulate effectively such systems, many enhanced sampling methods have been developed. An important class of such methods is based on an adaptive-bias approach and includes adaptive umbrella sampling,[1] metadynamics (MetaD),[2,3] and the recently developed on-the-fly probability enhanced sampling (OPES).[4−6] Adaptive-bias methods operate by adding to the system's energy $U(\mathbf{R})$ an external bias potential $V = V(\mathbf{s})$, which is a function of a set of collective variables (CVs), $\mathbf{s}$. The CVs, $\mathbf{s} = \mathbf{s}(\mathbf{R})$, depend on the atomic coordinates $\mathbf{R}$ and are meant to describe the slow modes associated with the rare event under study. They also define a free energy surface (FES), $F(\mathbf{s}) = -1/\beta \log P(\mathbf{s})$, where $\beta = (k_B T)^{-1}$ is the inverse Boltzmann factor and $P(\mathbf{s})$ the marginal $\mathbf{s}$ distribution, $P(\mathbf{s}) \propto \int e^{-\beta U(\mathbf{R})} \delta[\mathbf{s} - \mathbf{s}(\mathbf{R})] d\mathbf{R}$. The bias is periodically updated until it converges to a chosen form. A popular choice is to have it exactly offset the underlying FES, $V(\mathbf{s}) = -F(\mathbf{s})$, so that the resulting $\mathbf{s}$ distribution is uniform.

The main limitation of adaptive-bias methods is that finding good collective variables is sometimes difficult and a bad choice of CVs might not promote the desired transitions in an affordable computer time. In practical applications, one generally has to live with suboptimal CVs[7] that still can drive transitions but do not include some of the slow modes. In this case, applying a static bias cannot speed up the slow modes that are not accounted for, and thus transitions remain quite infrequent. It is sometimes possible to achieve a faster transition rate using a rapidly changing bias, which can push the system out of a metastable state through a high free energy pathway, different from the energetically favored one. However, unless one wishes to deal explicitly with out-of-equilibrium statistics,[8−10] it is not possible to obtain reliable information about the system while the bias changes in a nonadiabatic fashion. To estimate the FES and other observables, one must let the adaptive-bias method approach convergence, and as the bias becomes quasi-static, transitions inevitably become less frequent.

We refer to this situation as an exploration−convergence tradeoff that every adaptive-bias enhanced sampling method has to deal with when suboptimal CVs are used. Some methods, like OPES, focus more on quickly converging to a quasi-static bias potential and thus obtaining an efficiently reweighted FES, while others, like metadynamics, focus more on escaping metastable states and exploring the phase space. We will demonstrate this qualitative difference in some

Received: February 14, 2022
Published: May 26, 2022

prototypical systems. For simplicity, in the paper, we only consider the well-tempered variant of metadynamics,[3] but in the Supporting Information, we provide examples that use the original nontempered MetaD[2] and other popular variants, such as parallel-bias MetaD.[11]

We propose here a variant of OPES, named OPES-explore, which focuses on rapid exploration, rather than on fast convergence. It shares many features with the original OPES, and is designed to be an easy-to-use tool requiring few input parameters. To this end, we also introduce an adaptive bandwidth algorithm that can be used in both OPES variants and further reduces the number of input parameters that need to be specified. The detailed description of the adaptive bandwidth algorithm is given in the Supporting Information. All OPES simulations presented make use of this algorithm.

2. OPES METHOD

The enhanced sampling method OPES works by adding an adaptive-bias potential to the energy of the system so as to modify the Boltzmann probability distribution into the desired target one. Most adaptive-bias methods aim at sampling uniformly the CV space, but it has been shown that choosing a different target distribution could be advantageous.[12,13] There are two different classes of target distributions that can be sampled with OPES: metadynamics-like and replica-exchange-like. We will consider here only the former type, introduced in ref 4, but the interested reader can find information about OPES for replica-exchange-like sampling in ref 5.

To define a metadynamics-like target distribution, one has to choose a set of collective variables, $\mathbf{s} = \mathbf{s}(\mathbf{R})$. As stated in the introduction, the unbiased marginal probability along such CVs is $P(\mathbf{s}) \propto \int e^{-\beta U(\mathbf{R})} \delta[\mathbf{s} - \mathbf{s}(\mathbf{R})] d\mathbf{R}$, where $U(\mathbf{R})$ is the potential energy. The target distribution is then defined by requiring a specific marginal probability distribution over the CVs, $p^{tg}(\mathbf{s})$. Consequently, the desired bias potential is written as

$$V(\mathbf{s}) = -\frac{1}{\beta}\log\frac{p^{tg}(\mathbf{s})}{P(\mathbf{s})} \tag{1}$$

so that $\int e^{-\beta[U(\mathbf{R}) + V(\mathbf{s})]} \delta[\mathbf{s} - \mathbf{s}(\mathbf{R})] d\mathbf{R} \propto p^{tg}(\mathbf{s})$. A typical choice for $p^{tg}(\mathbf{s})$ is the well-tempered distribution[3]

$$p^{WT}(\mathbf{s}) \propto [P(\mathbf{s})]^{1/\gamma} \tag{2}$$

where the bias factor $\gamma > 1$ controls how much the original distribution is smoothed out. In the limit of $\gamma = \infty$, one targets a uniform distribution.

The core idea of OPES is to update self-consistently the estimate of the probability distributions and of the bias potential in an on-the-fly fashion similar to self-healing umbrella sampling.[14] The estimate of the unbiased probability is obtained via a weighted kernel density estimation so that at step n, one has

$$P_n(\mathbf{s}) = \frac{\sum_k^n w_k G(\mathbf{s}, \mathbf{s}_k)}{\sum_k^n w_k} \tag{3}$$

where the weights w_k are given by $w_k = e^{\beta V_{k-1}(\mathbf{s}_k)}$, and the Gaussian kernels $G(\mathbf{s}, \mathbf{s}') = h \exp[-1/2(\mathbf{s} - \mathbf{s}')^T \Sigma^{-1} (\mathbf{s} - \mathbf{s}')]$ have a diagonal covariance matrix $\Sigma_{ij} = \sigma_i^2 \delta_{ij}$ and fixed height $h = \prod_i (\sigma_i \sqrt{2\pi})^{-1}$. The number of kernels to represent $P_n(\mathbf{s})$ would grow linearly with simulation time, but this is avoided thanks to an on-the-fly kernel compression algorithm,[15] as described in detail in the supporting information of ref 4. The compression algorithm also allows for the bandwidth of the kernels to shrink over time, as the effective sample size $N_{eff}^{(n)} = (\sum_k^n w_k)^2 / \sum_k^n w_k^2$ grows. The idea is to start with a coarse estimate of $P(\mathbf{s})$ and then refine it as more data are available. The kernel bandwidth of the i-th CV at step n is

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

where d is the total number of CVs.

The instantaneous bias is based on the probability estimate $P_n(\mathbf{s})$; following eq 1 and using the approximation $p^{WT}(\mathbf{s}) \propto [P_n(\mathbf{s})]^{1/\gamma}$, one has

$$V_n(\mathbf{s}) = (1 - 1/\gamma)\frac{1}{\beta}\log\left(\frac{P_n(\mathbf{s})}{Z_n} + \varepsilon\right) \tag{5}$$

where ϵ is a regularization term that limits the maximum possible absolute value of the bias potential and Z_n can be understood as normalization of $P_n(\mathbf{s})$ over the CV space thus far explored, Ω_n

$$Z_n = \frac{1}{|\Omega_n|}\int_{\Omega_n} P_n(\mathbf{s})\, d\mathbf{s} \tag{6}$$

The integral is calculated approximately as a sum of P_n over the compressed kernels, as explained in the supporting information of ref 4. The intuitive idea is that new kernels are added to the compressed representation only when a new region of CV space is sampled (otherwise they are merged with the existing ones), thus the explored CV-space volume $|\Omega_n|$ is approximately proportional to the total number of compressed kernels.

The introduction of the Z_n term is one of the key innovations of OPES. In similar methods, once a new metastable state is found, one often sees a dramatic increase in the exit time, compared to the first one[16] (see Supporting Information, Figure S8). This exit time problem is present also when the CVs are optimal and should not be confused with the exploration−convergence tradeoff that is the primary concern of this paper. Other convergence-focused methods introduce extra parameters to tackle this problem, for example in transition-tempered metadynamics[17] prior knowledge of the position of all metastable states is required. Instead, OPES avoids the exit time problem by taking into account the expansion of the CV space via the Z_n term, which allows the bias to adjust more quickly when a new CV-space region is sampled.[4]

At the start of an OPES simulation, only a handful of parameters need to be chosen, namely the initial kernel bandwidth, the pace at which the bias is updated, and the approximate FES barrier height that needs to be overcome. From this last information, a prescription is given to automatically set the values γ and ϵ. The number of parameters can be reduced even further if one uses, as we shall do here, the adaptive bandwidth algorithm discussed in the Supporting Information.

3. EXPLORATORY OPES VARIANT

We present now a new OPES variant called OPES-explore which, compared to the original OPES formulation, leads to a faster exploration of the phase space at the cost of slower convergence. We have recalled that the Z_n term allows OPES to quickly adapt the bias when a metastable state is found in a

https://doi.org/10.1021/acs.jctc.2c00152
J. Chem. Theory Comput. 2022, 18, 3988−3996

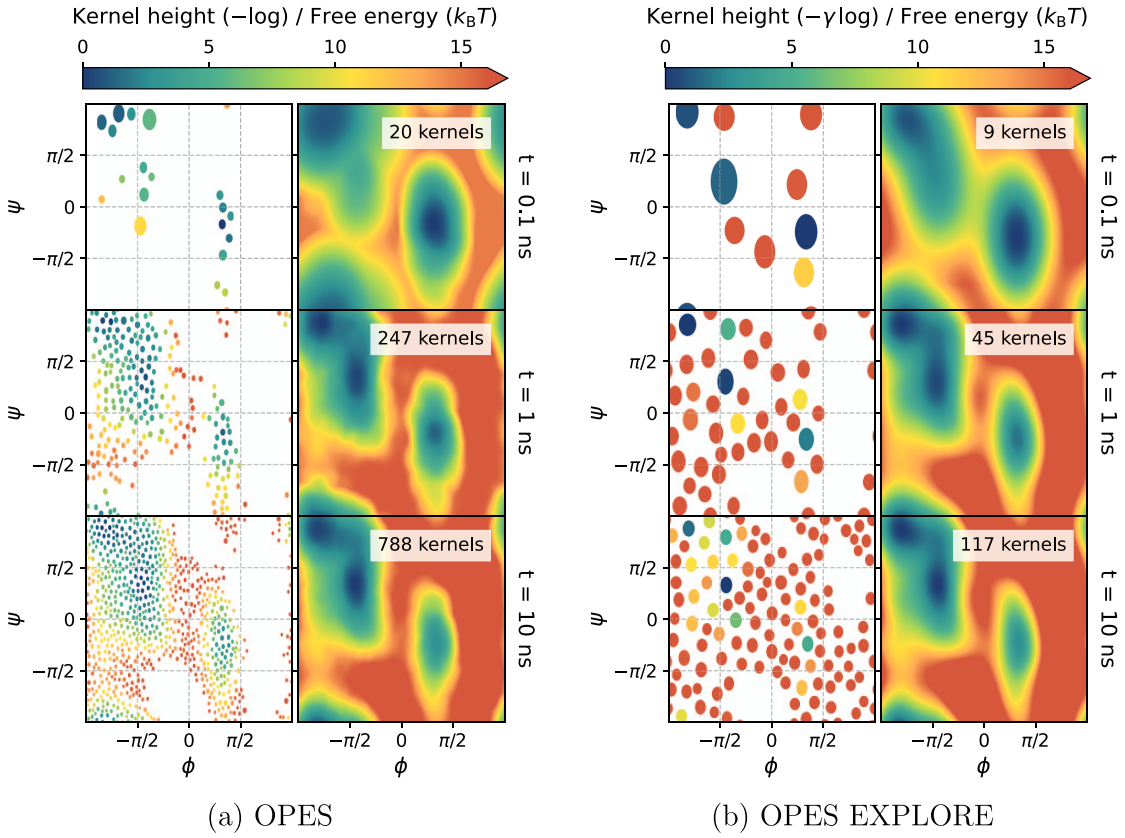

Figure 1. Time evolution of a typical simulation of alanine dipeptide in a vacuum using the two OPES variants with the dihedral angles ϕ and ψ as CVs. For each method, the compressed kernels are shown on the left with the point size indicating the adaptive bandwidth, and the corresponding free energy estimate $F_n(\phi, \psi)$ on the right. (a) In the original OPES, kernels make up the unbiased distribution estimate $P_n(\phi, \psi)$ and $F_n(\phi, \psi) = -1/\beta \log P_n(\phi, \psi)$, while (b) in OPES-explore, kernels make up the sampled distribution estimate $p_n^{WT}(\phi, \psi)$ and $F_n(\phi, \psi) = -\gamma\ 1/\beta \log p_n^{WT}(\phi, \psi)$. All $F_n(\phi, \psi)$ are shifted to have zero minimum. Notice how OPES-explore requires fewer kernels and visits higher FES regions.

previously unexplored region of CV space. However, if the CVs used are suboptimal, it may happen that a new metastable state is found in an already explored $\mathbf{s}$ region.[17,18] In such a case, the Z_n term remains constant and is therefore ineffective in accelerating the exit time. Instead, to encourage a rapid exit, one would need a method that allows the bias to significantly change shape again. Fortunately, it is possible to achieve this exploratory behavior simply by making a minimal change to the OPES protocol, which gives rise to the OPES-explore variant.

In formulating OPES-explore, we restrict ourselves to the case of using as a target the well-tempered distribution, $p^{tg}(\mathbf{s}) = p^{WT}(\mathbf{s})$, eq 2. In OPES, the bias is expressed as a function of $P_n(\mathbf{s})$, the on-the-fly estimate of the unknown equilibrium distribution $P(\mathbf{s})$. At the beginning of the simulation, this estimate is not reliable but it improves over time and converges in a self-consistent way. In OPES-explore instead, one builds the bias starting from the on-the-fly estimate of the distribution that is being sampled in the biased simulation

$$p_n^{WT}(\mathbf{s}) = \frac{1}{n}\sum_k^n G(\mathbf{s}, \mathbf{s}_k) \tag{7}$$

where $\mathbf{s}_k$ is the CV value sampled at step k. As the simulation converges, $p_n^{WT}(\mathbf{s})$ approaches the target well-tempered distribution $p^{WT}(\mathbf{s})$. Thus, analogously to Section 2, we use the approximation $P(\mathbf{s}) \propto [p_n^{WT}(\mathbf{s})]^\gamma$ and write the bias according to eq 1

$$V_n(\mathbf{s}) = (\gamma - 1)\frac{1}{\beta}\log\left(\frac{p_n^{WT}(\mathbf{s})}{Z_n} + \epsilon\right) \tag{8}$$

where ϵ and Z_n have been added for the same reasons as in eq 5. We notice that the expressions in eqs 3 and 7, which define the probability estimates used in the two OPES schemes, converge, respectively, to $P(\mathbf{s})$ and $p^{WT}(\mathbf{s})$ only within the self-consistent scheme where the simulation runs with a bias that is updated on-the-fly according to eqs 5 and 8, respectively. Both OPES variants are applications of the more general eq 1, but OPES estimates on-the-fly $P(\mathbf{s})$ and uses it to calculate the bias, while OPES-explore does the same but with $p^{WT}(\mathbf{s}) \propto [P(\mathbf{s})]^{1/\gamma}$.

The free energy surface as a function of the CVs can be estimated in two distinct ways, either directly from the probability estimate, $F_n(\mathbf{s}) = -\gamma\ 1/\beta \log p_n^{WT}(\mathbf{s})$, or via importance sampling reweighting, e.g., using a weighted kernel density estimation

$$F_n(\mathbf{s}) = -\frac{1}{\beta}\log\sum_k^n e^{\beta V_{k-1}(\mathbf{s}_k)} G(\mathbf{s}, \mathbf{s}_k) \tag{9}$$

In standard OPES, these two estimates are equivalent, while in OPES-explore (similarly to MetaD), they can differ significantly in the first part of the simulation until they eventually converge to the same estimate.

In Figure 1, we contrast an OPES and OPES-explore simulation of alanine dipeptide in a vacuum, which has become

https://doi.org/10.1021/acs.jctc.2c00152
J. Chem. Theory Comput. 2022, 18, 3988−3996

a standard test for enhanced sampling methods. Both simulations have the same input parameters and use the adaptive bandwidth scheme described in the Supporting Information. The bias is initially quite coarse, but the width of the kernels reduces as the simulation proceeds and the details of the FES are increasingly better described. It can clearly be seen that the OPES-explore variant employs fewer kernels compared to the original OPES. This is due to the fact that in OPES-explore, the kernel density estimation is used for $p^{\mathrm{WT}}(\mathbf{s}) \propto [P(\mathbf{s})]^{1/\gamma}$ that is a smoothed version of $P(\mathbf{s})$, and thus requires fewer details. This more compact representation can be useful, especially in higher dimensions, where the number of kernels can greatly increase despite the compression algorithm. However, as a drawback, it can result in a less accurate bias estimate, especially for large values of γ.

4. FEWER TRANSITIONS CAN LEAD TO BETTER CONVERGENCE

The difference in performance between OPES and OPES-explore cannot be judged from the alanine dipeptide example because in this case, the CVs chosen are extremely efficient. To highlight the difference between the two methods, we study a simple two-dimensional model potential that is known as the Müller potential,[20] see Figure 2a, using the x coordinate as a collective variable. This is a clear example of suboptimal CV since it can discriminate the metastable states but not the transition state.

For two-dimensional systems, the free energy along the CV, $F(x)$, can be computed precisely with numerical integration, Figure 2b. From $F(x)$, the free energy difference between the two metastable states can be calculated as

$$\Delta F = -\frac{1}{\beta} \log \frac{\int_0^1 e^{-\beta F(x)} dx}{\int_{-1.3}^0 e^{-\beta F(x)} dx} \tag{10}$$

While it is possible to distinguish better the two states using also the y coordinate, this does not result in a significant difference in the ΔF value (see Supporting Information, Sec. S3). On the other hand, x does a poor job of identifying the transition state, which is around $x \approx -0.7$ and $y \approx 0.6$, and not at $x \approx 0$ as it would seem from $F(x)$. As a consequence, it is not possible to significantly increase the transition rate between the states using a static bias that is a function of x only.

To show this, we consider the effect of adding to the system the converged well-tempered bias $V(x) = -(1 - 1/\gamma)F(x)$, with $\gamma = 20$. In Figure 2b, we can see the effect of the bias on the FES along x, which becomes almost completely flat. However, when we consider the full 2D landscape, as shown in Figure 2c, we can see that such a bias does not really remove the barrier between the two states. From the height of the barrier at the transition state, one can roughly estimate that adding $V(x)$ improves the transition rate by about one order of magnitude. Nevertheless, transitions remain quite rare, around one of every 10^6 uncorrelated samples (see Supporting Information, Sec. S3).

We want to compare the two OPES variants and well-tempered metadynamics in this challenging setting, where CVs are suboptimal and the total simulation time is not enough to reach full convergence. This type of situation is not uncommon in practical applications, and it is thus of great interest. Given enough time, all of the methods considered will converge to the same bias potential and sample the same target distribution, but we shall see that before reaching this limit, they behave very differently.

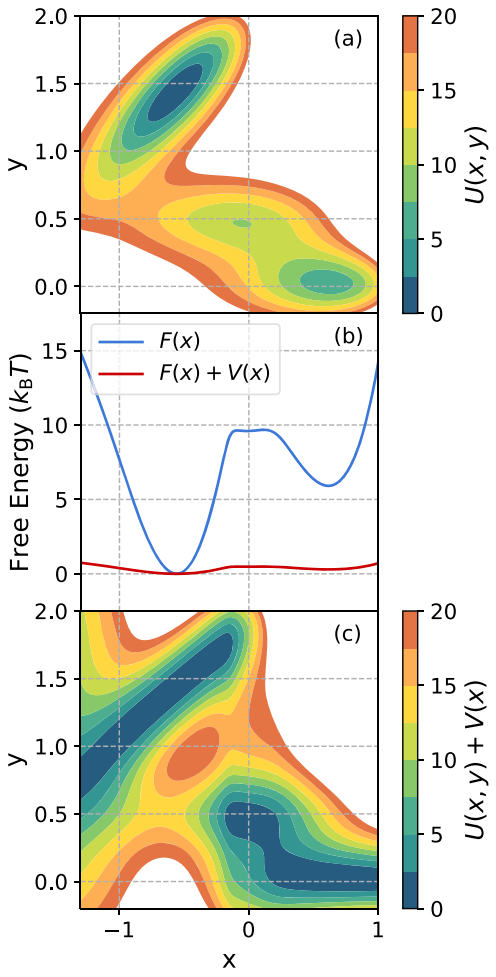

Figure 2. (a) Müller potential energy surface, $U(x, y)$. (b) Free energy surface along the x coordinate, $F(x)$, with and without the addition of the bias potential $V(x) = -(1 - 1/\gamma)F(x)$, where $\gamma = 20$. (c) Potential energy modified by the bias potential, $U(x, y) + V(x)$. All the energies reported are shifted to have zero minimum. It can be seen that despite the almost flat profile along x, the transition region between the states remains at high energy.

Figure 3a shows a typical run of the Müller potential obtained by biasing the x coordinate with OPES, OPES-explore, or MetaD. As a simple way to visualize the evolution of the bias, we also report in Figure 3b, the ΔF_n estimate obtained directly from the applied bias, using $F_n(x) = -(1 - 1/\gamma)^{-1}V_n(x)$ in eq 10. We can see a qualitative difference between OPES and the other two methods.

OPES reaches a quasi-static bias that is very close to the converged one but samples a distribution that is far from the well-tempered one, where the two basins would be about equally populated. On the other hand, the x distribution sampled by OPES-explore is closer to the target well-tempered one, but its bias is far from converged and makes ample oscillations around the correct value. Metadynamics behaves similarly to OPES-explore. This is the exploration−convergence tradeoff described in Section 1. Since the CV is suboptimal, even when using the converged bias $V(x)$ to see a transition occur, one has to wait for an average number of steps $\tau \approx 10^6$, which is more than the total length of the simulation.

https://doi.org/10.1021/acs.jctc.2c00152
J. Chem. Theory Comput. 2022, 18, 3988−3996

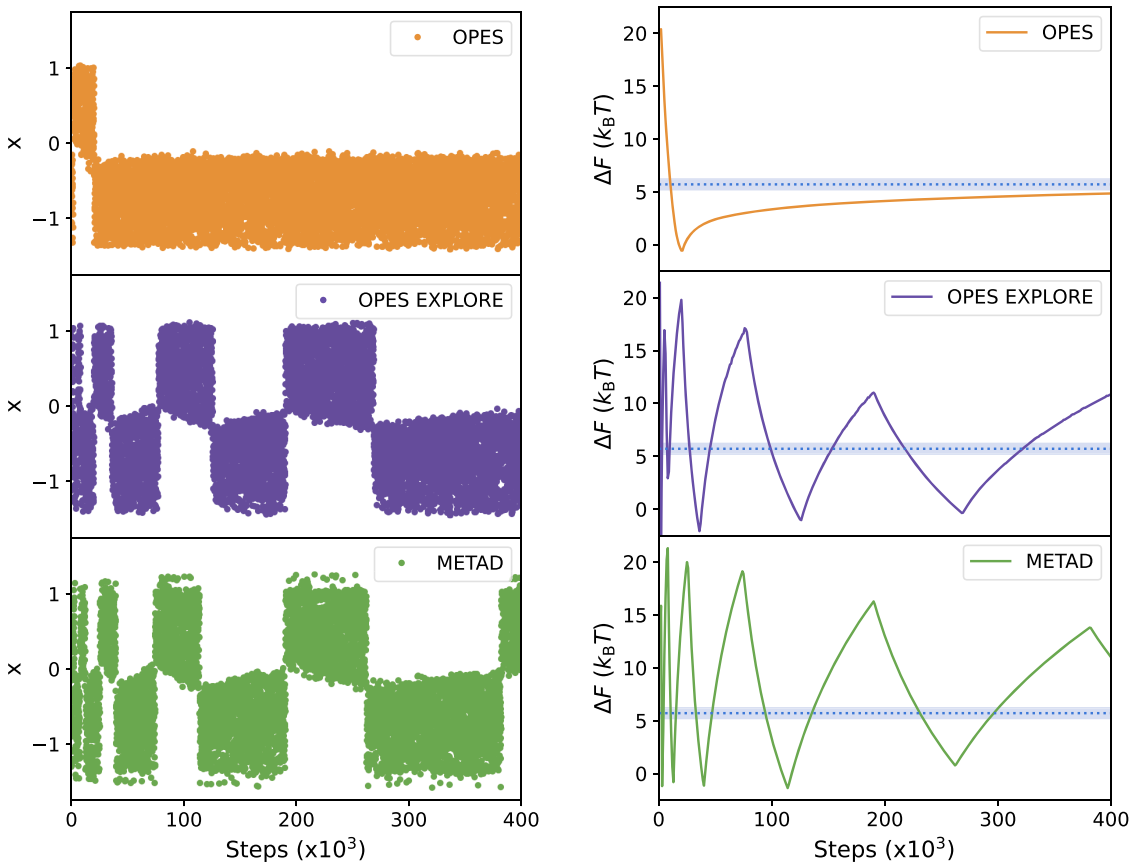

Figure 3. Typical simulations of the Müller potential using different methods for biasing the x coordinate. Given more time, the three methods will converge to the same bias potential and will sample the same target distribution. (a) Shows the trajectory along the CV and (b) shows the corresponding ΔF_n, eq 10, calculated using the FES estimate obtained directly from the applied bias, $F_n(x) = -(1 - 1/\gamma)^{-1}V_n(x)$. The correct ΔF value is highlighted by a blue stripe 1 k_BT thick.

However, it is possible to greatly accelerate transitions using a time-dependent bias that forces the system into higher energy pathways, which are not accessible at equilibrium.

In OPES-explore, the bias is based on the estimate of the sampled probability $p_n^{WT}(\mathbf{s})$ and pushes to make it similar to the almost flat well-tempered target. This means that to have a quasi-static bias, $p_n^{WT}(\mathbf{s})$ should both be almost flat and not change significantly as the simulation proceeds. Clearly, this cannot happen unless the simulation is longer than τ, otherwise, most of the time would be spent in the same basin and $p_n^{WT}(\mathbf{s})$ would be far from flat. On the contrary, in OPES, the bias is based on the reweighted estimate $P_n(\mathbf{s})$, and thus it can reach a quasi-static regime even before properly sampling the target distribution.

In Figure 4a, we show the ΔF_n estimate averaged over 25 independent runs, all starting from the main basin $x < 0$. We can see that on average, OPES provides the best ΔF_n estimate at any n in spite of the fact that it induces far fewer transitions. In fact, most of the time, only one full back-and-forth transition is observed (see Supporting Information, Figure S5a). One should notice that after a single transition, the ΔF_n estimate is far from being accurate (see Figure 3b) but, since the bias quickly becomes quasi-static, it is possible to collect equilibrium samples and reliably reweight them, and the average estimate becomes more accurate the more simulations are run. Instead, in OPES-explore and MetaD, despite starting from independent initial conditions, the runs are highly correlated due to the transitions being mostly driven by the strong changes in the bias rather than the natural fluctuations of the system. As a consequence of this, a systematic error is present in the average estimate, even if ΔF_n is further averaged over time, to remove the oscillatory behavior of OPES-explore and MetaD. Such a systematic error depends on the characteristic of the system and the chosen CVs and is hard to predict whether it will be relevant or small. Nevertheless, one can be sure that it reduces over time as the bias converges.[24]

Estimates of ΔF using different reweighting schemes are shown in Figure 4b. For OPES and OPES-explore, eq 9 has been used, while for MetaD, we consider two of the most popular reweighting schemes, namely last-bias reweighting[19,23] and bias-offset reweighting.[21,22] As expected, the reweighting estimate of OPES is virtually identical to the direct estimate obtained from the bias, while for the other two methods, the two estimates differ. The reweighting of OPES-explore has very small statistical uncertainty, which further highlights the presence of a systematic error in the free energy difference estimate. Like others before us,[22,25,26] we observe empirically that the last-bias reweighting for MetaD tends to be in agreement with the direct estimate, even when the simulation is far from converged, while the bias-offset reweighting provides a very unreliable estimate if the MetaD bias has not

https://doi.org/10.1021/acs.jctc.2c00152
J. Chem. Theory Comput. 2022, 18, 3988−3996

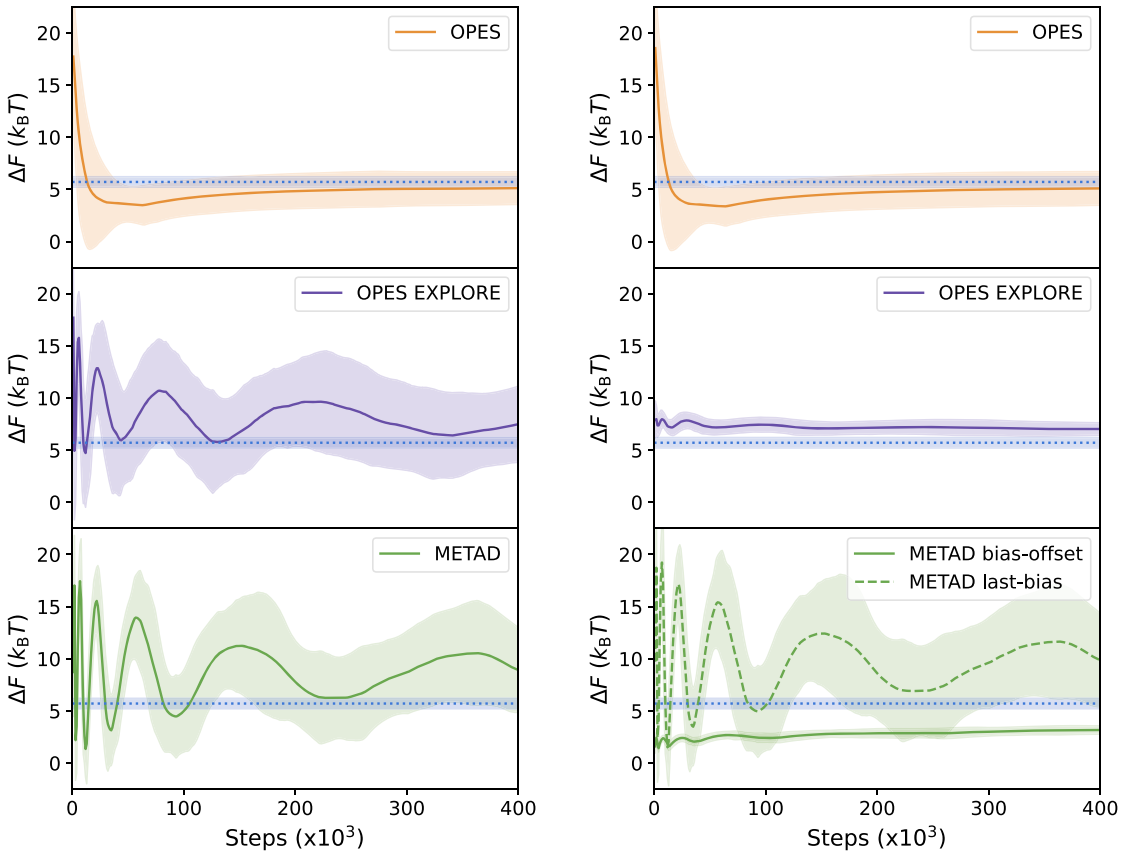

Figure 4. Estimate of the free energy difference ΔF for the Müller potential obtained by averaging 25 independent runs for each biasing method. The standard deviation is also shown for each estimate. Given more time, all of these estimates will converge to the correct ΔF. All simulations start from the main basin, $x < 0$, but with different initial conditions. (a) Shows the estimate obtained directly from the applied bias, as shown in Figure 3b, while (b) shows the corresponding estimate obtained via reweighting. For metadynamics, two different reweighting schemes are considered, bias-offset[21,22] and last-bias reweighting.[19,23] The correct ΔF value is highlighted by a blue stripe 1 k_BT thick.

reached a quasi-static regime and the initial part of the simulation is not discarded. Once again, it must be noted that the simulations considered here are not fully converged, otherwise all of the different estimates of the various methods would have yielded the correct result, without systematic errors. However, for most practical purposes, they behave very differently; thus, it is important to choose between an exploration-focused or a convergence-focused enhanced sampling method, depending on the specific aim of the simulation.

5. SOMETIMES EXPLORATION IS WHAT MATTERS

In the example of the previous section, it was shown that OPES converges to a quasi-static bias faster than OPES-explore and provides more accurate FES estimates. However, FES estimation is not the only goal of an enhanced sampling simulation. In complex systems where good CVs are not available, convergence can remain out of reach; still, one might be interested in exploring the phase space and finding all of the relevant metastable basins. In such a situation, OPES-explore can be a useful tool.

We consider here alanine tetrapeptide in a vacuum as a test system, as in ref 4. It has three ϕ dihedral angles, each of them can change from positive to negative values and vice versa with a relatively low probability. This leads to $2^3 = 8$ distinct metastable basins, each corresponding to a different combination of ϕ angles signs, as shown in Figure 5. Here, we are not interested in estimating the FES, but rather we want to compare the ability of different methods to explore this space and discover all metastable states.

Figure 6 shows the number of explored basins averaged over 10 independent simulations for each enhanced sampling method. The simulations in the top panel (Figure 6a) bias the ϕ angles, $V = V(\phi_1, \phi_2, \phi_3)$, which are good CVs, while in the bottom (Figure 6b), the suboptimal ψ angles are used, $V = V(\psi_1, \psi_2, \psi_3)$. In all methods, the exploration time increases approximately by two orders of magnitude when suboptimal CVs are used (please note the horizontal logarithmic scale). As expected, OPES and OPES-explore have similar exploration speeds when using good CVs, while with suboptimal CVs, OPES struggles to find all of the metastable basins. This is because the same region of CV space might correspond to two different metastable basins or to a basin and a transition state, as for the Müller potential.[18,19] In this situation, the previously estimated bias must change considerably for the simulation to escape the current metastable state quickly.

The exploration speed of MetaD depends critically on the input parameters and requires a trial-and-error tuning. We report here only the outcome of MetaD simulations in which a standard choice of the input parameters has been made. As can be seen in Figure 6, in these simulations, the exploration speed is roughly one order of magnitude slower than that of OPES-

https://doi.org/10.1021/acs.jctc.2c00152
J. Chem. Theory Comput. 2022, 18, 3988−3996

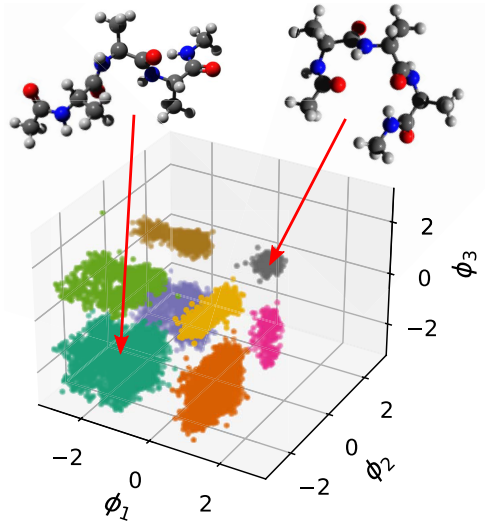

Figure 5. Eight metastable basins of alanine tetrapeptide in a vacuum sampled via OPES-explore by biasing the three ψ angles, a suboptimal set of CVs. Each basin is identified by the sign of the three ϕ angles for a total of 2^3 possible combinations. The most stable basin has ϕ_1, ϕ_2, $\phi_3 < 0$, while the least stable basin has ϕ_1, ϕ_2, $\phi_3 > 0$.

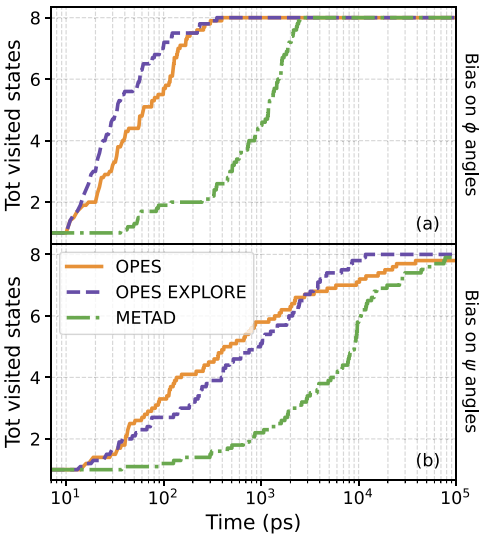

Figure 6. Exploration time of the eight metastable basins of alanine tetrapeptide over 100 ns. The lines are an average of over 10 independent runs for each method, showing the total number of visited basins. (a) Shows bias as a function of the three ϕ angles, $V = V(\phi_1, \phi_2, \phi_3)$, while (b) shows the three ψ angles are used, $V = V(\psi_1, \psi_2, \psi_3)$. See the Supporting Information for results with different input parameters and other MetaD variants, such as parallel-bias MetaD.[11]

explore. However, the performance of MetaD simulations can be improved using different settings, as shown in the Supporting Information. In the Supporting Information, we also report and briefly discuss results obtained with non-tempered metadynamics,[2] adaptive-Gaussians metadynamics,[23] and parallel-bias metadynamics.[11] None of these MetaD variants significantly improve the exploration speed, and some make it even worse.

Finally, in the Supporting Information, Sec. S5, we show how a preliminary OPES-explore run can be combined with a multithermal OPES simulation[5,6] to efficiently sample alanine tetrapeptide and reach a converged FES, even without explicitly biasing the ϕ angles.

6. CONCLUSIONS

With the help of model systems, we show that there is an exploration−convergence tradeoff in adaptive-bias methods when suboptimal CVs are used. This tradeoff should not be confused with the exit time problem, which is present also with optimal CVs, and is discussed in Section 2 and refs 16, 17. Contrary to the exit time problem, the exploration−convergence tradeoff cannot be solved. It is an intrinsic limitation of CV-based adaptive-bias methods, which comes as a consequence of suboptimal CVs. We believe that the best way to handle this tradeoff is to have separate methods that clearly focus on one or the other aspect so that they can be used depending on the application. In a convergence-focused method, the bias soon becomes quasi-static to allow for accurate reweighting and free energy estimation. However, with suboptimal CVs, this leads to a slow transition rate and a long time is required to sample the target distribution. As discussed, even if one knows the true $F(\mathbf{s})$ and directly applies the converged bias, one would not obtain a faster exploration. In an exploration-focused method, it is possible to improve the exploration speed by letting the bias change substantially even in a CV region that has already been visited. While this may increase the number of transitions, it comes at the cost of a slower convergence.

The original OPES method focuses on fast convergence to provide an accurate estimate of the free energy surface and reweighted observables. As a consequence, it is very sensitive to the quality of the CVs (see e.g., Figure 3a), and any improvement in the CVs results in a clear acceleration of the transition rate. This is a particularly useful property when developing machine learning-based CVs, and in fact, OPES has already been used several times in this context.[27−32]

In other situations, improving the CVs may first require a better exploration of the phase space.[28,33−35] Furthermore, one may be interested simply in exploring the metastable states of a system rather than estimating an accurate FES.[36−39] For this reason, we have introduced a variant of the OPES method, OPES-explore, which focuses on quickly sampling the target distribution and exploring the phase space.

We have shown that also well-tempered metadynamics is an exploration-focused method. One of the main advantages of OPES-explore over MetaD is that it is easier to use since it requires fewer input parameters and it has a more straightforward reweighting scheme (but more advanced ones can also be used[25,40]). Another important difference between the two methods is that OPES-explore, similarly to OPES, by default provides a maximum threshold to the applied bias potential, thus it avoids unreasonably high free energy regions. To obtain the same effect with MetaD, one typically has to define some *ad hoc* static bias walls by trial and error. This last feature of OPES-explore has been recently leveraged by Raucci et al. to systematically discover reaction pathways in chemical processes.[41]

Finally, we should clarify that OPES-explore, just as metadynamics, might not be able to exit any metastable state if the CVs are too poor,[19,42] and its improved exploration capability can only be harnessed if the CVs are close enough to the correct ones to make such transitions possible. The speed and the small number of input parameters of OPES-explore are

https://doi.org/10.1021/acs.jctc.2c00152
J. Chem. Theory Comput. 2022, 18, 3988−3996

extremely helpful for quickly testing several candidate CVs to find out which can drive transitions and discard the bad ones.

We believe that OPES-explore is an important addition to the OPES family of methods and will become a useful tool for researchers as it pushes forward the trend for more robust and reliable enhanced sampling methods.

■ ASSOCIATED CONTENT

Supporting Information

The Supporting Information is available free of charge at https://pubs.acs.org/doi/10.1021/acs.jctc.2c00152.

Description of the adaptive bandwidth algorithm, computational details regarding the Müller potential and further biased trajectories, exploration speed for alanine tetrapeptide using other methods, and description of a multithermal−multiumbrella simulation to improve upon the OPES-explore run (PDF)

■ AUTHOR INFORMATION

Corresponding Author

Michele Invernizzi − *Freie Universität Berlin, 14195 Berlin, Germany*; orcid.org/0000-0002-3328-6557; Email: michele.invernizzi@fu-berlin.de

Author

Michele Parrinello − *Italian Institute of Technology, 16163 Genova, Italy*

Complete contact information is available at: https://pubs.acs.org/10.1021/acs.jctc.2c00152

Notes

The authors declare no competing financial interest.

An open-source implementation of the OPES and OPES-explore methods is available in the enhanced sampling library PLUMED from version 2.8.[43] All of the data and input files needed to reproduce the simulations presented in this paper are available on PLUMED-NEST (www.plumed-nest.org), the public repository of the PLUMED consortium,[44] as plumID:22.003.

■ ACKNOWLEDGMENTS

The authors thank Valerio Rizzi and Umberto Raucci for their useful discussions. M.I. acknowledges support from the Swiss National Science Foundation through an Early Postdoc.Mobility fellowship. Calculations were carried out on the Euler cluster at ETH Zurich and on workstations provided by USI Lugano.

■ REFERENCES

(1) Mezei, M. Adaptive umbrella sampling: Self-consistent determination of the non-Boltzmann bias. *J. Comput. Phys.* **1987**, *68*, 237−248.

(2) Laio, A.; Parrinello, M. Escaping free-energy minima. *Proc. Natl. Acad. Sci. U.S.A.* **2002**, *99*, 12562−12566.

(3) Barducci, A.; Bussi, G.; Parrinello, M. Well-Tempered Metadynamics: A Smoothly Converging and Tunable Free-Energy Method. *Phys. Rev. Lett.* **2008**, *100*, No. 020603.

(4) Invernizzi, M.; Parrinello, M. Rethinking Metadynamics: From Bias Potentials to Probability Distributions. *J. Phys. Chem. Lett.* **2020**, *11*, 2731−2736.

(5) Invernizzi, M.; Piaggi, P. M.; Parrinello, M. Unified Approach to Enhanced Sampling. *Phys. Rev. X* **2020**, *10*, 41034.

(6) Invernizzi, M. OPES: On-the-fly Probability Enhanced Sampling Method. *Il Nuovo Cimento C* **2021**, *44*, 112.

(7) Invernizzi, M.; Parrinello, M. Making the Best of a Bad Situation: A Multiscale Approach to Free Energy Calculation. *J. Chem. Theory Comput.* **2019**, *15*, 2187−2194.

(8) Jarzynski, C. Nonequilibrium Equality for Free Energy Differences. *Phys. Rev. Lett.* **1997**, *78*, 2690−2693.

(9) Donati, L.; Keller, B. G. Girsanov reweighting for metadynamics simulations. *J. Chem. Phys.* **2018**, *149*, No. 072335.

(10) Bal, K. M. Reweighted Jarzynski Sampling: Acceleration of Rare Events and Free Energy Calculation with a Bias Potential Learned from Nonequilibrium Work. *J. Chem. Theory Comput.* **2021**, *17*, 6766−6774.

(11) Pfaendtner, J.; Bonomi, M. Efficient Sampling of High-Dimensional Free-Energy Landscapes with Parallel Bias Metadynamics. *J. Chem. Theory Comput.* **2015**, *11*, 5062−5067.

(12) Valsson, O.; Parrinello, M. Variational Approach to Enhanced Sampling and Free Energy Calculations. *Phys. Rev. Lett.* **2014**, *113*, No. 090601.

(13) White, A. D.; Dama, J. F.; Voth, G. A. Designing Free Energy Surfaces That Match Experimental Data with Metadynamics. *J. Chem. Theory Comput.* **2015**, *11*, 2451−2460.

(14) Marsili, S.; Barducci, A.; Chelli, R.; Procacci, P.; Schettino, V. Self-healing Umbrella Sampling: A Non-equilibrium Approach for Quantitative Free Energy Calculations. *J. Phys. Chem. B* **2006**, *110*, 14011−14013.

(15) Sodkomkham, D.; Ciliberti, D.; Wilson, M. A.; Fukui, K.-I.; Moriyama, K.; Numao, M.; Kloosterman, F. Kernel density compression for real-time Bayesian encoding/decoding of unsorted hippocampal spikes. *Knowledge-Based Syst.* **2016**, *94*, 1−12.

(16) Fort, G.; Jourdain, B.; Lelièvre, T.; Stoltz, G. Self-healing umbrella sampling: convergence and efficiency. *Stat. Comput.* **2017**, *27*, 147−168.

(17) Dama, J. F.; Rotskoff, G.; Parrinello, M.; Voth, G. A. Transition-tempered metadynamics: Robust, convergent metadynamics via on-the-fly transition barrier estimation. *J. Chem. Theory Comput.* **2014**, *10*, 3626−3633.

(18) Pietrucci, F. Strategies for the exploration of free energy landscapes: Unity in diversity and challenges ahead. *Rev. Phys.* **2017**, *2*, 32−45.

(19) Bussi, G.; Laio, A. Using metadynamics to explore complex free-energy landscapes. *Nat. Rev. Phys.* **2020**, *2*, 200−212.

(20) Müller, K.; Brown, L. D. Location of saddle points and minimum energy paths by a constrained simplex optimization procedure. *Theor. Chim. Acta* **1979**, *53*, 75−93.

(21) Tiwary, P.; Parrinello, M. A time-independent free energy estimator for metadynamics. *J. Phys. Chem. B* **2015**, *119*, 736−742.

(22) Valsson, O.; Tiwary, P.; Parrinello, M. Enhancing Important Fluctuations: Rare Events and Metadynamics from a Conceptual Viewpoint. *Annu. Rev. Phys. Chem.* **2016**, *67*, 159−184.

(23) Branduardi, D.; Bussi, G.; Parrinello, M. Metadynamics with Adaptive Gaussians. *J. Chem. Theory Comput.* **2012**, *8*, 2247−2254.

(24) Dama, J. F.; Parrinello, M.; Voth, G. A. Well-Tempered Metadynamics Converges Asymptotically. *Phys. Rev. Lett.* **2014**, *112*, No. 240602.

(25) Marinova, V.; Salvalaglio, M. Time-independent free energies from metadynamics via mean force integration. *J. Chem. Phys.* **2019**, *151*, No. 164115.

(26) Giberti, F.; Cheng, B.; Tribello, G. A.; Ceriotti, M. Iterative Unbiasing of Quasi-Equilibrium Sampling. *J. Chem. Theory Comput.* **2020**, *16*, 100−107.

(27) Bonati, L.; Rizzi, V.; Parrinello, M. Data-Driven Collective Variables for Enhanced Sampling. *J. Phys. Chem. Lett.* **2020**, *11*, 2998−3004.

(28) Bonati, L.; Piccini, G.; Parrinello, M. Deep learning the slow modes for rare events sampling. *Proc. Natl. Acad. Sci. U.S.A.* **2021**, *118*, No. e2113533118.

(29) Karmakar, T.; Invernizzi, M.; Rizzi, V.; Parrinello, M. Collective variables for the study of crystallisation. *Mol. Phys.* **2021**, *119*, No. e1893848.
(30) Trizio, E.; Parrinello, M. From Enhanced Sampling to Reaction Profiles. *J. Phys. Chem. Lett.* **2021**, *12*, 8621−8626.
(31) Rizzi, V.; Bonati, L.; Ansari, N.; Parrinello, M. The role of water in host-guest interaction. *Nat. Commun.* **2021**, *12*, No. 93.
(32) Ansari, N.; Rizzi, V.; Carloni, P.; Parrinello, M. Water-Triggered, Irreversible Conformational Change of SARS-CoV-2 Main Protease on Passing from the Solid State to Aqueous Solution. *J. Am. Chem. Soc.* **2021**, *143*, 12930−12934.
(33) Branduardi, D.; Gervasio, F. L.; Parrinello, M. From A to B in free energy space. *J. Chem. Phys.* **2007**, *126*, No. 054103.
(34) McCarty, J.; Parrinello, M. A variational conformational dynamics approach to the selection of collective variables in metadynamics. *J. Chem. Phys.* **2017**, *147*, No. 204109.
(35) Mendels, D.; Piccini, G.; Parrinello, M. Collective Variables from Local Fluctuations. *J. Phys. Chem. Lett.* **2018**, *9*, 2776−2781.
(36) Piaggi, P. M.; Parrinello, M. Predicting polymorphism in molecular crystals using orientational entropy. *Proc. Natl. Acad. Sci. U.S.A.* **2018**, *115*, 10251−10256.
(37) Capelli, R.; Carloni, P.; Parrinello, M. Exhaustive Search of Ligand Binding Pathways via Volume-Based Metadynamics. *J. Phys. Chem. Lett.* **2019**, *10*, 3495−3499.
(38) Ahlawat, P.; Hinderhofer, A.; Alharbi, E. A.; Lu, H.; Ummadisingu, A.; Niu, H.; Invernizzi, M.; Zakeeruddin, S. M.; Dar, M. I.; Schreiber, F.; Hagfeldt, A.; Grätzel, M.; Rothlisberger, U.; Parrinello, M. A combined molecular dynamics and experimental study of two-step process enabling low-temperature formation of phase-pure α-FAPbI 3. *Sci. Adv.* **2021**, *7*, No. eabe3326.
(39) Francia, N. F.; Price, L. S.; Salvalaglio, M. Reducing crystal structure overprediction of ibuprofen with large scale molecular dynamics simulations. *CrystEngComm* **2021**, *23*, 5575−5584.
(40) Carli, M.; Laio, A. Statistically unbiased free energy estimates from biased simulations. *Mol. Phys.* **2021**, *119*, No. e1899323.
(41) Raucci, U.; Rizzi, V.; Parrinello, M. Discover, Sample, and Refine: Exploring Chemistry with Enhanced Sampling Techniques. *J. Phys. Chem. Lett.* **2022**, *13*, 1424−1430.
(42) Bussi, G.; Branduardi, D. *Free-Energy Calculations with Metadynamics: Theory and Practice*; John Wiley & Sons, Inc., 2015; pp 1−49.
(43) Tribello, G. A.; Bonomi, M.; Branduardi, D.; Camilloni, C.; Bussi, G. PLUMED 2: New feathers for an old bird. *Comput. Phys. Commun.* **2014**, *185*, 604−613.
(44) The PLUMED consortium. Promoting transparency and reproducibility in enhanced molecular simulations. *Nat. Methods* **2019**, *16*, 670−673.

https://doi.org/10.1021/acs.jctc.2c00152
J. Chem. Theory Comput. 2022, 18, 3988−3996