

# Faster Molecular Dynamics with Neural Network Potentials via Distilled Multiple Time-Stepping and Nonconservative Forces

Nicolaï Gouraud, Côme Cattin, Thomas Plé,\* Olivier Adjoua, Louis Lagardère, and Jean-Philip Piquemal\*



Cite This: <https://doi.org/10.1021/acs.jctc.6c00653>



Read Online

ACCESS |



Metrics & More




Article Recommendations



Supporting Information

**ABSTRACT:** Following our previous work [Cattin, C. et al. *J. Phys. Chem. Lett.* **2026**, *17*(5), 1288–1295], we propose the DMTS-NC approach, a distilled multi-time-step (DMTS) strategy using nonconservative (NC) forces, to further accelerate atomistic molecular dynamics (MD) simulations using foundation neural network models such as FeNNix-Bio1. Therein, a dual-level reversible reference system propagator algorithm (RESPA) formalism couples a target accurate conservative potential to a simplified distilled representation optimized for the production of nonconservative forces. Despite being nonconservative, the distilled architecture is designed to enforce key physical priors, such as equivariance under rotation and cancellation of atomic force components. These choices facilitate the distillation process and therefore improve drastically the robustness of simulation, significantly limiting abnormal discrepancies between the two models, thus achieving excellent agreement with the forces data. Overall, the DMTS-NC scheme is found to be more stable and efficient than its conservative counterpart with additional speedups reaching 15–30% over DMTS. Requiring no fine-tuning steps, it is easier to implement and can be pushed to the limit of the system's physical resonances to maintain accuracy while providing maximum efficiency. We obtain additional speedup by combining hydrogen mass repartitioning (HMR) and High Hydrogen Friction (HHF) to further extend the largest time step up to 10 fs of our schemes while conserving stability and accuracy. As for DMTS, DMTS-NC is applicable to any neural network potential (NNPs) and can be applied to approaches that are computationally heavier than FeNNix-Bio1. We show a proof of principle applying the approach to the distillation of MACE-OFF23 with consequent speedups ranging from 3.66 to 5.64 compared to a single time step.



## 1. INTRODUCTION

Molecular dynamics (MD)<sup>1,2</sup> is a popular numerical tool aimed at inferring properties of matter from computer simulations at the atomistic level. With the steady development of efficient algorithms and computational power, “in silico” methods enable the simulation of increasingly complex systems with high accuracy, driving a variety of scientific applications in biology, chemistry, drug design, and material sciences.<sup>3</sup> Studying the time evolution of a system, i.e., resolving the associated equations of motion, requires a physical model encoding the interactions between particles. Ideally, at this level of description, this model should be grounded on quantum mechanics. However, solving numerically the electronic structure problem, with so-called *ab initio* methods,<sup>4</sup> remains computationally expensive, which limits their application to relatively small systems and on short time scales. A more coarse grained approach, realizing a good compromise between computational efficiency and modeling at the atomic scale, consists in using an empirical potential<sup>5–10</sup> with a given functional form, involving *a priori* unknown parameters that must be optimized to replicate a set of targeted thermodynamic or structural properties.

In recent years, a third approach has emerged, namely Neural Network potentials (NNPs),<sup>11–26</sup> which can be seen as

the application of machine-learning architectures to the analytic potential approach, learning an energy model from large *ab initio* computation databases. Nowadays, several libraries such as SchnetPack,<sup>27</sup> DeepMD-kit,<sup>28</sup> MLAtom,<sup>29</sup> or FeNNol<sup>30</sup> enable the implementation of NNPs, and large foundation models such as MACE<sup>31,32</sup> and FeNNix-Bio1<sup>33</sup> allow general-purpose simulations, covering complete areas of applications of molecular dynamics.<sup>24,31–34</sup> Several important advantages over traditional approaches can be highlighted. First, the large number of parameters, whose optimization is fully automated, make NNPs highly flexible and largely transferable across systems. Second, they inherently handle chemical reactions, which most classical empirical force fields do not. Finally, again thanks to the large number of parameters, NNPs can achieve close-to quantum-mechanical accuracy, only for a fraction of the computation cost of *ab initio* methods. However, they remain significantly more expensive to evaluate

Received: April 2, 2026

Revised: May 28, 2026

Accepted: May 28, 2026

than traditional empirical potentials, which motivates the development of algorithms increasing the speed of simulation based on NNPs without losing their high accuracy.

In a molecular dynamics simulation, the number of evaluations of the gradient of the potential energy per unit of physical time is inversely proportional to the time discretization step size  $\delta$ . The maximum value of this parameter which preserves stability and accuracy of the simulation is limited by the highest-frequency motions of the system, such as chemical bond vibrations. A well-known approach to increase the step size used in classical integrators such as BAOAB<sup>35,36</sup> is called multi-time-stepping<sup>37–39</sup> (MTS). When applied to an empirical force field model, it consists in treating the fast-varying short-range forces within a small step size  $\delta$ , while only applying the long-range, more regular forces at larger time step  $\Delta = n\delta$ . This effectively increases the speed of simulation by reducing the rate at which expensive long-range forces are computed. Note that although MTS simulations are stable and accurate for  $\Delta$  several times larger than  $\delta$ , its maximum value is also bounded, this time by resonance phenomena<sup>40,41</sup> resulting from a coupling of internal frequencies of the molecular system with nonphysical periodicity created by the various time steps, which can induce numerical instabilities for large values of  $\Delta$ . Despite these limitations, for which several alternatives exist,<sup>42–46</sup> MTS methods such as reference system propagator algorithm (RESPA)<sup>37</sup> or BAOAB-RESPA-1<sup>39</sup> are widely adopted.

MTS schemes are also applicable to *ab initio* MD,<sup>47</sup> and several recent works studied their use in machine-learning approaches.<sup>48–52</sup> Note that the principle of MTS as described above is not readily usable with NNPs, since there is no natural decomposition of the forces into “cheap” and “expensive” parts. Recently, we proposed<sup>52</sup> a strategy based on distillation<sup>53,54</sup> of the FeNNix-Bio1(M) foundation model<sup>33</sup> which was performed via the FeNNol<sup>30</sup> library. In this strategy, a small, less precise but fast-to-evaluate model is trained on data labeled with the FeNNix-Bio1(M) model instead of Density Functional Theory. Then, following the MTS procedure, this small model is applied in several iterations of an inner loop, before being corrected by an application of the difference between the large and small force models in an external loop. Therefore, by reducing the number of computations of the large expensive model, simulation speed is increased, enabling longer stable simulations.

As proposed in several previous works,<sup>21,55–59</sup> another path to accelerating simulation speed with ML models is to use nonconservative (NC) forces, i.e., forces that do not necessarily derive from a potential. Indeed, directly predicting interatomic forces enables to bypass the costly backpropagation step required to compute conservative forces, thus improving computational efficiency at fixed model architecture. However, Bigi et al.<sup>60</sup> identified severe simulation artifacts induced by this approach and alternatively proposed to use nonconservative forces in an MTS scheme. In their framework, a single model is trained to output both conservative and nonconservative forces: the latter are used at every step of the simulation, while conservative forces are only evaluated every few steps. In this work, we follow a similar direction and extend our Distilled Multi-time-Step (DMTS) procedure to the use of nonconservative forces, which we denote DMTS-NC. Compared to the unified model proposed in ref 60, the distillation step significantly reduces the computational cost of the model used in the MTS internal steps, while the absence of

differentiation procedures make both training and evaluation faster than the distilled conservative models used previously in ref 52. In addition, by imposing several physical priors such as equivariance under rotation and cancellation of atomic force components (which are properties that are satisfied by conservative forces), the distilled model is able to closely fit the reference one. We show that the improved agreement between the models, with training reaching a mean absolute error near 1 kcal/mol/Å on a diverse data set labeled with FeNNix-Bio1(M), allows to use larger external time steps, further enhancing the speed and robustness of DMTS schemes. Importantly, like the original DMTS framework, DMTS-NC is a model-agnostic acceleration strategy, and can be applied to any neural network potential, independently of its specific architecture.

The rest of this article is organized as follows. In Section 2, we describe the architecture and training of the non-conservative model, the associated DMTS-NC integrator, as well as several techniques used to further stabilize the dynamics, namely hydrogen mass repartitioning (HMR), high hydrogen friction (HHF), and a rewind procedure preventing occasional disagreements between the models. Section 3 is devoted to numerical experiments and validation of the improved stability and speed of DMTS-NC on various molecular systems. In addition to simulations based on FeNNix-Bio1, we highlight the generality of DMTS-NC by applying the method to the MACE-OFF23 foundation model.

## 2. METHODS

### 2.1. Model Architecture

The architecture of the nonconservative model is a modification of FeNNix-Bio1<sup>33</sup> in which the neural network learns to directly predict forces instead of energies. By removing the requirement of having conservative forces, i.e., derived from a potential, certain properties such as Newton's third law are not guaranteed anymore. However, our architecture is designed to still enforce some physical priors such as equivariance under rotation and cancellation of atomic force components.

Let us first quickly recall the general architecture of FeNNix-Bio1. In a first step, each atom is embedded in a  $N_f$ -dimensional vector space, describing electronic structure and charge information. This initial geometry-independent embedding is then updated to include local geometric information via two message-passing layers of an equivariant transformer. The first layer considers a high-resolution description of a very short-range (within a  $R_c^{(sr)} = 3.5$  Å cutoff) geometry, while the second layer incorporates both short-range and medium-range messages up to a cutoff of  $R_c^{(lr)} = 7.5$  Å. Then, atomic energies are obtained from mixture-of-experts multilayer perceptrons that take the final embeddings as an input, with routing depending on the chemical group of each atom. Finally an explicit screened nuclear repulsion term is added. We refer to ref 33 for more details.

We now turn to the nonconservative model. Its embedding architecture is similar to the previous one, except that only short-range messages are considered in both layers. This allows the model to mainly focus on local structure, as it is specifically designed to handle short-range forces and the associated high-frequency motions within the inner loops of an MTS integrator. At the end of the embedding layers, each atom  $i$  is characterized by a scalar embedding  $x_i \in \mathbb{R}^{N_f}$  and a tensorial embedding  $\hat{V}_i \in \mathbb{R}^{n_1 \times (\lambda_{\max} + 1)^2}$ . The latter is a collection, along  $n_1$  channels, of geometric tensors of orders up to  $\lambda_{\max}$ . We then use  $n_s$  scalar attention heads and  $n_t$  tensor heads, one for each first-order geometric tensor channel. For every atom  $i$  and each attention head  $h \in \{1, \dots, n_s + n_t\}$ , vectors  $q_{ih}$  and  $k_{ih}$  are computed via a linear projection of the scalar embedding  $x_i$  on a subspace of dimension  $n_c$ .

$$q_{ih} = W_{qh}x_i$$

$$k_{ih} = W_{kh}x_i,$$

where the weight matrices  $W_{qh}$  and  $W_{kh}$  are optimized during the training. Let  $\mathcal{N}(i)$  denote the neighbor list of atom  $i$ , i.e., the set of atoms  $j \neq i$  within the short-range cutoff distance  $R_c^{(sr)}$ . For each head  $h \in \{1, \dots, n_s + n_l\}$  and each neighbor  $j \in \mathcal{N}(i)$ , the scaled dot product coefficient  $c_{ijh}$  is defined as

$$c_{ijh} = \frac{1}{\sqrt{n_c}} \sum_{k=1}^{n_c} q_{ihk} k_{jhk}$$

Let  $\vec{R}_{ij}$  be the vector pointing from  $i$  to  $j$  and  $(B_h(r_{ij}))_{h \in \{1, \dots, n_s\}}$  a radial basis vector corresponding to the decomposition of  $r_{ij} = |\vec{R}_{ij}|$  in a basis of  $n_s$  Bessel functions, as described in ref 61. The three-dimensional vector  $F_{ij}$  is then constructed as

$$F_{ij} = \left( \sum_{h=1}^{n_s} c_{ijh} B_h(r_{ij}) \right) \left( \frac{\vec{R}_{ij}}{r_{ij}} + \sum_{h=1}^{n_l} c_{ij(n_s+h)} \vec{V}_{ih} \right) f_c(r_{ij})$$

where  $\vec{V}_{ih}$  is the vector part of the tensor embedding  $\hat{V}_{ih}$  (i.e., the irreducible representations of order  $\lambda = 1$ ) and  $f_c(r_{ij})$  is a polynomial cutoff function going smoothly to zero at the short-range cutoff distance  $R_c^{(sr)}$ , as defined in ref 61. The final force vector  $F_i$  acting on atom  $i$  is given by the antisymmetric sum of the  $F_{ij}$  to which is added the short-range screened nuclear repulsion term  $F_{rep} = -\nabla E_{rep}$  that is used in FeNNix-Bio1

$$F_i = \sum_{j \in \mathcal{N}(i)} (F_{ij} - F_{ji}) + F_{rep,i} \in \mathbb{R}^3$$

The expression of  $E_{rep}$  follows the analytic NHL parametrization from 62

$$E_{rep}(r_{ij}) = \frac{Z_i Z_j}{4\pi\epsilon_0 r_{ij}} \sum_{n=1}^3 a_{nij} e^{-b_{nij} r_{ij}}$$

where  $Z_i$  is the atomic number of atom  $i$ , and  $a_{nij}$ ,  $b_{nij}$  are parameters, obtained via fitting to *ab initio* reference data in ref 62, that depend on the unordered pair of species  $Z_\nu Z_\rho$ .

Note that  $F_i$  explicitly depends on the tensorial embedding  $\hat{V}_\nu$ , while in the original architecture of FeNNix-Bio1, it is only used as an internal descriptor in the embedding phase, and not given as an input to the perceptrons that compute the final energies. In particular, here the functional form of  $F_i$  ensures that the sum of the forces over all atoms is zero, and more generally, the sum of the  $F_i$  over any connected component of the interaction graph of the atoms is also zero. In addition, all of the  $F_{ij}$  and therefore the  $F_i$  too, are equivariant under rotation.

We conclude this paragraph with the parameter values. In FeNNix-Bio1(M),  $N_f = 176$ ,  $n_c = 16$ ,  $\lambda_{max} = 3$ ,  $n_l = 4$ , and  $n_s = 10$ . In the small NC model, we have  $N_f = 64$ ,  $n_c = 8$ ,  $\lambda_{max} = 1$ ,  $n_l = 2$ , and  $n_s = 8$ . For the sake of completeness, let us mention that in the initial geometry-independent embedding, the number of charge equilibration channels is  $N_Q = 32$  in FeNNix-Bio1(M), and  $N_Q = 8$  in the NC model, the species embedding neural network involves one layer of 256 neurons in FeNNix-Bio1(M) and only 64 in the NC model. In the message-passing phase, the attention mechanism involves  $h_s = 16$  scalar attention heads in FeNNix-Bio1(M) and  $h_s = 4$  in the NC model, and a two hidden-layers neural network, with respectively 352 and 176 neurons in FeNNix-Bio1(M) and a one-hidden-layer neural network with 64 neurons in the NC model. Finally, the polynomial  $f_c(r_{ij})$  is of order 8 in FeNNix-Bio1(M), and 5 in the NC model. This results in a total number of 286,736 parameters in the NC model, compared to 9,526,855 for FeNNix-Bio1(M).

## 2.2. Model Training

Following the principle of knowledge distillation,<sup>53,54</sup> as applied in ref 52, the nonconservative force model is trained as a distilled version of the larger FeNNix-Bio1(M)<sup>33</sup> NNP. A subset of conformations of the

SPICE2 data set<sup>63,64</sup> is evaluated with FeNNix-Bio1(M), which then constitutes the training data set for the nonconservative force model. This data set contains a variety of small organic molecules and biologically relevant complexes, thus covering a wide chemical space.

The training was done for 2200 epochs with 1000 batches of 128 conformations per epoch. We used the Muon optimizer,<sup>65</sup> with an initial learning rate of  $1.0 \times 10^{-5}$ , linearly increasing up to  $5.0 \times 10^{-4}$  after a warm-up phase, then decreasing down to a final rate of  $1.0 \times 10^{-6}$  according to a cosine one-cycle learning-rate schedule. The total training, performed on a Nvidia A100 40GB GPU, took 28 h and 47 min. The resulting model, which directly aims at fitting the force vectors produced by FeNNix-Bio1(M) instead of the energies, ends up achieving excellent agreement with the data, reaching a final mean absolute error of MAE = 1.46 kcal/mol and a root-mean-square error of RMSE = 2.33 kcal/mol, a significantly lower value than for the small conservative model of ref 52, for which MAE = 3.44 kcal/mol and RMSE = 5.53 kcal/mol.

## 2.3. Multi-time-Step Integrator

Denote  $U: \mathbb{R}^d \rightarrow \mathbb{R}$  a potential energy function, and  $\mathcal{L}$  the infinitesimal generator of associated Langevin dynamics. Most single-time-step (STS) numerical schemes rely on so-called splitting of the dynamics. For instance, the BAOAB method<sup>35,36</sup> consists in splitting the Langevin process into three parts A, B, and O, corresponding respectively to free transport, acceleration, and fluctuation/dissipation, which comes down to splitting the generator as  $\mathcal{L} = \mathcal{L}_A + \mathcal{L}_B + \mathcal{L}_O$  and simulating successively the parts B-A-O-A-B, thanks to the Trotter/Strang formula

$$e^{t\mathcal{L}} = e^{t/2\mathcal{L}_B} e^{t/2\mathcal{L}_A} e^{t\mathcal{L}_O} e^{t/2\mathcal{L}_A} e^{t/2\mathcal{L}_B} + \mathcal{O}(t^3)$$

Multi-time-Step (MTS) methods such as BAOAB-RESPA push the splitting of the dynamics one step further, by decomposing the force vector field  $-\nabla U = F_S + F_L$  and considering separately the associated accelerations  $\mathcal{L}_{B_S}$  and  $\mathcal{L}_{B_L}$ . In a classical force field,  $F_S$  typically gathers the fast-varying, easy to compute short-range forces and  $F_L$  includes more regular many-body or long-range forces that are more expensive to evaluate.

Given a time step  $\delta > 0$ , denote  $Q_\delta$  one transition of the BAOAB chain associated with the force  $F_S$ , namely

$$Q_\delta = e^{\delta/2\mathcal{L}_{B_S}} e^{\delta/2\mathcal{L}_A} e^{\delta\mathcal{L}_O} e^{\delta/2\mathcal{L}_A} e^{\delta/2\mathcal{L}_{B_S}}$$

Then, by denoting  $\Delta = n\delta$  some multiple of  $\delta$ , the BAOAB-RESPA MTS scheme relies on the following approximation of the continuous dynamics

$$e^{\Delta\mathcal{L}} \approx e^{\Delta/2\mathcal{L}_{B_L}} (Q_\delta)^n e^{\Delta/2\mathcal{L}_{B_L}}$$

In other words, letting  $x$  and  $v$  denote the position and velocity vectors and  $M$  the mass matrix, one transition of the integrator consists of the following steps

1.  $v \leftarrow v - \frac{\Delta}{2} M^{-1} F_L(x)$ .
2. Perform  $n$  times the BAOAB scheme associated with  $F_S$  with time step  $\delta$ .
3.  $v \leftarrow v - \frac{\Delta}{2} M^{-1} F_L(x)$ .

Note that in this method,  $F_S$  and  $F_L$  are not required to be conservative forces: they only need to satisfy  $F_S + F_L = -\nabla U$ . In DMTS-NC,  $F_S$  is the smaller, faster to evaluate machine-learned model approximating  $\nabla U$  that we described above and  $F_L = -\nabla U - F_S$ . To sum up, the full procedure is given in Algorithm 1.

## 2.4. Hydrogen Mass Repartitioning

The resonance phenomenon<sup>40,41</sup> that limits the maximum usable value of  $\Delta$  can be mitigated in various ways. A standard procedure, applied for instance in ref 39, and which does not significantly impair sampling rates is called Hydrogen Mass Repartitioning (HMR).<sup>66</sup> It consists in increasing the mass of hydrogen atoms (in our simulations, by 3.0 Da) and redistributing the deficit over heavier atoms in order to preserve the total mass of each molecule. This has the effect of red-

**Algorithm 1** DMTS-NC Integration Step

**Input:**  $F_{\text{NN}}(x)$  the reference force field evaluated at configuration  $x$ ,  $F_{\text{NNX}}(x)$  the cheaper non-conservative force model,  $M$  the mass matrix,  $\Delta$  the external time-step,  $n$  the number of iterations of the internal loop.

```

1: if first_step then
2:    $F_{\text{NC}} \leftarrow F_{\text{NNX}}(x)$ 
3:    $F \leftarrow F_{\text{NN}}(x)$ 
4: end if
5:  $v \leftarrow v + \frac{\Delta}{2} M^{-1}(F - F_{\text{NC}})$ 
6: for  $i = 1$  to  $n$  do
7:    $v \leftarrow v + \frac{\Delta}{2n} M^{-1} F_{\text{NC}}$ 
8:    $x \leftarrow x + \frac{\Delta}{2n} v$ 
9:    $v \leftarrow \text{Thermostat} \left( v, \frac{\Delta}{n} \right)$ 
10:   $x \leftarrow x + \frac{\Delta}{n} v$ 
11:   $F_{\text{NC}} \leftarrow F_{\text{NNX}}(x)$ 
12:   $v \leftarrow v + \frac{\Delta}{2n} M^{-1} F_{\text{NC}}$ 
13: end for
14:  $F \leftarrow F_{\text{NN}}(x)$ 
15:  $v \leftarrow v + \frac{\Delta}{2} M^{-1}(F - F_{\text{NC}})$ 

```

shifting the frequency of intramolecular bond and angle motions, allowing in turn to increase the external time step  $\Delta$ , without biasing configurational sampling (that only depends on potential energy, which is unchanged), nor the temperature (whose distribution does not depend on the mass). In addition, as shown below in Section 3 and already noted in refs 39,52, by preserving the mass of each molecule, dynamic properties such as diffusion coefficients are not significantly impacted either, contrary to other methods.<sup>43–45</sup>

### 2.5. High Hydrogen Friction

To sample the canonical measure, a standard choice is to simulate (underdamped) Langevin dynamics, parametrized by a friction coefficient  $\gamma$  measuring the rate of velocity dissipation. This process can be generalized by replacing the scalar  $\gamma$  by any (possibly position-dependent) symmetric positive-definite matrix  $\Gamma$ . While determining an optimal  $\Gamma$  remains a nontrivial question (see ref 67 for results in this direction), a simple possibility is to take a constant diagonal matrix, assigning a specific friction coefficient to each atomic species.

Taking a higher friction on hydrogen (a method to which we refer as HHF in the following) has the effect of damping bond oscillations, without impacting the diffusion coefficient (see Section 3) as much as a fully overdamped regime on all atoms, while also being much simpler to implement and to parametrize than colored-noise thermostats based on the generalized Langevin equation.<sup>42</sup>

In addition, and in combination with HHF, the intrinsically diffusive behavior caused by momentum flips in overdamped regimes can be addressed by replacing the standard Langevin thermostat (or rather its species-dependent-friction variant) by the Fast Forward Langevin (FFL) thermostat.<sup>68</sup> It consists in applying the following velocity update in the thermostat part of the integrator

$$v \leftarrow |e^{-\gamma\delta} v + \sqrt{\beta^{-1}(1 - e^{-2\delta\gamma})} M^{-1/2} G| \frac{v}{|v|}$$

where  $G$  is a standard multidimensional Gaussian vector,  $M$  the mass matrix of the system,  $\beta$  the inverse temperature,  $\gamma$  the (species-dependent) friction coefficient, and  $|\cdot|$  the Euclidean norm.

### 2.6. Preventing Rare Model Disagreements (Rewind Procedure)

In addition to intrinsic resonance problems associated with multi-time-stepping, DMTS faces a specific pitfall: the possibility that on rare occasions,  $F_S$  and  $-\nabla U$  disagree significantly, which can make the dynamics suddenly diverge because of the abnormally strong applied force  $F_L = -\nabla U - F_S$ . Unlike resonances, which are detectable early in a simulation, these failures caused by “hallucinations” of the small force field are difficult to anticipate, since the dynamics can remain stable for a long period before encountering a disagreement.

One way of addressing this issue, proposed in ref 52, relies on an active learning procedure. It consists in fine-tuning the small model based on data frames, collected automatically during an initial simulation, for which the difference between the models exceeds some unrealistic threshold. In addition to being system-specific, collecting the pathological data can take a long time if disagreements are too rare, in which case extremely long simulations are also required after the fine-tuning to validate that the new model is fully stable. As we

show in Section 3, the excellent overall agreement between models, allowed by the nonconservative architecture focusing on forces instead of energies, results in much fewer “holes” than in the original DMTS of ref 52. Therefore, even at larger values of  $\Delta$ , the dynamics remains robust without relying on system-specific fine-tuning.

Instead, in DMTS-NC, a generic on-the-fly procedure is implemented: at each time step, if the difference between the two force models (or some other indirect indicator) exceeds a threshold, the dynamics rewinds until the last saved frame and MTS is replaced by STS for a given number of steps (typically a few picoseconds of physical time), before going back to multi-time-stepping. Given the rarity of those rewinds, occurring on average every 10 ns at the maximum stable  $\Delta$ , the overall simulation speed is almost not impacted by the procedure.

### 2.7. Computation of the Virial Term in NPT

Several barostats are currently implemented in FeNNol for simulations in the isothermal–isobaric (NPT) ensemble: Monte Carlo,<sup>69,70</sup> Berendsen,<sup>71</sup> Langevin piston,<sup>72</sup> and Bussi’s stochastic velocity rescaling.<sup>73</sup>

A common issue of multi-time-step methods in NPT is that the small discretization bias on the potential energy, while insignificant in NVT for most observables of interest, has a non-negligible impact on the estimation of the virial term  $q \cdot \nabla U(q)$ , which, in turn, induces a problematic bias on the volume distribution. In the Tinker molecular dynamics software,<sup>74</sup> this virial bias is corrected by adding the average of the intermediate virials at the end of each internal loop to the final estimator, computed from the configuration and volume at the end of the external step.

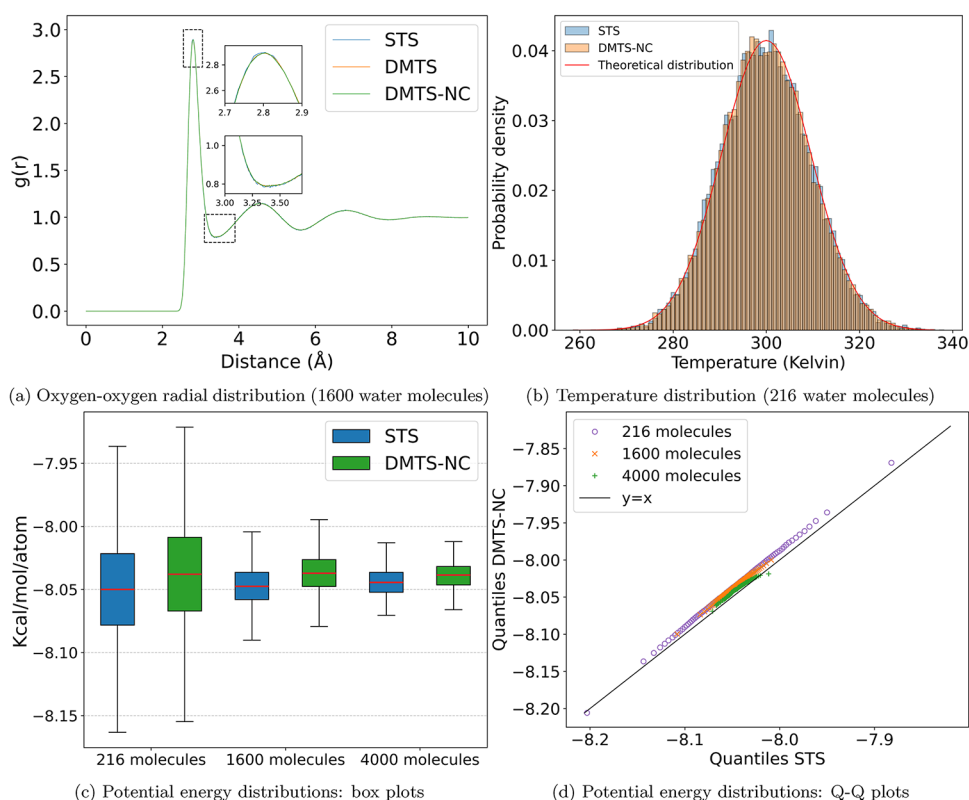
Here, we apply a similar method, however, since nonconservative forces are applied in the internal loops, the intermediate virials are not well-defined. Instead, the correction term added to the final estimator of the virial is the average, over all internal loops, of the proxy term  $-\sum_{i=1}^N F_i \cdot q_i$ , where  $F_i$  is the force acting on atom  $i$ . To resolve the coordinate ambiguity of  $q_i$  regarding the choice of the periodic image under periodic boundary conditions, this term can be rewritten as a function of the direction vectors  $r_{ij} = q_j - q_i$  with minimum image convention, and the pairwise force terms  $F_{ij}$  and  $F_{\text{rep},ij}$  defined in the model architecture section. By using the fact that the repulsion terms are antisymmetrical and that neighbor lists are symmetrical, a simple computation shows that

$$-\sum_{i=1}^N F_i \cdot q_i = \sum_{i=1}^N \sum_{j \in \mathcal{N}(i)} \left( F_{ij} + \frac{1}{2} F_{\text{rep},ij} \right) \cdot r_{ij}$$

In Section 1 of Supporting Information, we show that applying this correction term removes the bias on pressure and density distributions.

## 3. NUMERICAL RESULTS

The goal of this section is to assess performance, stability, and sampling accuracy of the DMTS-NC method, on bulk water boxes of various sizes, two solvated proteins and a collection of solvated small molecules. Simulations were performed in periodic boundary conditions, applying the Langevin thermostat at  $T = 300$  K temperature and  $\gamma = 1$  ps<sup>-1</sup> friction. In the HHF variant, we tested both Langevin and Fast Forward Langevin thermostats with a higher friction of  $\gamma_H = 10$  ps<sup>-1</sup> on hydrogen atoms and  $\gamma = 1$  ps<sup>-1</sup> on other species. Single time step simulations used the BAOAB integrator. In all DMTS and DMTS-NC simulations, HMR was systematically activated with an additional 3.0 Da on hydrogen atoms. Isobaric–isothermal simulations applied Bussi’s barostat,<sup>73</sup> targeting a 1.0 atm pressure, with a  $\tau_P = 2.0$  ps volume relaxation time. Training and MD used the FeNNol<sup>30</sup> Python library coupled to the Tinker-HP package<sup>75,76</sup> within its Deep-HP module<sup>48</sup> on a single NVIDIA A100 40GB GPU card.



**Figure 1.** Sampling properties of bulk water simulations. All data was collected from 2 ns simulations, with one frame saved per picosecond. Panel (a) shows that radial distributions are not affected by the use of DMTS-NC compared to DMTS or STS. Panel (b) shows that temperatures fit the equilibrium theoretical distribution, and samples of potential energies, represented on a boxplot (panel (c)) and a quantile-quantile plot (panel (d)), highlight the low-bias identical-variance deviation of the distributions from STS to DMTS-NC, as expected for a multi-time-step method in a stable regime.

### 3.1. Speed and Sampling Performances

**3.1.1. Bulk Water.** We start by comparing, for various sizes of bulk water boxes and various time steps, the STS BAOAB integrator, the BAOAB-RESPA DMTS integrator with the small conservative model used in ref 52, and DMTS-NC, the version using the nonconservative force model. To assess sampling properties, Figure 1 shows oxygen–oxygen radial distributions (computed with FeNNol), as well as distributions of temperatures and potential energies across simulations. In particular, empirical distributions of temperature (known to be sensitive to numerical instabilities due to multi-time-stepping) are plotted along the equilibrium theoretical distribution, given by  $\frac{T}{d}\chi^2(d)$ , where  $T$  is the equilibrium average temperature (here, 300 K),  $d$  is the total configurational dimension of the system (i.e., three times the number of atoms), and  $\chi^2(d)$  is the Chi-squared distribution with  $d$  degrees of freedom. In addition, Table 3 sums up the average potential energies and temperatures.

Table 1 shows the speed of the algorithms expressed in nanoseconds of simulation per day, highlighting the significant acceleration of DMTS-NC with respect to conservative model DMTS and STS. Note that in addition to being faster, thanks to a better fit to the training data, DMTS-NC exhibits fewer numerical instabilities, and the maximum usable time step is limited by classical MTS resonances instead of the problem of abnormal model disagreements, as described in Section 2. Simulations show that DMTS-NC provides an acceleration ranging from 2.94 to 4.31 with respect to STS and from 30.8 to 55.7% with respect to conservative DMTS.

**Table 1. Simulation Performances for Bulk Water Systems, Expressed in Nanoseconds per Day<sup>a</sup>**

integrator	number of atoms			numerical instabilities
	648	4800	12,000	
STS	36.89	9.08	3.63	no
DMTS 1–5 fs	69.54	30.17	11.95	no
DMTS 1–6 fs	80.62	34.34	NaN	data holes
DMTS 1–7 fs	NaN	NaN	NaN	
DMTS-NC 1–5 fs	91.13	32.69	12.39	no
DMTS-NC 1–6 fs	103.67	36.15	13.67	no
DMTS-NC 1.625–6.5 fs	108.29	41.10	15.64	no
DMTS-NC 1–7 fs	114.11	39.72	13.56	high temp. (resonances)
max. acc. DMTS-NC/DMTS	55.7%	36.2%	30.8%	
max. acc. DMTS-NC/STS	2.94	4.53	4.31	

<sup>a</sup>In DMTS-NC, the stability limit is reached around an external time step of  $\Delta = 6.5$  fs, after what resonances start to appear. In contrast, the main limitation with the (generic) conservative model is the presence of “data holes” already at  $\Delta = 6$  fs.

**3.1.2. Solvated Proteins.** We next turn to two protein systems: solvated phenol-lysosyme protein–ligand complex (16,221 atoms) and solvated dihydrofolate reductase protein (DHFR, 23,558 atoms). Simulation speeds are summarized in Table 2, showing that DMTS-NC provides an acceleration factor with respect to STS of 3.40 for phenol-lysosyme and

**Table 2. Speed Performances for the Two Ligand-Protein Systems<sup>a</sup>**

integrator	speed (ns/day)		numerical instabilities
	phenol-lyso.	DHFR	
STS	2.55	1.85	no
DMTS 1.75–3.5	6.59	4.46	no
DMTS 1–4	NaN	NaN	
fine-tuned DMTS 2–4	7.45	4.97	no
DMTS-NC 1–4	7.20	4.89	no
DMTS-NC 1–5	8.42	5.57	no
DMTS-NC 1.25–5	8.67	5.72	no
DMTS-NC 1.1–5.5	NaN	NaN	
max. acc. DMTS-NC/generic DMTS	31.6%	28.3%	
max. acc. DMTS-NC/fine-tuned DMTS	16.4%	15.1%	
max. acc. DMTS-NC/STS	3.40	3.09	

<sup>a</sup>Similarly to bulk water simulations, DMTS-NC allows to use a larger external time step than conservative DMTS, without relying on active learning.

3.09 for DHFR, an acceleration with respect to the generic conservative DMTS of 31.6% for phenol-lysosyme and 28.3% for DHFR, and finally an acceleration with respect to the fine-tuned (after a system-specific active learning procedure) conservative DMTS of 16.4% for phenol-lysosyme and 15.1% for DHFR. Note that in all simulations, DMTS-NC uses a generic, nonfine-tuned model.

Regarding potential energy and average temperature, statistics were gathered from 5 ns of simulation and shown in Table 3. Similarly to simulations of bulk water, the low-bias identical-variance deviation of the distributions from STS to DMTS-NC, expected for multi-time-stepping in a stable regime, indicates that using an external time step of  $\Delta = 5$  fs yields stable and accurate simulations, while resonance problems start to arise for larger values of  $\Delta$ . Finally, in Figure 2, we show the time evolution of the protein backbone RMSD and the ligand's Distance to Bound Configuration (DBC), computed using the Colvars library,<sup>77</sup> during a 20 ns simulation of the solvated phenol-lysosyme complex, with frames collected every picosecond.

### 3.1.3. Hydration Free Energies of Small Molecules.

We then compare STS and DMTS-NC for calculating the hydration free energies (HFEs) of a collection of 44 solvated small molecules, both in NVT and NPT ensembles. Free

energy is computed using the alchemical lambda-ABF method,<sup>78</sup> following the alchemical parametrization for the FeNNix-Bio1 architecture described in ref 33.

Each simulation was done for 6 ns, with a time step of 1 fs on STS and 1 fs–6 fs on DMTS-NC. Results are illustrated in Figure 3, where error bars result from three runs for each molecule and each setup. Thanks to the increased speed of DMTS-NC, simulation cost is greatly reduced with respect to STS, while preserving its high accuracy: the average mean absolute error between the HFE predicted by STS and DMTS-NC is 0.110 kcal/mol in both NVT and NPT simulations, and the root-mean-square error (RMSE) is 0.142 kcal/mol in NVT and 0.135 kcal/mol in NPT. We refer to the Supporting Information for a complete list of the molecules and the corresponding HFE values.


**3.1.4. HHF Increased Speed.** By damping hydrogen atom movements, the HHF method described in Section 2 with a  $\gamma_{\text{H}} = 10 \text{ ps}^{-1}$  friction on hydrogen and  $\gamma_{\text{O}} = 1 \text{ ps}^{-1}$  on oxygen allowed to push further the external time step to  $\Delta = 10$  fs on water boxes systems,  $\Delta = 7$  fs on DHFR, and  $\Delta = 6$  fs on phenol-lysosyme, yielding an additional increase in speed, mostly on bulk water. Simulation speed and accelerations are gathered in Table 4, and statistics on temperature and potential energy are given in Table 5. In bulk water systems, HHF allows for an acceleration between 25.92 and 32.61% with respect to regular DMTS-NC, and from 7.96 to 14.57% on solvated proteins. When comparing with STS, the HHF simulations yield an acceleration factor ranging from 3.45 (on DHFR) to 5.28 (on the 4800 atoms water box). While temperature, potential energy, and overall stability are well preserved in this setup, the acceleration provided by HHF comes at the cost of degraded diffusion coefficients, as discussed in Section 3.2.

**3.1.5. Distillation of MACE-OFF23.** The DMTS and DMTS-NC procedures are generic and can be applied to any NNP. To illustrate this, we applied DMTS-NC to the MACE-OFF23(S) foundation model.<sup>32</sup> Its distillation (based on the same data set) into a nonconservative model yielded a good agreement with the data (MAE = 1.484 kcal/mol, RMSE = 2.363 kcal/mol), very similar to the distillation of FeNNix-Bio1(M). We compared the speed of simulation between STS and DMTS-NC integrators on the same systems (three water boxes and two solvated proteins) and gathered the results in Table 6. Since MACE-OFF is more expensive to evaluate than FeNNix-Bio1 (see ref 33 for a performance comparison), while the distilled nonconservative model has exactly the same architecture as previously, the acceleration factor is greater

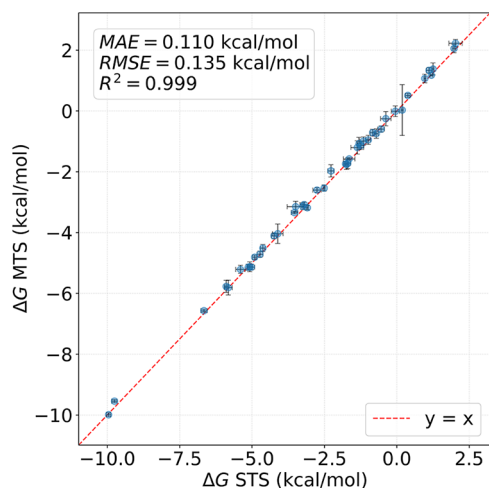
**Table 3. Average Temperatures and Potential Energies ( $\pm$  Empirical Standard Deviations)<sup>a</sup>**

system	integrator	temperature (K)	potential energy (kcal/mol/atom)		
water boxes	216 mol.	STS	299.84 $\pm$ 9.57	-8.049 $\pm$ 0.042	
		DMTS-NC	299.87 $\pm$ 9.56	-8.038 $\pm$ 0.043	
	1600 mol.	STS	300.09 $\pm$ 3.60	-8.047 $\pm$ 0.017	
		DMTS-NC	299.93 $\pm$ 3.54	-8.037 $\pm$ 0.016	
solvated proteins	phenol-lyso.	4000 mol.	STS	300.61 $\pm$ 2.93	-8.043 $\pm$ 0.014
			DMTS-NC	299.69 $\pm$ 2.22	-8.039 $\pm$ 0.011
	DHFR	STS	300.15 $\pm$ 1.99	-7.869 $\pm$ 0.009	
		DMTS-NC	301.51 $\pm$ 1.99	-7.850 $\pm$ 0.009	
		STS	299.97 $\pm$ 1.64	-7.819 $\pm$ 0.008	
		DMTS-NC	302.26 $\pm$ 1.64	-7.796 $\pm$ 0.007	

<sup>a</sup>DMTS-NC simulations were performed at the highest stable time steps that we found, i.e.,  $\delta = 1.25 \text{ fs} - \Delta = 5 \text{ fs}$  for the protein systems, and  $\delta = 1.2 \text{ fs} - \Delta = 6 \text{ fs}$  for bulk water.



**Figure 2.** Time evolution of the protein backbone RMSD (panel (a)) and the ligand's Distance to Bound Configuraition (panel (b)) during a 20 ns simulation of the solvated lysosyme-phenol. The DMTS-NC integrator uses  $\delta = 1.25$  fs– $\Delta = 5$  fs time steps.



**Figure 3.** Hydration free energy of 44 small molecules in NPT. Error bars are computed from three independent lambda-ABF runs for each integrator and each molecule.

**Table 4. HHF Speed Performances and Acceleration Compared to DMTS-NC and STS**

system size/name		water boxes		solvated prot.	
		1600 mol.	4000 mol.	phenol-lys.	DHFR
speed (ns/day)	STS 1 fs	9.08	3.63	2.55	1.85
	DMTS-NC	36.2	13.7	8.42	5.57
	HHF	47.9	17.2	9.09	6.38
HHF accel. w.r.t.	DMTS-NC	32.61%	25.92%	7.96%	14.57%
	STS	5.28	4.74	3.56	3.45

here than for DMTS based on FeNNix-Bio1, now ranging from 3.66 to 5.64.

**Table 5. HHF Average Temperatures and Potential Energies**

system	time steps (fs)	temperature (K)	pot. energy (kcal/mol/atom)	
			DMTS-NC	STS
water boxes	1600 mol.	300 ± 4	−8.02 ± 0.02	−8.05 ± 0.02
	4000 mol.	300 ± 2	−8.02 ± 0.01	−8.04 ± 0.01
solvated prot.	phenol-lys.	297 ± 2	−7.87 ± 0.01	−7.87 ± 0.01
	DHFR	298 ± 2	−7.81 ± 0.01	−7.82 ± 0.01


### 3.2. Mitigation of Resonances

In MTS methods, the maximum value of the external time step is usually limited by resonance phenomena<sup>40,41</sup> resulting from a coupling of internal frequencies of the molecular system with nonphysical periodicity created by the various time steps, which can induce numerical instabilities such as abnormal fluctuations of kinetic energy for large values of  $\Delta$ . This phenomenon is well highlighted on velocity autocorrelation spectra, as shown on Figure 4 (computed from the MD trajectory at the level of the smallest time step using the Wiener-Khinchin theorem) for various external time steps. In this spectral representation, the initial peaks represent physical frequencies, such as intramolecular angle bending (around 1000  $\text{cm}^{-1}$ ) and bond stretching (the two peaks around 2000  $\text{cm}^{-1}$ ). Higher-frequency peaks are nonphysical artifacts resulting from MTS. As the external time step  $\Delta$  increases, the main artifact shifts toward lower frequencies, eventually coupling with the physical stretching peaks when  $\Delta$  becomes too large.


To mitigate this problem and push further the stability limit, various alternative methods have been proposed<sup>42–46,66</sup> some of which can be difficult to parametrize, or have a negative impact on certain dynamic properties. As explained in Section 2 and illustrated in Section 3.1, we explored three such strategies, namely Hydrogen Mass Repartitioning (HMR), High Hydrogen-Friction (HHF), and Fast Forward Langevin (FFL), all of which are straightforward to implement and to parametrize. When analyzing MTS strategies, or any method aiming at accelerating simulation speed, a standard sanity check consists in computing the diffusion coefficient, a dynamic property that measures the tendency of particles to drift through the system. A significant drop of this quantity can indicate a reduced effective sampling rate, which in turn offsets the computational gain, as longer trajectories become necessary to keep same amount of sampling. As explained in

Table 6. DMTS-NC Performances with MACE-OFF23(S)

system size/name	water boxes			solvated prot.	
	216 mol.	1600 mol.	4000 mol.	phenol-lys.	DHFR
speed (ns/day)	7.40	1.04	0.38	0.31	0.22
	DMTS-NC	37.59	5.88	1.14	1.00
acceleration	5.08	5.64	5.45	3.66	4.62



**Figure 4.** Hydrogen velocity autocorrelation spectra for various external time steps, in a system of 1600 water molecules. From left to right, the first peak corresponds to molecule libration and slow fluctuations in the hydrogen bond network, the second peak to angle bending oscillations and the third and fourth peaks to bond stretching motions. The ones on the inset graph are nonphysical artifacts. As the external time step increases, these peaks shift to the left, eventually coupling with the bond-oscillation peak at the stability limit.



**Figure 5.** Mean-square displacement (panel (a)) and diffusion coefficients (panel (b)) for various time steps and thermostats.

ref 39, the acceleration rate should ideally be higher than the corresponding loss in diffusion.

Here, we computed diffusion coefficients using the Einstein formula (see, for instance, ref 79), which relates the diffusion coefficient  $D$  to the growth rate of Mean-Square-Displacement (MSD). By denoting  $X_t$  the position of the system at time  $t$ ,  $N$  the number of atoms and  $\mathbb{E}$  the expectation with respect to realizations of the stochastic dynamics, the Einstein formula states that

$$D = \frac{1}{6N} \lim_{t \rightarrow \infty} \frac{d}{dt} \mathbb{E}[|X_t - X_0|^2]$$

In other words, the diffusion coefficient is proportional to the long time limit of the slope of the MSD, which we plot in Figure 5 for a 4800 atoms water box simulated during 20 ns, with a frame saved every picosecond. Quantitative results, gathered in Table 7, are the following. In all explored setups, the percentage of loss in diffusion with respect to STS is significantly lower than the associated increase in speed: for instance, regular DMTS-NC with 1–6 fs time steps only loses 12.6% in diffusion while accelerating simulation by 4.53 (see Section 3.1). In this case, the reduction in diffusion is largely explained by the use of HMR, as a similar drop is observed (12.0% loss) by using HMR on the STS integrator. In HHF simulations, as expected, the increased friction on hydrogen

Table 7. Diffusion Coefficients for Various MTS Strategies<sup>a</sup>

integrator	diffusion coeff. (Å <sup>2</sup> /ps)	loss w.r.t. STS (%)
STS	0.1672 ± 0.0017	
STS-HMR	0.1472 ± 0.0012	12.0
DMTS 1–5	0.1525 ± 0.0007	8.8
DMTS-NC 1–6	0.1461 ± 0.0005	12.6
HHF 1–10 Langevin	0.1016 ± 0.0003	39.2
HHF 1–10 FFL	0.1126 ± 0.0007	32.6

<sup>a</sup>The loss represents the relative decrease in diffusion compared to the STS reference.

atoms results in a larger decrease in diffusion, although replacing Langevin thermostat by Fast Forward Langevin thermostat (see Section 2) reduces the loss from 39.2 to 32.6%. While the impact on diffusion coefficients is indeed higher in HHF, the acceleration in simulation speed is also greater: on the 4800 atoms water system, HHF simulations with 1–10 fs time steps yield 5.28-fold acceleration with respect to standard STS. Overall, the combined use of HMR, HHF, and FFL helps extending the stability limit of DMTS-NC integrators, with a relatively low impact on diffusion coefficient compared to the associated additional speedup.

#### 4. CONCLUSIONS

The DMTS-NC approach marks an improvement in atomistic molecular dynamics with neural network potentials by seamlessly merging model distillation with nonconservative (NC) forces and multi-time-stepping (MTS) strategies. By leveraging a distilled architecture that strictly enforces physical priors—such as equivariance under rotation and the cancellation of atomic force components—this method secures a 15–30% speedup over its conservative predecessors and provides up to a 5.6-fold acceleration over single-time-step methods. Remarkably, the framework remains architecture-agnostic and robust, delivering near-*ab initio* accuracy across foundation models like FeNNix-Biol and MACE-OFF23 without requiring the labor-intensive burden of system-specific fine-tuning. When integrated with stability-extending techniques like Hydrogen Mass Repartitioning and High Hydrogen Friction, DMTS-NC successfully narrows the long-standing performance gap between high-accuracy neural network potentials and efficient classical force fields, setting a new standard for high-throughput, reliable simulations. These advancements represent a further step toward closing the performance gap with classical force fields. Future work will focus on optimizing our implementation to handle more complex systems and investigate subtle sampling properties. As a future direction, we plan to explore further acceleration frameworks, including splitting schemes inspired by RESPA-1 and RESPA-2<sup>38,39</sup> as well as various resonance-mitigating strategies such as isokinetic phase-space algorithm<sup>44</sup> and velocity jumps.<sup>46</sup>

#### ■ ASSOCIATED CONTENT

##### SI Supporting Information

The Supporting Information is available free of charge at <https://pubs.acs.org/doi/10.1021/acs.jctc.6c00653>.

Pressure and density distributions; free energies of solvation: list of studied molecules (NVT and NPT) (PDF)

#### ■ AUTHOR INFORMATION

##### Corresponding Authors

Thomas Plé – Sorbonne Université, Laboratoire de Chimie Théorique, UMR 7616 CNRS, 75005 Paris, France; Email: [thomas.ple@sorbonne-universite.fr](mailto:thomas.ple@sorbonne-universite.fr)

Jean-Philip Piquemal – Qubit Pharmaceuticals, Advanced Research Department, 75014 Paris, France; Sorbonne Université, Laboratoire de Chimie Théorique, UMR 7616 CNRS, 75005 Paris, France; [orcid.org/0000-0001-6615-9426](https://orcid.org/0000-0001-6615-9426); Email: [jean-philip.piquemal@sorbonne-universite.fr](mailto:jean-philip.piquemal@sorbonne-universite.fr)

##### Authors

Nicolaï Gouraud – Qubit Pharmaceuticals, Advanced Research Department, 75014 Paris, France; [orcid.org/0009-0004-0722-037X](https://orcid.org/0009-0004-0722-037X)

Côme Cattin – Sorbonne Université, Laboratoire de Chimie Théorique, UMR 7616 CNRS, 75005 Paris, France

Olivier Adjoua – Sorbonne Université, Laboratoire de Chimie Théorique, UMR 7616 CNRS, 75005 Paris, France

Louis Lagardère – Qubit Pharmaceuticals, Advanced Research Department, 75014 Paris, France; Sorbonne Université, Laboratoire de Chimie Théorique, UMR 7616 CNRS, 75005 Paris, France; [orcid.org/0000-0002-7251-0910](https://orcid.org/0000-0002-7251-0910)

Complete contact information is available at:

<https://pubs.acs.org/10.1021/acs.jctc.6c00653>

##### Author Contributions

N.G., T.P., L.L., and J.-P.P. conceived and designed the research and numerical experiments. N.G. performed the numerical experiments using FeNNix-Bio1. C.C. performed the distillation and tests using MACE-OFF23. N.G. and O.A. wrote the code. All authors analyzed the data. N.G., T.P., and J.-P.P. wrote the paper with the input of all other authors.

##### Notes

The authors declare the following competing financial interest(s): Louis Lagardère and Jean-Philip Piquemal are shareholders and co-founders of Qubit Pharmaceuticals. All the remaining authors declare no conflict of interest.

Calculations were performed using the FeNNol library: <https://github.com/FeNNol-tools/FeNNol> and the Tinker-HP package <https://github.com/TinkerTools/tinker-hp>. Distance to Bound Configuration (DBC) trajectories were computed using the Colvars library <https://github.com/Colvars/colvars>. Pretrained models are available on Github at: <https://github.com/FeNNol-tools/FeNNol-PMC>.

#### ■ ACKNOWLEDGMENTS

This work has received funding from the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation program (grant agreement No 810367), project EMC2 (J.-P.P.). Computations have been performed at IDRIS (Jean Zay) on GENCI Grants: no. A0150712052 (J.-P.P.) and A0180716167 (N.G.).

#### ■ REFERENCES

- (1) Leimkuhler, B.; Matthews, C. *Molecular Dynamics*; Springer, 2015.
- (2) Allen, M. P.; Tildesley, D. J. *Computer Simulation of Liquids*; Oxford University Press, 2017.
- (3) Ciccotti, G.; Dellago, C.; Ferrario, M.; Hernández, E. R.; Tuckerman, M. E. Molecular simulations: past, present, and future (a

Topical Issue in EPJB. *Eur. Phys. J. B* **2022**, *95* (1), No. 3,

DOI: 10.1140/epjb/s10051-021-00249-x.

(4) Jensen, F. *Introduction to Computational Chemistry*; Wiley, 2017.

(5) Ren, P.; Ponder, J. W. Polarizable Atomic Multipole Water Model for Molecular Mechanics Simulation. *J. Phys. Chem. B* **2003**, *107* (24), 5933–5947.

(6) Reif, M. M.; Hünenberger, P. H.; Oostenbrink, C. New interaction Parameters for Charged Amino Acid Side Chains in the GROMOS Force Field. *J. Chem. Theory Comput.* **2012**, *8* (10), 3705–3723.

(7) Maier, J. A.; Martinez, C.; Kasavajhala, K.; Wickstrom, L.; Hauser, K. E.; Simmerling, C. ff14SB: Improving the Accuracy of Protein Side Chain and Backbone Parameters from ff99SB. *J. Chem. Theory Comput.* **2015**, *11* (8), 3696–3713.

(8) Robertson, M. J.; Tirado-Rives, J.; Jorgensen, W. L. Improved Peptide and Protein Torsional Energetics with the OPLS-AA Force Field. *J. Chem. Theory Comput.* **2015**, *11* (7), 3499–3509.

(9) Huang, J.; Rauscher, S.; Nawrocki, G.; et al. CHARMM36m: an improved force field for folded and intrinsically disordered proteins. *Nat. Methods* **2017**, *14*, 71–73.

(10) MacKerell, A. D. Empirical Force Fields for Proteins: Current Status and Future Directions. In *Annual Reports in Computational Chemistry*; Elsevier, 2005; Chapter 7, Vol. 1, pp 91–102.

(11) Behler, J.; Parrinello, M. Generalized neural-network representation of high-dimensional potential-energy surfaces. *Phys. Rev. Lett.* **2007**, *98* (14), No. 146401.

(12) Bartók, A. P.; Payne, M. C.; Kondor, R.; Csányi, G. Gaussian approximation potentials: The accuracy of quantum mechanics, without the electrons. *Phys. Rev. Lett.* **2010**, *104* (13), No. 136403.

(13) Shakouri, K.; Behler, J.; Meyer, J.; Kroes, G.-J. Accurate neural network description of surface phonons in reactive gas-surface dynamics: N<sub>2</sub>+ Ru (0001). *J. Phys. Chem. Lett.* **2017**, *8* (10), 2131–2136.

(14) Smith, J. S.; Isayev, O.; Roitberg, A. E. ANI-1: an extensible neural network potential with DFT accuracy at force field computational cost. *Chem. Sci.* **2017**, *8* (4), 3192–3203.

(15) Schütt, K. T.; Kindermans, P.-J.; Sauceda, H. E.; Chmiela, S.; Tkatchenko, A.; Müller, K.-R. In *SchNet: A Continuous-Filter Convolutional Neural Network for Modeling Quantum Interactions*, Proceedings of the 31st International Conference on Neural Information Processing Systems. NIPS'17; Curran Associates Inc.: Long Beach, California, USA, 2017; pp 992–1002.

(16) Chmiela, S.; Sauceda, H. E.; Müller, K.-R.; Tkatchenko, A. Towards exact molecular dynamics simulations with machine-learned force fields. *Nat. Commun.* **2018**, *9* (1), No. 3887.

(17) Grisafi, A.; Ceriotti, M. Incorporating long-range physics in atomic-scale machine learning. *J. Chem. Phys.* **2019**, *151* (20), No. 204105.

(18) Zubatyuk, R.; Smith, J. S.; Leszczynski, J.; Isayev, O. Accurate and transferable multitask prediction of chemical properties with an atoms-in-molecules neural network. *Sci. Adv.* **2019**, *5* (8), No. eaav6490.

(19) Devereux, C.; Smith, J. S.; Huddleston, K. K.; et al. Extending the Applicability of the ANI Deep Learning Molecular Potential to Sulfur and Halogens. *J. Chem. Theory Comput.* **2020**, *16* (7), 4192–4202.

(20) Unke, O. T.; Chmiela, S.; Gastegger, M.; Schütt, K. T.; Sauceda, H. E.; Müller, K.-R. SpookyNet: Learning force fields with electronic degrees of freedom and nonlocal effects. *Nat. Commun.* **2021**, *12* (1), No. 7273.

(21) Gasteiger, J.; Becker, F.; Günnemann, S. In *Gemnet: Universal Directional Graph Neural Networks for Molecules*, Advances in Neural Information Processing Systems 34, 2021; pp 6790–6802.

(22) Grisafi, A.; Nigam, J.; Ceriotti, M. Multi-scale approach for the prediction of atomic scale properties. *Chem. Sci.* **2021**, *12* (6), 2078–2090.

(23) Musaelian, A.; Batzner, S.; Johansson, A.; et al. Learning local equivariant representations for large-scale atomistic dynamics. *Nat. Commun.* **2023**, *14* (1), No. 579.

(24) Plé, T.; Lagardère, L.; Piquemal, J.-P. Force-field-enhanced neural network interactions: from local equivariant embedding to atom-in-molecule properties and long-range effects. *Chem. Sci.* **2023**, *14* (44), 12554–12569.

(25) Unke, O. T.; Stöhr, M.; Ganscha, S.; et al. Biomolecular dynamics with machine-learned quantum-mechanical force fields trained on diverse chemical fragments. *Sci. Adv.* **2024**, *10* (14), No. eadn4397.

(26) Kabylda, A.; Frank, J. T.; Suárez-Dou, S.; et al. Molecular Simulations with a Pretrained Neural Network and Universal Pairwise Force Fields. *J. Am. Chem. Soc.* **2025**, *147* (37), 33723–33734.

(27) Schütt, K. T.; Hessmann, S. S. P.; Gebauer, N. W. A.; Lederer, J.; Gastegger, M. SchNetPack 2.0: A neural network toolbox for atomistic machine learning. *J. Chem. Phys.* **2023**, *158* (14), No. 144801.

(28) Zeng, J.; Zhang, D.; Lu, D.; et al. DeePMD-kit v2: A software package for deep potential models. *J. Chem. Phys.* **2023**, *159* (5), No. 054801, DOI: 10.1063/5.0155600.

(29) Dral, P. O.; Ge, F.; Hou, Y. F.; et al. MLatom 3: A platform for machine learning-enhanced computational chemistry simulations and workflows. *J. Chem. Theory Comput.* **2024**, *20* (3), 1193–1213.

(30) Plé, T.; Adjoua, O.; Lagardère, L.; Piquemal, J.-P. FeNNol: An efficient and flexible library for building force-field-enhanced neural network potentials. *J. Chem. Phys.* **2024**, *161* (4), No. 042502.

(31) Batatia, I.; Kovacs, D. P.; Simm, G.; Ortner, C.; Csányi, G. In *MACE: Higher Order Equivariant Message Passing Neural Networks for Fast and Accurate Force Fields*, Advances in Neural Information Processing Systems 35, 2022; pp 11423–11436.

(32) Kovács, D. P.; Moore, J. H.; Browning, N. J.; et al. MACE-OFF: Short-Range Transferable Machine Learning Force Fields for Organic Molecules. *J. Am. Chem. Soc.* **2025**, *147* (21), 17598–17611.

(33) Plé, T.; Adjoua, O.; Benali, A. et al. A Foundation Model for Accurate Atomistic Simulations in Drug Design *ChemRxiv* 2025 DOI: 10.26434/chemrxiv-2025-flhgn-v4.

(34) Benali, A.; Plé, T.; Adjoua, O. et al. Pushing the Accuracy Limit of Foundation Neural Network Models with Quantum Monte Carlo Forces and Path Integrals. 2025, arXiv:2504.07948. arXiv.org e-Print archive. <https://arxiv.org/abs/2504.07948>.

(35) Leimkuhler, B.; Matthews, C. Robust and efficient configurational molecular sampling via Langevin dynamics. *J. Chem. Phys.* **2013**, *138* (17), No. 174102.

(36) Leimkuhler, B.; Matthews, C. Rational Construction of Stochastic Numerical Methods for Molecular Sampling. *Appl. Math. Res. Express* **2012**, *2013* (1), 34–56.

(37) Tuckerman, M.; Berne, B. J.; Martyna, G. J. Reversible multiple time scale molecular dynamics. *J. Chem. Phys.* **1992**, *97* (3), 1990–2001.

(38) Zhou, R.; Harder, E.; Xu, H.; Berne, B. Efficient multiple time step method for use with Ewald and particle mesh Ewald for large biomolecular systems. *J. Chem. Phys.* **2001**, *115* (5), 2348–2358.

(39) Lagardère, L.; Aviat, F.; Piquemal, J.-P. Pushing the Limits of Multiple-Time-Step Strategies for Polarizable Point Dipole Molecular Dynamics. *J. Phys. Chem. Lett.* **2019**, *10* (10), 2593–2599.

(40) Biesiadecki, J. J.; Skeel, R. D. Dangers of Multiple Time Step Methods. *J. Comput. Phys.* **1993**, *109* (2), 318–328.

(41) Ma, Q.; Izaguirre, J. A.; Skeel, R. D. Verlet-I/R-RESPA/Impulse is Limited by Nonlinear Instabilities. *SIAM J. Sci. Comput.* **2003**, *24* (6), 1951–1973.

(42) Morrone, J. A.; Markland, T. E.; Ceriotti, M.; Berne, B. J. Efficient multiple time scale molecular dynamics: Using colored noise thermostats to stabilize resonances. *J. Chem. Phys.* **2011**, *134* (1), No. 014103.

(43) Albaugh, A.; Tuckerman, M. E.; Head-Gordon, T.; Combining, I. Iteration-Free Polarization with Large Time Step Stochastic-Isokinetic Integration. *J. Chem. Theory Comput.* **2019**, *15* (4), 2195–2205.

(44) Leimkuhler, B.; Margul, D. T.; Tuckerman, M. E. Stochastic, resonance-free multiple time-step algorithm for molecular dynamics

- with very large time steps. *Mol. Phys.* **2013**, *111* (22–23), 3579–3594.
- (45) Margul, D. T.; Tuckerman, M. E. A Stochastic, Resonance-Free Multiple Time-Step Algorithm for Polarizable Models That Permits Very Large Time Steps. *J. Chem. Theory Comput.* **2016**, *12* (5), 2170–2180.
- (46) Gouraud, N.; Lagardère, L.; Adjoua, O.; Plé, T.; Monmarché, P.; Piquemal, J.-P. Velocity Jumps for Molecular Dynamics. *J. Chem. Theory Comput.* **2025**, *21* (6), 2854–2866.
- (47) Liberatore, E.; Meli, R.; Rothlisberger, U. A Versatile Multiple Time Step Scheme for Efficient ab Initio Molecular Dynamics Simulations. *J. Chem. Theory Comput.* **2018**, *14* (6), 2834–2842.
- (48) Inizan, T. J.; Plé, T.; Adjoua, O.; et al. Scalable hybrid deep neural networks/polarizable potentials biomolecular simulations including long-range effects. *Chem. Sci.* **2023**, *14* (20), 5438–5452.
- (49) Fu, X.; Musaelian, A.; Johansson, A.; Jaakkola, T.; Kozinsky, B. In *Learning Interatomic Potentials at Multiple Scales*, NeurIPS 2023 AI for Science Workshop, 2023.
- (50) Osadchey, R.; Nam, K.; Cui, Q. Toward Improving Multiple Time Step QM/MM Simulations with  $\Delta$ -Machine Learning. *J. Phys. Chem. B* **2025**, *129* (40), 10451–10466.
- (51) Mouvet, F.; Browning, N. J.; Baudin, P.; Liberatore, E.; Rothlisberger, U. Machine learning-enhanced multiple time-step ab initio molecular dynamics. *J. Chem. Phys.* **2025**, *163* (18), No. 184115.
- (52) Cattin, C.; Plé, T.; Adjoua, O.; Gouraud, N.; Lagardère, L.; Piquemal, J.-P. Accelerating Molecular Dynamics Simulations with Foundation Neural Network Models Using Multiple Time Steps and Distillation. *J. Phys. Chem. Lett.* **2026**, *17* (5), 1288–1295.
- (53) Hinton, G.; Vinyals, O.; Dean, J. Distilling the Knowledge in a Neural Network. 2015, arXiv:1503.02531. arXiv.org e-Print archive. <https://arxiv.org/abs/1503.02531>.
- (54) Gou, J.; Yu, B.; Maybank, S. J.; Tao, D. Knowledge distillation: A survey. *Int. J. Comput. Vis.* **2021**, *129* (6), 1789–1819.
- (55) Hu, W.; Shuaibi, M.; Das, A. et al. Forcenet: A Graph Neural Network for Large-Scale Quantum Calculations. 2021, arXiv:2103.01436. arXiv.org e-Print archive. <https://arxiv.org/abs/2103.01436>.
- (56) Liao, Y.-L.; Wood, B.; Das, A.; Smidt, T. Equiformerv2: Improved Equivariant Transformer for Scaling to higher-Degree Representations. 2023, arXiv:2306.12059. arXiv.org e-Print archive. <https://arxiv.org/abs/2306.12059>.
- (57) Neumann, M.; Gin, J.; Rhodes, B. et al. Orb: A Fast, Scalable Neural Network Potential. 2024, arXiv:2410.22570. arXiv.org e-Print archive. <https://arxiv.org/abs/2410.22570>.
- (58) Rhodes, B.; Vandenhaute, S.; Šimkus, V. et al. Orb-v3: Atomistic Simulation at Scale. 2025, arXiv:2504.06231. arXiv.org e-Print archive. <https://arxiv.org/abs/2504.06231>.
- (59) Eissler, M.; Korjakow, T.; Ganscha, S.; Unke, O. T.; Müller, K.-R.; Gugler, S. How Simple Can You Go? An Off-the-Shelf Transformer Approach to Molecular Dynamics. 2025, arXiv:2503.01431. arXiv.org e-Print archive. <https://arxiv.org/abs/2503.01431>.
- (60) Bigi, F.; Langer, M. F.; Ceriotti, M. In *The Dark Side of the Forces: Assessing Non-Conservative Force Models for Atomistic Machine Learning*, Forty-Second International Conference on Machine Learning, 2025.
- (61) Gasteiger, J.; Groß, J.; Günnemann, S. Directional Message Passing for Molecular Graphs. 2003, arXiv:2003.03123. arXiv.org e-Print archive. <https://arxiv.org/abs/2003.03123>.
- (62) Nordlund, K.; Lehtola, S.; Hobler, G. Repulsive interatomic potentials calculated at three levels of theory. *Phys. Rev. A* **2025**, *111* (3), No. 032818.
- (63) Eastman, P.; Behara, P. K.; Dotson, D. L.; et al. SPICE, A Dataset of Drug-like Molecules and Peptides for Training Machine Learning Potentials. *Sci. Data* **2023**, *10* (1), No. 11.
- (64) Eastman, P.; Pritchard, B. P.; Chodera, J. D.; Markland, T. E. Nutmeg and SPICE: Models and Data for Biomolecular Machine Learning. *J. Chem. Theory Comput.* **2024**, *20* (19), 8583–8593.
- (65) Jordan, K.; Jin, Y.; Boza, V. et al. *Muon: An Optimizer for Hidden Layers in Neural Networks*, 2024.
- (66) Feenstra, K. A.; Hess, B.; Berendsen, H. J. C. Improving efficiency of large time-scale molecular dynamics simulations of hydrogen-rich systems. *J. Comput. Chem.* **1999**, *20* (8), 786–798.
- (67) Chak, M.; Kantas, N.; Lelièvre, T.; Pavliotis, G. A. Optimal friction matrix for underdamped Langevin sampling. *ESAIM: Math. Modell. Numer. Anal.* **2023**, *57* (6), 3335–3371.
- (68) Hijazi, M.; Wilkins, D. M.; Ceriotti, M. Fast-forward Langevin dynamics with momentum flips. *J. Chem. Phys.* **2018**, *148* (18), No. 184109.
- (69) Chow, K.-H.; Ferguson, D. M. Isothermal-isobaric molecular dynamics simulations with Monte Carlo volume sampling. *Comput. Phys. Commun.* **1995**, *91* (1), 283–289.
- (70) Åqvist, J.; Wennerström, P.; Nervall, M.; Bjelic, S.; Brandsdal, B. O. Molecular dynamics simulations of water and biomolecules with a Monte Carlo constant pressure algorithm. *Chem. Phys. Lett.* **2004**, *384* (4), 288–294.
- (71) Berendsen, H. J. C.; Postma, J. P. M.; van Gunsteren, W. F.; DiNola, A.; Haak, J. R. Molecular dynamics with coupling to an external bath. *J. Chem. Phys.* **1984**, *81* (8), 3684–3690.
- (72) Feller, S. E.; Zhang, Y.; Pastor, R. W.; Brooks, B. R. Constant pressure molecular dynamics simulation: The Langevin piston method. *J. Chem. Phys.* **1995**, *103* (11), 4613–4621.
- (73) Bussi, G.; Zykova-Timan, T.; Parrinello, M. Isothermal-isobaric molecular dynamics using stochastic velocity rescaling. *J. Chem. Phys.* **2009**, *130* (7), No. 074101.
- (74) Rackers, J. A.; Wang, Z.; Lu, C.; et al. Tinker 8: Software Tools for Molecular Design. *J. Chem. Theory Comput.* **2018**, *14* (10), 5273–5289.
- (75) Lagardère, L.; Jolly, L. H.; Lipparini, F.; et al. Tinker-HP: a massively parallel molecular dynamics package for multiscale simulations of large complex systems with advanced point dipole polarizable force fields. *Chem. Sci.* **2018**, *9* (4), 956–972.
- (76) Adjoua, O.; Lagardère, L.; Jolly, L. H.; et al. Tinker-hp: Accelerating molecular dynamics simulations of large complex systems with advanced point dipole polarizable force fields using gpus and multi-gpu systems. *J. Chem. Theory Comput.* **2021**, *17* (4), 2034–2053.
- (77) Fiorin, G.; Klein, M. L.; Héning, J. Using collective variables to drive molecular dynamics simulations. *Mol. Phys.* **2013**, *111* (22–23), 3345–3362.
- (78) Lagardère, L.; Maurin, L.; Adjoua, O.; et al. Lambda-ABF: Simplified, Portable, Accurate, and Cost-Effective Alchemical Free-Energy Computation. *J. Chem. Theory Comput.* **2024**, *20* (11), 4481–4498.
- (79) Tuckerman, M. *Statistical Mechanics: Theory and Molecular Simulation*; Oxford Graduate Texts: Oxford, 2010.