

Quantum machine learning corrects classical forcefields: Stretching DNA base pairs in explicit solvent

Cite as: J. Chem. Phys. **157**, 064107 (2022); <https://doi.org/10.1063/5.0094727>

Submitted: 05 April 2022 • Accepted: 13 July 2022 • Published Online: 09 August 2022

 Joshua T. Berryman,  Amirhossein Taghavi, Florian Mazur, et al.



View Online



Export Citation



CrossMark

ARTICLES YOU MAY BE INTERESTED IN

[Automatic derivation of many-body theories based on general Fermi vacua](#)

The Journal of Chemical Physics **157**, 064111 (2022); <https://doi.org/10.1063/5.0097858>

[Atomistic global optimization X: A Python package for optimization of atomistic structures](#)

The Journal of Chemical Physics **157**, 054701 (2022); <https://doi.org/10.1063/5.0094165>

[Orbital pair selection for relative energies in the domain-based local pair natural orbital coupled-cluster method](#)

The Journal of Chemical Physics **157**, 064102 (2022); <https://doi.org/10.1063/5.0100010>

Lock-in Amplifiers
up to 600 MHz



Zurich
Instruments



Quantum machine learning corrects classical forcefields: Stretching DNA base pairs in explicit solvent

Cite as: J. Chem. Phys. 157, 064107 (2022); doi: 10.1063/5.0094727

Submitted: 5 April 2022 • Accepted: 13 July 2022 •

Published Online: 9 August 2022



Joshua T. Berryman,^{a)} Amirhossein Taghavi,^{b)} Florian Mazur,^{b)} and Alexandre Tkatchenko^{b)}

AFFILIATIONS

Department of Physics and Materials Science, University of Luxembourg, L-1511 Luxembourg City, Luxembourg

^{a)} Author to whom correspondence should be addressed: josh.berryman@uni.lu

^{b)} Also at: Laboratoire de Physique et Chimie Théoriques, Université de Lorraine, Faculté des Sciences et Technologies, Boulevard des Aiguillettes, B.P. 70239, F-54506 Vandoeuvre-les-Nancy, France.

ABSTRACT

In order to improve the accuracy of molecular dynamics simulations, classical forcefields are supplemented with a kernel-based machine learning method trained on quantum-mechanical fragment energies. As an example application, a potential-energy surface is generalized for a small DNA duplex, taking into account explicit solvation and long-range electron exchange–correlation effects. A long-standing problem in molecular science is that experimental studies of the structural and thermodynamic behavior of DNA under tension are not well confirmed by simulation; study of the potential energy vs extension taking into account a novel correction shows that leading classical DNA models have excessive stiffness with respect to stretching. This discrepancy is found to be common across multiple forcefields. The quantum correction is in qualitative agreement with the experimental thermodynamics for larger DNA double helices, providing a candidate explanation for the general and long-standing discrepancy between single molecule stretching experiments and classical calculations of DNA stretching. The new dataset of quantum calculations should facilitate multiple types of nucleic acid simulation, and the associated Kernel Modified Molecular Dynamics method (KMMD) is applicable to biomolecular simulations in general. KMMD is made available as part of the AMBER22 simulation software.

© 2022 Author(s). All article content, except where otherwise noted, is licensed under a Creative Commons Attribution (CC BY) license (<http://creativecommons.org/licenses/by/4.0/>). <https://doi.org/10.1063/5.0094727>

I. INTRODUCTION

The dominant long-range interactions in molecular systems can be treated classically with great efficiency as a pairwise Coulomb term ($U_c \propto 1/r$) plus a pairwise fluctuating dipole–dipole attraction (a dispersion interaction, $U_d \propto 1/r^6$), with atoms represented as spheres having fitted non-integer charge and attraction parameters, both of which in theory can be derived from *ab initio* quantum mechanics.^{1–3} Limitations to the accuracy of this approach are especially evident when atomic polarizability is anisotropic or when the atoms are part of a flexible molecule in which the atomic polarizability and the partial charge can couple to the molecular conformation. These two limitations apply very much to nucleic acids where the charged backbone, anisotropically delocalized electrons in the aromatic bases, polar solvent, and close localization of cations

give a strongly many-body character to the interatomic non-bonded forces.^{4,5}

Although the limitations of classical forcefields are understood, with large discrepancies noted in less-standard conformations, such as single stranded DNA,⁶ due to the complexity of nucleic acid molecules and the important role of dispersion interactions, it has been quite difficult to arrive at solid quantum benchmark calculations to improve these models. Substantial progress has been made in recent years by the calculation of dispersion interactions in an explicit many-body manner, working with tensors of directional atom-centered polarizabilities (for example, in the many-body dispersion method⁷). This approach was found to be quantitative for DNA stacking energies⁸ and benchmarked favorably against other methods over a dataset of “challenging” non-covalent complexes, including DNA bases.⁹

Evaluation of the potential energies of static molecular conformations has limited relevance to the highly dynamic finite-temperature behavior of biomolecules: ideally, in order to study large nucleic acid molecules in explicit solvent and salt, a potential energy surface (PES) should be defined over the nuclear coordinates. The PES must be differentiable and inexpensive to compute, scaling better than $\mathcal{O}(N^2)$ in the number of atoms. In order to integrate the accuracy of the *ab initio* calculations to the existing efficient infrastructure of classical forcefields, a Kernel Modified Molecular Dynamics (KMMD) approach is devised here.

The KMMD approach, described in detail below (Sec. II), uses a kernel machine to learn not the full PES, but only the part of the quantum correction to the classical PES susceptible to prediction by selected internal degrees of freedom of the nucleic acid chain. Solvent-solvent interactions are thus calculated as usual for the chosen classical forcefield, and solvent-DNA interactions are modified only when they can be predicted by the DNA conformation. This is a form of coarse-graining: degrees of freedom that are not explicitly corrected may still be corrected on average. DNA-DNA interactions are only modified insofar as can be predicted from the training data, for instance, because information relating to bond lengths was excluded from the training data, bond lengths are only improved in the KMMD as far as they are driven by correlations with the explicitly treated torsions. The role of solvent is significant for nucleic acids, so it is important to understand that although the KMMD correction presented is not a correction to solvent interactions, because solvent was present in the training data, it is nonetheless an implicit correction to interactions in a solvated system.

Having implemented the KMMD method, simulations were made of a GG-CC DNA complex (with Watson-Crick base pairing artificially imposed), under constant force at the O3' and O5' termini, before and after quantum corrections. The observed small reduction in the work to extend with base pairing preserved was consistent with the large long-standing difference between single-molecule stretching experiments on DNA polymers and classical simulation (comparing simulations^{10,11} against a review of experimental data,¹² work in experiment is 50%–70% of work in classical simulations). While a change in work to extend single steps should imply a change of the same sign for large polymers, for the present prototype study, only a small GG-CC complex was studied, and direct experimental measurement of single-step stretching free energies in order to make a quantitative comparison is non-trivial. Experimental measurement of single base stacking energy *has* been made for a comparable proxy system, a nicked duplex.¹³ This setup is not exactly equivalent to the calculation geometry; however, the experimental energies are closer to the KMMD corrected calculation than to the uncorrected (*vide infra*: Sec. III).

Various alternative hypotheses exist in relation to the discrepancy between experimental and calculated DNA extension dynamics, and not all are mutually inconsistent. It is possible that experimental systems contain chemical defects not accounted for in classical simulation or by KMMD: nicking of the backbone due to some form of tension-accelerated hydrolysis, for example, or occasional base-pair mismatches due to failure of polymerase enzymes. On the simulation side, apart from the forcefield errors documented here, there is the problem of scaling up results from a simulation on a microseconds \times nanometers scale to experiments in the world of seconds \times millimeters and from seconds \times millimeters

to the abstraction of a fully adiabatic, quasi-equilibrium stretching. Advanced non-equilibrium data analysis using the Jarzynski equality is often applied to make the scaling-up in time, even when starting from the relatively long timescales of experiment.¹⁴ A corresponding recovery of equilibrium stretching behavior from continuous stretching simulations requires care and expertise;¹⁵ in DNA, there are certain very rare dynamical fluctuations, such as base breathing, which might well have an impact on the stretching process but which are expected to manifest only on microsecond timescales,¹⁶ thus requiring many microseconds for effective sampling. Finite-size thermodynamics arising through the generalization from small to large systems are equally subtle and difficult to treat for a complex molecule, such as DNA.¹⁷ Finally, although the present study addresses collective electronic effects in the DNA, collective electron fluctuations in the solvent are not treated beyond the nearest few molecules, and these have been shown to be important for biomolecules up to the 20 Å scale.¹⁸

II. METHODS

A. Generation of DNA training data

69 GG-CC DNA dyad structures were sampled from existing classical atomistic simulations of 24bp (24 base pair) DNA duplexes at varying extensions documented in previous work.¹⁹ Hydrogens were added to the O5' and O3' ends. Each dyad structure was solvated with 2446 TIP3P waters²⁰ and two sodium ions (neutralizing the system), and the solvent box was then allowed to equilibrate while holding the dyads in stiff restraints for 10 ns. This classical simulation and processing was done using AmberTools.²¹ The OL15 forcefield²² was used. Dynamics were classical (no nuclear quantum effects were treated), allowing all atomic masses to be set to 12 amu.

From the initial 69 structures, fresh conformations were bootstrapped from long KMMD simulations with a total time of 1 μ s (sampling one per ns) to make a training set of 1000 solvated 2bp duplexes. The bootstrapping process was designed as a carousel of 69 independent runs, each restrained to have roughly constant extension O3'–O5' and each restarting after a 10 ns window, to use fresh training data generated by all previous runs. When a total of 1 μ s dynamics was achieved (1000 frames at one per nanosecond), the bootstrapping process was stopped.

For each training snapshot, a structure of the DNA with 20 water molecules and two sodium ions surrounding it in a cluster was created. The water molecules were selected based on the 20 closest to any solute atom. This number was found to be sufficient such that backbone charges and the solute-solvent hydrogen bonds were always coordinated. Energies were found for the cluster structures using density functional theory (DFT) calculations with the hybrid PBE0 functional^{23–25} for semi-local electron exchange and correlation together with the MBD method^{7,26} for long-range electron correlation interactions (including van der Waals forces). The FHI-aims code was used to converge the electron density,²⁷ with the associated atom-centered basis sets represented at the “intermediate” level where available (H, C, N, O) but at the “tight” level when intermediate basis sets were not available for the given element (Na, P). Memory requirement for a neutralized dyad plus 20 water molecules was \sim 12 GB, with convergence taking \sim 12 h on 8 cores. Two structures that did not converge within 72 h were

rejected as having unphysical geometry, leaving 998 initial reference DNA structures for use in the first stage of the machine learning calculation.

B. Kernel method for generalization to PES

The training data points $\{\Phi_i(\vec{x}_i)\}$ were defined as the residuals of the quantum and classical energies: $\Phi_i = H_Q(\vec{x}_i) - H_{AMBER}(\vec{x}_i)$. Evaluations of $\{H_{AMBER}\}$ are made at the start of each run for convenient re-use of the quantum data across forcefields. Feature extraction was carried out in order to reduce the degrees of freedom of the problem by defining a vector of dihedral angles $\vec{\phi}_i(\vec{x}_i)$ based on the Cartesian coordinates. The canonical DNA torsions $(\alpha, \beta, \gamma, \delta, \epsilon, \zeta, \chi, \nu_0, \nu_1, \nu_2, \nu_3, \nu_4)$ were collected, with one occurrence per strand of the backbone torsions $\alpha, \beta, \gamma, \epsilon, \zeta$ and two occurrences per strand of the sugar torsions δ, ν_{0-4} and the base twisting angle χ , giving a total of 38 angular degrees of freedom to summarize the state of a base-pair step. These were then (reversibly) expanded to vectors \vec{p}_i of 76 Cartesian degrees of freedom by taking the sin and cos of each angular degree of freedom: $\vec{p}_i = (\sin \vec{\phi}_i, \cos \vec{\phi}_i)$.

The feature space of transformed angles forms a toroidal manifold embedded in \mathbb{R}^{76} . The maximum Euclidean distance between two points is $r_{\max} = 2\sqrt{38} \approx 12.3$, and the average distance between two random points is $\sqrt{76} \approx 8.7$. For small angular differences, the distance in feature space is very close to the imaged distance in angular space (perturbing a single angle by 0.1 rad implies $r = 0.0996$ in the feature space, 1 rad implies $r = 0.959$, and π rad implies 2). The PES correction at a feature space point \vec{q} is estimated from the training set using a single-parameter Gaussian kernel,

$$\Phi(\vec{q}) = \frac{\sum_i \Phi_i \exp(-\|\vec{q} - \vec{p}_i\|^2 / \sigma^2)}{\sum_i \exp(-\|\vec{q} - \vec{p}_i\|^2 / \sigma^2)}, \quad (1)$$

where $\sigma^2 := 0.1$ for all results discussed here. In the limit of small σ , the PES becomes a Voronoi diagram in the feature space (losing differentiability). In the limit of large σ , the PES becomes flat. For any σ , the PES deals gracefully with duplicate training points, simply averaging their contributions: this is valuable because it permits degrees of freedom not represented in the feature space to be averaged over; however, the caveat is introduced that if multiple copies of training points exist, untreated degrees of freedom should preferably be sampled using the Gibbs measure at the target temperature for simulation. The denominator in Eq. (1) is treated as an empirical confidence value for a given point evaluation. This kernel method, a minor variation of the extant family of kernel methods, is referred to as a “Normalized Radial Basis Function” (“nRBF”) kernel machine or, equivalently, as an nRBF network.²⁸

To calculate the force, first define $p_{i,d}$ as the displacement of point i in dimension d in the Cartesian feature space from the reference position of training point i , having D treated dihedral angles. For even $d \in [0, 2, \dots, 2D - 2]$, $p_{i,d} = \cos \phi_{i,d/2} - \cos \phi_{i,d/2}^0$; for odd $d \in [1, 3, \dots, 2D - 1]$, $p_{i,d} = \sin \phi_{i,(d-1)/2} - \sin \phi_{i,(d-1)/2}^0$. $R_i^2 = \sum_{d \in [0, 2D-1]} p_{i,d}^2$.

We write a “weight” for the i th training point, w_i , as

$$w_i = \exp -R_i^2 / \sigma^2.$$

The normalization term follows

$$Z = \sum_i w_i.$$

The energy delta is

$$\Phi = \frac{\sum_i \Phi_i w_i}{Z}.$$

Differentiating, using a primed notation for the partial derivative of a function in dimension d for point i : $f' = \frac{\partial f}{\partial p_{i,d}}$,

$$w' = -2p_{i,d} w_i / \sigma^2, \\ Z' = w'.$$

Following the quotient rule,

$$\Phi' = \frac{\Phi_i w' Z - Z' \sum_i \Phi_i w_i}{Z^2}.$$

Substituting $Z' = w' = -2p_{i,d} w_i / \sigma^2$,

$$\Phi' = -2p_{i,d} w_i \frac{\Phi_i Z - \sum_i \Phi_i w_i}{\sigma^2 Z^2}.$$

Substituting $\Phi = \sum_i \Phi_i w_i / Z$,

$$\Phi' = -2p_{i,d} w_i \frac{(\Phi_i - \Phi)}{\sigma^2 Z}. \quad (2)$$

The Cartesian feature space has dimension $2D$, where D is the number of torsions studied. Forces are not, in general, tangent to the subspace of the Cartesian space, which can be mapped directly back to the space of the angles. Therefore, they are projected onto the subspace. By construction, the force gradients are linear with respect to individual training points and angles but must be treated in pairs for the purpose of projection back to the subspace of dihedral angles. The 2D unit radius vector corresponding to angle $\phi_{i,d}$ for $d \in [1, D]$ is already known as

$$\hat{u}_{i,d} = \cos \phi_{i,d} \hat{e}_1 + \sin \phi_{i,d} \hat{e}_2, \\ \hat{u}_{i,d} = p_{i,2d} \hat{e}_1 + p_{i,2d+1} \hat{e}_2.$$

To project the orthogonal to a unit vector, we subtract a term scaled by the inner product,

$$\vec{\Phi}'_d \leftarrow \vec{\Phi}'_d - (\hat{u}_{i,d} \cdot \vec{\Phi}'_d) \hat{u}_{i,d}. \quad (3)$$

This vector can then be converted directly into an angular force.

C. Validation of machine learning approach

Validation of the KMMD approach was carried out by subsampling a set of 1232 frames and then evaluating the trained system against the remaining unseen frames. Mean Absolute Error (MAE) values shown are averages over five independent selections of training/test sets. Noise is greater for larger training sets due to the smaller validation sets. The KMMD was robust at predicting energies regardless of the σ^2 parameter and made more efficient use

of training data than the other, more sophisticated, approaches. The “normalized Radial Basis Function Network” (nRBF) has an identical functional form to that of the KMMD [Eq. (1)], except that Φ_i are not training point energies but arbitrary values fixed by a least-squares fit to the training data.²⁹ In theory, this fitting should allow performance improvement; however, with a relatively modest number of training points (as in this example, where the number of centers is still small enough to be computationally cheap at evaluation time), there is no gain but potentially some minor loss due to the under-constraint of the fitting problem. Relaxing the constraint that the weights should be normalized [removing the denominator of Eq. (1)] converts the system to a standard RBF network. This removal of a constraint gives even greater freedom in fitting; however, no benefit was observed.

Kernel Ridge Regression (KRR) is a generalization of the RBF network, which has been used successfully to fit atomization energies of small molecules *in vacuo* using a feature space constructed around the Coulomb matrix of atom–atom distances.³⁰ The Laplacian kernel $e^{-|x|}$ recommended in this related example was tested in the present feature space and found to perform about as well as the Gaussian nRBF for some choices of the KRR regularization parameter α . Although the KMMD performed best on the present dataset, all methods have the potential for optimization and for efficiency gains against larger datasets. In the present examples, all training points became nodes of evaluation. For larger training sets, it is possible to merge or prune training points or even to use non-physically constructed points as RBF centers, gaining efficiency at evaluation time in exchange for a small or zero loss of accuracy (Fig. 1).

D. Constant force simulations

In order to monitor the influence of extensional force on the DNA dynamics, 36 parallel simulation replicates were prepared with the four terminus atoms G,C@O_{5'},O_{3'} subjected to constant-force

restraints of strengths (15, 20, ..., 185, 190) pN, oriented toward fixed sites at the top and bottom of the simulation box. Each replicate was equilibrated in a box of 2446 TIP3P water molecules and two sodium ions for 50 ns. During equilibration, frames evaluated with a low confidence [small denominator of Eq. (1)] were saved, and those points with the lowest confidence were used to augment the training set to 1232 frames. In production, force-extension time series were collected for 1500 ns per replicate, with a 1 fs time step.

The calculation was run with DNA bases restrained to Watson–Crick pairing, following standard B-DNA geometry,³¹ with the intention to therefore focus on the behavior of stack interactions. In detail, the restraints were such that there was no energy penalty for the proton–acceptor pairs G@N₁H–C@N₃, G@NH₂–C@O, and C@H₂–G@O to be within distances ± 0.2 Å of their equilibrium values, 1.94, 1.85, and 2.95 Å, respectively. Beyond these ranges, harmonic restraints were imposed with a spring constant of 20 kcal mol^{−1} Å^{−2}, for a further 0.5 Å, after which restraints became constant-force. For the two stronger hydrogen bonds, G@N₁–C@N₃ and C@N–G@O donor–acceptor distance restraints were also added in the same pattern, at equilibrium distances of 2.5 and 2.51 Å, respectively, thus (weakly) enforcing planarity of the base pairs.

E. Equilibrium potential of mean force calculations

As an examination of the equilibrium behavior of the KMMD calculation, Potentials of Mean Force (PMF) over ϵ, ζ, χ dihedral angles were calculated using sets of multiple parallel simulations with and without KMMD. Sampling was enhanced using the technique of Hamiltonian Replica Exchange (HREM) within the sets of 120 simulations, which individually sampled umbrellas spaced 3° apart in the torsion for which the given PMF was desired. Umbrella torsion restraints were harmonic, with the spring constant 250 kcal mol^{−1} rad^{−2}. The two chains were kept together

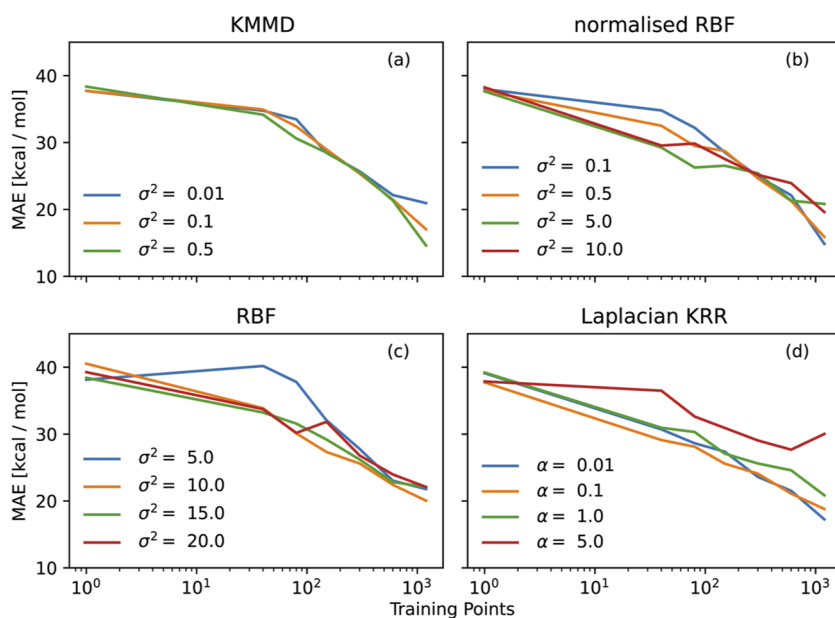


FIG. 1. Validation of KMMD implementation (a) against related machine learning methods, including a conventional nRBF (b) with weights Φ_i as trained parameters. (c) and (d) Performance is tested against unseen validation data for increasing amounts of seen training data. The basic KMMD approach is stable with respect to the parameters and performs well.

using a weak flat-bottomed distance restraint between the two phosphorous atoms, with zero force at distances between 20 and 30 Å but a harmonic force with a spring constant of 20 kcal mol⁻¹ Å⁻² thereafter. Exchange flux between replicas was verified as leading to a physically credible free energy distribution; however, the PMF was calculated using the Weighted Histogram Analysis Method (WHAM), implemented using the utility distributed by Grossfield.³²

Replica exchange runs were chunked into sections of 10 ps (10 000 steps of 1 fs each). Each chunk had 2000 cycles of exchange attempts (one per 5 MD steps), with an average acceptance rate $r \approx 0.3$. The first chunk for each run was discarded as equilibration, and the subsequent two were used as production runs. A Langevin thermostat with a temperature of 300 K and a coupling rate of 1 ps⁻¹ was applied to accelerate divergence between the 120 replicas in each set. Implicit solvent was applied using a generalized Born solvation model.³³ No cutoff was applied to the non-bonded interactions. When error bars were calculated with the assumption of a decorrelation time (per replica) equal to 100 fs or 20 cycles of replica exchange, these appeared as almost too small to be visible. Production runs for each equilibrium calculation comprised in total 120 × 20 ps, or 2.4 ns, of dynamics with HREMD-accelerated sampling.

III. RESULTS

A. Correction vs multiple forcefields

Stretching simulations were made against the OL15 forcefield, which is considered to be among the best classical DNA models at the time of writing.³⁴ However, it was found prudent to make a survey of multiple forcefields in order to confirm that the deficiency in treating stacking is general (Fig. 2). The bsc1 forcefield³⁵ was, therefore, compared against OL15 by post-processing snapshots taken from the KMMD simulations. The DESRES refinement of non-bonded interaction parameters for nucleobases³⁶ was also checked using the OL15 bonded and other non-specified parameters (“OL15DES”).

First, snapshots were ordered by the length of the average vector between chain-adjacent bases, measured at the N atoms where the chain joins the backbone [Fig. 2(a)]. In this case, the behavior around the minimum remained unchanged by adding KMMD, although the potential became apparently stiffer for large separation of the bases. The minimum appeared very sharp with ff99, whether corrected or not. This effect arises because all frames were generated using OL15+KMMD and are, therefore, close to the energy minima or valleys as defined by that PES; other forcefields may have similar features but are subject to slight geometry shifts, therefore giving apparently stiffer potentials. When lengths x were measured using the terminus atoms O5′–O3′ [Fig. 2(b)], an increase in the equilibrium length of the duplex was observed. This shift of ≈ 0.25 Å was independent of the base classical forcefield and was associated with a continued softening of the potential with respect to further extension. All three forcefields performed similarly for snapshots having small end-to-end length x and at large x . The bsc1 forcefield had the smallest over-estimate of pulling energy. The differences between the forcefields seem to be largely independent of the KMMD correction: all were altered by the KMMD by quite similar amounts.

B. Modulations of collective dynamics

The purpose of the KMMD correction to the forcefield is not primarily to moderate individual dihedral potentials but to treat collective interactions in molecules with multiple soft torsions. Figure 3 compares structural correlations in the DNA duplex under stretch with and without the ML correction. A basic probe of correlation is to collect the proportion of frames in which the χ angles adjacent on the chain were both in the same energy basin (of three available, as determined by the tetrahedral C1′ atom at which the base joins the sugar). As expected, the C3:C4 stacked pair of bases has an overall smaller chance for adjacent χ angles to be in the same minimum than does the G1:G2 stack due to the smaller size of the C base. For both stacks, there is a changeover in the effect of the KMMD with increasing extension: at short extension, KMMD destabilizes the base stack relative to standard MD with the OL15 forcefield, while at higher




FIG. 2. Postprocessing frames with different forcefields shows that the differences between them are largely independent of the larger errors fixed by using the KMMD Machine Learning correction (marked as “ML”). (a): Base–base distance (measured N9–N9 or N1–N1) shows moderate effects, with ML slightly increasing stiffness but only when far from the minimum. (b): Extension measured between the O5′ and O3′ termini is noticeably softer after correction and has a shifted minimum, regardless of the starting forcefield, in a seeming conflict with the tendency of bases to resist having large separation.

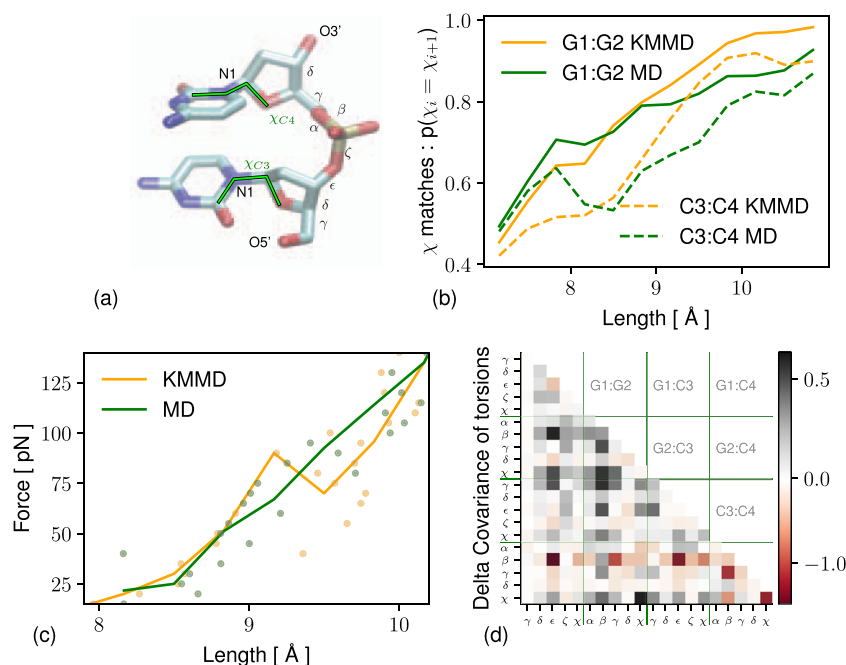


FIG. 3. (a) The CC chain of the GG-CC duplex, with backbone and χ torsions labeled. The four atoms defining each χ are shown with green highlighted lines. (b) Base–base interactions vs extension: proportion of χ angles on the same chain to be in the same basin. At short extension, KMMD gives less correlation between adjacent χ . At long extension, KMMD strengthens the correlation. (c) Force–extension. Points are independent constant-force simulations; lines are a moving average. (d): Change in covariance of dihedrals, KMMD-MD, over all extensions. Significant signals for χ and for torsions adjacent to the central phosphate.

extension, ordering of bases increases for all systems but especially so with the KMMD [Fig. 3(b)]. This signal is consistent with stabilization of edge–edge interactions by the KMMD when under imposed extension, as investigated in Figs. 5 and 6.

As simulations were run at constant force, it is proper to show the force–extension curve [Fig. 3(c)]. Although this observable is noisy and subject to finite size effects, it does support a reduction of the work to extend the DNA base step, reducing the discrepancy against the lower work to force-melt macroscopic DNA typically seen in experiment relative to simulation. Although the KMMD force–extension trace appears visually to be flatter than the uncorrected trace, an outlier point at 90 pN has the effect of giving almost equal values for the work to extend from 8 to 10 Å when a simple integration is performed using the trapezium rule (1.69 kcal mol^{−1} with KMMD vs 1.72 kcal mol^{−1} without). If the offending data point is removed, the corrected work becomes 1.56 kcal mol^{−1}, giving a difference that is still quite small in absolute terms but significant as a fraction of the total work to stretch the base-pair step. The experimental value given by Yakovchuk *et al.*¹³ for the base pair stacking energy of GG-CC at 300 K is salt-dependent, giving 1.5 ± 0.1 kcal mol^{−1} at a salt concentration of 45 mM, equivalent to the selected 2Na⁺ in 2446 water.

Figure 3(d) gives a sketch of the pairwise torsion angle correlations most affected by the KMMD correction. Covariance $C_{1,2}$ of a pair of dihedrals θ_1, θ_2 is defined here as $(\theta_1 - \bar{\theta}_1)(\theta_2 - \bar{\theta}_2)$ in units radians squared [the means $\bar{\theta}_i$ are found as $\arctan2(\cos \bar{\theta}_i, \sin \bar{\theta}_i)$]. The (signed) differences in the absolute values of the covariances $|C_{KMMD}| - |C_{MD}|$ are shown in Fig. 3(d) as a heat map. Overall, the KMMD correction promotes disorder (negative delta covariance, orange); however, covariance between some torsions is (more weakly) increased (positive delta, gray), in particular, involving χ and ζ angles. The β angle linking G bases becomes more covariant (gray),

while the β angle linking the C bases becomes much less covariant and more disordered.

C. Conformational analysis

The strength of the ML method is that second-order and higher correlations between torsion angles are treated; however, it is still instructive to project onto individual torsions in order to see which conformations of the duplex are stabilized/destabilized relative to classical forcefields. In order to compare the ML corrected forcefield with traditional nucleic acid forcefields, a series of replica exchange calculations were made (REMD) at zero imposed force and without hydrogen bond restraints, generating equilibrium potentials of mean force (PMFs). Three of the most important dihedral angles were examined: the ϵ and ζ joining two G bases and the χ angle of the first G. The qualitative picture was instructive: from the point of view of individual dihedral angles (as opposed to correlations between them), the KMMD correction is a step backward through time. Qualitatively, a one-dimensional view of the KMMD correction seems to show it as operating in the opposite direction to a recently proposed correction to the AMBER ff10 forcefield,³⁷ modifying the shapes of the free energy landscapes over the relevant dihedrals with the opposite sign relative to the original ff10^{38,39} (ff10 is a recent ancestor of the OL15²² starting point for the present study). The correction of Aytenfisu *et al.* did not alter the functional form of the forcefield but did re-fit with a much wider range of nucleic acid construct structures, with the novelty being that many structures drawn from crystallographic data were used, as well as the typical systematic rotations of dihedral angles in otherwise canonical structures. This use of crystallographic data selectively well-fits ordered structures. In the KMMD-corrected simulation, with training data drawn from highly disordered non-equilibrium structures, the less-dominant *syn* conformation for the χ angle was relatively




FIG. 4. Potentials of Mean Force (PMF) for key dihedrals, sampled without any stretching or WC bonding imposed.

stabilized, individually promoting disorder in the base conformation even though it is clear from Fig. 3 that collective order/disorder in base stacking without stretch is not much altered by KMMD. Training set bias is therefore less important when using KMMD because it is possible for a given dihedral to be both stabilized and destabilized at the same angle, depending on conformations of other dihedrals elsewhere in the system. The 1D PMFs for ϵ and ζ appear, in contrast to χ , to be order-stabilizing in the KMMD without stretch, and this is consistent with the variance-covariance plot of Fig. 3(d), showing small or negative changes in the individual variances of these angles under stretch with the correction applied. This is accentuated by a positive (gray) covariance between the adjacent ϵ and ζ angles: decreased individual variance coupled to increased covariance indicates an overall reduction of entropy associated with the backbone degrees of freedom without imposed stretch. The behavior of the χ – χ couplings vs stretch in Fig. 3(b) indicates, however, a slight increase in base-stack entropy at low or no imposed stretch, so we can state that the classical forcefields have at equilibrium a slight under-stiffening of the backbone and a slight over-emphasis on base–base interactions or on the stiffness of the χ angles, which also control base positioning (Fig. 4).

These data do not imply a simple revision of the existing backbone torsion potentials, and a re-fitting of individual dihedral angle potentials such that they could agree with the KMMD 1D-PMFs would land close to the AMBER ff10, which is, after more than a decade now, regarded as relatively obsolete even among traditional forcefields. What is shown, however, is that conflicting sets of training data (the dihedral scans that generated ff10 and the




FIG. 6. A conformation highly favored by the ML relative to the uncorrected OL15 forcefield. Hydrogen bonding is preserved, and bases on the same strand remain (sometimes/partly) in contact despite a large vertical stretch.

crystallographic structures that contributed to the Aytenfisu *et al.* correction) can be reconciled by considering a larger number of degrees of freedom together. Collective dihedral angle corrections can indirectly address errors arising from other terms of the forcefield: the example of the steric clash between the hydrogen attached to the Guanine C8 (or to any purine C8 Hydrogen) and the pentose O4' is investigated by permuting the angles labeled in Fig. 5(a). Although the Pauli exclusion that drives the sharp increase of energy as atoms approach overlap is not addressed directly by the KMMD correction, because molecular geometry is determined predominantly by dihedral angles, it is possible to indirectly address a clash that the standard classical treatment misses.

D. Conformations of greatest deviation

Analyzing the individual conformations with the greatest differences between KMMD and classical treatment interrogates the limitations of classical methods and helps to motivate the operation of the method. The top 100 structures with large favorable corrections following KMMD are dominated by conformations showing roughly 90° base–base interactions (the single most-favored conformation is shown in Fig. 6).

We refer to stretched duplex geometries with mostly preserved hydrogen bonding and partly preserved stacking (subject to tilt and slide) as “ τ -steps.”¹⁹ These have been put forward among various candidate geometries for stretched DNA since the earliest molecular

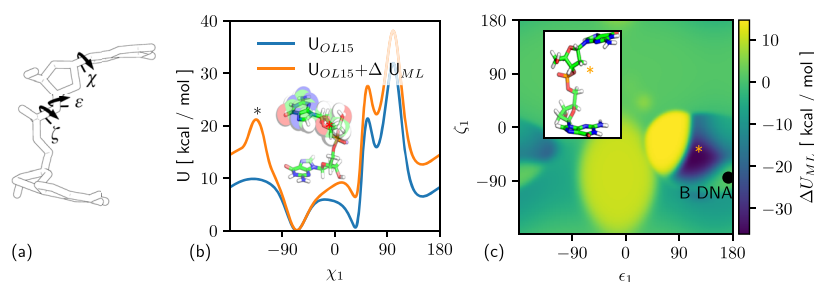


FIG. 5. 1D and 2D cuts holding other angles fixed. (a) Strand 1 (GG) of the dyad, showing angles permuted in the conformation examined. (b) A steric clash (red and white spheres, inset foreground) at $\chi = -140^\circ$ is penalized more harshly by the ML method. (c) A highly extended conformation (*) and inset is strongly favored by the ML.

mechanics studies⁴⁰ and linked specifically to stretched GG-CC pairs in recent classical MD work.⁴¹ Our primary structural observation into DNA under tension, therefore, is that the posited τ -step conformation is confirmed and that this is much more favorable than found when modeled classically. The strength of this interaction geometry is not inconsistent with the chemical understanding of 90° interactions between aromatic molecules with hydrogen bonding groups at the edges, in general, discussed as “remarkably strong,”⁴² although the canonical T-shape aromatic geometry is typically presented as edge–face rather than edge–edge. The lack of directionality imposed by a point-like treatment of classical atoms leaves little hope of correctly treating the highly directional combination of hydrogen bonding and aromatic stacking in the framework of current mainstream forcefields.

E. Outlook

Distillation of *ab initio* quantum mechanical calculations on fragments into a smooth potential energy surface was successful in this initial example, demonstrating integration of machine-learned physics for short-range interactions with classical modeling of longer-range Coulomb and pairwise dispersion physics, as well as classical treatment of solvent and of stiffer short-range degrees of freedom (bond lengths and angles) not strongly affecting the geometry. Generalization to long chains as a sum of overlapping fragments does not seem, in principle, to be a major challenge, nor does generalization to a wider array of treated repeating units, such as amino acids or RNA. Classical MD simulations corrected with KMMD were able to provide a qualitatively correct alteration to the force-extension behavior of a small DNA construct, bringing results more closely in line with experiment. KMMD does not, as applied here, treat chemical effects, collective electron dynamics in the solvent, or convergence errors due to rare fluctuations; however, by application of Occam’s razor following the results presented here, we may reduce our concern in relation to these potential sources of error.

The KMMD concept as described here does ultimately have well-defined limitations, arising, in particular, from the fact that non-bonded interactions are treated indirectly only. It is entirely feasible to construct an augmented feature space containing two-body or N-body distance information as well as the angular information treated in the present implementation. Indeed, this has been done for kernel methods targeting smaller molecules;³⁰ however, there are always costs for adding complexity. Development of accuracy and generality for kernel-modified MD is likely to continue in the near future by considered design of the feature space in which the kernel operates and the thoughtful integration of physics, which can be calculated cheaply and well from first principles with physics, which is best treated through machine learning.

KMMD was shown to modify DNA dynamics with a qualitatively correct reduction of the force to extend a small two-base pair construct in water, of the correct sign, such that it can now be hoped that simulations of larger systems will manifest thermodynamics (and perhaps structure) in line with experiment. The putative biologically relevant conformations for many base pair DNA, which have been suggested from experimental data in relation to stretching^{19,43} and also to hydrophobic denaturants,⁴⁴ were not susceptible to be confirmed or denied by this study.

In particular, the hypothetical Σ phase of DNA in which the symmetry is broken with a three-bp periodicity is obviously not susceptible to direct study by simulations of 2bp only.

ACKNOWLEDGMENTS

The experiments presented in this paper were carried out using the HPC facilities of the University of Luxembourg,⁴⁵ supported by the Luxembourg National Research Fund (FNR) (Grant No. C20/MS/14769845/BroadApp).

AUTHOR DECLARATIONS

Conflict of Interest

The authors have no conflicts to disclose.

Author Contributions

Joshua T. Berryman: Conceptualization (lead); Data curation (lead); Funding acquisition (equal); Investigation (lead); Methodology (lead); Project administration (equal); Resources (equal); Software (lead); Supervision (lead); Validation (equal); Visualization (lead); Writing – original draft (lead); Writing – review & editing (lead). **Amirhossein Taghavi:** Conceptualization (supporting); Investigation (supporting); Writing – original draft (supporting). **Florian Mazur:** Conceptualization (equal); Investigation (equal); Methodology (equal); Visualization (equal); Writing – original draft (supporting). **Alexandre Tkatchenko:** Conceptualization (supporting); Funding acquisition (equal); Project administration (equal); Resources (equal); Supervision (supporting); Writing – review & editing (supporting).

DATA AVAILABILITY

Software to carry out the KMMD and other calculations is available from ambermd.org as part of the AmberTools 22 package.²¹ The dataset of quantum energy evaluations generated for this work is available from the NOMAD database under a Document Object Identifier: <https://dx.doi.org/10.17172/NOMAD/2022.04.25-110.17172/NOMAD/2022.04.25-1>.

REFERENCES

- ¹F.-Y. Dupradeau, A. Pigache, T. Zaffran, C. Savineau, R. Lelong, N. Grivel, D. Lelong, W. Rosanski, and P. Cieplak, *Phys. Chem. Chem. Phys.* **12**, 7821 (2010).
- ²P. Bleiziffer, K. Schaller, and S. Riniker, *J. Chem. Inf. Model.* **58**, 579 (2018).
- ³A. Tkatchenko and M. Scheffler, *Phys. Rev. Lett.* **102**, 073005 (2009).
- ⁴J. Šponer, J. Leszczyński, and P. Hobza, *J. Phys. Chem.* **100**, 5590 (1996).
- ⁵J. Šponer, K. E. Riley, and P. Hobza, *Phys. Chem. Chem. Phys.* **10**, 2595 (2008).
- ⁶C. Nganou, S. D. Kennedy, and D. W. McCamant, *J. Phys. Chem. B* **120**, 1250 (2016).
- ⁷A. Tkatchenko, R. A. DiStasio, Jr., R. Car, and M. Scheffler, *Phys. Rev. Lett.* **108**, 236402 (2012).
- ⁸R. A. DiStasio, Jr., O. A. von Lilienfeld, and A. Tkatchenko, *Proc. Natl. Acad. Sci. U. S. A.* **109**, 14791 (2012).
- ⁹Y. S. Al-Hamdani and A. Tkatchenko, *J. Chem. Phys.* **150**, 010901 (2019).
- ¹⁰S. A. Harris, Z. A. Sands, and C. A. Laughton, *Biophys. J.* **88**, 1684 (2005).
- ¹¹D. R. Roe and A. M. Chaka, *J. Phys. Chem. B* **113**, 15364 (2009).

- ¹²C. Bustamante, S. B. Smith, J. Liphardt, and D. Smith, *Curr. Opin. Struct. Biol.* **10**, 279 (2000).
- ¹³P. Yakovchuk, E. Protozanova, and M. D. Frank-Kamenetskii, *Nucleic Acids Res.* **34**, 564 (2006).
- ¹⁴J. Liphardt, S. Dumont, S. B. Smith, I. Tinoco, and C. Bustamante, *Science* **296**, 1832 (2002).
- ¹⁵S. Park, F. Khalili-Araghi, E. Tajkhorshid, and K. Schulten, *J. Chem. Phys.* **119**, 3559 (2003).
- ¹⁶R. Galindo-Murillo, D. R. Roe, and T. E. Cheatham, *Biochim. Biophys. Acta, Gen. Subj.* **1850**, 1041 (2015), recent developments of molecular dynamics.
- ¹⁷M. Manghi, J. Palmeri, and N. Destainville, *J. Phys.: Condens. Matter* **21**, 034104 (2008).
- ¹⁸M. Stöhr and A. Tkatchenko, *Sci. Adv.* **5**, eaax0024 (2019).
- ¹⁹A. Taghavi, P. van der Schoot, and J. T. Berryman, *Q. Rev. Biophys.* **50**, e15 (2017).
- ²⁰W. L. Jorgensen, J. Chandrasekhar, J. D. Madura, R. W. Impey, and M. L. Klein, *J. Chem. Phys.* **79**, 926 (1983).
- ²¹D. A. Case, H. M. Aktulga, K. Belfon, I. Y. Ben-Shalom, J. T. Berryman, S. R. Brozell, D. S. Cerutti, T. E. Cheatham III, G. A. Cisneros, V. W. D. Cruzeiro, T. A. Darden, R. E. Duke, G. Giambasu, M. K. Gilson, H. Gohlke, A. W. Goetz, R. Harris, S. Izadi, S. A. Izmailov, K. Kasavajhala, M. C. Kaymak, E. King, A. Kovalenko, T. Kurtzman, T. S. Lee, S. LeGrand, P. Li, C. Lin, J. Liu, T. Luchko, R. Luo, M. Machado, V. Man, M. Manathunga, K. M. Merz, Y. Miao, O. Mikhailovskii, G. Monard, H. Nguyen, K. A. O'Hearn, A. Onufriev, F. Pan, S. Pantano, R. Qi, A. Rahnamoun, D. R. Roe, A. Roitberg, C. Sagui, S. Schott-Verdugo, A. Shajan, J. Shen, C. L. Simmerling, N. R. Skrynnikov, J. Smith, J. Swails, R. C. Walker, J. Wang, J. Wang, H. Wei, R. M. Wolf, X. Wu, Y. Xiong, Y. Xue, D. M. York, S. Zhao, and P. A. Kollman, Amber 2022, University of California, San Francisco, 2022; available at ambermd.org.
- ²²M. Zgarbová, A. M. Rosnik, F. J. Luque, C. Curutchet, and P. Jurečka, *J. Comput. Chem.* **36**, 1874 (2015).
- ²³J. P. Perdew, M. Ernzerhof, and K. Burke, *J. Chem. Phys.* **105**, 9982 (1996).
- ²⁴C. Adamo and V. Barone, *J. Chem. Phys.* **110**, 6158 (1999).
- ²⁵X. Ren, P. Rinke, V. Blum, J. Wieferink, A. Tkatchenko, A. Sanfilippo, K. Reuter, and M. Scheffler, *New J. Phys.* **14**, 053020 (2012).
- ²⁶A. Ambrosetti, A. M. Reilly, R. A. DiStasio, Jr., and A. Tkatchenko, *J. Chem. Phys.* **140**, 18A508 (2014).
- ²⁷V. Blum, R. Gehrke, F. Hanke, P. Havu, V. Havu, X. Ren, K. Reuter, and M. Scheffler, *Comput. Phys. Commun.* **180**, 2175 (2009).
- ²⁸J. Moody and C. J. Darken, *Neural Comput.* **1**, 281 (1989).
- ²⁹G. Bugmann, *Neurocomputing* **20**, 97 (1998).
- ³⁰K. Hansen, G. Montavon, F. Biegler, S. Fazli, M. Rupp, M. Scheffler, O. A. von Lilienfeld, A. Tkatchenko, and K.-R. Müller, *J. Chem. Theory Comput.* **9**, 3404 (2013).
- ³¹W. Saenger, *Principles of Nucleic Acid Structure* (Springer, 1984).
- ³²A. Grossfield, "WHAM: the weighted histogram analysis method," Version XXXX, http://membrane.urmc.rochester.edu/wordpress/?page_id=126.
- ³³H. Nguyen, D. R. Roe, and C. Simmerling, *J. Chem. Theory Comput.* **9**, 2020 (2013).
- ³⁴P. D. Dans, I. Ivani, A. Hospital, G. Portella, C. González, and M. Orozco, *Nucleic Acids Res.* **45**, 4217 (2017).
- ³⁵I. Ivani, P. D. Dans, A. Noy, A. Pérez, I. Faustino, A. Hospital, J. Walther, P. Andrio, R. Goñi, A. Balaceanu *et al.*, *Nat. Methods* **13**, 55 (2016).
- ³⁶D. Tan, S. Piana, R. M. Dirks, and D. E. Shaw, *Proc. Natl. Acad. Sci. U. S. A.* **115**, E1346 (2018).
- ³⁷A. H. Aytenfisu, A. Spasic, A. Grossfield, H. A. Stern, and D. H. Mathews, *J. Chem. Theory Comput.* **13**, 900 (2017).
- ³⁸A. Pérez, F. J. Luque, and M. Orozco, *J. Am. Chem. Soc.* **129**, 14739 (2007).
- ³⁹M. Zgarbová, M. Otyepka, J. Šponer, A. Mládek, P. Banáš, T. E. Cheatham, and P. Jurečka, *J. Chem. Theory Comput.* **7**, 2886 (2011).
- ⁴⁰A. Lebrun and R. Lavery, *Nucleic Acids Res.* **24**, 2260 (1996).
- ⁴¹J. W. Shepherd, R. J. Greenall, M. I. J. Probert, A. Noy, and M. C. Leake, *Nucleic Acids Res.* **48**, 1748 (2020).
- ⁴²L. R. Rutledge and S. D. Wetmore, *J. Chem. Theory Comput.* **4**, 1768 (2008).
- ⁴³N. Bosaeus, A. Reymer, T. Beke-Somfai, T. Brown, M. Takahashi, P. Wittung-Stafshede, S. Rocha, and B. Nordén, *Q. Rev. Biophys.* **50**, E11 (2017).
- ⁴⁴B. Feng, R. P. Sosa, A. K. F. Mårtensson, K. Jiang, A. Tong, K. D. Dorfman, M. Takahashi, P. Lincoln, C. J. Bustamante, F. Westerlund, and B. Nordén, *Proc. Natl. Acad. Sci. U. S. A.* **116**, 17169 (2019).
- ⁴⁵S. Varrette, P. Bouvry, H. Cartiaux, and F. Georgatos, in *Proceedings of the 2014 International Conference on High Performance Computing and Simulation (HPCS 2014)* (IEEE, Bologna, Italy, 2014), pp. 959–967.