JCIM JOURNAL OF CHEMICAL INFORMATION AND MODELING

pubs.acs.org/jcim Article

Automated Force Field Developer and Optimizer Platform: Torsion Reparameterization

Alejandro Blanco-Gonzalez, William Betancourt, Ryan Michael Snyder, Shi Zhang, Timothy J. Giese, Zeke A. Piskulich, Andreas W. Götz,* Kenneth M. Merz, Jr.,* Darrin M. York,* Hasan Metin Aktulga,* and Madushanka Manathunga*

Cite This: J. Chem. Inf. Model. 2026, 66, 3206−3219

ABSTRACT: General force fields such as General Amber Force Field (GAFF) have been designed for broad applicability and are widely used in protein−ligand binding simulations in structure-based drug discovery. However, the force field parameters are not always transferable across ligand molecules, and custom reparameterization is sometimes necessary for accurate binding free energy simulations. This is especially true for torsion parameters, which are highly dependent on stereoelectronic and steric effects. Here, we report a novel, flexible, and user-friendly computational tool called the Automated Force Field Developer and Optimizer (AFFDO) platform that allows generating accurate, tailored GAFF2 torsion parameters for drug-like molecules. For a given ligand, AFFDO selects the most important torsions, carries out GPU-accelerated density functional theory calculations to collect reference data and fits torsion terms using a fast gradient-based optimizer that leverages automated differentiation. We benchmark AFFDO by parametrizing a series of drug-like molecules and carrying out protein−ligand relative binding free energy (RBFE) simulations. The results show that AFFDO can significantly improve GAFF2 torsion parameters against QM reference data, which in some cases translates into better agreement with experimental RBFE values within a reasonable computational time.

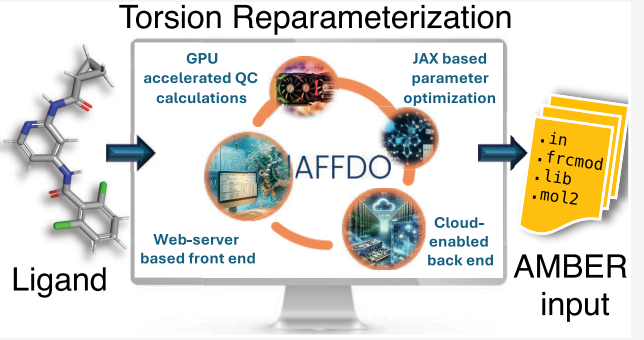

1. INTRODUCTION

Predicting binding free energy between a ligand and a target protein is an important task in structure-based drug discovery.[1,2] This is usually done using computational methods such as free energy perturbation and thermodynamic integration, where molecular mechanics force fields serve as the main workhorse.[3] The accuracy of the force field used is crucial for reliable predictions of free energy, thus for identifying promising drug candidates. Over the past few decades, several highly successful force fields such as AMBER,[4] CHARMM,[5] OPLS,[6] and OpenFF[7,8] have been developed for biochemical simulations. These force fields contain parameter sets for certain subsets of the chemical space. For example, in the widely used AMBER force field, ff19SB[9] and OL21[10] provide parameters for protein and DNA simulations, respectively. When a molecule falls outside of these subsets (e.g., an organic ligand), obtaining the necessary parameters can involve a complex, lengthy, and computationally demanding parametrization process.

Existing force fields provide a general parameter set (e.g., GAFF[11,12] parameter set in case of AMBER) that covers a broad range of molecules by defining new atom types, adding more chemical groups and utilizing a heuristic pattern matching algorithm. However, a generic parameter set may fail to deliver accurate parameters, specifically for torsion angles and nonbonded interactions. Among bonded interaction terms, torsion parameters are particularly sensitive to the local environment of the target molecule, as various stereoelectronic and steric effects must also be considered in their optimization. Additionally, interactions such as resonance effects among aromatic rings can influence torsional profiles, even when modifications occur in nonlocal areas that may not be detected through chemical perception.[13] These factors make torsion parameters less transferable compared to other valence parameters.

For reasons outlined above, torsion parameters are often subjected to reparameterization. A common approach to reparameterization is to start with a general force field and refine parameters for individual molecules, mainly by fitting against quantum chemical (QC) data. Multiple tools centered around this approach have been reported to date.[14−22] Although

Received: February 18, 2026
Revised: February 26, 2026
Accepted: February 27, 2026
Published: March 9, 2026

https://doi.org/10.1021/acs.jcim.6c00528
J. Chem. Inf. Model. 2026, 66, 3206−3219

torsion reparameterization based on QC data fitting can improve the accuracy of the free energy simulations, the process of fitting the parameters to numerous torsions is labor-intensive and may incur significant computational costs. One way to reduce computational costs is the use of machine learning methods such as the Accurate NeurAl networK engINe for Molecular Energies (ANI)[23,24] and modern semiempirical methods like xtb,[25] which is a compromise between the accuracy and the computational cost to compute the reference data. Parameterize,[26] a tool that allows the reparameterization of GAFF2/OpenFF torsion parameters, is a good example in this context.

In this paper, we report Automated Force Field Developer and Optimizer (AFFDO), a new GPU-accelerated software platform for force field parametrization. The current version of AFFDO is available at https://www.attmosdiscovery.com/tools as a web service free of charge. This free version allows reparameterization of GAFF torsion terms for molecules containing up to 90 atoms.

AFFDO optimizes the parameters in a force field by first generating high-quality QC data and then fitting against these data through a highly optimized end-to-end workflow; the implementation reported in this manuscript is specific to AMBER and targets torsion parameters only. However, extending AFFDO to generate parameter files compatible with other MD simulation packages such as GROMACS is planned for future development. Computational optimizations in AFFDO span multiple steps of the workflow. First, for a given ligand, significant torsions are identified by sampling the QC potential energy surface (PES); torsional scans are performed only for these significant torsions, eliminating unnecessary costs. Second, torsional scans are carried out using a low-cost semiempirical (xtb) method and are later refined using massively parallel, GPU-accelerated QC calculations. Finally, optimizations of force field parameters against QC scans are performed using a fast gradient-based optimizer that leverages automated differentiation for rapid convergence.

Using the AFFDO platform, custom torsion parameter optimizations can be completed for target molecules in a relatively short time (compared to the time it takes to carry out the corresponding free energy computations) and without user intervention. AFFDO allows step-by-step monitoring of the process and modifications by the user can be applied to any step of the workflow. We demonstrate the utility, accuracy, and performance of AFFDO by generating torsion parameters for a series of drug molecules. We then show how the reparameterized force field performs in predicting the relative binding free energy (RBFE) values of several protein−ligand systems reported in a widely used public data set.[27] The next sections of this paper are organized as follows. In Section 2, we provide methodological details and describe the implementation of the AFFDO platform. Section 3 is devoted to results and discussion where we present the results of benchmark simulations, including timing data. Finally, we conclude this study and provide future directions in Section 4.

2. AFFDO WORKFLOW

For a set of N atoms described by their coordinates in 3D Cartesian space $\mathbb{R}^{3\times N}$, the functional form for the AMBER force field, which is largely shared by CHARMM and OPLS force fields as well, is given by eq 1.

$$V(\mathbb{R}^{3\times N}) = \sum_{i\in\text{bonds}} k_{b_i}(l_i - l_i^0)^2 + \sum_{i\in\text{angles}} k_{a_i}(\theta_i - \theta_i^0)^2 + \sum_{i\in\text{torsions}} \sum_n \frac{1}{2} V_i^n[1 + \cos(n\omega_i - \gamma_i^n)] + \sum_{i<j} \left[s_{ij}^{\text{LJ}} \varepsilon_{ij} \left(\left(\frac{r_{\text{min},ij}}{r_{ij}} \right)^{12} - 2 \left(\frac{r_{\text{min},ij}}{r_{ij}} \right)^{6} \right) + s_{ij}^{\text{C}} \frac{q_i q_j}{4\pi\varepsilon_0 r_{ij}} \right] \quad (1)$$

Here, the bond (first term) and angle (second term) energies are approximated as harmonic oscillators, where l_i^0 and θ_i^0 denote the equilibrium bond length and bond angle, respectively, and k_{b_i} and k_{a_i} are the corresponding force constants. The torsional contribution (third term) models the energy associated with rotation about a bond and is expressed as a Fourier series to account for multiple periodic contributions arising from chemical effects such as bond order, conjugation, and neighboring substituents. V_i^n is the barrier height of the n-th Fourier component, ω_i is the dihedral angle, n is the periodicity, and γ_i^n is the associated phase offset.

Nonbonded interactions (final term) describe short-range Pauli repulsion and van der Waals dispersion via a Lennard−Jones potential, together with electrostatic interactions between partial charges. The nonbonded sum runs over atom pairs $i < j$ that are not excluded by molecular topology, following standard AMBER conventions (e.g., exclusion of 1−2 and 1−3 interactions). Interactions between atoms separated by three bonds (1−4 pairs) are included with topology-dependent scaling, using distinct scaling factors for the Lennard−Jones and electrostatic contributions (SCNB and SCEE, respectively), encoded in the pair-dependent scaling coefficients s_{ij}^{LJ} and s_{ij}^{C} appearing in eq 1, while all more distant pairs (1−5 and beyond) interact without additional scaling. Long-range electrostatic interactions are evaluated using particle mesh Ewald−type methods.[28]

Parameterization of a new molecule involves determining the optimal values for terms in eq 1. While generic force fields such as GAFF can serve as a good model for organic molecules and pharmaceutical compounds in general, it is unreasonable to expect high fidelity from a generic force field consistently, given the practically infinite size of the set of organic molecules and pharmaceutical compounds. Creating force field descriptions customized to a target molecule has been the main theme behind efforts such as OpenFF BespokeFit[14] and OPLS3e,[29] too. Much like OpenFF BespokeFit and OPLS3e, the current customization targets in AFFDO are the torsion parameters because they are crucial for accurate structural representation of molecules. Horton et al.[14] describe the general methods and software tools used for customizing OpenFF's torsion parameters; Roos et al.[29] describe a significant expansion of the OPLS3e torsion angle database to improve the fidelity of OPLS along with a tool that can help users create their own custom OPLS torsion parameters.

As discussed above, QM-based torsion fitting is an established approach; accordingly, AFFDO's contribution is not the concept of torsion fitting itself, but the specific integration and automation choices described below. AFFDO's end-to-end workflow design emphasizes throughput and fully automatic

https://doi.org/10.1021/acs.jcim.6c00528
J. Chem. Inf. Model. 2026, 66, 3206−3219

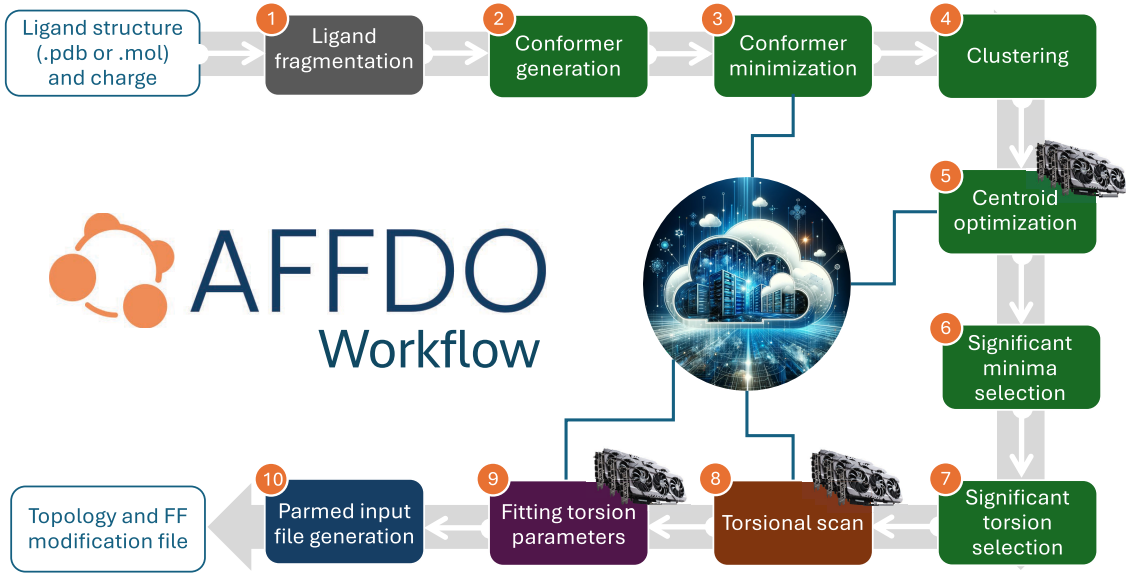

Figure 1. Parameterization procedure of AFFDO workflow. Steps 3, 5, 8, and 9 involve CPU or GPU parallelization and are carried out on a user-specified backend. Different tile colors gray, green, brown, violet, navy blue represent the five main tasks of the workflow: ligand fragmentation (step 1), identification of significant torsions (step 2−7), generation of reference data (step 8), torsion fitting (step 9), and the generation of final output files (step 10).

operation without user involvement. In addition to this general motivation, AFFDO introduces several workflow-level developments that go beyond torsion fitting. First, AFFDO includes a torsion-focused fragmentation module (FragMentor) that operates as an initial, fully integrated step of the pipeline: it evaluates whether the molecule contains potentially reparameterizable torsions and uses these targets to decide whether QM scans should be performed on the full molecule or on automatically constructed fragments, only when this reduces QM cost without compromising the relevant local torsional environment. Second, AFFDO implements an automated procedure to identify and prioritize chemically significant torsions from conformer ensembles, avoiding unnecessary fitting. Third, torsion fitting is performed using a gradient-based optimizer that leverages gradient information (L-BFGS-B; SciPy/JAX implementation) to efficiently refine torsion parameters. In addition, several implementation-level advances were required to support high-throughput execution (with GPU acceleration where applicable) in practice, including automated torsion scans and constrained optimizations in the QM workflow (DL-Find + QUICK integration) and RESP-charge computation within QUICK. These implementation details are described in later sections.

For optimizing the torsion parameters for a molecule, it is necessary to determine the optimal barrier height, phase, and periodicity terms for each torsion in the molecule. The underlying torsion parametrization procedure of AFFDO is summarized in Figure 1. This multistep procedure achieves five main tasks: torsion-focused fragmentation (step 1), identification of significant torsions (step 2−7), generation of reference data (step 8), torsion fitting (step 9) and generation of the final output files (step 10).

2.1. Ligand Fragmentation

The parametrization of large ligands can be challenging due to computational cost and potential steric clashes during torsional scans. This can be avoided by fragmenting the ligand, parametrizing the fragments separately, and transferring torsion parameters from fragments to the original ligand at the end of the procedure. AFFDO is equipped with a dedicated fragmentation tool, FragMentor, specifically designed for torsion parametrization. Its primary novelty lies in its fully automatic decision-making, which autonomously determines whether ligand fragmentation is necessary based on predefined criteria. This ensures that fragmentation is only applied when beneficial, aligning with AFFDO's streamlined and automated workflow.

FragMentor strategically selects cleavage sites to target torsions suitable for parametrization while preserving their original chemical environments. The fragmentation decision and fragment sizes are guided by the number of *backbone-connected* torsions and by ensuring that torsions arising from distinct molecular regions are not forced into a single fragment, preventing excessive fragmentation while ensuring that chemically meaningful torsional regions are properly parametrized. Here, *backbone-connected* torsions refer to the subset of significant torsions whose central bonds lie on the longest non-hydrogen chain of the ligand (identified via a graph-theoretic longest-path filter before fragmentation). FragMentor operationally preserves the local chemical environment by partitioning significant torsions into connected *torsion regions* (consecutive torsions on the molecular graph) and generating a separate fragment for each region; fragments are retained only if they contain the atoms defining the region. Torsions within rings, end-of-chain torsions, and functional groups are explicitly excluded from cleavage, preserving key molecular substructures. The resulting fragments contain unique torsions, ensuring that the derived torsion profiles remain representative of the parent molecule. FragMentor also employs RDKit-based algorithms[30] to break bonds and cap fragment ends with methyl groups. Additionally, detailed atom-index mappings between fragments and the original ligand (ligand-fragment atom maps) are systematically stored in JSON format, enabling accurate reconstruction and parameter transfer after optimization. Overall, FragMentor, integrated within AFFDO's fully auto-

https://doi.org/10.1021/acs.jcim.6c00528
J. Chem. Inf. Model. 2026, 66, 3206−3219

mated pipeline, reduces manual intervention and improves computational efficiency.

We note that related fragmentation tooling has been developed to reduce the cost of QM torsion drives, including OpenFF-Fragmenter.[13] In its recommended WBO-sensitive mode, OpenFF-Fragmenter follows a bond-centric strategy in which fragments are grown for a *driven* rotatable bond until the Wiberg bond order (WBO) of that bond matches the parent molecule within a user-defined tolerance (i.e., one fragment is constructed per targeted torsion-drive bond). In contrast, FragMentor is implemented as the entry-point of AFFDO and is optimized for end-to-end GAFF torsion reparameterization: it operates on the *set of torsions selected by AFFDO* and targets workflow-level efficiency by minimizing the number of fragments/AFFDO jobs, while preserving the local chemical context required for fitting (as described above). We emphasize that both tools share the goal of reducing torsion-drive cost, but they operate at different granularity (bond-by-bond WBO-driven growth versus workflow-integrated torsion-region coverage) and are optimized for different refitting workflows.

2.2. Identification of Significant Torsions

An important step of the torsion parametrization is the identification of chemically important torsions of a given molecule. The goal here is to eliminate the parametrization of unnecessary torsions and reduce the computational cost. To achieve this task we developed the following novel methodology. For a given molecular structure, AFFDO generates a large number of conformers using the ETKDG method[31] implemented in RDKit.[30] Generated conformer structures are minimized using the GFN2-xTB method[32] implemented in xTB.[25] To filter redundant conformations and improve efficiency, minimized geometries are clustered using SciPy's hierarchical clustering approach.[33,34] The algorithm groups conformers based on geometric and energetic similarity, iteratively merging those with minimal Euclidean distances (RMSD and energy) to ensure a structurally diverse yet nonredundant selection. In our implementation, each conformer is represented by a two-feature vector consisting of its heavy-atom RMSD to the lowest-energy conformer and its relative energy (ΔE) from the conformer-minimization step. To combine these features in a single Euclidean distance without an ad hoc unit conversion, we L2-normalize each feature across the conformer set and apply Ward hierarchical clustering to the resulting normalized feature vectors (SciPy[33,34] linkage, method = "ward").

AFFDO then identifies the lowest-energy centroid and eliminates other centroids above a configurable energy window intended as a pragmatic screening criterion to retain low-energy conformers while discarding high-energy outliers (default ΔE_{keep} = 3.5 kcal/mol). If a higher-level reference is requested (e.g., DFT in this work), the retained centroid conformers are reoptimized at that level prior to torsion identification; otherwise, the xTB-minimized centroids from the prescreening stage are used directly for the torsion-dispersion analysis. Specifically, for each candidate torsion, AFFDO computes the wrapped dihedral angles across the retained centroid set and retains a torsion if the maximum centroid-to-centroid dihedral difference exceeds a configurable cutoff (default 30°). The resulting torsion list can contain rotations about methyl/terminal groups and multiple torsion definitions arising from the same central bond. Such torsions are pruned using chemically motivated heuristics. In particular, torsions involving terminal atoms (atoms with only one heavy-atom neighbor) are excluded by default because their large centroid-to-centroid angle differences often arise from quasi-free rotations that are not representative of the conformational changes targeted for reparameterization; this filter is configurable. Torsions entirely contained within rings are also excluded by default (independent of ring size) because many ring torsions are conformationally constrained and typically exhibit limited angle variation across the centroid ensemble; this choice may be revisited in future versions to better accommodate more flexible cyclic systems such as medium-sized rings (8−11-membered) and macrocycles (commonly ≥ 12-membered), where torsional reparameterization can be more impactful. When multiple torsion definitions arise from the same bond, AFFDO identifies the longest non-hydrogen chain of the molecule using a graph algorithm and prioritizes the most "backbone-like" torsion based on how deeply the terminal atoms lie within the heavy-atom framework.[26] The pruned torsion list from this procedure results in the ultimate torsions (hereafter, *significant torsions*) which will be used in later stages.

2.3. Generating Reference Data

Starting from each centroid structure identified in the previous step, 360° torsional scans are carried out for each significant torsion using the GFN2-xTB method on a uniform dihedral grid with a default spacing of 20° between consecutive scan points (18 grid points over 360°). The grid spacing is a configurable AFFDO setting (default 20° in this work). This is achieved by using the CPU-parallel `xtb` program.[25] In principle, one could use a single centroid structure (e.g., the structure corresponding to the most stable minimum) for all torsional scans. However, due to steric clashes, torsional scans initiated from certain structures can result in abrupt spikes or shifts in the energy profiles. Given the low computational cost of GFN2-xTB, we mitigate this issue by starting scans from all centroids and selecting the smoothest energy profile for each torsion. In AFFDO, this selection happens automatically by ranking the candidate scans using a composite score that (i) favors profiles that are well-described by a low-order Fourier representation and (ii) penalizes unphysical roughness in the energy curve, particularly discontinuous jumps between consecutive scan points and excessive high-frequency jaggedness/curvature. The highest-ranked profile is retained for refinement and subsequent fitting. Next, selected energy profiles are refined using density functional theory (DFT) employing the PBE0-D3BJ functional[35] with the 6−31G* basis set for neutral and cationic molecules, and the 6−31+G* basis set for anionic molecules.[36,37] The GPU-accelerated QUICK QC program[38−43] is employed for these constrained optimization calculations. Note that the constraint is applied as an explicit *dihedral-angle constraint* on the scanned torsion (rather than Cartesian positional constraints on the four atoms) during the constrained optimization. To achieve this task, we enabled constrained optimization capability in QUICK during this work.

We then recompute each profile using GAFF2 as a preliminary step to refine force field parameters through comparison between the MM energy profiles and higher-level reference curves. This process is facilitated by a new implementation added to AmberTools during this work. Specifically, we interfaced the sander program with the DL-FIND software[44] to allow for constrained geometry optimizations using several algorithms and coordinate systems. The interface is made by directly linking the DL-FIND library with

https://doi.org/10.1021/acs.jcim.6c00528
J. Chem. Inf. Model. 2026, 66, 3206−3219

the sander executable. This greatly improves the computational performance because it avoids the use of system calls to launch and reinitialize the sander program at each optimization step to obtain the energy and forces. The default setting for geometry optimization is the limited-memory Broyden-Fletcher-Goldfarb-Shanno (L-BFGS) algorithm[45] and a hybrid delocalized coordinate (HDLC) system.[46] The coordinate system is constructed using the residue definitions in the AMBER topology, which we use to partition the system into fragments for the HDLC representation. The interface, included in the release of AmberTools25, allows select bond length, bond angle, and torsion angle values to be constrained during optimization.

2.4. Torsion Fitting

As our current implementation is specific to AMBER, AFFDO constructs the initial force field for a given ligand using GAFF/GAFF2 parameters[11,12] and AM1-BCC,[47,48] ABCG2,[49] and RESP[38,50,51] atomic charges. Specifically, the Antechamber program[52] in AmberTools[53] is used to assign atom types and parameters for the target ligand. The charges are computed for all centroid structures, according to the model charge selection, and final atomic charges are obtained by a Boltzmann averaging procedure. Specifically, AFFDO computes Boltzmann weights from the relative centroid energies $\Delta E_k = E_k - E_{\min}$ at $T = 298$ K, $w_k \propto \exp(-\Delta E_k/(k_B T))$ (normalized over the retained centroids), and then forms the per-atom averaged charges as $\overline{q}_i = \sum_k w_k\, q_i^{(k)}$. This is because most of these charge models are highly dependent on the target molecular conformations.

To fit the torsion parameters, we developed a novel methodology that uses a quasi-Newton optimization algorithm that only requires gradient information, specifically the L-BFGS-B method implemented in the SciPy package[33,34] coupled with Google's JAX library.[54] At a high level, BFGS and similar Quasi-Newton methods minimize an objective function using an iterative approximation of the function's Hessian, B. For a given iteration k, search direction p_k, step size generated by line search α_k, and corresponding parameter update s_k, we have the following eq 2

$$
\begin{aligned}
p_k &= -B_k \Delta f(x_k) \\
x_{k+1} &= x_k + \alpha_k p_k \text{ and } s_k = x_{k+1} - x_k \\
y_k &= \Delta f(x_{k+1}) - \Delta f(x_k) \\
B_{k+1} &= B_k + \frac{y_k y_k^T}{y_k^T s_k} - \frac{B_k s_k s_k^T B_k^T}{s_k^T B_k s_k}
\end{aligned}
\tag{2}
$$

These methods are more robust than simple gradient descent but notably still require the gradient of the objective function $\Delta f(x_k)$. In this case, we define our objective function as the sum of squared error between QC and GAFF torsional energy profiles. In the present implementation, the objective targets agreement of torsional energy profiles; adding geometric terms is a natural extension for future development. Both profiles are generated by performing constrained minimizations on a uniform 20° dihedral grid, where the *scanned torsion itself* is enforced via an AMBER dihedral-angle restraint at each target angle (force constant 500 kcal/mol/rad^2 with a ± 0.5° tolerance band around the target). We find that existing GAFF2 parameters generally serve as good initial guesses and ensure rapid convergence. The main challenge then remains rapid evaluation of the gradients of the loss function. We perform the torsional scans and loss function evaluations using an implementation of the AMBER force field written with Google's JAX library, which is a set of function transformations for Python that are an extension of NumPy. Through its support for automatic differentiation (i.e., Autograd), we can obtain the gradient of the objective function to machine precision with no additional implementation overhead. These function transformations also enable acceleration on GPUs for energy evaluations and constrained geometry optimizations.[54] In the context of a method like L-BFGS-B, this means that we can avoid the imprecision and massive overhead of finite difference or complex step approximations of the gradient of energy profiles. In this work, the ligand-specific fitting strategy used throughout the QM-benchmark and RBFE simulations refines torsional parameters by adjusting the barrier height, phase, periodicity, and the 1−4 electrostatic and van der Waals scaling factors (`scee` and `scnb`). These scaling factors modulate Coulombic and dispersion interactions between atoms three bonds apart, helping reproduce the local electronic environment more realistically and preventing overestimation of nonbonded contributions. AFFDO also provides a streamlined torsion-only variant in which only the torsional barrier heights are optimized while all other terms, including `scee` and `scnb`, remain fixed to their standard GAFF values. This alternative mode retains the standard GAFF nonbonded parametrization and provides a robust option when such ligand-specific tuning is not desired.

2.5. Generating Output Files

In this task, AFFDO prepares a series of files that can be directly used as input to Production Free Energy Simulation Setup and Analysis Framework (ProFESSA)[55] to run binding free energy simulations. ProFESSA is an automated, flexible, end-to-end workflow for performing alchemical free energy simulations using AMBER's GPU-accelerated MD engine[56−59] and is centered around the Drug Discovery Boost package,[60] which provides enhanced sampling features and analysis tools.[61−63] Since the parameter fitting described above is performed by running the fitting algorithm for each significant torsion separately, it results in multiple topology files, each containing a set of modifications corresponding to the optimized parameters. Furthermore, if the parametrization were performed using ligand fragments rather than the full ligand, one would end up with multiple topology files significantly different from each other. Therefore, it is necessary to merge all the changes and construct a final topology file for the full ligand, but such a topology file will be inconvenient to use in ProFESSA. This is because during the setup of free energy simulations, the ligands are embedded in condensed phases (solvent and protein environments), and new topology files are created. To circumvent this issue, AFFDO processes these topology files and generates an input file for the parmed program[64] of AmberTools.[53] Furthermore, force field library (`.lib`), force field modification (`.frcmod`), and a coordinate file that also contains atomic charges (`.mol2`) are also generated. If the parametrization was performed using ligand fragments, ligand-fragment atom maps are taken into consideration during the preparation of above files. The resulting files can then be used to setup binding free energy simulations. The parmed input file can be used to update the solvent−ligand and protein−ligand topology files with new torsion parameters. Analysis of free energy simulation results were performed using the FE-ToolKit package.[65]

https://doi.org/10.1021/acs.jcim.6c00528
J. Chem. Inf. Model. 2026, 66, 3206−3219

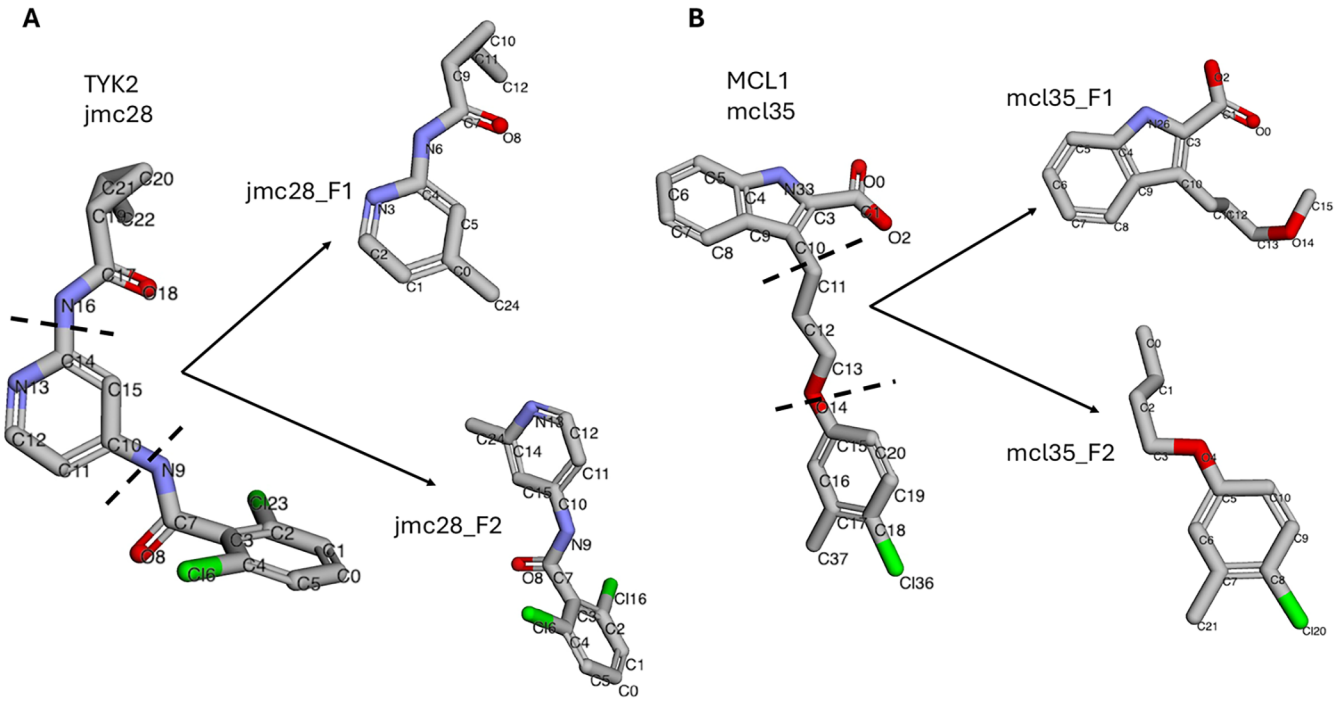

Figure 2. General fragmentation pattern of the representative ligands. Fragmentation of jmc28 ligand from TYK2 (A) and L35 ligand from MCL1 (B) protein systems are shown here. This AFFDO integrated feature streamlines the torsion parametrization process.

3. BENCHMARKING STUDIES

We benchmarked AFFDO by reparameterizing torsions for ligands drawn from the widely used Wang et al. data set.[27] Specifically, we considered ligand series associated with the TYK2 and MCL1 systems, which comprise closely related analogs differing by modest functional-group substitutions, thereby mimicking a lead-optimization campaign.

Two benchmark components were evaluated. First, we assessed torsion-profile agreement by comparing GAFF2 versus AFFDO-optimized torsional potential energy surfaces against QC constrained dihedral scans. This torsion-profile benchmark includes 16 TYK2 ligands and 40 MCL1 ligands. Second, we evaluated the downstream impact of torsion reparameterization on RBFE predictions using the RBFE networks defined in the same data set; these RBFE results are reported for 15 TYK2 transformations and 14 MCL1 transformations (see Table S2 of the Supporting Information). We also report the timing and computational cost of torsion parameter optimization with AFFDO.

For these benchmarking tests, AFFDO was installed on a Linux server with Intel Core i7−12700K CPU (12 core) and 64 GB RAM. A cloud instance consisting of 4 NVIDIA RTX4090 GPUs, 36 AMD EPYC 7282 CPU cores and 64 GB memory was reserved from Vast.ai and configured as the backend of AFFDO. AmberTools/24, QUICK/24.03 and xTB/6.6.0 were installed on the cloud instance. For AmberTools and QUICK compilations, CUDA/12.0, GCC/9.3.0 and OpenMPI/4.0.3 were used.

3.1. Torsion Reparameterization Using AFFDO

Selected ligand structures were obtained from published data and used as input to the AFFDO workflow. Ligands were fragmented using FragMentor (Figure 2), and significant torsions were identified. The threshold energy value for conformer selection and the torsion angle difference between a pair of conformers for significant torsion identification were set to 3.5 kcal/mol and 30°, respectively. The resulting number of significant torsions for the representative ligands, jmc28 and L35, were 4 and 5, respectively.

To generate the QC reference data, torsional scans of TYK2 ligands were carried out at the PBE0-D3BJ/6−31G* level of theory. For anionic MCL1 ligands, PBE0-D3BJ/6−31+G* level of theory was used. An interval of 20° was used for all scans. The default settings were used for parameter fitting (GAFF2 torsion parameters as initial guess, 0.05 step size, 0.001 kcal/mol RMSD convergence criteria). The torsional scan energy profiles were also computed using standard GAFF2, as well as using parameters optimized by AFFDO ("customized GAFF2"). The computed torsional scan energy profiles using QC, GAFF2 and AFFDO-optimized parameters of the most significant torsions are depicted in Figure 3. The plots for the remaining torsions are reported in Figure S1.

As can be seen in Figure 3A,3B, standard GAFF2 poorly describes the selected torsions of the jmc28 ligand. Specifically, in the former (6−7−9−11 torsion), GAFF2 results in lower barrier heights and a phase shift with respect to the Reference QC energy profile. In the latter (7−9−10−11 torsion), the phase of GAFF2 is consistent with QC, but the barrier heights are significantly different and the minima are not correctly reproduced. The energy profiles of L35 (Figure 3C,D) suggest that GAFF2 describes the corresponding torsions properly. The minima are consistent with those of the QC PES although barrier heights are slightly different. As clearly evident from all cases, the reparameterization improves the torsional scan plots significantly, as customized GAFF2 correctly describes the selected torsions.

To quantitatively assess the improvements achieved by AFFDO, we computed root-mean-square error (RMSE), mean absolute error (MAE), and correlation metrics (Pearson and Spearman coefficients) for both TYK2 and MCL1 ligand parametrization. The results, presented in Table 1, compare

https://doi.org/10.1021/acs.jcim.6c00528
J. Chem. Inf. Model. 2026, 66, 3206−3219

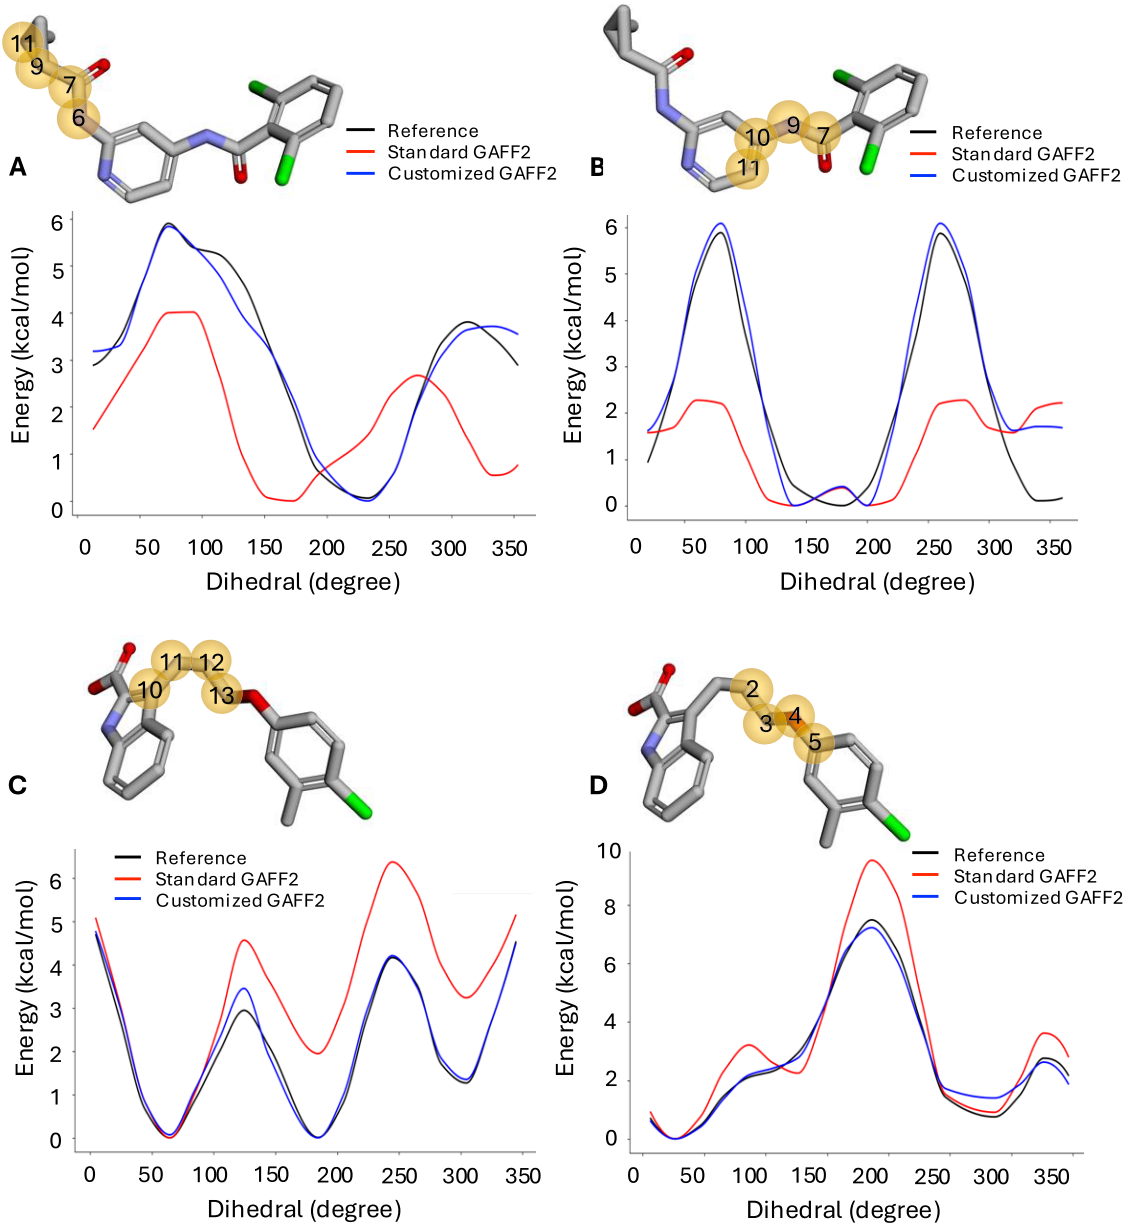

Figure 3. Torsional scan energy profiles of the most significant torsions of jmc28 ligand from TYK2 (A, B) and L35 ligand from MCL1 protein system (C, D). The corresponding torsions are highlighted in the molecular structures by yellow spheres.

GAFF2 and AFFDO performance relative to the QC reference data within each system. For clarity, these statistics quantify agreement with QC torsion-scan energy profiles evaluated over all dihedral-scan grid points and should not be interpreted as RBFE predictive accuracy. Metrics are computed from energy differences across the scanned dihedral angles using the QC (DFT) profile as reference. Additionally, RMSD values quantify deviations in optimized geometries relative to QC-optimized structures, providing further validation of the impact of torsion reparameterization on molecular conformations.

AFFDO demonstrates substantial improvements in both systems, with the enhancements being particularly pronounced in TYK2. The RMSE reduction for TYK2 (1.92 → 0.24 kcal/mol) indicates a significant refinement in torsional parametrization accuracy, leading to a much closer match with QC energy profiles. Similarly, Pearson and Spearman correlation coefficients show a notable increase, confirming the improved consistency between AFFDO-optimized torsion potentials and the reference data. While the improvements in the MCL1 (1.08 → 0.69 kcal/mol) system are less pronounced, they remain consistent with the observed torsional energy profiles (Figures 3 and S1). In this case, the primary contribution of AFFDO lies in correcting the energy barrier heights rather than major phase shifts or minima realignments, which were already reasonably well captured by GAFF2. This analysis provides a direct measure of AFFDO's effectiveness in refining ligand torsion potentials by demonstrating its capability to systematically reduce deviations from quantum mechanical reference data. These improvements highlight AFFDO's ability to enhance force field accuracy, particularly in cases where standard GAFF2 struggles to correctly capture torsional energy landscapes.

3.2. Relative Binding Free Energy Simulations

To further investigate the practical implications of AFFDO-optimized parameters, we assess their impact on relative binding free energy (RBFE) simulations by evaluating the fidelity of the

https://doi.org/10.1021/acs.jcim.6c00528
J. Chem. Inf. Model. 2026, 66, 3206−3219

Table 1. Comparison of GAFF2 and AFFDO Torsion-Profile Accuracy for the TYK2 and MCL1 Ligand Sets[c]

	TYK2		MCL1	
Metric	GAFF2	AFFDO	GAFF2	AFFDO
MAE[a] (kcal/mol)	1.56 ± 0.21	0.20 ± 0.06	0.81 ± 0.06	0.51 ± 0.06
RMSE[a] (kcal/mol)	1.92 ± 0.29	0.24 ± 0.07	1.08 ± 0.07	0.69 ± 0.08
Pearson (r) (unitless)	0.57	0.96	0.90	0.94
Spearman (ρ) (unitless)	0.52	0.95	0.90	0.94
Max RMSD[b] (Å)	0.44	0.39	0.63	0.63

[a]Uncertainties reported as 95% confidence interval, calculated analytically. [b]RMSD values measure conformational deviations relative to QC-optimized structures. [c]Metrics quantify agreement with QC torsional potential energy surfaces obtained from constrained dihedral scans evaluated on a 20° angular grid (i.e., scan-point energies) and should not be interpreted as RBFE predictive accuracy. The benchmark includes 16 TYK2 ligands and 40 MCL1 ligands.

customized GAFF2 force field with reoptimized torsion parameters in a series of ligand transformations targeting TYK2 and MCL1 protein systems. The initial protein structures were obtained from published data[27] and simulations were prepared using the ProFESSA workflow with the AMBER FF19SB force field for proteins, GAFF2 for ligands, and TIP4P-Ew[66] for water molecules. A second set of simulations were prepared by updating the AMBER topology files with optimized torsion parameters obtained from AFFDO (using the torsion + `scee/scnb` fitting mode described above). All remaining topology terms, including the ligand partial charges (AM1-BCC as assigned by ProFESSA), were kept identical between the GAFF2 and AFFDO cases. Accordingly, any differences in RBFE performance reported here are attributable to the modified torsional terms rather than a change in the charge model. This was done by executing `parmed` with the AFFDO-generated input file and the corresponding aqueous- and complex-phase topology files. All simulations were executed on the aforementioned cloud instance with 4 RTX4090 GPUs. Three independent trials were performed for each RBFE edge; each trial used 21 λ windows with 2.4 ns equilibration and 5 ns production per window (see Section S3 in the Supporting Information for full RBFE protocol details).

In Figure 4, we report RBFE values computed using standard GAFF2 and reparameterized GAFF2 along with experimental data. Experimental and computed RBFEs for *all* transformations (standard GAFF2 and AFFDO-parametrized GAFF2), together with statistical uncertainties, are reported in Tables S2 and S3 (Supporting Information, Section S4). Three key insights emerge from these data. First, as shown in Figure 4A,B, torsion reparameterization improves the RBFE predictions of most ligand transformations. Notably, most points in blue (AFFDO) are closer to the dashed lines than the red points (GAFF2), indicating reduced deviation from experimental values. Across all transformations, including both improved and unchanged cases, AFFDO reduces the MAE of standard GAFF2 by approximately ∼0.4 kcal/mol. Among the subset of transformations (ca. 80%) where reparameterization is beneficial, the average reduction in absolute error is ∼0.5 kcal/mol, with the highest observed improvement reaching 2.45 kcal/mol (see Table S3 for this subset). Second, while it is reasonable to expect improvements in RBFE predictions for the TYK2 system due to substantial changes in torsional energy profiles after reparameterization (Figure 3A,3B), the consistent improvements observed for the MCL1 system were more surprising, given that the torsional energy barriers for its ligands were only moderately altered (Figure 3C,D). This suggests that even subtle modifications to the torsional landscape can enhance RBFE accuracy. Third, the uncertainty associated with RBFE estimates is consistently lower in reparameterized GAFF2, indicating improved convergence in free-energy calculations. For direct comparisons between standard GAFF2 and reparameterized GAFF2, uncertainties for both force fields are explicitly shown in Figure 4 (horizontal and vertical error bars, respectively), ensuring a statistically balanced comparison. These findings collectively reinforce the role of torsional refinement in enhancing force field performance, though its effectiveness varies across different ligand transformations due to factors beyond torsion parameters alone.

To elucidate how torsion reparameterization impacts RBFE accuracy, we analyzed the torsional sampling patterns from molecular dynamics trajectories of representative ligand transformations (see detailed analysis in Section S5 of the SI). We found that customized torsion parameters significantly redefined the ligand potential energy surfaces, influencing conformational accessibility across both end-states. For example, in the transformation of the TYK2 system jmc23-ejm55 (Figure 5), the repair of the torsion substantially increased the heights of the energy barrier and refined the minima regions, resulting in a narrower and more coherent conformational sampling between the ligands. This improved sampling overlap facilitated smoother alchemical transformations and directly enhanced the reliability of RBFE predictions compared to standard GAFF2. A similar effect was observed in the MCL1 system for the L49−L67 transformation (Figure S4 of SI). However, these improvements were not universal. For instance, in the ejm49−ejm50 (TYK2) (Figure S3 of SI) and L67−L27 (MCL1) (Figure S4 of SI) transformations, the optimized parameters did not substantially alter the conformational landscape sampled by standard GAFF2, resulting in RBFE estimates comparable to those obtained without reparameterization.

Overall, this analysis demonstrates that torsion reparameterization improves RBFE estimates by significantly redefining the ligand potential energy surface in agreement with high-quality quantum mechanical reference data, thereby modulating conformational accessibility in both end-states and enhancing sampling quality. For the investigated cases, the reparameterization also leads to increased overlap between ligand conformational distributions, which facilitates sampling, leading to smoother alchemical free-energy pathways, thus reducing error bars of RBFE predictions. However, these conclusions are drawn exclusively from the present data set involving the TYK2 and MCL1 systems. Further validation across additional ligand transformations and protein systems will be required to draw more general conclusions. In particular, while good torsional parameters are essential, there are other factors that influence the quality of RBFE estimates, including nonbonded interactions, charge models, and error cancellation. Importantly, AFFDO's value extends beyond RBFE to other free energy estimations, such as absolute binding free energies (ABFE) and solvation free energies, which could directly benefit from enhanced PES descriptions of individual ligands.

https://doi.org/10.1021/acs.jcim.6c00528
J. Chem. Inf. Model. 2026, 66, 3206−3219

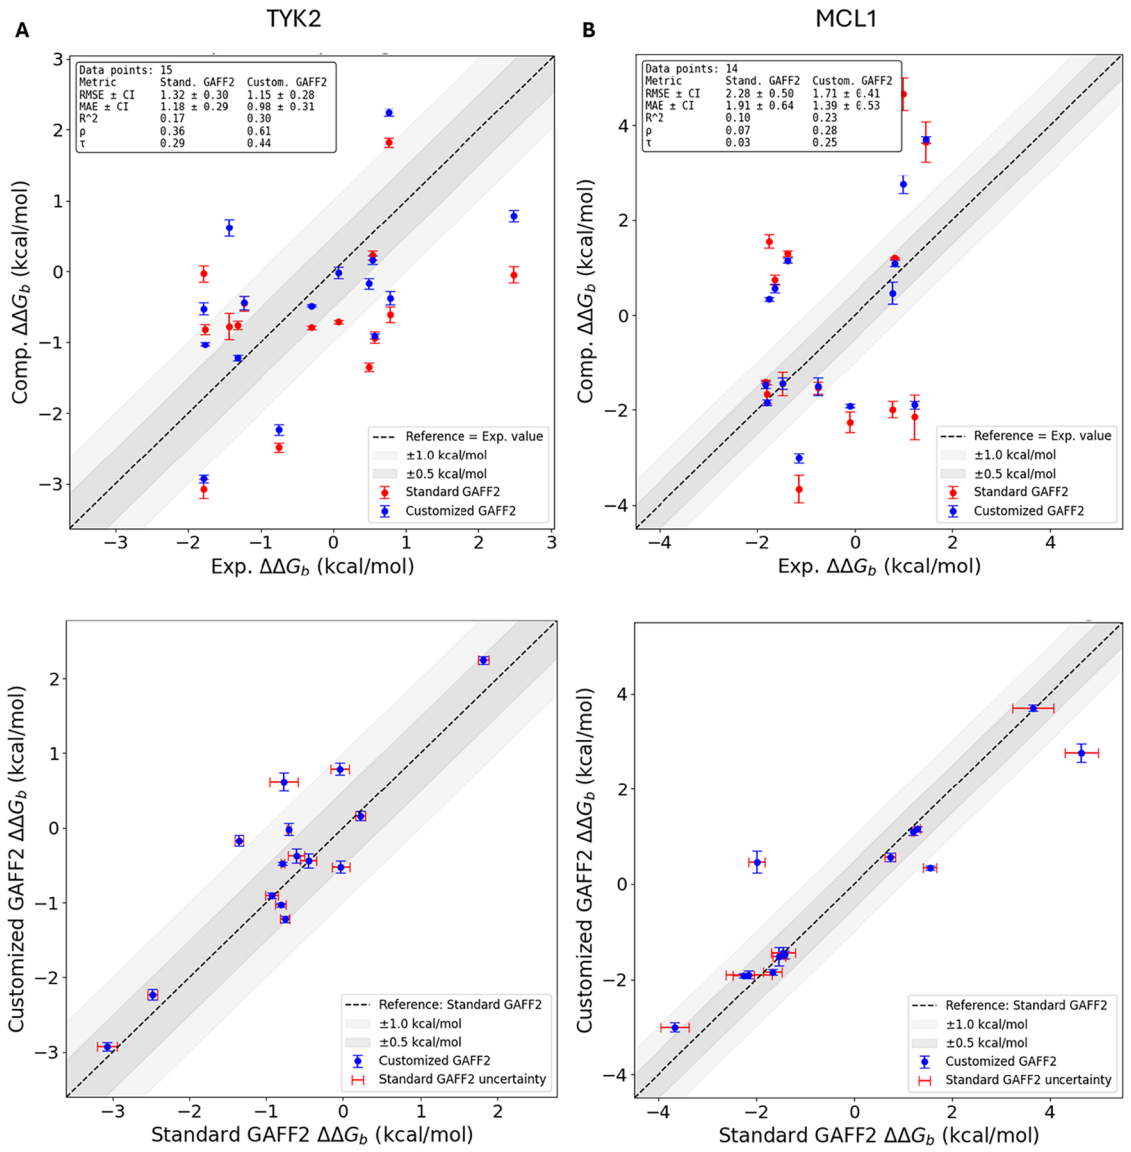

Figure 4. Computed versus experimental relative binding free energies (RBFEs) for the TYK2 (top left, panel A) and MCL1 (top right, panel B) ligand sets. Top panels compare experimental RBFEs to predictions from standard GAFF2 (red) and customized GAFF2 (blue). Bottom panels compare customized GAFF2 against standard GAFF2 predictions for each transformation. Error bars denote statistical uncertainty from three independent RBFE trials using identical analysis protocols (see Supporting Information); in the bottom panels, horizontal (red) and vertical (blue) error bars correspond to standard GAFF2 and customized GAFF2 uncertainties, respectively. The dashed line indicates perfect agreement. Insets report summary statistics for each force field (RMSE, MAE, R^2, Spearman ρ, Kendall τ, and N).

3.3. Timing and Computational Cost of Torsion Optimization with AFFDO

In Table 2, we report wallclock times taken for each step of the AFFDO workflow using a fragment (out of the two available) from jmc28 and L35 ligand parametrization. The most time-consuming steps include centroid optimization (due to QC computations), torsional scans (again due to QC), and parameter fitting (due to a series of torsional scans needed during gradient-based optimization).

Although the computational cost of torsion optimization using AFFDO can be high (47 min for jmc28_F1 and 7.3 h for L35_F1), RBFE computations themselves are significantly more time-consuming (approximately 20−24 h per MCL1 transformation and over 48 h per TYK2 transformation on identical GPU hardware); these runtimes correspond to the full triplicate campaign, i.e., three independent RBFE trials per transformation. The additional cost for torsion optimization is thus easily justifiable and further is inherently parallelizable across scanned points. Estimating the total time for reparameterizing a molecule based purely on its size or number of torsions is challenging since QC torsional scan durations depend significantly on the complexity of the conformational energy landscape, convergence of the constrained geometry optimization and SCF procedures, which can vary considerably.

Furthermore, the associated wallclock times can be reduced in a number of ways. In general, providing more CPU cores for xtb calculations and employing a multilevel parallelism can be helpful. Similarly, providing more GPUs for QC calculations would reduce the wall time spent on torsional scans. For xtb runs reported in this study, we used all available CPU cores in the cloud instance but torsional scan calculations were launched one after another. Instead, launching all scans simultaneously with multiple CPU cores for each calculation could lead to better

https://doi.org/10.1021/acs.jcim.6c00528
J. Chem. Inf. Model. 2026, 66, 3206−3219

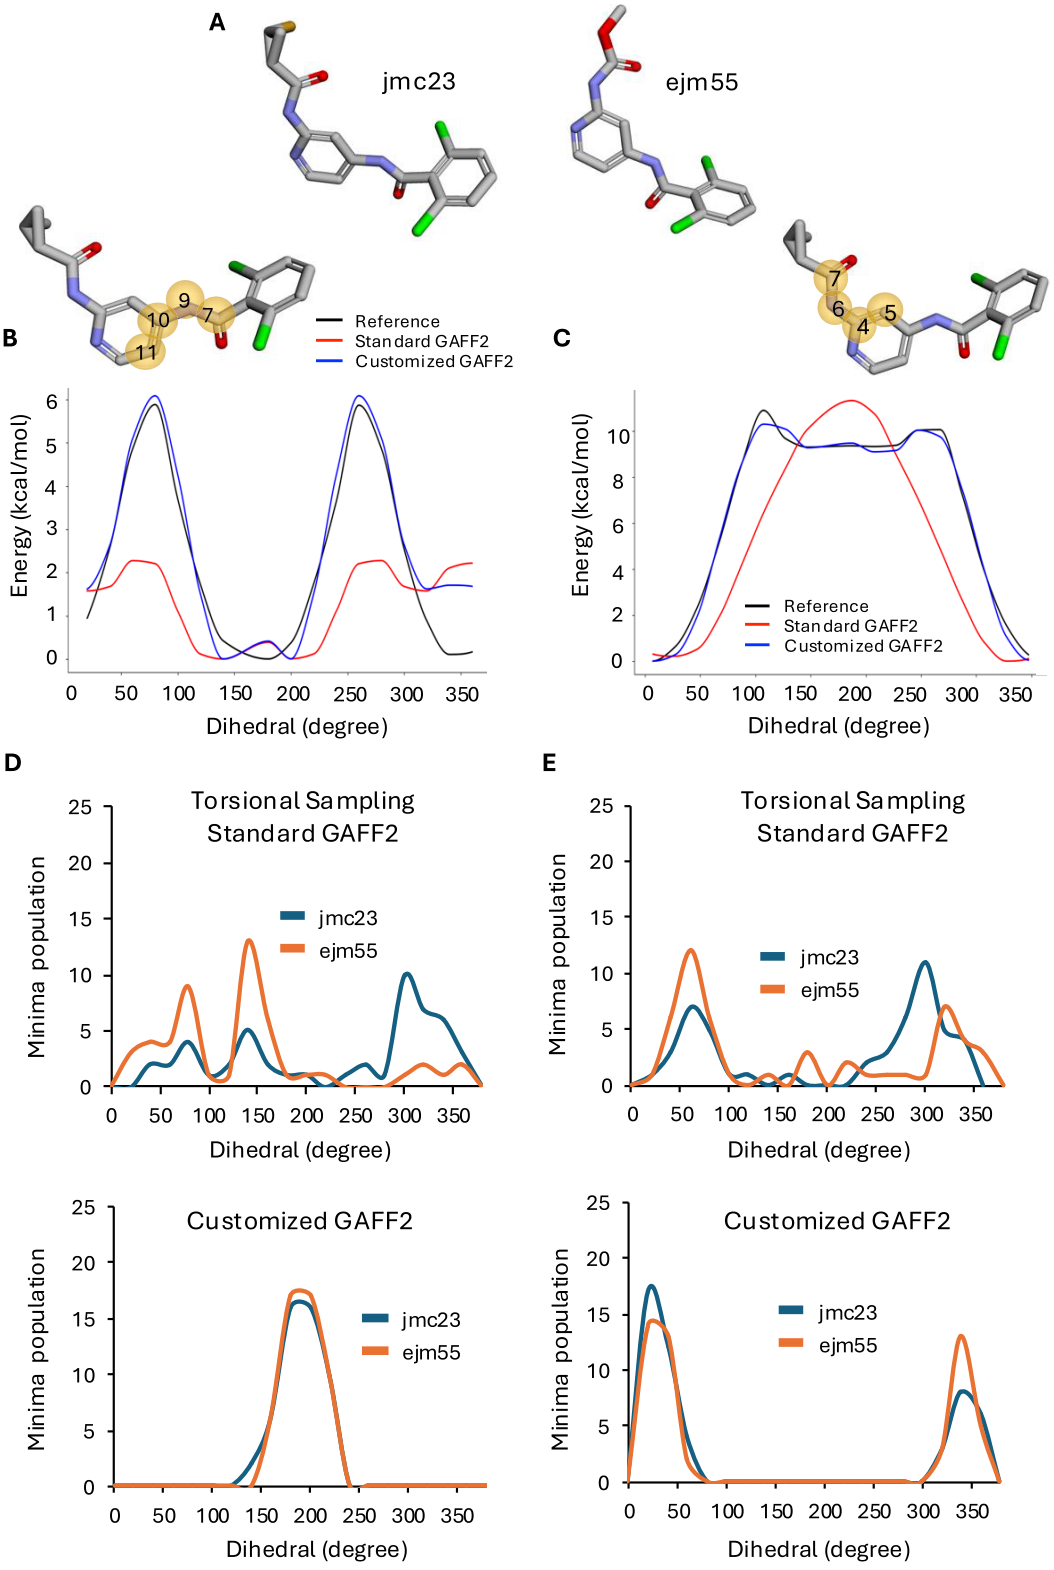

Figure 5. Comparative torsional sampling analysis between standard GAFF2 and customized GAFF2 for the jmc23-ejm55 ligand transformation (TYK2 system) in aqueous solution, an example where customized GAFF2 improves RBFE accuracy. (A) Chemical structures of ligands jmc23 and ejm55 show structural variations. (B–C) Reference torsional energy profiles (black) and corresponding profiles generated with standard GAFF2 (red) and customized GAFF2 (blue) for torsions 7–9–10–11 (B) and 7–6–4–5 (C). Insets show the reparameterized ligand structure indicating the torsions analyzed. (D, E) Torsional sampling distributions of torsions 7–9–10–11 (D) and 7–6–4–5 (E) comparing standard GAFF2 (top panel) and customized GAFF2 (bottom panel) for ligands jmc23 and ejm55, illustrating the population distributions across different PES minima. All panels report the dihedral angle scale in degrees (*X*-axis) using the absolute torsion values, with the minimum located at its corresponding dihedral angle.

https://doi.org/10.1021/acs.jcim.6c00528
J. Chem. Inf. Model. 2026, 66, 3206–3219

Table 2. Wall Time (Seconds) Taken from AFFDO Workflow Steps during Torsion Reparameterization of jmc28_F1 and L35_F1 Ligand Fragments (See Figure 2) on a Cloud Instance with 36 AMD EPYC 7282 CPU cores, 4 NVIDIA RTX4090 GPUs, and 64 GB Memory

Step	Wall time (s)	
	jmc28_F1 (1 torsion)	L35_F1 (3 torsions)
1. Ligand fragmentation	<1	<1
2. Conformer generation	4	4
3. Conformer minimization (XTB)	27	148
4. Clustering	<1	<1
5. Centroid optimization (QC and MM)	900	2600
6. Significant minima selection	<1	<1
7. Significant torsion selection	<1	<1
8. Torsional scan (XTB and QC)	1742	22995
9. Fitting torsion parameters	133	623
10. Parmed input file generation	<1	<1
Total run time	2807	26370

performance overall. QC calculations were launched in batches of 4, utilizing all 4 GPUs for 4 different calculations. Availability of more GPUs will allow increasing this batch size and therefore speed up the procedure. Additionally, in this study, we used conservative cutoffs for QC calculations (10^{-7} *a.u.* density matrix cutoff for self-consistent field (SCF) convergence, 10^{-12} *a.u.* integral cutoff, 10^{-8} exchange correlation cutoff, and default geometry optimization cutoffs) and 20° intervals for torsional scans. Loosening these cutoffs and carrying out scans with more relaxed angular intervals (e.g., 30°), which is configurable in AFFDO, can lead to significant performance gains.

4. CONCLUSIONS AND FUTURE DIRECTIONS

We have developed a new computational infrastructure called AFFDO that facilitates force field development and optimization. In particular, the current version enables the reparameterization of GAFF/GAFF2 torsion parameters for drug-like molecules. Benchmarks demonstrated that AFFDO improves agreement with QC torsion-scan reference data (by ∼0.8 kcal/mol RMSE across the benchmark torsion-scan data set[27]) in an automated manner and within a reasonable computational time using cloud computing resources. RBFE simulations using standard GAFF2 versus reparameterized GAFF2 showed that AFFDO-optimized torsions reduced overall prediction errors for MCL1 (by ∼0.5 kcal/mol RMSE), whereas the aggregate improvement for TYK2 was smaller (by ∼0.2 kcal/mol RMSE) and not statistically significant at the full-network level, indicating that for the TYK2 benchmark, RBFE accuracy may be limited by model deficiencies beyond torsion descriptions.

AFFDO can be further improved in terms of usability and performance. The overall runtime of the reparameterization process is currently dominated by reference data generation; this cost could be reduced by enabling support for multiple cloud instances, which is not possible at the moment. Runtime could also be substantially reduced by leveraging fast AI-based models for torsional-scan reference data generation, such as the QDπ2 model,[67−69] along with storing reference data in a database[70] so that it can be reused when reparameterizing new ligands. Finally, we will endeavor to expand the scope of our force field parameter optimization infrastructure to include improved nonbonded (e.g., Lennard-Jones), electrostatic, and linear-response models.

■ ASSOCIATED CONTENT

Data Availability Statement

AFFDO is available as a web service at https://www.attmosdiscovery.com/tools free of charge, where users can submit jobs and access results via a job ID (including an in-browser report with torsion plots) and downloadable outputs. Input files used for the benchmarks in this work (both AFFDO runs and the RBFE workflow) are provided in the Supporting Information.

Supporting Information

The Supporting Information is available free of charge at https://pubs.acs.org/doi/10.1021/acs.jcim.6c00528.

Torsion plots, performance evaluation of AFFDO torsional parametrization across additional systems, technical details of RBFE simulations, RBFE predictions, and sampling investigation (PDF)

Input files for AFFDO torsion reparameterization runs and the RBFE simulations reported in this work (ZIP)

■ AUTHOR INFORMATION

Corresponding Authors

Andreas W. Götz − *ATTMOS Inc., East Lansing, Michigan 48823, United States; San Diego Supercomputer Center, University of California San Diego, La Jolla, California 92093-0505, United States;* orcid.org/0000-0002-8048-6906; Email: agoetz@sdsc.edu

Kenneth M. Merz, Jr. − *ATTMOS Inc., East Lansing, Michigan 48823, United States; Department of Chemistry and Department of Biochemistry and Molecular Biology, Michigan State University, East Lansing, Michigan 48824-1322, United States;* orcid.org/0000-0001-9139-5893; Email: merz@chemistry.msu.edu

Darrin M. York − *ATTMOS Inc., East Lansing, Michigan 48823, United States; Laboratory for Biomolecular Simulation Research, Center for Integrative Proteomics Research and Department of Chemistry and Chemical Biology, Rutgers University, Piscataway, New Jersey 08854-8087, United States;* orcid.org/0000-0002-9193-7055; Email: Darrin.York@rutgers.edu

Hasan Metin Aktulga − *Department of Computer Science and Engineering, Michigan State University, East Lansing, Michigan 48824-1322, United States; ATTMOS Inc., East Lansing, Michigan 48823, United States*; Email: hma@msu.edu

Madushanka Manathunga − *ATTMOS Inc., East Lansing, Michigan 48823, United States;* orcid.org/0000-0002-3594-8112; Email: madu@attmosdiscovery.com

Authors

Alejandro Blanco-Gonzalez − *ATTMOS Inc., East Lansing, Michigan 48823, United States;* orcid.org/0000-0001-7379-2588

William Betancourt − *Department of Computer Science and Engineering, Michigan State University, East Lansing, Michigan 48824-1322, United States*

Ryan Michael Snyder − *Laboratory for Biomolecular Simulation Research, Center for Integrative Proteomics Research and Department of Chemistry and Chemical Biology, Rutgers University, Piscataway, New Jersey 08854-8087, United States*

https://doi.org/10.1021/acs.jcim.6c00528
J. Chem. Inf. Model. 2026, 66, 3206−3219

Shi Zhang − *Laboratory for Biomolecular Simulation Research, Center for Integrative Proteomics Research and Department of Chemistry and Chemical Biology, Rutgers University, Piscataway, New Jersey 08854-8087, United States;* orcid.org/0000-0002-0281-9314

Timothy J. Giese − *Laboratory for Biomolecular Simulation Research, Center for Integrative Proteomics Research and Department of Chemistry and Chemical Biology, Rutgers University, Piscataway, New Jersey 08854-8087, United States;* orcid.org/0000-0002-0653-9168

Zeke A. Piskulich − *Laboratory for Biomolecular Simulation Research, Center for Integrative Proteomics Research and Department of Chemistry and Chemical Biology, Rutgers University, Piscataway, New Jersey 08854-8087, United States;* orcid.org/0000-0003-0304-305X

Complete contact information is available at: https://pubs.acs.org/10.1021/acs.jcim.6c00528

Notes

The authors declare the following competing financial interest(s): Andreas W. Goetz, Darrin M. York, Kenneth M. Merz, Hasan Metin Aktulga, and Madushanka Manathunga are co-founders of ATTMOS, Inc. Zeke A. Piskulich, Ryan Snyder, and Timothy J. Giese are consultants for ATTMOS, Inc.

■ ACKNOWLEDGMENTS

Research reported in this publication was supported by the National Institute of General Medical Sciences of the National Institutes of Health under Award No. R44GM150347 and R01GM130641. The content is solely the responsibility of the authors and does not necessarily represent the official views of the National Institutes of Health. MM and HMA thank MSU Research Foundation, Lansing Regional SmartZone, and East Lansing Technology Innovation Center. Computational resources were provided by the Advanced Cyberinfrastructure Coordination Ecosystem: Services and Support (ACCESS) program, which is supported by the National Science Foundation grants No. 2138259, 2138286, 2138307, 2137603, and 2138296 (supercomputer Delta at NCSA through allocation CHE240135). This material is partly based on work supported by the National Science Foundation under Award Nos. 2209717, 2209718, and 2435622. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the National Science Foundation.

■ REFERENCES

(1) York, D. M. Modern Alchemical Free Energy Methods for Drug Discovery Explained. *ACS Phys. Chem. Au* **2023**, *3*, 478−491.

(2) Cournia, Z.; Chipot, C.; Roux, B.; York, D. M.; Sherman, W. *Free Energy Methods in Drug Discovery—Introduction*; Armacost, K.; Thompson, D. C., Eds.; American Chemical Society, 2021; Vol. *1397*, pp 1−38.

(3) Cole, D. J.; Horton, J. T.; Nelson, L.; Kurdekar, V. The future of force fields in computer-aided drug design. *Future Med. Chem.* **2019**, *11*, 2359−2363.

(4) Cornell, W. D.; Cieplak, P.; Bayly, C. I.; Gould, I. R.; Merz, K. M.; Ferguson, D. M.; Spellmeyer, D. C.; Fox, T.; Caldwell, J. W.; Kollman, P. A. A Second Generation Force Field for the Simulation of Proteins, Nucleic Acids, and Organic Molecules. *J. Am. Chem. Soc.* **1995**, *117*, 5179−5197.

(5) MacKerell, A. D., Jr.; Bashford, D.; Bellott, M.; Dunbrack, R. L., Jr.; Evanseck, J. D.; Field, M. J.; Fischer, S.; Gao, J.; Guo, H.; Ha, S.; Joseph-McCarthy, D.; Kuchnir, L.; Kuczera, K.; Lau, F. T. K.; Mattos, C.; Michnick, S.; Ngo, T.; Nguyen, D. T.; Prodhom, B.; Reiher, W. E.; Roux, B.; Schlenkrich, M.; Smith, J. C.; Stote, R.; Straub, J.; Watanabe, M.; Wiórkiewicz-Kuczera, J.; Yin, D.; Karplus, M. All-Atom Empirical Potential for Molecular Modeling and Dynamics Studies of Proteins. *J. Phys. Chem. B* **1998**, *102*, 3586−3616.

(6) Jorgensen, W. L.; Tirado-Rives, J. The OPLS [optimized potentials for liquid simulations] potential functions for proteins, energy minimizations for crystals of cyclic peptides and crambin. *J. Am. Chem. Soc.* **1988**, *110*, 1657−1666.

(7) Qiu, Y.; Smith, D. G. A.; Boothroyd, S.; Jang, H.; Hahn, D. F.; Wagner, J.; Bannan, C. C.; Gokey, T.; Lim, V. T.; Stern, C. D.; Rizzi, A.; Tjanaka, B.; Tresadern, G.; Lucas, X.; Shirts, M. R.; Gilson, M. K.; Chodera, J. D.; Bayly, C. I.; Mobley, D. L.; Wang, L.-P. Development and Benchmarking of Open Force Field v1.0.0—the Parsley Small-Molecule Force Field. *J. Chem. Theory Comput.* **2021**, *17*, 6262−6280.

(8) Boothroyd, S.; Behara, P. K.; Madin, O. C.; Hahn, D. F.; Jang, H.; Gapsys, V.; Wagner, J. R.; Horton, J. T.; Dotson, D. L.; Thompson, M. W.; Maat, J.; Gokey, T.; Wang, L.-P.; Cole, D. J.; Gilson, M. K.; Chodera, J. D.; Bayly, C. I.; Shirts, M. R.; Mobley, D. L. Development and Benchmarking of Open Force Field 2.0.0: The Sage Small Molecule Force Field. *J. Chem. Theory Comput.* **2023**, *19*, 3251−3275.

(9) Tian, C.; Kasavajhala, K.; Belfon, K. A. A.; Raguette, L.; Huang, H.; Migues, A. N.; Bickel, J.; Wang, Y.; Pincay, J.; Wu, Q.; Simmerling, C. ff19SB: Amino-Acid-Specific Protein Backbone Parameters Trained against Quantum Mechanics Energy Surfaces in Solution. *J. Chem. Theory Comput.* **2020**, *16*, 528−552.

(10) Zgarbová, M.; Šponer, J.; Otyepka, M.; Cheatham, T. E. I.; Galindo-Murillo, R.; Jurečka, P. Refinement of the Sugar−Phosphate Backbone Torsion Beta for AMBER Force Fields Improves the Description of Z- and B-DNA. *J. Chem. Theory Comput.* **2015**, *11*, 5723−5736.

(11) Wang, J.; Wolf, R. M.; Caldwell, J. W.; Kollman, P. A.; Case, D. A. Development and testing of a general amber force field. *J. Comput. Chem.* **2004**, *25*, 1157−1174.

(12) Wang, J.; Wolf, R. M.; Caldwell, J. W.; Kollman, P. A.; Case, D. A. Development and testing of a general amber force field. *J. Comput. Chem.* **2005**, *26*, 114.

(13) Stern, C. D.; Bayly, C. I.; Smith, D. G. A.; Fass, J.; Wang, L.-P.; Mobley, D. L.; Chodera, J. D. Capturing non-local through-bond effects in molecular mechanics force fields I: Fragmenting molecules for quantum chemical torsion scans *bioRxiv* 2022.

(14) Horton, J. T.; Boothroyd, S.; Wagner, J.; Mitchell, J. A.; Gokey, T.; Dotson, D. L.; Behara, P. K.; Ramaswamy, V. K.; Mackey, M.; Chodera, J. D.; Anwar, J.; Mobley, D. L.; Cole, D. J. Open Force Field BespokeFit: Automating Bespoke Torsion Parametrization at Scale. *J. Chem. Inf. Model.* **2022**, *62*, 5622−5633.

(15) Horton, J. T.; Allen, A. E. A.; Dodda, L. S.; Cole, D. J. QUBEKit: Automating the Derivation of Force Field Parameters from Quantum Mechanics. *J. Chem. Inf. Model.* **2019**, *59*, 1366−1381.

(16) Ringrose, C.; Horton, J. T.; Wang, L.-P.; Cole, D. J. Exploration and validation of force field design protocols through QM-to-MM mapping. *Phys. Chem. Chem. Phys.* **2022**, *24*, 17014−17027.

(17) Morado, J.; Mortenson, P. N.; Verdonk, M. L.; Ward, R. A.; Essex, J. W.; Skylaris, C.-K. ParaMol: A Package for Automatic Parameterization of Molecular Mechanics Force Fields. *J. Chem. Inf. Model.* **2021**, *61*, 2026−2047.

(18) Sami, S.; Menger, M. F.; Faraji, S.; Broer, R.; Havenith, R. W. A. Q-Force: Quantum Mechanically Augmented Molecular Force Fields. *J. Chem. Theory Comput.* **2021**, *17*, 4946−4960.

(19) Betz, R. M.; Walker, R. C. Paramfit: Automated optimization of force field parameters for molecular dynamics simulations. *J. Comput. Chem.* **2015**, *36*, 79−87.

(20) Huang, L.; Roux, B. Automated Force Field Parameterization for Nonpolarizable and Polarizable Atomic Models Based on Ab Initio Target Data. *J. Chem. Theory Comput.* **2013**, *9*, 3543−3556.

(21) Mayne, C. G.; Saam, J.; Schulten, K.; Tajkhorshid, E.; Gumbart, J. C. Rapid parameterization of small molecules using the force field toolkit. *J. Comput. Chem.* **2013**, *34*, 2757−2770.

(22) Kumar, A.; MacKerell, A. D., Jr. FFParam-v2.0: AComprehensive Tool for CHARMM Additive and Drude Polarizable Force-Field Parameter Optimization and Validation. *J. Phys. Chem. B* **2024**, *128*, 4385−4395.
(23) Devereux, C.; Smith, J. S.; Huddleston, K. K.; Barros, K.; Zubatyuk, R.; Isayev, O.; Roitberg, A. E. Extending the Applicability of the ANI Deep Learning Molecular Potential to Sulfur and Halogens. *J. Chem. Theory Comput.* **2020**, *16*, 4192−4202.
(24) 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*, 3192−3203.
(25) Bannwarth, C.; Caldeweyher, E.; Ehlert, S.; Hansen, A.; Pracht, P.; Seibert, J.; Spicher, S.; Grimme, S. Extended tight-binding quantum chemistry methods. *WIREs Comput. Mol. Sci.* **2021**, *11*, No. e1493.
(26) Galvelis, R.; Doerr, S.; Damas, J. M.; Harvey, M. J.; De Fabritiis, G. A Scalable Molecular Force Field Parameterization Method Based on Density Functional Theory and Quantum-Level Machine Learning. *J. Chem. Inf. Model.* **2019**, *59*, 3485−3493.
(27) Wang, L.; Wu, Y.; Deng, Y.; Kim, B.; Pierce, L.; Krilov, G.; Lupyan, D.; Robinson, S.; Dahlgren, M. K.; Greenwood, J.; Romero, D. L.; Masse, C.; Knight, J. L.; Steinbrecher, T.; Beuming, T.; Damm, W.; Harder, E.; Sherman, W.; Brewer, M.; Wester, R.; Murcko, M.; Frye, L.; Farid, R.; Lin, T.; Mobley, D. L.; Jorgensen, W. L.; Berne, B. J.; Friesner, R. A.; Abel, R. Accurate and Reliable Prediction of Relative Ligand Binding Potency in Prospective Drug Discovery by Way of a Modern Free-Energy Calculation Protocol and Force Field. *J. Am. Chem. Soc.* **2015**, *137*, 2695−2703.
(28) Darden, T.; York, D.; Pedersen, L. Particle mesh Ewald: An N·log(N) method for Ewald sums in large systems. *J. Chem. Phys.* **1993**, *98*, 10089−10092.
(29) Roos, K.; Wu, C.; Damm, W.; Reboul, M.; Stevenson, J. M.; Lu, C.; Dahlgren, M. K.; Mondal, S.; Chen, W.; Wang, L.; Abel, R.; Friesner, R. A.; Harder, E. D. OPLS3e: Extending Force Field Coverage for Drug-Like Small Molecules. *J. Chem. Theory Comput.* **2019**, *15*, 1863−1874.
(30) RDKit Cheminformatics Software RDKit: Open-Source Cheminformatics, 2024, https://www.rdkit.org. Accessed February, 2025.
(31) Riniker, S.; Landrum, G. A. Better Informed Distance Geometry: Using What We Know To Improve Conformation Generation. *J. Chem. Inf. Model.* **2015**, *55*, 2562−2574.
(32) Bannwarth, C.; Ehlert, S.; Grimme, S. GFN2-xTB—An Accurate and Broadly Parametrized Self-Consistent Tight-Binding Quantum Chemical Method with Multipole Electrostatics and Density-Dependent Dispersion Contributions. *J. Chem. Theory Comput.* **2019**, *15*, 1652−1671.
(33) Byrd, R. H.; Lu, P.; Nocedal, J.; Zhu, C. A Limited Memory Algorithm for Bound Constrained Optimization. *SIAM J. Sci. Comput.* **1995**, *16*, 1190−1208.
(34) Virtanen, P.; Gommers, R.; Oliphant, T. E.; Haberland, M.; Reddy, T.; Cournapeau, D.; Burovski, E.; Peterson, P.; Weckesser, W.; Bright, J.; van der Walt, S. J.; Brett, M.; Wilson, J.; Millman, K. J.; Mayorov, N.; Nelson, A. R. J.; Jones, E.; Kern, R.; Larson, E.; Carey, C. J.; Polat, İ.; Feng, Y.; Moore, E. W.; VanderPlas, J.; Laxalde, D.; Perktold, J.; Cimrman, R.; Henriksen, I.; Quintero, E. A.; Harris, C. R.; Archibald, A. M.; Ribeiro, A. H.; Pedregosa, F.; van Mulbregt, P. SciPy 1.0 Contributors SciPy 1.0: Fundamental Algorithms for Scientific Computing in Python. *Nat. Methods* **2020**, *17*, 261−272.
(35) Adamo, C.; Barone, V. Toward Reliable Density Functional Methods without Adjustable Parameters: The PBE0Model **1999** *110* 6158 6170 .
(36) Petersson, G. A.; Bennett, A.; Tensfeldt, T. G.; Al-Laham, M. A.; Shirley, W. A.; Mantzaris, J. A Complete Basis Set Model Chemistry. I. The Total Energies of Closedshell Atoms and Hydrides of the First-row Elements. *J. Chem. Phys.* **1988**, *89*, 2193−2218.
(37) Petersson, G. A.; Al-Laham, M. A. A Complete Basis Set Model Chemistry. II. Openshell Systems and the Total Energies of the First-row Atoms. *J. Chem. Phys.* **1991**, *94*, 6081−6090.
(38) Manathunga, M.; O'Hearn, K. A.; Shajan, A.; Smith, J.; Miao, Y.; He, X.; Ayers, K.; Brothers, E.; Palos, E.; Tripathy, V.; Götz, A. W.; Merz, K. M. QUICK-25.03, University of California, San Diego, CA and Michigan State University, East Lansing, MI, 2025 https://github.com/merzlab/QUICK. accessed February 2025.
(39) Miao, Y.; Merz, K. M., Jr. Acceleration of High Angular Momentum Electron Repulsion Integrals and Integral Derivatives on Graphics Processing Units. *J. Chem. Theory Comput.* **2015**, *11*, 1449−1462.
(40) Manathunga, M.; Miao, Y.; Mu, D.; Götz, A. W.; Merz, K. M., Jr. Parallel Implementation of Density Functional Theory Methods in the Quantum Interaction Computational Kernel Program. *J. Chem. Theory Comput.* **2020**, *16*, 4315−4326.
(41) Manathunga, M.; Jin, C.; Cruzeiro, V. W. D.; Miao, Y.; Mu, D.; Arumugam, K.; Keipert, K.; Aktulga, H. M.; Merz, K. M., Jr; Götz, A. W. Harnessing the Power of Multi-GPU Acceleration into the Quantum Interaction Computational Kernel Program. *J. Chem. Theory Comput.* **2021**, *17*, 3955−3966.
(42) Manathunga, M.; Aktulga, H. M.; Götz, A. W.; Merz, K. M., Jr. Quantum Mechanics/Molecular Mechanics Simulations on NVIDIA and AMD Graphics Processing Units. *J. Chem. Inf. Model.* **2023**, *63*, 711−717.
(43) Shajan, A.; Manathunga, M.; Götz, A. W.; Merz, K. M., Jr. Geometry Optimization: A Comparison of Different Open-Source Geometry Optimizers. *J. Chem. Theory Comput.* **2023**, *19*, 7533−7541.
(44) Kästner, J.; Carr, J. M.; Keal, T. W.; Thiel, W.; Wander, A.; Sherwood, P. DL-FIND: An open-source geomtery optimizer for atomistic simulations. *J. Phys. Chem. A* **2009**, *113*, 11856−11865.
(45) Nocedal, J. Updating quasi-Newton matrices with limited storage. *Math. Comput.* **1980**, *35*, 773−782.
(46) Billeter, S. R.; Turner, A. J.; Thiel, W. Linear scaling geometry optimization and transition state search in hybrid delocalised internal coordinates. *Phys. Chem. Chem. Phys.* **2000**, *2*, 2177−2186.
(47) Jakalian, A.; Bush, B. L.; Jack, D. B.; Bayly, C. I. Fast, efficient generation of high-quality atomic charges. AM1-BCC model: I. Method. *J. Comput. Chem.* **2000**, *21*, 132−146.
(48) Jakalian, A.; Jack, D. B.; Bayly, C. I. Fast, efficient generation of high-quality atomic charges. AM1-BCC model: II. Parameterization and validation. *J. Comput. Chem.* **2002**, *23*, 1623−1641.
(49) He, X.; Man, V. H.; Yang, W.; Lee, T.-S.; Wang, J. A fast and high-quality charge model for the next generation general AMBER force field. *J. Chem. Phys.* **2020**, *153*, No. 114502.
(50) Cornell, W. D.; Cieplak, P.; Bayly, C. I.; Kollman, P. A. Application of RESP charges to calculate conformational energies, hydrogen bond energies, and free energies of solvation. *J. Am. Chem. Soc.* **1993**, *115*, 9620−9631.
(51) Bayly, C. I.; Cieplak, P.; Cornell, W. D.; Kollman, P. A. A well-behaved electrostatic potential based method using charge restraints for determining atom-centered charges: the RESP model. *J. Phys. Chem. A* **1993**, *97*, 10269−10280.
(52) Wang, J.; Wang, W.; Kollman, P. A.; Case, D. A. Automatic atom type and bond type perception in molecular mechanical calculations. *J. Mol. Graphics Modell.* **2006**, *25*, 247−260.
(53) Case, D. A.; Aktulga, H. M.; Belfon, K.; Cerutti, D. S.; Cisneros, G. A.; Cruzeiro, V. W. D.; Forouzesh, N.; Giese, T. J.; Götz, A. W.; Gohlke, H.; Izadi, S.; Kasavajhala, K.; Kaymak, M. C.; King, E.; Kurtzman, T.; Lee, T.-S.; Li, P.; Liu, J.; Luchko, T.; Luo, R.; Manathunga, M.; Machado, M. R.; Nguyen, H. M.; O'Hearn, K. A.; Onufriev, A. V.; Pan, F.; Pantano, S.; Qi, R.; Rahnamoun, A.; Risheh, A.; Schott-Verdugo, S.; Shajan, A.; Swails, J.; Wang, J.; Wei, H.; Wu, X.; Wu, Y.; Zhang, S.; Zhao, S.; Zhu, Q.; Cheatham, T. E. I.; Roe, D. R.; Roitberg, A.; Simmerling, C.; York, D. M.; Nagan, M. C.; Merz, K. M., Jr. AmberTools. *J. Chem. Inf. Model.* **2023**, *63*, 6183−6191.
(54) Bradbury, J.; Frostig, R.; Hawkins, P.; Johnson, M. J.; Leary, C.; Maclaurin, D.; Necula, G.; Paszke, A.; VanderPlas, J.; Wanderman-Milne, S.; Zhang, Q. JAX: composable transformations of Python +NumPy programs, 2018, https://github.com/jax-ml/jax. accessed February 2025.

https://doi.org/10.1021/acs.jcim.6c00528
J. Chem. Inf. Model. 2026, 66, 3206−3219

(55) Ganguly, A.; Tsai, H.-C.; Fernández-Pendás, M.; Lee, T.-S.; Giese, T. J.; York, D. M. AMBER Drug Discovery Boost Tools: Automated Workflow for Production Free- Energy Simulation Setup and Analysis (ProFESSA). *J. Chem. Inf. Model.* **2022**, *62*, 6069−6083.
(56) Götz, A. W.; Williamson, M. J.; Xu, D.; Poole, D.; Le Grand, S.; Walker, R. C. Routine Microsecond Molecular Dynamics Simulations with AMBER on GPUs. 1. Generalized Born. *J. Chem. Theory Comput.* **2012**, *8*, 1542−1555.
(57) Salomon-Ferrer, R.; Götz, A. W.; Poole, D.; Le Grand, S.; Walker, R. C. Routine Microsecond Molecular Dynamics Simulations with AMBER on GPUs. 2. Explicit Solvent Particle Mesh Ewald. *J. Chem. Theory Comput.* **2013**, *9*, 3878−3888.
(58) Lee, T.-S.; Cerutti, D. S.; Mermelstein, D.; Lin, C.; LeGrand, S.; Giese, T. J.; Roitberg, A.; Case, D. A.; Walker, R. C.; York, D. M. GPU-Accelerated Molecular Dynamics and Free Energy Methods in Amber18: Performance Enhancements and New Features. *J. Chem. Inf. Model.* **2018**, *58*, 2043−2050.
(59) Case, D. A.; Cerutti, D. S.; Cruzeiro, V. W. D.; Darden, T. A.; Duke, R. E.; Ghazimirsaeed, M.; Giambaşu, G. M.; Giese, T. J.; Götz, A. W.; Harris, J. A.; Kasavajhala, K.; Lee, T.-S.; Li, Z.; Lin, C.; Liu, J.; Miao, Y.; Salomon-Ferrrer, R.; Shen, J.; Snyder, R.; Swails, J.; Walker, R. C.; Wang, J.; Wu, X.; Zeng, J.; Cheatham, T. E., III; Roe, D. R.; Roitberg, A.; Simmerling, C.; York, D. M.; Nagan, M. C.; Merz, K. M. J. Recent Developments in Amber Biomolecular Simulations. *J. Chem. Inf. Model.* **2025**, *65*, 7835−7843.
(60) Lee, T.-S.; Lin, Z.; Allen, B. K.; Lin, C.; Radak, B. K.; Tao, Y.; Tsai, H.-C.; Sherman, W.; York, D. M. *Free Energy Methods in Drug Discovery: Current State and Future Directions*; American Chemical Society, 2021; Chapter 7 pp, 161−187.
(61) Zhang, S.; Giese, T. J.; Lee, T.-S.; York, D. M. Alchemical Enhanced Sampling with Optimized Phase Space Overlap. *J. Chem. Theory Comput.* **2024**, *20*, 3935−3953.
(62) Lee, T.-S.; Tsai, H.-C.; Ganguly, A.; York, D. M. ACES: Optimized Alchemically Enhanced Sampling. *J. Chem. Theory Comput.* **2023**, *19*, 472−487.
(63) Tsai, H.-C.; Xu, J.; Guo, Z.; Yi, Y.; Tian, C.; Que, X.; Giese, T.; Lee, T.-S.; York, D. M.; Ganguly, A.; Pan, A. Improvements in Precision of Relative Binding Free Energy Calculations Afforded by the Alchemical Enhanced Sampling (ACES) Approach. *J. Chem. Inf. Model.* **2024**, *64*, 7046−7055.
(64) Shirts, M. R.; Klein, C.; Swails, J. M.; Yin, J.; Gilson, M. K.; Mobley, D. L.; Case, D. A.; Zhong, E. D. Lessons learned from comparing molecular dynamics engines on the SAMPL5 dataset. *J. Comput.-Aided Mol. Des.* **2017**, *31*, 147−161.
(65) Giese, T. J.; Snyder, R.; Piskulich, Z.; Barletta, G. P.; Zhang, S.; McCarthy, E.; Solen, E.; York, D. M. FE-ToolKit: A Versatile Software Suite for Analysis of High- Dimensional Free Energy Surfaces and Alchemical Free Energy Networks. *J. Chem. Inf. Model.* **2025**, *65*, 5273−5279.
(66) Horn, H. W.; Swope, W. C.; Pitera, J. W.; Madura, J. D.; Dick, T. J.; Hura, G. L.; Head-Gordon, T. Development of an improved four-site water model for biomolecular simulations: TIP4P-Ew. *J. Chem. Phys.* **2004**, *120*, 9665−9678.
(67) Zeng, J.; Tao, Y.; Giese, T. J.; York, D. M. QDπ: A Quantum Deep Potential Interaction Model for Drug Discovery. *J. Chem. Theory Comput.* **2023**, *19*, 1261−1275.
(68) Giese, T. J.; Zeng, J.; Lerew, L.; McCarthy, E.; Tao, Y.; Solen, E.; York, D. M. Software Infrastructure for Next-Generation QM/MM−ΔMLP Force Fields. *J. Phys. Chem. B* **2024**, *128*, 6257−6271.
(69) Zeng, J.; Giese, T. J.; Zhang, D.; Wang, H.; York, D. M. DeePMD-GNN: ADeePMD-kit Plugin for External Graph Neural Network Potentials. *J. Chem. Inf. Model.* **2025**, *65*, 3154−3160.
(70) Zeng, J.; Giese, T. J.; Götz, A. W.; York, D. M. The QDπ dataset, training data for drug-like molecules and biopolymer fragments and their interactions. *Sci. Data* **2025**, *12*, No. 693.

https://doi.org/10.1021/acs.jcim.6c00528
J. Chem. Inf. Model. 2026, 66, 3206−3219