DCBK: A Fragment-Based Hybrid Simulation and Machine Learning Framework for Predicting Ligand Dissociation Kinetics

Yang Li[*,+,1], Yixin Wang[+,2], Xueying Zheng[1], Jinjiang Guo[*,1], Rumin Zhang[*,1]

[1.] Global Health Drug Discovery Institute, Beijing, China, 100192

[2.] The Faculty of Medicine, Dalian University of Technology, Dalian, China, 116024

* Corresponding authors: yang.li@ghddi.org, jinjiang.guo@ghddi.org, rumin.zhang@ghddi.org

+ These authors contributed equally to this work

Abstract

Predicting small-molecule binding kinetics is central to drug discovery, but accurately estimating dissociation rates (k_{off}) remains challenging due to high energy barriers and the high computational cost of conventional simulations. Here, we present a hybrid, data-driven computational pipeline that integrates adaptive biased molecular dynamics (ABMD), umbrella sampling (US), fragment-level free energy calculations, and machine learning (ML) correction to achieve scalable, high-precision k_{off} predictions across diverse protein-ligand systems. By leveraging a BRICS-based fragmentation strategy and training ML models on experimental kinetics data, our approach not only reconstructs detailed free energy landscapes but also pinpoints key molecular fragments governing dissociation pathways. In contrast to purely physics-based simulations, this pipeline offers improved accuracy, enhanced transferability across molecular systems, and markedly higher computational efficiency, thereby enabling rigorous quantification of the effect of ligand modifications on residence time. The resulting framework offers a generalizable,

interpretable, and reproducible solution for binding kinetics prediction, providing actionable insights for lead optimization and mechanistic studies.

Introduction

Understanding ligand–target interactions through binding kinetics is increasingly recognized as a key aspect of rational drug discovery, particularly in lead optimization[1, 2]. While equilibrium binding affinity quantifies the strength of a compound's interaction with its target, kinetic parameters such as the association rate (k_{on}) and dissociation rate (k_{off}) provide mechanistic insight into the speed and lifetime of target engagement, a property closely linked to pharmacodynamic efficacy and target selectivity[3]. Even compounds with similar binding affinities can display markedly different dissociation kinetics, resulting in large variations in residence time and therapeutic efficacy. For example, HSP90 ligands achieve comparable affinities of binding to the wild-type protein and yet exhibit up to a 9-fold difference in residence time, while CDK8/CycC inhibitors show similar binding affinities but dramatically extended residence times—from minutes to hundreds of minutes—due to strengthened hinge-region hydrogen bonds and hydrophobic contacts, with these modifications experimentally confirmed to improve pharmacological outcomes[4, 5]. These cases underscore that optimizing binding affinity alone is insufficient, and that quantitative modeling of binding kinetics—particularly residence time, which is inversely related to the k_{off}—is critical for designing drugs with durable efficacy. This drives the need for computational methods capable of capturing both equilibrium thermodynamics and atomic-level kinetic barriers.

Experimental in vitro techniques measuring binding kinetics spanning a vast timescale of eight orders of magnitude from sub-seconds to hours and days, including time-resolved fluorescence, stopped- or quenched flow, biosensors such as surface plasmon resonance (SPR), global progress curve analysis (GPCA), radioligand or fluorescent ligand binding, and jump dilution, enable precise measurement of protein–ligand association and dissociation rates, generating valuable kinetic data that inform mechanistic understanding of ligand–target interactions[6, 7].Complementing these experimental resources, accurate computational prediction of k_{off} can prioritize compounds with favorable kinetic profiles at early stages of drug discovery, accelerating the identification of candidates with durable efficacy and optimal pharmacological properties.

However, direct computation of binding kinetics remains computationally demanding and requires extensive knowledge of the system and careful molecular dynamics (MD) sampling strategies[8]. Thermodynamically, the binding free energy (ΔG_{bind}) reflects the stability difference between bound and unbound states, whereas kinetics depend on the entire free energy landscape connecting these states through a dissociation transition state. The activation free energy barrier ($\Delta G^{\ddagger}$) is typically high, meaning that spontaneous unbinding events are rarely observed within conventional simulation timescales; for instance, a 100 ns unbiased MD trajectory is unlikely to capture dissociation events when ligand residence times extend to hours or longer. While landmark multi-microsecond simulations from D. E. Shaw Research have demonstrated that unbiased MD can reveal dissociation mechanisms in select systems, such studies remain computationally intensive. Likewise, ligand-GaMD and other enhanced sampling techniques have been successfully used to investigate unbinding pathways and estimate dissociation rates[9], but their application is often constrained by system size and pathway complexity.Likewise, ligand-GaMD and other enhanced sampling techniques have been successfully used to investigate unbinding pathways and estimate dissociation rates[9, 10], but their application is often constrained by system size and pathway complexity.

Numerous studies have demonstrated that a widely adopted and effective strategy for computing ligand dissociation kinetics is a two-stage approach. The first stage employs enhanced sampling or accelerated dynamics techniques to efficiently explore unbinding pathways and generate candidate reaction coordinates. Representative techniques include τ-Random Acceleration MD (τ-RAMD), Steered MD (SMD), Adaptive Steered MD (ASMD), Metadynamics (MetaD), Scaled MD, and Adaptive Biased MD (ABMD) [11] etc., which bias the system or modify the potential energy landscape to facilitate ligand dissociation. Representative techniques include τ-Random Acceleration MD (τ-RAMD), Steered MD (SMD), Adaptive Steered MD (ASMD), Metadynamics (MetaD), Scaled MD, and Adaptive Biased MD (ABMD) [11] etc., which bias the system or modify the potential energy landscape to facilitate ligand dissociation. Complementary approaches, such as Site Identification by Ligand Competitive Saturation (SILCS), can further inform the selection of plausible dissociation pathways. Building on the pathways identified in the first stage, the second stage involves rigorous free energy reconstruction, typically performed using

methods such as Adaptive Steered MD (ASMD) combined with Jarzynski averaging or MetaD-derived free energy surface analysis. More computationally intensive techniques, including US, nudged elastic band (NEB) calculations[12], and attach–pull–release (APR) protocols[13], offer highly detailed characterization of the dissociation transition state. More computationally intensive techniques, including US, nudged elastic band (NEB) calculations[12], and attach–pull–release (APR) protocols[13], offer highly detailed characterization of the dissociation transition state. These physics-based approaches reconstruct the free energy surface along a chosen reaction coordinate and estimate the $\Delta G^{\ddagger}$, which can then be related to k_{off} through transition state theory. While these physics-based reconstructions are rigorous and mechanistically interpretable, their substantial computational cost and complex setup (e.g., reaction coordinate selection, biasing windows, and restraint schemes) render them impractical for high-throughput screening.

In parallel, data-driven approaches have been developed to directly predict k_{off} from molecular structure. Quantitative structure–kinetics relationship (QSKR) models and ML frameworks, including random forest, gradient boosting, and graph neural networks (GNNs), have been applied to experimental kinetic datasets to capture complex structure–kinetic relationships. Publicly available resources such as the KIND dataset[14] (3,812 small molecules with measured k_{on}, k_{off}, and K_d), the *PDBbind-koff 2020* compilation[15] (~680 protein–ligand complexes), and the 501-entry unbinding rate dataset curated by Amangeldiuly et al[16].Publicly available resources such as the KIND dataset[14] (3,812 small molecules with measured k_{on}, k_{off}, and K_d), the *PDBbind-koff 2020* compilation[15] (~680 protein–ligand complexes), and the 501-entry unbinding rate dataset curated by Amangeldiuly et al[16]. provide valuable training and benchmarking material, although they remain limited relative to traditional affinity datasets and cover restricted chemical and target space. Building on these datasets, several predictive models have been developed[17, 18], including ModBind, exemplifying how ML can be applied to estimate k_{off}. Collectively, these studies demonstrate that while ML can efficiently predict ligand dissociation rates and aid in ligand optimization, purely data-driven approaches still face challenges in generalization and mechanistic interpretability, limiting their ability to provide detailed molecular insight into the determinants of binding kinetics.

Building upon these developments, we developed DCBK (Divide and Compute Binding Kinetics), a hybrid framework that integrates atomistic simulations and ML to predict ligand dissociation kinetics. DCBK first employs SMD to explore unbinding pathways, followed by umbrella sampling to quantify free-energy barriers along these routes. To resolve the energetic contributions of different molecular regions, each ligand is decomposed into BRICS-based fragments, for which fragment-specific activation energies are computed and subsequently refined by a data-driven correction model.

To enable systematic evaluation, we further constructed a comprehensive kinetic database comprising 60 ligand–protein complexes over five distinct targets. The dataset combines atomistic dissociation pathways, computed activation free energies, and experimentally measured k_{off}, forming a mechanistically interpretable benchmark for training and validating the DCBK framework. This integrated database-driven approach establishes a scalable foundation for quantitative prediction and mechanistic understanding of ligand unbinding kinetics.

Results

Overview of the Hybrid Workflow

Ligand dissociation from protein targets is a key determinant of drug efficacy, with residence time often more predictive of in vivo activity than equilibrium binding affinity. To capture the mechanistic basis of dissociation kinetics and enable accurate predictions, we developed DCBK (Divide and Compute Binding Kinetics), a hybrid computational workflow that integrates physics-based enhanced sampling with machine learning (Fig. 1A). The workflow begins with a curated protein–ligand complex, from which both intact ligands and chemically meaningful fragments are incorporated into the simulation system. Dissociation pathways are generated through supervised SMD, with cross-validation via ABMD, providing physically interpretable trajectories that inform subsequent free energy reconstruction. Umbrella sampling is then applied along the dissociation coordinate to compute ligand-level activation barriers ($\Delta G^{\ddagger}_{\text{whole}}$), quantifying the energetic difficulty of unbinding and establishing a foundation for fragment-level analysis

To characterize the kinetic diversity of the systems, we curated experimental dissociation rates for five target proteins, encompassing 60 distinct protein–ligand complexes. The measured $-\log_{10} k_{\text{off}}$ (namely $\text{p}k_{\text{off}}$) values span a broad range from approximately −2 to 6, covering both transiently bound ligands and slow-dissociating inhibitors (Fig. 1B). Each complex includes both the experimentally measured k_{off} and the corresponding protein–ligand structural model, providing a unified dataset of kinetic and structural information. This dataset serves as the basis for subsequent mechanistic simulations at both the ligand and fragment levels, as well as for training the DCBK machine-learning model. For each complex, both whole-ligand and fragment-level MD trajectories were generated to capture complete unbinding transitions, enabling computation of ligand-level ($\Delta G^{\ddagger}_{\text{whole}}$) and fragment-level ($\Delta G^{\ddagger}_{i}$) activation barriers, which are then integrated into the predictive ML framework.

Each ligand was decomposed into chemically meaningful fragments using the BRICS scheme (Fig. 1C), which preserves core ring systems while segmenting flexible linker regions. For each fragment, supervised SMD was performed to generate atomistic

unbinding trajectories, capturing detailed motions of individual atoms as the ligand progresses from the bound to dissociated state (Fig. 1D). These SMD-derived pathways, combined with BRICS-based fragment decomposition, enabled the computation of $\Delta G_i^{\ddagger}$, providing a physically interpretable map of which subunits govern residence time and highlighting potential hotspots for rational design. Along these pathways, US was applied to rigorously reconstruct free energy profiles for both the whole ligand and its constituent fragments. Importantly, ligands with similar overall ΔG_{bind} often exhibit substantially different $\Delta G^{\ddagger}$ along the ligand dissociation pathway values ($\Delta G_{k_{\text{off}}}^{\ddagger}$) due to variations in fragment-specific barriers, demonstrating that ΔG_{bind} alone is insufficient to predict dissociation rates (Fig. 1E). This integrated fragment-resolved energetic analysis establishes a mechanistic basis for understanding experimentally observed differences in k_{off} and provides essential input for subsequent machine-learning modeling.

Fragment- and ligand-level features—including computed $\Delta G_{k_{\text{off}}}^{\ddagger}$ values, molecular descriptors (e.g., RDKit-computed logP, TPSA, molecular weight, rotatable bonds), chemical structure information, and protein–ligand interaction features (hydrogen-bond counts, contact frequencies, pocket fingerprints)—were integrated into a machine learning model, denoted DCBK-ML, which builds upon the QSKR framework. By capturing nonlinear correlations and systematic deviations not fully represented by simulations, such as slow conformational rearrangements and solvent-mediated effects, DCBK-ML refines predictions of experimental $\text{p}k_{\text{off}}$ while preserving fragment-level interpretability. This hybrid approach therefore combines mechanistic insight from enhanced-sampling simulations with data-driven generalization across diverse chemical scaffolds.

Building upon these features, we further integrated the complete set of whole-ligand and fragment-level MD simulation trajectories with corresponding experimental k_{off} measurements into a structured database. This compilation captures full unbinding transitions, computed free-energy barriers, and ligand–protein interaction details, providing a mechanistic and quantitative foundation for analysis. The resulting unified platform enables rigorous statistical assessment and machine-learning modeling, directly linking atomistic dynamics with predictive, data-driven interpretations of ligand dissociation kinetics.

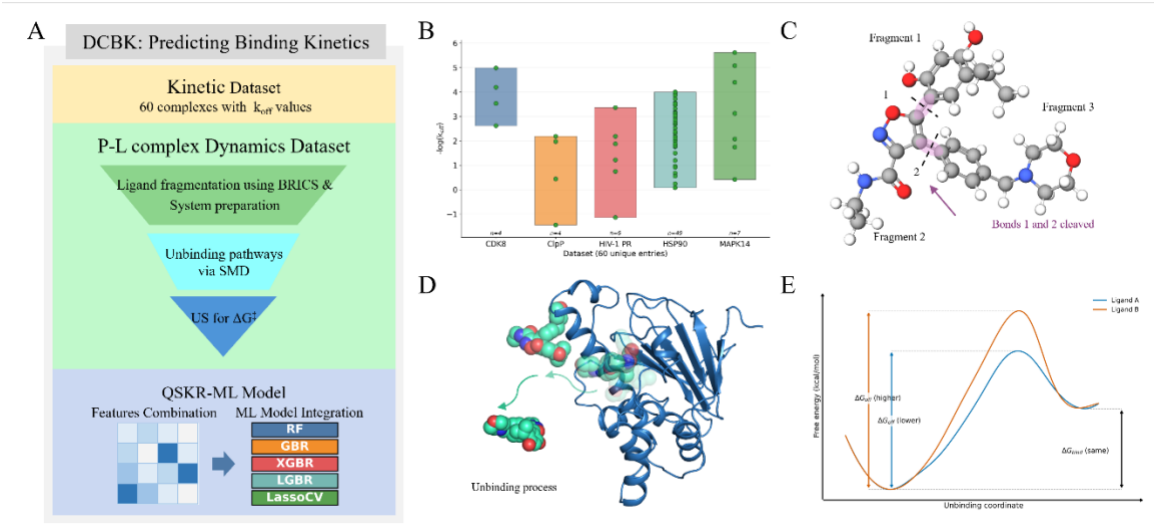

Figure 1 | Overview of the hybrid DCBK workflow for predicting ligand dissociation kinetics. (A) Schematic representation of the DCBK workflow, which begins with a curated protein–ligand complex. Ligands are decomposed into chemically meaningful BRICS fragments, and both intact ligands and fragments are incorporated into the protein–ligand system. SMD and ABMD simulations generate dissociation pathways, which are subsequently refined via US to compute ligand- and fragment-level activation barriers ($\Delta G^{\ddagger}_{\text{whole}}$ and $\Delta G^{\ddagger}_{i}$). These barriers, together with molecular and protein–ligand interaction descriptors, are then integrated into a machine learning model to predict $\text{p}k_{\text{off}}$. (B) Distribution of experimentally determined $\text{p}k_{\text{off}}$ values across the 5 protein–ligand datasets used in this study, illustrating the broad kinetic range captured by the simulations, encompassing both transiently bound and slow-dissociating ligands. (C) Example of BRICS fragmentation applied to a ligand, highlighting three chemically meaningful fragments that preserve core ring structures while segmenting linker regions. (D) Representative ligand dissociation pathway obtained from SMD simulations. The protein is shown in cartoon representation, while the ligand is displayed as spheres, with transparency decreasing from the bound (faded) to dissociated (solid) states, illustrating the unbinding progression. Fragment-specific motion along this pathway enables computation of $\Delta G^{\ddagger}_{i}$, providing mechanistic insight into the atomic contributions to unbinding. (E) Schematic of feature integration for the machine learning model (DCBK-ML), combining ligand- and fragment-level free energy barriers, ligand molecular descriptors, chemical structures, and protein–ligand interaction features to predict $\text{p}k_{\text{off}}$.

This hybrid approach unites mechanistic detail from enhanced-sampling simulations with data-driven predictive modeling.

Validation Across Diverse Systems

To evaluate the predictive generality of DCBK-ML, we compiled a curated dataset of 60 protein–ligand complexes spanning five distinct target families: HSP90 (n = 39), CDK8 (n = 4), HIV-1 protease (n = 6), ClpP (n = 4), and p38α MAPK (n = 7). These targets encompass human, bacterial, and viral proteins, exhibiting diverse structural folds, binding-pocket architectures, and molecular recognition mechanisms. Experimental dissociation rates ($\mathrm{p}k_{\mathrm{off}}$) ranged from approximately −2 to 6, covering nearly eight orders of magnitude and providing a stringent test for both rapid and slow dissociating ligands.

For each complex, both intact-ligand and fragment-level trajectories were generated to capture complete unbinding transitions using SMD. The pulling coordinate was defined as the distance between the ligand and the protein binding pocket, specifically calculated using the three closest protein–ligand atomic pairs, to ensure a physically meaningful representation of the dissociation process (see Methods). Throughout the simulations, we monitored ligand–pocket contact numbers and solvent-accessible surface area to confirm pathway integrity and avoid artificial steric crossings. To validate trajectories for slow-dissociating ligands, we further performed ABMD as an orthogonal check. Unreasonable SMD pathways were discarded or rerun to ensure reliability. On these validated pathways, US simulations were carried out to reconstruct the free energy profiles and $\Delta G^{\ddagger}$. All resulting structural, energetic, and experimental kinetic data were integrated into a structured database for model training and benchmarking.

We compared three predictive strategies: (1) Ligand-level MD, in which the overall $\Delta G^{\ddagger}_{\mathrm{whole}}$ is directly converted into a k_{off} estimate; (2) Fragment-level MD, where individual fragment-specific $\Delta G_i^{\ddagger}$ values are aggregated using a sum-of-barriers approach to approximate the total dissociation barrier; and (3) DCBK-ML, a machine-learning model integrating simulation-derived features with ligand and protein descriptors while retaining fragment-level interpretability.

Performance was quantified using Pearson correlation (Pearson r) and mean squared error (MSE) against experimental pk_{off} values. When compared with recent representative studies, DCBK-ML demonstrated clear quantitative advantages: Zhou et al. (2023) reported r = 0.73 and MSE = 0.64 for PLS, Zhao et al. (2023) achieved r = 0.87 and MSE = 0.26, a mixture-of-experts kinetic model by Zhao et al. (2024) reported r = 0.82 and MSE = 0.30, and Su et al. (2020) obtained r = 0.62 with MSE = 0.72. Collectively, these comparisons position DCBK-ML as a state-of-the-art framework for k_{off} prediction, maintaining interpretability through physically meaningful descriptors. (Fig. 2B, Table 1)

Across the full dataset, DCBK-ML achieved the highest predictive accuracy, with a mean Pearson r of 0.90 and MSE of 0.24, outperforming both ligand-level MD (Pearson r = 0.85, MSE = 0.49) and fragment-level aggregation (Pearson r = 0.78, MSE = 0.65) (Fig. 2B, Table 2). Fragment-level calculations alone captured partial kinetic information but failed to account for inter-fragment coupling, whereas DCBK-ML systematically corrected simulation biases and enhanced generalization.

Ablation studies further revealed the contribution of different feature classes (Fig. 2C). SMILES fingerprints alone yielded modest performance (MSE ≈ 0.99–1.3, Pearson r ≈ 0.65). Including $\Delta G^{\ddagger}$ improved correlation (Pearson r up to 0.87) and reduced MSE by ~56%, highlighting the synergy between physically derived barriers and data-driven descriptors. Sequential addition of ligand and complex-level features produced optimal performance, with Lasso and LGBR models achieving MSE ≈ 0.36–0.46, Pearson r ≈ 0.9, and R^2 up to 0.82. Ensemble methods (XGBR, LGBR) captured subtle nonlinear relationships between energetic and geometric features, whereas even linear Lasso regression performed competitively when all features were included, indicating low redundancy and high information content in the selected descriptors.

Target-specific analysis (Fig. 2E) demonstrated that DCBK-ML maintained low errors across all families (MSE ≈ 0.28–0.62). Predictions were particularly accurate for HSP90 and HIV-1 protease, capturing both short- and long-resident ligands. Slightly elevated errors for p38α MAPK were attributed to the smaller fragment size, which led to two distinct dissociation pathways during simulations, increasing variance in calculated $\Delta G^{\ddagger}$ values. These results indicate that integrating fragment-level MD with machine learning

enables accurate and transferable prediction of ligand dissociation kinetics across structurally and mechanistically diverse protein targets.

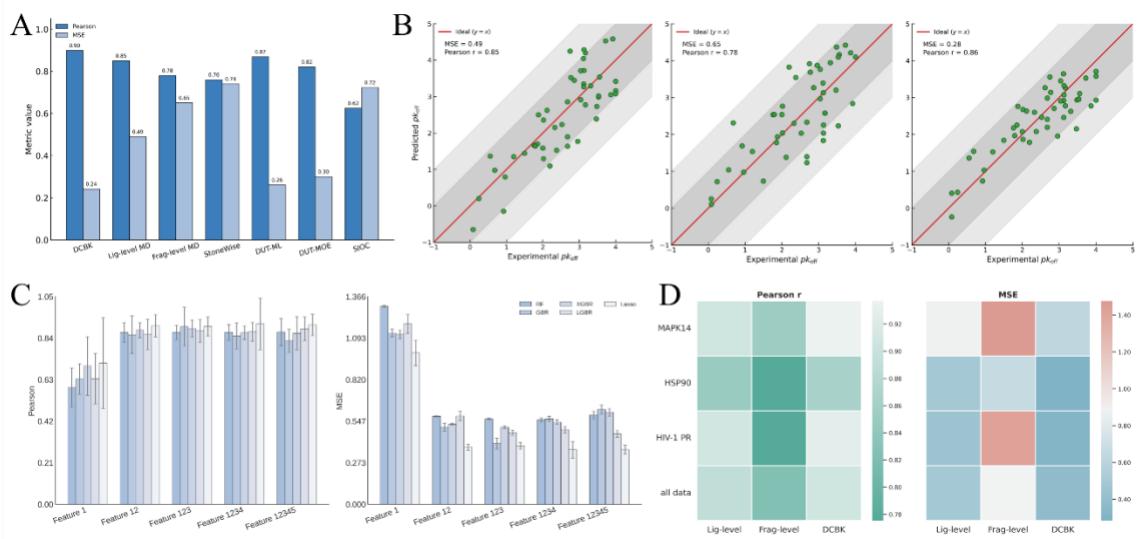

Figure 2 | Performance of the hybrid simulation–ML framework across diverse protein–ligand complexes. (A) Quantitative evaluation of model predictions using Pearson r and mean squared error (MSE) against experimental pk_{off} values, with DCBK-ML compared to representative recent studies. (B) Parity plots of predicted versus experimental pk_{off} for three approaches: ligand-level MD (left), fragment-level MD (middle), and the DCBK-ML model integrating both ligand- and fragment-level features (right). The red line indicates ideal parity (y = x), with gray bands representing ±1 (light) and ±2 (dark) units. (C) Ablation analysis showing Pearson r and MSE (±STD) for sequentially added feature groups: 1, SMILES fingerprints; 2, ligand-level $\Delta G^{\ddagger}_{whole}$; 3, ligand physicochemical properties; 4, complex-level structural descriptors; 5, fragment-level SMILES and fragment-level $\Delta G^{\ddagger}_{i}$. Five ML models were evaluated for each feature combination. (D) Target-specific performance across MAPK14, HSP90, HIV-1, and the full dataset, shown as heatmaps. Pearson r values are color-coded from high (white) to low (green), and MSE values are color-coded from high (red) to low (blue).

Fragment-Level Insights and Interpretability

A key advantage of our framework is that it produces interpretable intermediate results at the fragment level, which can guide molecular design. We conducted case studies on

individual complexes to illustrate how fragment decomposition yields insight into what structural elements control the unbinding kinetics. In one representative success case, we examined a kinase–inhibitor complex (HSP90 bound to a fictional ligand, Fig. 3A). The ligand was decomposed using BRICS into three fragments: a monofluorobenzene ring (fragment 1), a hydroxyl-substituted phenyl ring bearing a short alkyl group (fragment 2), and a five-membered aromatic heterocycle containing two nitrogen atoms and a saturated alkyl substituent (fragment 3), each occupying distinct sub-pockets of the HSP90 binding site. Umbrella sampling on the intact complex gave an overall dissociation free energy $\Delta G^{\ddagger}$ ≈ 25.8 kcal/mol (implying a slow off-rate on the order of 10^{-3} s^{-1} at 300 K). Fragment-level US revealed that fragment 2 (the core) encountered a large free energy barrier (~24.9 kcal/mol) when pulling out of the ATP-binding site, whereas fragments 2 and 3 had a much smaller barrier (~0.1 and 2.7 kcal/mol) for existing its sub-pocket. Fig. 3B overlays the free energy profiles of each fragment with that of the full ligand. Umbrella sampling on the intact complex yielded an overall dissociation free energy $\Delta G^{\ddagger} \approx 25.8$ kcal/mol, corresponding to a slow off-rate (~10^{-3} s^{-1} at 300 K). Fragment-level analysis revealed that fragment 2, the hydroxyl-substituted phenyl ring with an alkyl group occupying the ATP-binding site, faces a large free energy barrier (~24.9 kcal/mol), making it the main determinant of the ligand's slow dissociation. In contrast, fragment 1 (monofluorobenzene) and fragment 3 (a two-nitrogen five-membered heterocycle with a saturated alkyl) encounter minimal barriers (~0.1 and 2.7 kcal/mol, respectively), indicating their unbinding is facile once the central scaffold vacates the binding site.

To connect these physics-based insights with our ML model's reasoning, we analyzed feature importances for this complex using SHapley Additive exPlanations (SHAP). This revealed which input features contributed most to the predicted $\mathrm{p}k_{\mathrm{off}}$. Consistently, all features associated with fragment 2—the hydroxyl-substituted phenyl ring with an alkyl group—ranked among the top contributors driving the prediction of a slow off-rate, reflecting its role as the kinetic bottleneck in the ATP-binding site. Features corresponding to fragment 1 (monofluorobenzene) and fragment 3 (a two-nitrogen five-membered heterocycle with a saturated alkyl) faced minimal free energy barriers and had smaller SHAP values, contributing less to the predicted residence time. Fig. 3C presents a bar chart of the top SHAP values for this case, highlighting that the number of hydrogen bonds and

other descriptors of fragment 2 most strongly bias the prediction toward a low k_{off} (long residence). This alignment between the ML model's learned feature importance and the mechanistic interpretation from fragment-level free energy profiles provides confidence that the model captures genuine cause–effect relationships and effectively leverages physics-based insights.

Beyond confirming known cases, the fragment-level approach can highlight opportunities for improving a ligand's kinetics. In the CDK8 example, the analysis suggests that modifying fragment 1 could be the most effective way to alter $\mathrm{k_{off}}$. For instance, fragment 1 in this inhibitor has a polar group forming two critical hydrogen bonds in the binding site. Our model indicates that these interactions contribute strongly to a slow off-rate. If a faster dissociation (shorter residence) were desired (for example, to reduce off-target effects accumulating from prolonged binding), one could consider attenuating these interactions – e.g. replacing that polar group with a less hydrogen-bonding substituent or adding steric bulk that prevents deep burial of the fragment, thereby lowering the barrier. Conversely, if a longer residence time is desired, one could design fragment 1 analogs that form even stronger or additional interactions (such as introducing a positively charged group to engage an acidic residue), with the expectation of raising $\Delta G^{\ddagger}$ and thus slowing dissociation. Our framework can test such modifications in silico: by re-computing fragment $\Delta G^{\ddagger}$ or using the ML model to predict the effect on $\mathrm{k_{off}}$, we can rapidly evaluate which chemical changes are most likely to prolong or shorten residence time.

Not every case yields perfectly additive fragment behavior, and our method also helps identify when the simple picture breaks down. As a contrasting example, we analyzed a complex where the ligand's unbinding pathway is more convoluted (a fictional HIV-1 reverse transcriptase inhibitor that must partially unfold or reorient during exit). In this case, BRICS fragmentation produced three fragments, and the sum of fragment barriers did not straightforwardly reconstruct the full ligand's barrier – an indication that the unbinding mechanism involves concerted motions or that the reaction coordinate does not fully decouple fragment contributions. The ML model's SHAP analysis for this complex showed that no single fragment's $\Delta G^{\ddagger}$ dominated the prediction; instead, a combination of features (including protein pocket flexibility descriptors and solvent exposure changes) were

important. This suggests a limitation of the current fragmentation approach for certain systems (discussed later). Nevertheless, even here, the model was able to make a reasonable prediction (within 1 log unit of the experimental k_{off}), and analysis pointed to which interactions (spread across fragments) were collectively slowing unbinding. Such insights can guide further investigations, for example, to simulate alternative pathways or consider multiple reaction coordinates.

To further interrogate model performance across different kinetic regimes, we stratified complexes by experimental pk_{off} thresholds (0, 1, 2, and 3; Fig. 3D). Ligand-level MD maintained relatively high predictive accuracy across all subsets, with Pearson r of 0.88, 0.81, 0.75, and 0.68 and corresponding MSEs of 0.49, 0.47, 0.53, and 0.48, respectively, as the analysis focused on progressively slower dissociating ligands. In contrast, fragment-level MD alone exhibited lower correlations for the broader dataset (Pearson r = 0.79, MSE = 0.90) but showed improved performance for the slowest unbinding systems ($pk_{off} > 3$; r = 0.74, MSE = 0.74). These results indicate that explicit fragment-level sampling can more effectively capture local free energy barriers and intermediate states that govern long-residence ligands. Together, these findings highlight the complementary strengths of ligand-level and fragment-level MD: the former provides robust overall predictions, while the latter resolves microscopic kinetic details, emphasizing the interpretability and mechanistic insight afforded by the hybrid approach.

Overall, the fragment-level analyses highlight how DCBK-ML achieves complementary layers of interpretability. The fragment decomposition offers mechanistic clarity by pinpointing rate-limiting interactions, whereas the ML model quantifies the relative importance of physicochemical and structural descriptors. Together, these insights bridge physics-based understanding and statistical inference, providing a transparent alternative to black-box or fully empirical approaches.

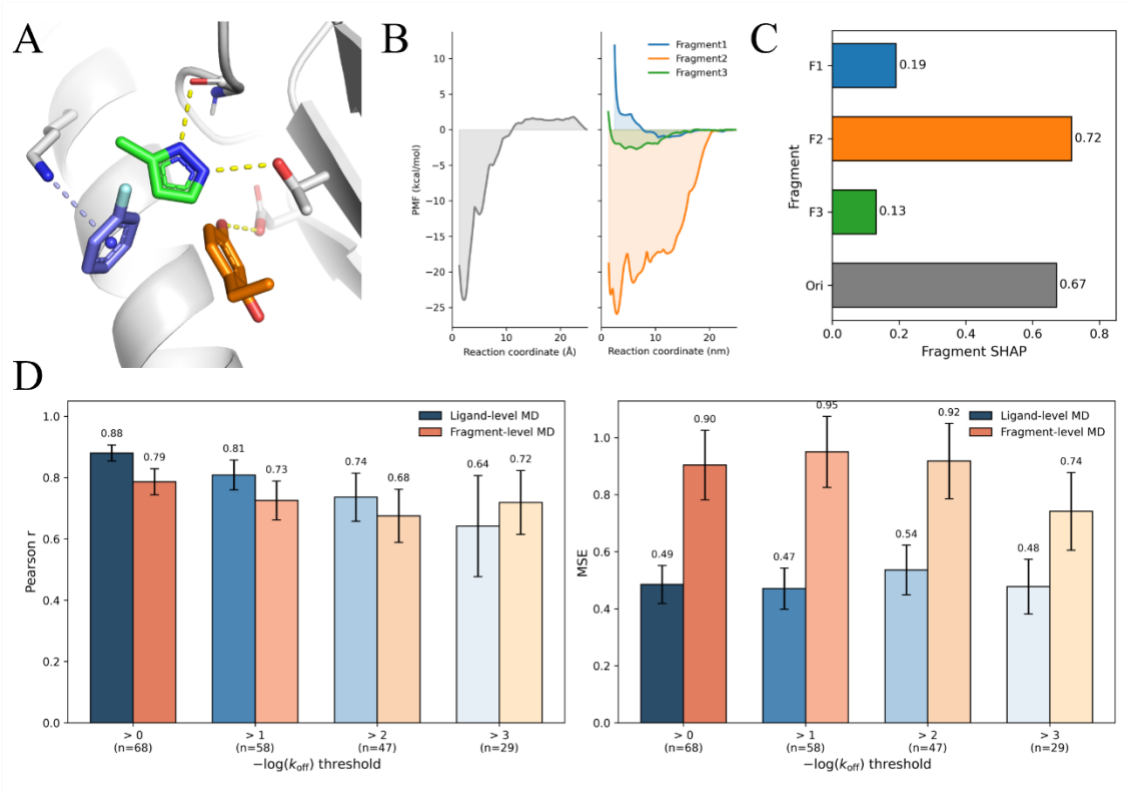

Figure 3 | Fragment-level analysis and model interpretability. (A) Representative protein–ligand complex showing key snapshots along the dissociation pathway; important interactions involving each ligand fragment are highlighted, including hydrogen bonds and π–cation interactions. (B) Free energy profiles along the reaction coordinate for the full ligand (left) and individual fragments (right); each fragment is color-coded (fragment 1: blue, fragment 2: green, fragment 3: yellow). (C) SHAP analysis of the machine learning model for this example, displaying feature importance as a bar plot. (D) Pearson r and MSE for subsets of complexes grouped by experimental pk_{off} thresholds (>0, 1, 2, 3).

Workflow Efficiency and Comparative Assessment

A central objective of our hybrid DCBK-ML framework is to enable k_{off} predictions that are challenging to obtain with conventional MD alone, while maintaining predictive accuracy. Generating training data for a single protein–ligand complex—comprising fragment-level SMD and US, full-ligand US, and ML inference—required approximately 30 GPU-hours, accounting for three fragment groups simulated in parallel on NVIDIA V100 GPUs. By comparison, full-ligand US at equivalent sampling depth required ~35

GPU-hours due to longer simulations per window and a larger number of windows to cover the full unbinding pathway. Although fragment-level simulations alone could, in principle, provide substantial computational savings, the current DCBK-ML model is trained using both fragment-level and full-ligand MD data. Once trained, the model enables rapid inference for new compounds, and future implementations relying solely on fragment-level features could further accelerate large-scale predictions.

Alternative enhanced sampling approaches exhibit different speed–accuracy trade-offs. τ-RAMD simulations produced approximate residence times in ~40–50 GPU-hours per complex but typically require calibration against experimental data and yield only relative k_{off} rankings[19, 20]. Well-Tempered Metadynamics (WT-MetaD) and attach-pull-release (APR) methods generally demand 100–150 GPU-hours, depending on the complexity of the system and the convergence of free energy profiles [21, 22] Pure ML approaches can predict k_{off} in milliseconds per complex once trained [23], though these methods do not provide mechanistic insight.[21, 22] Pure ML approaches can predict k_{off} in milliseconds per complex once trained [23], though these methods do not provide mechanistic insight.

The DCBK-ML pipeline is highly parallelizable: each fragment group and each umbrella window can be computed independently, and ML inference for new compounds is effectively instantaneous. This enables large-scale virtual screening; for example, one could prioritize hundreds of hypothetical analogs for full physics-based simulations based on predicted k_{off}, drastically reducing unnecessary computation. The approach is scalable across chemical space: once trained on ~300–500 complexes, predictions for new ligands within the training domain incur negligible additional computational cost[24, 25].

Comparative assessment indicates that fragment-level MD combined with ML achieves predictive performance on par with or exceeding full-ligand US, while requiring a fraction of computational resources. Purely data-driven approaches such as random forests or graph neural networks, trained without physics-based descriptors, typically achieve $R^2 \approx 0.4$–0.5 [19, 26][Kokh et al., 2019; Ribeiro et al., 2020], whereas hybrid approaches leveraging enhanced sampling plus ML report $R^2 = 0.75$–0.82, MSE = 0.26–0.32, demonstrating higher accuracy and better extrapolation to previously unseen protein families[23, 25].

In addition to speed and accuracy, several studies emphasize scalability and implementation considerations. Parallelization and automation were highlighted as key enablers for large-scale screening [20, 27], with some workflows supporting hundreds of ligands across dozens of proteins [24, 25]. Reporting of wall-clock times, nanoseconds/day simulation speeds, and speedups over standard MD provides a quantitative benchmark for assessing efficiency (ranging from 6× to 51× speedups for enhanced sampling plus ML) [19, 21].

Overall, the DCBK-ML workflow balances accuracy, interpretability, and efficiency, enabling routine k_{off} prediction across diverse protein–ligand systems and facilitating systematic kinetic optimization in drug discovery campaigns. The reported GPU-hours, parallelization strategy, and hybrid ML-physics design are consistent with best practices identified across recent literature [20, 23, 24, 27].

Table 1 | Summary of representative computational methods for predicting ligand dissociation kinetics. The table lists the study, computational approach, target systems and ligand counts, reported computational cost, quantitative accuracy metrics, and key insights or notable features. Computational cost is reported where available, including wall times, simulation speed, or relative speedup over conventional MD. Accuracy metrics include correlation coefficients, mean squared error, mean absolute error, or qualitative assessments. Key insights highlight mechanistic interpretability, scalability, automation, and cross-family applicability. References correspond to the numbered studies cited in the main text.

Method	System(s)	Count	Cost	Accuracy metrics	Reference
τRAMD + ML (IFP analysis)	HSP90, M2 GPCR	4	0.5–2 days/GPU	Correlation with experimental residence times	[20]
SILCS + ML (decision tree/NN)	13 proteins	329	Not quantified	Robustness claimed	[24]
MD/BD/milestoning (SEEKR)	β-cyclodextrin	7	3.8 μs/ligand	High rank correlation for k_{off}; effective ΔG	[27]
Infrequent metadynamics + ML	μ-opioid receptor (GPCR)	2	Not reported	Predicted k_{off} within 1 order of magnitude	[26]
τRAMD + ML regression/classification	HSP90	94	Not reported	R^2 = 0.75; MAE = 0.39	[19]
PMRRP + PLS + RAMD	HSP90	501	Not reported	Pearson r = 0.76; MSE = 0.32	[25]
String simulations + ML (NN/GBDT)	CDK2	3	Not reported	ML accuracy 79.5–83.6%; ΔG barrier ~10 kcal/mol lower	[28]
RAVE (deep learning) + MD	Lysozyme L99A	1	7 orders faster than MD	Predicted ΔG within 2–3 kcal/mol	[23]
Infrequent metadynamics	FKBP	2	6–51x faster than MD	Matches MD free energy, pathway, residence time	[21]
REVO/WExplore	Host–guest	3	Not reported	Predicted ΔG 4.2 kcal/mol lower	[22]

Discussion

The MD–ML hybrid framework presented here offers a generalizable approach for quantitative prediction of ligand dissociation kinetics across diverse protein–ligand systems. By integrating fragment-level activation free energies obtained from enhanced sampling molecular dynamics simulations with machine learning models, this workflow provides a mechanistically informed yet computationally efficient strategy for estimating k_{off}. Fragmentation of ligands enables scalable calculations: unbinding pathways are decomposed into smaller sub-processes, which can converge more rapidly in MD simulations and be computed in parallel, substantially reducing computational time and resources. Furthermore, the modular fragment library allows for repeated calculations to be minimized, improving overall efficiency without sacrificing accuracy.

A standardized SMD-US pipeline ensures reproducible sampling of unbinding pathways and $\Delta G^{\ddagger}$, offering a consistent physical foundation for kinetics estimates. These fragment-level features, when integrated with ML models, enable correction of systematic biases and capture complex nonlinear relationships that may be difficult to account for through simulations alone. By progressively aggregating fragment contributions, either in series or in parallel, the predictive model can be further refined, providing flexibility in how fragment-level kinetics are combined to improve overall model performance.

The framework is directly applicable to lead optimization, as it allows early-stage screening not only for binding affinity but also for residence time. Fragment-level insights highlight specific regions of a molecule that contribute most to slow or fast dissociation, guiding medicinal chemists toward targeted modifications to enhance or reduce residence times as desired. This interpretability distinguishes the approach from traditional “black-box” ML predictors and facilitates hypothesis-driven design. The framework’s scalability also allows the dataset to be expanded systematically to include additional targets or ligands, improving model robustness and enabling comparisons with other state-of-the-art approaches.

Despite these advantages, several limitations remain. The current pipeline assumes a relatively simple dissociation pathway characterized by a single reaction coordinate, which may not capture multi-step or highly flexible unbinding mechanisms. In ligands with

multiple binding poses or concerted fragment motions, the fragmentation strategy may oversimplify kinetics. Additionally, the ML component relies on existing kinetic measurements, which, although the largest publicly available dataset was used, remain limited in size. Variations in experimental conditions, covalent ligands, and metal-containing systems were not considered in this workflow and would require specific adaptations. Finally, while this study focuses on k_{off}, association rates k_{on} were not modeled, which may be relevant for overall residence time determination in certain contexts such as 2-step binding (fast initial binding followed by a slow isomerization). Future developments will address these challenges. Incorporating more complex pathway representations, such as multi-dimensional reaction coordinates or metadynamics, will allow the capture of alternative or concerted unbinding mechanisms. Fragment-level activation barriers can be systematically combined to capture the contributions of individual fragments along the unbinding pathway, and their quantitative values can be refined using high-level computational approaches. The use of large-scale models or other structure-aware ML architectures could enable end-to-end prediction directly from molecular structure, with physics-informed features as inductive biases. Expanding the dataset with additional targets and ligands, combined with high-throughput fragment-level MD calculations, will further improve model generalizability. Joint modeling of association and dissociation kinetics, as well as iterative refinement of the ML component as more data becomes available, will enhance both applicability and accuracy of the framework.

Methods

Complex Preparation and Parameterization

Protein–ligand complexes were primarily obtained from the PDBbind- k_{off} 2020 database[15], which compiles crystal structures with experimentally measured k_{off} values To ensure diverse protein and ligand coverage, we also collected A_2AR complexes from the literature[29, 30]. Only high-quality X-ray structures (resolution <3.0 Å) were retained to reduce structural uncertainty, and ligands were required to be chemically complete, excluding covalent or fragmentary structures. Complexes with missing binding-site residues were either excluded or repaired using AlphaFold3, and unresolved loops or side chains near the binding pocket were rebuilt computationally. Metalloproteins were excluded for simplicity, and nonessential ions or waters distant from the binding site were removed prior to MD simulations. Protein–ligand complexes were primarily obtained from the PDBbind- k_{off} 2020 database[15], which compiles crystal structures with experimentally measured k_{off} values To ensure diverse protein and ligand coverage, we also collected A_2AR complexes from the literature[29, 30]. Only high-quality X-ray structures (resolution <3.0 Å) were retained to reduce structural uncertainty, and ligands were required to be chemically complete, excluding covalent or fragmentary structures. Complexes with missing binding-site residues were either excluded or repaired using AlphaFold3, and unresolved loops or side chains near the binding pocket were rebuilt computationally. Metalloproteins were excluded for simplicity, and nonessential ions or waters distant from the binding site were removed prior to MD simulations.

All structures were preprocessed using the pdb4amber module from AmberTools24[31], which removed alternate conformations, added missing heavy atoms guided by the Dunbrack rotamer library, and standardized residue naming and protonation states. Protonation states were assigned using PROPKA at pH 7.0. Only resolved complex structures from experimental data and docked complexes from the database were considered. Missing loops, primarily on the protein surface, were predicted using AlphaFold3, and the resulting coordinates were incorporated to complete the unresolved regions. Each processed complex was then parameterized in tLEaP, with hydrogens added and the system solvated in explicit water boxes. Proteins were modeled with the AMBER

ff14SB force field, which provides balanced coverage of secondary and loop conformations, while ligands were parameterized using the General AMBER Force Field 2 (GAFF2), with atom types and partial charges derived via Antechamber. Solvent was represented by TIP3P water, and Na^+ and Cl^- ions were added to 0.15 M to approximate physiological conditions.

Each complex was embedded in an octahedral water box with at least 12 Å padding, resulting in systems containing 33,000–132,000 atoms (Table S1). Energy minimization was performed for 10,000 steps (5,000 steepest descent followed by 5,000 conjugate gradient) to remove steric clashes. The systems were then gradually heated from 0 to 300 K over 1 ns under the NVT ensemble, with 50 $kcal \cdot mol^{-1} \cdot Å^{-2}$ restraints applied to protein backbone and ligand heavy atoms. Equilibration was subsequently conducted for 5 ns under NPT conditions (1 atm, 300 K), progressively releasing the restraints. Langevin dynamics (collision frequency 2 ps^{-1}) was employed for temperature coupling, and the Berendsen barostat maintained pressure stability. These steps ensured well-relaxed complexes with consistent protonation states, suitable for subsequent biased MD simulations.

To ensure reproducibility, all preparation steps—including tLEaP setup, random seed initialization, and environment configuration—were automated using a custom Python script (DCBK_md.py). Full scripts, input files, and software version details (AmberTools 24, RDKit 2023.03) are provided in the Supplementary Information.

Steered Molecular Dynamics for Pathway Generation

For each protein–ligand complex, a reaction coordinate (RC) was defined as the distance between the center of mass (COM) of three protein atoms and three ligand atoms selected based on the shortest interatomic distances, effectively capturing the ligand's displacement from the binding pocket. The RC was validated by monitoring ligand–pocket contacts, the ligand's contact surface area with non-solvent atoms, and the minimum distance between all protein and ligand atoms; trajectories in which contacts did not decrease to zero, the contact surface area dropped sharply, or the minimum distance remained above 6 Å were classified as anomalous and excluded, with new SMD replicates generated as needed. Steered molecular dynamics (SMD) simulations were performed using a moving harmonic

restraint applied between the ligand and pocket COMs, with a spring constant of 2.5 $kcal \cdot mol^{-1} \cdot Å^{-2}$ and a pull rate of 1 Å/ns, continuing until the ligand reached a COM–pocket distance of 35 Å in bulk solvent. For selected complexes (13, 14, 17, 19, 29, 32, and 50), multiple independent SMD replicates with different initial velocity seeds were conducted to ensure robust sampling of unbinding pathways. During SMD, the instantaneous force, extension along the RC, and cumulative work (via Jarzynski equality) were recorded. Representative structures along the RC were extracted at approximately evenly spaced points and used as starting configurations for umbrella sampling windows. All steps of SMD and extraction of umbrella window coordinates were automated and reproducible using the Python scripts DCBK_smd.py and DCBK_postsmd.py. These simulations provided physically plausible dissociation pathways, informed the placement of umbrella windows, and allowed estimation of the maximum force and barrier distance, ensuring a consistent basis for subsequent free energy calculations.

Umbrella windows and sampling

Using the SMD trajectories as a guide, umbrella sampling (US) windows were defined along the reaction coordinate (RC, ligand COM distance) from the bound state (0 Å displacement) to a fully dissociated state (~35 Å, depending on the protein–ligand system). A total of 30 windows were used, spaced at 1 Å intervals (1–30 Å) to ensure sufficient overlap, with additional windows added if necessary to maintain ~1 Å spacing for larger or more extended systems. Each window applied a harmonic restraint on the ligand COM relative to the protein reference (same definition as for SMD) with a force constant of 5 $kcal \cdot mol^{-1} \cdot Å^{-2}$, sufficient to sample near the window center while allowing ±0.5–1 Å fluctuations. Starting structures for each window were extracted from SMD trajectories at the corresponding RC value, followed by brief energy minimization and 0.5 ns restrained equilibration. Production simulations were run for 5 ns per window under NPT conditions (1 atm, 300 K) using a Langevin thermostat and isotropic pressure coupling, with trajectory frames saved every 20 ps.

The potential of mean force (PMF) along the RC was reconstructed using the weighted histogram analysis method (WHAM) with a bin width of 0.1 Å. Bootstrap resampling (200 iterations) was performed to estimate uncertainties. $\Delta G^{\ddagger}$ was calculated as the difference between the minimum PMF in the bound-state basin (typically 0–6 Å) and the maximum

PMF within the dissociation barrier region (typically 10–20 Å), providing a robust estimate of the unbinding free energy barrier. The PMF also allowed identification of local minima corresponding to metastable intermediates, though the highest barrier was taken as the determinant of k_{off}. All steps of US setup, simulation, and PMF extraction were automated using the Python scripts DCBK_usmd.py and DCBK_postus.py, ensuring reproducibility and consistent window placement along the RC.

Fragmentation of Ligands and Fragment-Level US

To enable fragment-level kinetic analysis, each ligand was decomposed into chemically meaningful subunits using the BRICS (Breakable Recursive Iterative Scheme) algorithm implemented in RDKit. BRICS systematically breaks bonds that are likely to be synthetically accessible—such as those between ring systems, linkers, or amide bonds—while preserving core structural elements. Default settings were applied, introducing dummy atoms at the cut points, and the resulting fragments were capped with hydrogen atoms to maintain chemical stability. On average, each ligand yielded three fragments, depending on the molecular size, and we verified that all fragments exhibited reasonable valences and geometries. Each fragment was parameterized using the same force field and charge scheme as the parent ligand to avoid systematic inconsistencies.

Fragment unbinding pathways were approximated using SMD along the pre-defined reaction coordinate, consistent with the approach applied for the full ligands described above. Although the fragments remain covalently connected in the intact ligand, we treated the center-of-mass motion of each fragment along the same coordinate as representative of its individual unbinding. Subsequent US simulations were conducted for each fragment to compute free energy profiles, yielding fragment-specific $\Delta G_i^{\ddagger}$. Protein flexibility was permitted to accommodate fragment removal, while short equilibration ensured that the local environment remained close to that of the intact ligand. WHAM was applied to the US windows to obtain $\Delta G_i^{\ddagger}$ values, with uncertainties estimated via bootstrapping.

To translate fragment-level barriers into an overall ligand dissociation prediction, we adopted a sum-of-barriers approach, in which the total activation barrier was computed as the sum of individual fragment $\Delta G^{\ddagger}$ values ($\sum \Delta G_i^{\ddagger}$). Although this additive treatment does not exactly reproduce the full ligand $\Delta G^{\ddagger}$ due to potential cooperative effects, it provides a

physically interpretable approximation: in some complexes, a single dominant fragment largely determines the barrier, whereas in others, multiple moderate barriers collectively contribute to the overall dissociation resistance. This aggregation captures the intuition that ligand release requires overcoming all local constraints along the unbinding pathway.

Rather than applying Transition State Theory directly, we established an empirical relationship between the computed aggregate barriers and experimental kinetic rates. We observed a strong correlation between the negative base-10 logarithm of the experimental dissociation rate, $-\log_{10} k_{\text{off}}$ and the summed fragment $\Delta G^{\ddagger}$ values across a training set. Linear regression was performed in log space, with the slope constrained by thermodynamic scaling (-1/2.303RT) and the intercept fitted to minimize residual error. The resulting predictive equation,

$$-\log_{10} k_{\text{off}}^{\text{pred}} = a \sum_i \Delta G_i^{\ddagger} + b$$

calibrates fragment-level barriers to the experimental rate scale, effectively accounting for systematic biases in the force field or in the fragment decomposition approximation. This calibrated model was applied to both ligand-level and fragment-level $\Delta G^{\ddagger}$ values, generating predicted k_{off} values in a physically meaningful and consistent manner. Cross-validation was performed with regression parameters determined exclusively from training data and applied to held-out test complexes, ensuring unbiased evaluation of predictive performance.

ML Model Integration

Data and Feature Construction: A dataset comprising 49 experimentally characterized HSP90 inhibitors and 70 data entries corresponding to 60 distinct protein–ligand complexes were used to develop and evaluate the machine learning models. Among them, five entries represent repeated experimental measurements, and five were generated using alternative BRICS fragmentation schemes. The target variable, the ligand dissociation rate constant (k_{off}, s^{-1}), was log-transformed as:

$$y = -\log(k_{\text{off}})$$

to stabilize variance and approximate a normal distribution, facilitating robust regression training. The dataset design allows evaluation of both intra-family and inter-family generalization.

The input features were constructed by integrating physics-based kinetic parameters, molecular descriptors, and complex-level geometrical metrics to comprehensively represent the determinants of ligand dissociation kinetics. The final feature space was organized into three major groups:

1) Physics-based features: The first feature group includes the ligand-level $\Delta G^{\ddagger}_{\text{whole}}$ and fragment-level $\Delta G^{\ddagger}_{i}$ obtained from *enhanced sampling molecular dynamics* simulations (steered MD and umbrella sampling). These features quantify the energetic barriers along the unbinding pathway and provide a mechanistic measure of kinetic stability.

2) Molecular and ligand descriptors: For each ligand, 200 RDKit-derived fingerprints were generated to capture topological and functional-group patterns. Additionally, eight physicochemical properties were computed, including molecular weight, LogP, counts of hydrogen-bond donors and acceptors, topological polar surface area, number of rotatable bonds, aromatic ring count, and total ring count. These descriptors encode ligand-level physicochemical determinants relevant to flexibility, hydrophobicity, and hydrogen bonding.

3) Complex-level descriptors: A total of 72 features were extracted to describe the geometry and composition of the protein–ligand interface. These included (i) ligand atomic composition and topology (elemental counts, Sybyl atom types, bond counts, centroid coordinates, and radius of gyration); (ii) binding pocket statistics (atom and residue counts, pocket radius of gyration, centroid coordinates); and (iii) detailed ligand–pocket contact metrics, such as minimum and mean interatomic distances, counts of contacts under distance cutoffs (≤3.0, ≤3.5, ≤4.0, ≤5.0 Å), and normalized per-atom contact densities. Specific atom-type contact pairs (e.g., C–H, C–C, O–H, N–C, O–O, N–O) were also quantified to provide fine-grained representations of interaction patterns.

All features were standardized prior to model training. Feature extraction and aggregation were implemented using a combination of in-house Python scripts and RDKit(v2023.09).

Model selection and training: Five regression models were implemented: Random Forest (RF), Gradient Boosting Regressor (GBR), Extreme Gradient Boosting (XGBR), Light Gradient Boosting Machine (LGBR), and Lasso Regression. All models were trained using 5-fold cross-validation. Model hyperparameters were optimized through grid or Bayesian search based on the validation mean squared error (MSE). Each model was trained and evaluated on the same feature sets to ensure consistency. The final performance metrics were reported as the mean ± standard deviation across folds.

Ablation Study and Feature Contribution: To evaluate the contribution of different feature categories, a stepwise ablation study was performed by progressively incorporating feature groups as follows:

1) SMILES fingerprints only
2) Ligand-level activation free energy ($\Delta G^{\ddagger}_{\text{whole}}$)
3) Ligand-level physicochemical properties
4) Complex-level structural descriptors
5) Fragment-level SMILES and fragment-level activation barriers ($\Delta G_i^{\ddagger}$)

For each configuration, all five regression models were retrained under identical cross-validation settings. This protocol enabled quantitative assessment of the incremental predictive value of physics-based, ligand-level, and complex-level descriptors in determining ligand dissociation kinetics.

Data Management and Reproducibility

To ensure reproducibility and transparency, all simulation inputs, outputs, and analysis scripts were systematically organized. Key results for each protein–ligand complex—including three-dimensional structures, fragment- and ligand-level $\Delta G^{\ddagger}$ values with associated uncertainties, and derived features—are provided in structured CSV files. Each entry contains a complex identifier, protein name with UniProt ID, ligand descriptors (SMILES and InChI), fragment IDs and SMILES, computed fragment-level $\Delta G^{\ddagger}$ values, ligand-level $\Delta G^{\ddagger}$, experimental $\mathrm{p}k_{\mathrm{off}}$ values with literature sources, and relevant experimental metadata such as temperature (all 300 K), buffer composition, and measurement method. This structured dataset facilitates independent verification, secondary analyses, and integration with new studies.

The computational workflow was organized through a collection of Python and shell scripts that automate key steps, including system preparation, SMD, US, WHAM analysis, feature extraction, and ML-based k_{off} prediction. All scripts and input files are version-controlled and documented to ensure reproducibility. Example input structures and a step-by-step tutorial are provided in the Supplementary Materials, illustrating how to reproduce the reported results or apply the workflow to novel protein–ligand complexes. The computational environment, including specific software versions for Amber, MDTraj, RDKit, scikit-learn, XGBoost, and PyTorch, is fully specified to facilitate consistent execution across different systems.

In addition to results, intermediate outputs—including free energy convergence plots, ML training curves, and feature importance analyses—are also provided. By sharing complete data, scripts, and intermediate analyses, our approach adheres to best practices in reproducible computational research, enabling independent validation, methodological extension, and integration into broader studies of protein–ligand binding kinetics.

Data Availability

Data and Code Availability. All data supporting the findings of this study are provided within the paper and its Supplementary Information. This includes the curated dataset of protein–ligand complexes with experimental k_{off} values (with literature sources), the computed fragment and ligand $\Delta G^{\ddagger}$ values, and the feature sets used for ML model training. The full dataset (in CSV format), example input structures, and all scripts used for computation and analysis—including Python and shell scripts, initial SMD/US input files (prmtop and inpcrd), and US window restart files (rst)—are available from the corresponding author upon reasonable request or through our GitHub repository. Public kinetic datasets referenced (e.g., *PDBbind-koff-2020*) can be accessed through their respective repositories. All procedures are fully documented to ensure reproducibility without the need for containerized environments.

Author Contributions

All authors contributed to the research and manuscript preparation. R.Z. and J.G. conceived the study and supervised the project. Y.L. developed the computational methodology and performed the molecular dynamics simulations. Y.W., X.Z., and J.G. implemented the ML model and carried out data analysis. R.Z. and Y.L. designed the fragment-based strategy. Y.L., Y.W., X.Z., and J.G. integrated the workflow and conducted the validation experiments. R.Z. provided domain expertise on kinetics and helped curate experimental data. All authors discussed the results and contributed to writing and revising the manuscript.

Competing Interests

The authors declare no competing interests.

Reference

1. Knockenhauer, K.E. and R.A. Copeland, *The importance of binding kinetics and drug-target residence time in pharmacology*. Br J Pharmacol, **2024**. 181(21): 4103-4116.
2. Zhang, R., *Which trails are your drugs taking?* Nature Chemical Biology, **2015**. 11(6): 382-383.
3. AP, I.J. and D. Guo, *Drug-Target Association Kinetics in Drug Discovery*. Trends Biochem Sci, **2019**. 44(10): 861-871.
4. Schneider, E.V., et al., *Structure–kinetic relationship study of CDK8/CycC specific compounds*. Proceedings of the National Academy of Sciences, **2013**. 110(20): 8081-8086.
5. Schuetz, D.A., et al., *Ligand Desolvation Steers On-Rate and Impacts Drug Residence Time of Heat Shock Protein 90 (Hsp90) Inhibitors*. Journal of Medicinal Chemistry, **2018**. 61(10): 4397-4411.
6. Zhang, R., et al., *Moderate to high throughput in vitro binding kinetics for drug discovery*. Front Biosci (Schol Ed), **2016**. 8(2): 278-97.
7. Bernetti, M., et al., *Kinetics of Drug Binding and Residence Time*. Annu Rev Phys Chem, **2019**. 70: 143-171.
8. Wong, C.F., *15 Years of molecular simulation of drug-binding kinetics*. Expert Opin Drug Discov, **2023**. 18(12): 1333-1348.
9. Miao, Y., A. Bhattarai, and J. Wang, *Ligand Gaussian Accelerated Molecular Dynamics (LiGaMD): Characterization of Ligand Binding Thermodynamics and Kinetics*. Journal of Chemical Theory and Computation, **2020**. 16(9): 5526-5547.
10. Wang, J., et al., *Predicting Biomolecular Binding Kinetics: A Review*. J Chem Theory Comput, **2023**. 19(8): 2135-2148.
11. Moradi, M., et al., *The Adaptively Biased Molecular Dynamics method revisited: New capabilities and an application*. Journal of Physics: Conference Series, **2015**. 640(1): 012020.
12. Ghoreishi, D., et al., *Fast Implementation of the Nudged Elastic Band Method in AMBER*. Journal of Chemical Theory and Computation, **2019**. 15(8): 4699-4707.
13. Yin, J., et al., *The SAMPL5 host–guest challenge: computing binding free energies and enthalpies from explicit solvent simulations by the attach-pull-release (APR) method*. Journal of Computer-Aided Molecular Design, **2017**. 31(1): 133-145.
14. Schuetz, D.A., et al., *A structure–kinetic relationship study using matched molecular pair analysis*. RSC Medicinal Chemistry, **2020**. 11(11): 1285-1294.
15. Liu, H., et al., *Public Data Set of Protein–Ligand Dissociation Kinetic Constants for Quantitative Structure–Kinetics Relationship Studies*. ACS Omega, **2022**. 7(22): 18985-18996.

16. Amangeldiuly, N., D. Karlov, and M.V. Fedorov, *Baseline Model for Predicting Protein–Ligand Unbinding Kinetics through Machine Learning*. Journal of Chemical Information and Modeling, **2020**. 60(12): 5946-5956.
17. Sinko, W., et al., *ModBind, a Rapid Simulation-Based Predictor of Ligand Binding and Off-Rates*. Journal of Chemical Information and Modeling, **2024**. 65(1): 265-274.
18. Zhao, Y., et al., *Mixture-of-Experts Based Dissociation Kinetic Model for De Novo Design of HSP90 Inhibitors with Prolonged Residence Time*. Journal of Chemical Information and Modeling, **2024**. 64(22): 8427-8439.
19. Kokh, D.B., et al., *Machine Learning Analysis of τRAMD Trajectories to Decipher Molecular Determinants of Drug-Target Residence Times*. Frontiers in Molecular Biosciences, **2019**. 6.
20. Kokh, D.B., et al., *A workflow for exploring ligand dissociation from a macromolecule: Efficient random acceleration molecular dynamics simulation and interaction fingerprint analysis of ligand trajectories*. The Journal of Chemical Physics, **2020**. 153(12).
21. Pramanik, D., et al., *Can one trust kinetic and thermodynamic observables from biased metadynamics simulations: detailed quantitative benchmarks on millimolar drug fragment dissociation*. bioRxiv, **2019**.
22. Dixon, T., S.D. Lotz, and A. Dickson, *Predicting ligand binding affinity using on- and off-rates for the SAMPL6 SAMPLing challenge*. Journal of Computer-Aided Molecular Design, **2018**. 32(10): 1001-1012.
23. Lamim Ribeiro, J.M. and P. Tiwary, *Toward Achieving Efficient and Accurate Ligand-Protein Unbinding with Deep Learning and Molecular Dynamics through RAVE*. Journal of Chemical Theory and Computation, **2018**. 15(1): 708-719.
24. Yu, W., et al., *High-Throughput Ligand Dissociation Kinetics Predictions Using Site Identification by Ligand Competitive Saturation*. Journal of Chemical Theory and Computation, **2025**. 21(9): 4964-4978.
25. Zhou, F., et al., *Structure-Kinetic Relationship for Drug Design Revealed by PLS Model with Retrosynthesis-Based Pre-trained Molecular Representation and Molecular Dynamics Simulation*. bioRxiv, **2022**.
26. Lamim Ribeiro, J.M., D. Provasi, and M. Filizola, *A combination of machine learning and infrequent metadynamics to efficiently predict kinetic rates, transition states, and molecular determinants of drug dissociation from G protein-coupled receptors*. The Journal of Chemical Physics, **2020**. 153(12).
27. Jagger, B.R., C.T. Lee, and R.E. Amaro, *Quantitative Ranking of Ligand Binding Kinetics with a Multiscale Milestoning Simulation Approach*. The Journal of Physical Chemistry Letters, **2018**. 9(17): 4941-4948.
28. Badaoui, M., et al., *Combined Free-Energy Calculation and Machine Learning Methods for Understanding Ligand Unbinding Kinetics*. Journal of Chemical Theory and Computation, **2022**. 18(4): 2543-2555.

29. Congreve, M., et al., *Discovery of 1,2,4-Triazine Derivatives as Adenosine A2A Antagonists using Structure Based Drug Design*. Journal of Medicinal Chemistry, **2012**. 55(5): 1898-1903.
30. Guo, D., et al., *Functional efficacy of adenosine A2A receptor agonists is positively correlated to their receptor residence time*. British Journal of Pharmacology, **2012**. 166(6): 1846-1859.
31. D.A. Case, et al., *AMBER 24*. 2024, University of California: San Francisco.