Research Square

Preprints are preliminary reports that have not undergone peer review.
They should not be considered conclusive, used to inform clinical practice,
or referenced by the media as validated information.

ITRG: Enables Controllable Biomolecular Complex Ensemble Generation through Interface Trajectory Conditioning and Repulsive Guidance

Xiaojun Yao

x.jyao@mpu.edu.mo

Macao Polytechnic University

Kai Xu

Yanan Tian

Jieying Zang

Bing Hu

https://orcid.org/0000-0002-5488-9728

Huanxiang Liu

Article

Keywords:

Posted Date: September 28th, 2026

DOI: https://doi.org/10.21203/rs.3.rs-10925626/v1

License: This work is licensed under a Creative Commons Attribution 4.0 International License.
Read Full License

Additional Declarations: There is **NO** Competing Interest.

ITRG: Enables Controllable Biomolecular Complex Ensemble Generation through Interface Trajectory Conditioning and Repulsive Guidance

Kai Xu[1], Yanan Tian[1], Jieying Zang[1], Bing Hu[2] Huanxiang Liu[1,*], Xiaojun Yao[1,*]

[1] Faculty of Applied Sciences, Macao Polytechnic University, Macao 999078, China

[2] Dalhousie University Faculty of Computer Science Halifax, NS, Canada

Corresponding authors *:

Huanxiang Liu Email: hxliu@mpu.edu.mo

Xiaojun Yao Email: xjyao@mpu.edu.mo

Abstract

Generative biomolecular structure models can predict diverse molecular complexes, yet their default inference procedures provide limited control over which interface states are explored and offer little mechanistic organization of the resulting ensembles. Here, we introduce ITRG, an interface-centered inference-time framework that treats the receptor–partner interface itself as the common target of perturbation, trajectory construction, and diversity control. Unlike global modulation of latent pair representations or generic guidance in predefined collective-variable spaces, ITRG selectively perturbs contact-associated protein–partner pair features, organizes the resulting feature states through contact-guided Markov chain Monte Carlo sampling, and applies descriptor-specific repulsive guidance to independently diversify ligand conformer, pose, and local binding environment. In representative protein–small-molecule systems, ITRG produced strength- and system-dependent contact redistribution, expanded accessible contact modes, and revealed distinct contact-retentive and contact-switching interface trajectories that were not captured by RMSD alone. Repulsive guidance further broadened targeted diversity dimensions without widespread steric deterioration. The same receptor–partner formulation extended naturally to protein–nucleic-acid and protein–peptide complexes, enabling base/backbone-resolved and peptide-interface analyses. Together, ITRG establishes a general strategy for manipulating and interpreting alternative molecular-recognition landscapes encoded by generative biomolecular models.

Introduction

Deep-learning-based structure prediction and generative modeling have transformed the modeling of biomolecular structures and interactions. AlphaFold2 established near-experimental accuracy for protein structure prediction[1], and recent all-atom or multi-component models, including AlphaFold3, RoseTTAFold All-Atom, Chai-1 and Boltz-family models, have extended this capability to protein-small-molecule, protein-nucleic-acid and protein-peptide complexes[2–6]. In parallel, diffusion- and generative-model-based approaches have enabled ligand pose prediction, protein-ligand complex generation and ligand-induced conformational modeling[7–10]. More recently, generative models for protein equilibrium ensembles have further emphasized the need to move beyond single-structure prediction toward ensemble-level structural modeling[11]**.** These advances create new opportunities for structure-based molecular design and biomolecular interface analysis, while also raising a key question: how can alternative protein-partner interface states be deliberately controlled, organized and interpreted during generation?

Many biomolecular recognition processes cannot be adequately described by a single dominant complex structure. Receptor flexibility, induced fit, conformational selection and local rearrangement of binding-site contacts are central to molecular recognition and interface adaptation[12–15]. Recent AlphaFold3-derived sampling methods and related conformational-ensemble prediction approaches have further highlighted that learned structure models can encode alternative conformational states, but also that default prediction pipelines often remain biased toward dominant or ground-state structures[16–18]. Across protein-small-molecule, protein-nucleic-acid and protein-peptide systems, alternative interface states may involve residue-level contact redistribution, shifts in interface participating patches, base- or backbone-mediated nucleic-acid recognition, or peptide-associated rearrangement of the binding surface. Protein-ligand interaction profiling, interaction fingerprint analysis, protein-DNA specificity modeling, protein-nucleic-acid docking and peptide-protein complex prediction further emphasize that interface-level residue contacts and local binding environments are critical objects of analysis[19–28]. Thus, generative complex modeling should ideally move beyond producing independent stochastic samples and provide explicit, interpretable control over interface-level conformational variation.

A major challenge is that interface remodeling is not always reflected by global coordinate displacement. Two generated complexes can remain close in root mean square deviation (RMSD)

while exhibiting distinct residue-level contact patterns, whereas larger coordinate shifts may still preserve much of the original interaction footprint. This issue is particularly important for compact or threshold-sensitive interfaces, where small coordinate changes near a contact cutoff can produce discrete changes in contact topology. Recent evaluations of AI-based docking and cofolding methods further indicate that RMSD-based assessment should be complemented by steric and geometric validity checks[29–32]. Interpreting generated complex ensembles therefore requires contact probability maps, contact mode organization, patch overlap analysis, and interface-distance and steric diagnostics in addition to conventional RMSD-based metrics.

Several strategies have attempted to access alternative conformational states from deep-learning structure models. Multiple-sequence-alignment (MSA) perturbation, sequence clustering, subsampling and stochastic inference have been used to expose alternative states from AlphaFold2-derived models[33–38]. These methods show that learned structure predictors can encode information about multiple conformational basins, but they often rely on heuristic input perturbations, large sampling budgets or post hoc selection. More recent inference-time steering strategies provide more direct access to structural priors encoded in pretrained biomolecular diffusion models. Boltz-sample and Metadiffusion represent two recent inference-time strategies for steering Boltz-2 conformational sampling, but they operate at distinct levels of the model[39,40]. Boltz-sample rescales the conditioning-encoder-derived latent pair tensor at the Pairformer input, after MSA processing and before Pairformer refinement, and applies this scalar modulation at each recycle step. This repeated global rescaling modulates the effective strength of residue-residue pairwise couplings across the full pair representation, thereby steering Boltz-2 toward alternative protein conformational states. Metadiffusion instead adds meta-energy gradients during denoising, enabling pretrained models to optimize, steer or explore user-defined collective variables without retraining and to connect generated ensembles with experimental observables such as small-angle X-ray scattering (SAXS) and nuclear magnetic resonance (NMR). Related diffusion-guidance and particle-guidance strategies have also shown that inference-time gradients or inter-sample repulsion can steer pretrained diffusion models toward user-defined objectives or more diverse sample distributions[41–44]. These inference-time strategies demonstrate that pretrained biomolecular diffusion models can be steered without retraining, but they do not directly define the receptor-partner interface itself as the joint target of feature perturbation, trajectory construction and diversity control. ITRG was designed to fill this gap by making the receptor-partner interface the

explicit target of both perturbation and diversity steering. Whereas Metadiffusion applies meta-energy gradients during coordinate denoising and implements repulsion as a generic inter-sample bias in user-defined collective-variable or global descriptor spaces, ITRG defines repulsive guidance in three task-specific descriptor spaces for biomolecular complex generation: ligand internal-distance geometry, protein-aligned ligand placement and dynamic-shell local binding environment. In this formulation, similar samples are pushed apart only with respect to the selected interface-relevant diversity mode, allowing conformer-, pose- and environment-level variation to be controlled separately. Thus, the repulsion in ITRG is not a general CV-level diversity bias, but a mode-specific interface-diversification mechanism designed to separate ligand-internal geometry, receptor-relative placement and chemically informed binding-context variation. In contrast to Boltz-sample, which globally rescales the latent pair tensor at the Pairformer input during each recycle step, ITRG restricts perturbation to contact-associated protein-partner pair-feature entries. This interface-localized design is intended to promote contact-driven remodeling while reducing the risk of globally altering the structural scaffold. ITRG acts on downstream Boltz-2 pair features after feature extraction and selectively perturbs only contact-associated protein-partner pair-feature entries defined by the interface contact matrix. This localized design avoids repeated global scaling of all residue-pair couplings and links inference-time steering directly to residue-level contact redistribution, contact-resolved trajectories and mode-specific interface diversity.

This distinction highlights a remaining gap in current inference-time steering strategies: the receptor-partner interface has not yet been treated as the primary object of localized pair-feature perturbation, contact-resolved trajectory construction and descriptor-specific diversity analysis. Existing collective-variable guidance primarily operates through predefined structural descriptors, whereas global pair-representation scaling modulates pairwise coupling strength across the system. In complex generation, however, alternative binding states are not only reflected by global conformational changes, but also by residue-level contact redistribution, patch drift, ligand pose shifts and changes in the local protein environment surrounding the partner molecule. These interface-level changes require more than global conformational steering or whole-representation scaling. An interface-specific mechanism is therefore needed to select perturbation targets, organize generated states into contact-resolved paths and separately control conformer-, pose- and environment-level sources of diversity during complex generation.

Here, we developed ITRG, an interface-centered perturbation framework with complementary path-generation and diversity-steering modules for Boltz-2-generated biomolecular complexes. The framework is organized around receptor-partner interface representations rather than global structural descriptors alone.

First, β-parameterized perturbation selectively modifies contact-associated protein-partner pair-feature entries, enabling controlled redistribution of interface contact patterns. β denotes a dimensionless interface-steering coefficient that controls the magnitude and direction of contact-aware pair-feature modulation; it is used as an inference-time hyperparameter.

Second, contact-guided MCMC sampling generates an ordered sequence of interface-associated feature states in the pair-feature space, following the general Metropolis-Hastings sampling principle [45,46]. Each state is independently used to condition Boltz-2 generation, producing a local conformational ensemble from which a representative structure is selected. The ordered representatives then form a contact-resolved interface trajectory, enabling local stepwise transitions and anchor-referenced remodeling to be interpreted using RMSD, contact overlap and patch overlap.

Third, we introduced descriptor-specific repulsion-guided sampling to control diversity during generation. Instead of applying a single generic diversity objective, we defined three complementary repulsive modes: conformer-guided sampling to expand ligand internal-distance geometry, pose-guided sampling to separate ligand placements in a protein-aligned coordinate frame, and environment-guided sampling to diversify the local binding context surrounding the ligand. The environment-guided mode uses a dynamic-shell binding-environment descriptor that integrates radial shell statistics, protein atom-type composition and protein-ligand contact features, allowing repulsion to act on chemically informed local binding-context variation.

We primarily applied ITRG to two representative protein-small-molecule systems: MfnG O-methyltransferase bound to L-tyrosine (PDB 7UX8) and T4 lysozyme L99A bound to toluene (PDB 4W53)[47,48]. We further tested whether the same receptor-partner formulation could be extended to protein-nucleic-acid and protein-peptide interfaces using human TATA-binding protein bound to TATA-sequence DNA (PDB 1TGH), the Fox-1 RNA-binding domain bound to UGCAUGU RNA (PDB 2ERR), and the Abl-SH3 domain bound to a designed high-affinity

peptide ligand (PDB 1BBZ)[49–51]. In protein-small-molecule systems, contact-aware perturbation reshaped contact-mode landscapes in a perturbation-strength- and system-dependent manner, and MCMC path analysis distinguished contact-retentive remodeling from cutoff-sensitive contact switching that would be obscured by RMSD-based comparisons alone. Repulsion-guided sampling further showed that conformer- and environment-level diversity can be selectively expanded, with the strongest effects observed for ligand-internal geometry and binding-environment diversification. Extending the same receptor-partner formulation to protein-DNA, protein-RNA and protein-peptide complexes enabled base/backbone-resolved analysis for nucleic-acid interfaces and total-interface analysis for peptide-associated rearrangements. Together, this framework provides a contact-aware, path-resolved and descriptor-specific strategy for controlling alternative interface states in Boltz-2-generated biomolecular complexes.

Results

Overview of the ITRG Framework

ITRG was designed to convert Boltz-2 complex generation from independent stochastic sampling into contact-aware exploration of receptor-partner interface states (**Fig. 1**). The framework contains three linked components: contact-aware perturbation of Boltz-2 pair features, MCMC-based interface path construction and descriptor-specific repulsion-guided diversity steering.

ITRG first represents different biomolecular complexes in a unified receptor-partner format, covering protein-small-molecule, protein-peptide and protein-nucleic-acid systems (**Fig. 1A**). For each input complex, Boltz-2 pair features are extracted as the initial representation Z. In parallel, the corresponding 3D complex structure is used to derive an interfacial contact matrix M, which encodes receptor-partner contact topology in the same matrix space (**Fig. 1B**). This contact matrix is then used as an interface mask to selectively perturb contact-associated pair-feature entries while leaving non-contact regions unchanged, yielding an anchor feature state Z_0. The perturbation strength is controlled by β, enabling controlled modulation of contact-associated pair-feature entries before downstream structure generation (**Fig. 1C**).

Starting from the anchor pair feature Z_0, contact-guided MCMC sampling uses the contact matrix M to restrict perturbations to selected pairwise regions. Each proposed pair-feature update is

scored and accepted or rejected according to the Metropolis-Hastings (MH) rule. After burn-in and thinning, the retained sampled pair features $\{Z_0, Z_1 \cdots Z_r\}$ are used for downstream conformational ensemble generation (**Fig. 1D**). Each sampled state is independently used to condition Boltz-2 structure generation, producing a local conformational ensemble for that state (**Fig. 1E**). A representative structure is then selected from each ensemble, and the ordered representatives are connected according to the MCMC sequence to form a contact-resolved interface trajectory (**Fig. 1G**). This design makes it possible to analyze both local adjacent-state transitions and cumulative anchor-referenced remodeling using RMSD, contact overlap and patch overlap.

In parallel, ITRG introduces descriptor-specific repulsion-guided sampling during diffusion generation (**Fig. 1F**). Three complementary diversity modes are defined: conformer-guided sampling expands ligand internal geometry, pose-guided sampling separates ligand placements in a protein-aligned frame, and environment-guided sampling diversifies the local binding context surrounding the ligand. In particular, the environment-guided mode uses a binding-environment descriptor to capture chemically informed changes in the local protein environment, providing a diversity objective beyond ligand-internal or pose-level displacement.

This framework defines three experimentally evaluated questions. First, we tested whether contact-aware β perturbation changes ensemble-level interface contact topology relative to the unmodified Boltz-2 baseline. Second, we asked whether MCMC sampling of perturbed pair features can organize these changes into contact-resolved interface trajectories. Third, we examined whether descriptor-specific repulsive guidance can selectively expand conformer-, pose- and environment-level diversity during generation. We then tested whether the same receptor-partner formulation extends beyond protein-small-molecule systems to protein-nucleic-acid and protein-peptide interfaces.

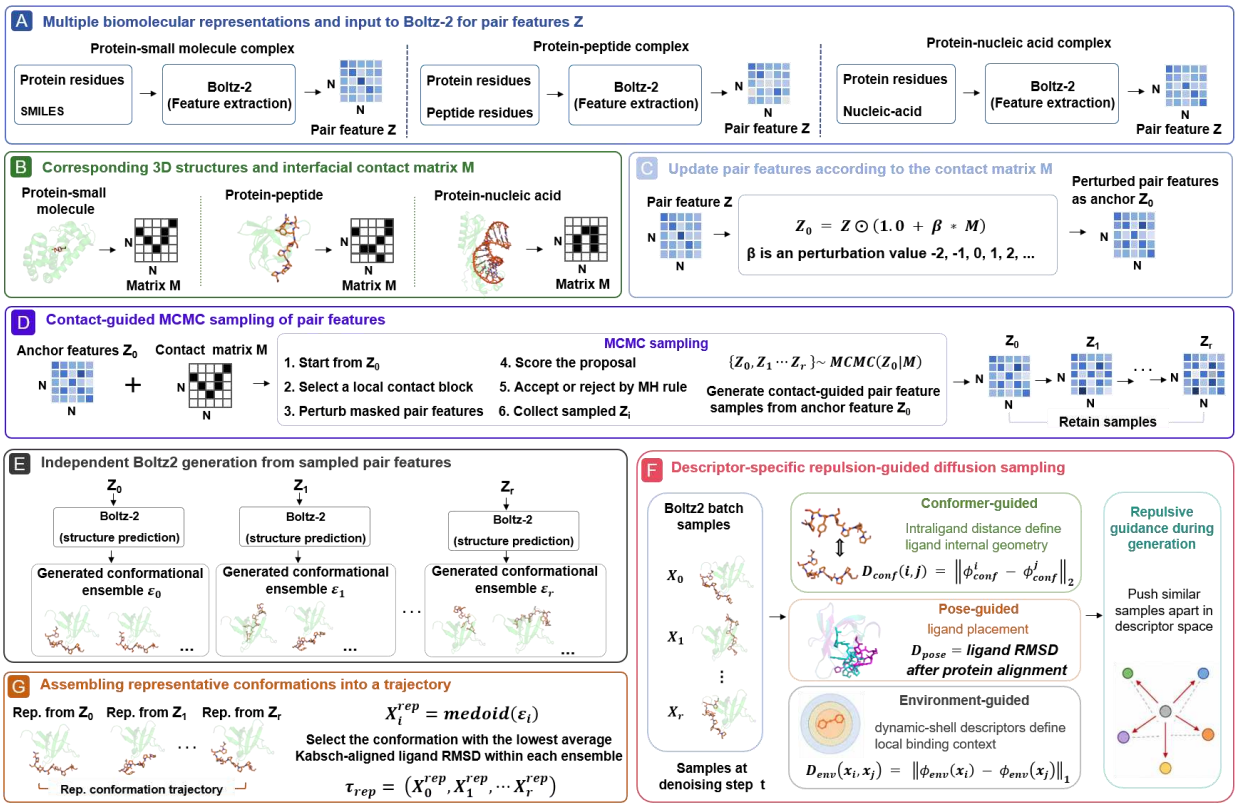


Figure 1. Overview of the ITRG framework. **A**, Protein-small-molecule, protein-peptide and protein-nucleic-acid complexes are represented in a unified receptor-partner format and processed by Boltz-2 to obtain pair features Z. **B**, The corresponding 3D structures are used to derive an interfacial contact matrix M, which encodes receptor-partner contact topology. **C**, Contact-aware perturbation selectively modifies pair-feature entries according to contact matrix M, producing perturbed anchor features Z_0. The perturbation is controlled by the interpolation parameter β, $\odot$ denotes element-wise. **D**, Starting from Z_0, MCMC sampling uses the contact matrix M to restrict local perturbations to selected pairwise regions. Each proposal is scored and accepted or rejected by the Metropolis-Hastings rule, yielding an ordered sequence of sampled pair-feature states,$\{Z_0,\ Z_1\ \cdots Z_r\}$. **E**, Each sampled state independently conditions Boltz-2 structure generation to produce a local conformational ensemble for that state. **F**, Descriptor-specific repulsion-guided diffusion sampling provides an additional diversity-control mechanism during Boltz-2 generation. Conformer-guided sampling uses intraligand-distance descriptors to diversify ligand internal geometry, pose-guided sampling uses protein-aligned ligand pose descriptors to diversify ligand placement, and environment-guided sampling uses dynamic-shell descriptors to diversify the local binding context. In each case, similar samples are pushed apart in the selected descriptor space during generation. **G**, Representative conformations selected from each generated ensemble are assembled in MCMC order to form a contact-resolved interface trajectory. The framework combines contact-aware interface perturbation, MCMC-based interface path construction and descriptor-specific repulsion-guided diversity control for interpretable exploration of alternative biomolecular interface states.

Contact-Aware β Perturbation Broadens Protein-Small-Molecule Contact-Mode Landscape

We first tested whether contact-aware β perturbation could broaden the contact landscapes of Boltz-2-generated protein-small-molecule complexes. Two representative systems, 7UX8 and 4W53, were analyzed under two perturbation strengths, $\beta = -1$ and $\beta = -2$, and compared with the unperturbed Boltz-2 baseline. For each complex and perturbation setting, 100 conformations were generated and evaluated using residue-level protein-ligand contact profiles, Δ-contact profiles relative to baseline, contact-map similarity, contact-mode clustering and steric clash statistics (**Fig. 2**). Detailed per-system contact redistribution and steric/interface-distance analyses are provided in **Supplementary Figs. S1-S4**.

Active contact profiles were constructed by retaining residues with contact probabilities ≥ 0.10 in either the baseline or perturbed ensembles, or residues with an absolute Δ-contact probability ≥ 0.05 relative to the baseline. These profiles showed that interface perturbation altered residue-level protein-ligand contact patterns in both systems. In 7UX8, the $\beta = -1$ ensemble retained partial similarity to the baseline contact topology, whereas $\beta = -2$ induced stronger contact redistribution. This trend was reflected by a decrease in contact-map Jaccard similarity from 0.50 at $\beta = -1$ to 0.32 at $\beta = -2$. Although the new-contact rate remained comparable between $\beta = -1$ and $\beta = -2$, the lower Jaccard similarity at $\beta = -2$ indicated a stronger rearrangement of the baseline contact topology. In 4W53, the response was more pronounced. The $\beta = -1$ ensemble largely preserved the baseline contact map, with a Jaccard similarity of 0.92, whereas $\beta = -2$ caused a marked contact-mode shift, with Jaccard similarity dropping to 0.00 and the new-contact rate increasing to 0.33. Thus, contact-aware perturbation reshaped protein-ligand contact landscapes in a perturbation-strength-dependent and system-dependent manner.

Contact-mode clustering further showed that perturbation expanded the accessible contact-state space, particularly under the stronger perturbation condition. In 7UX8, $\beta = -1$ slightly reduced the effective number of contact modes from 5 to 4 and marginally decreased the occupancy entropy from 1.43 to 1.39, suggesting moderate contact rearrangement without clear mode expansion. By contrast, $\beta = -2$ increased the effective number of contact modes from 5 to 9 and raised the occupancy entropy from 1.43 to 2.77, consistent with a broader and more evenly distributed contact-mode landscape. In 4W53, the baseline ensemble was concentrated in a single dominant

contact mode, with an effective mode count of 1 and near-zero entropy. Perturbation increased the effective number of modes to 3 at $\beta = -1$ and 6 at $\beta = -2$, with entropy increasing to 0.96 and 2.35, respectively. These results indicate that stronger interface perturbation promotes multi-modal contact sampling in both protein-small-molecule systems.

We next examined whether this expanded contact exploration was accompanied by severe steric deterioration. Across both complexes and perturbation strengths, the mean number of severe steric clashes remained zero. The mean minimum protein-ligand interface distances also remained within a reasonable range. For 7UX8, the mean minimum distance changed from 2.65 Å in the baseline to 2.70 Å at $\beta = -1$ and 2.57 Å at $\beta = -2$. For 4W53, the corresponding values changed from 3.40 Å in the baseline to 3.01 Å and 2.89 Å under $\beta = -1$ and $\beta = -2$, respectively. Although perturbation slightly reduced the minimum interface distance in some settings, no severe clashes were introduced, suggesting that the observed contact diversification was not primarily driven by gross steric failure.

These findings indicate that contact-aware pair-feature perturbation broadens protein-small-molecule contact landscapes by redistributing residue-level interfacial contacts and promoting alternative contact modes, particularly under stronger perturbation. The observed changes are therefore better interpreted as β-dependent shifts in contact topology rather than random geometric distortions, and they occurred without evidence of severe steric deterioration.

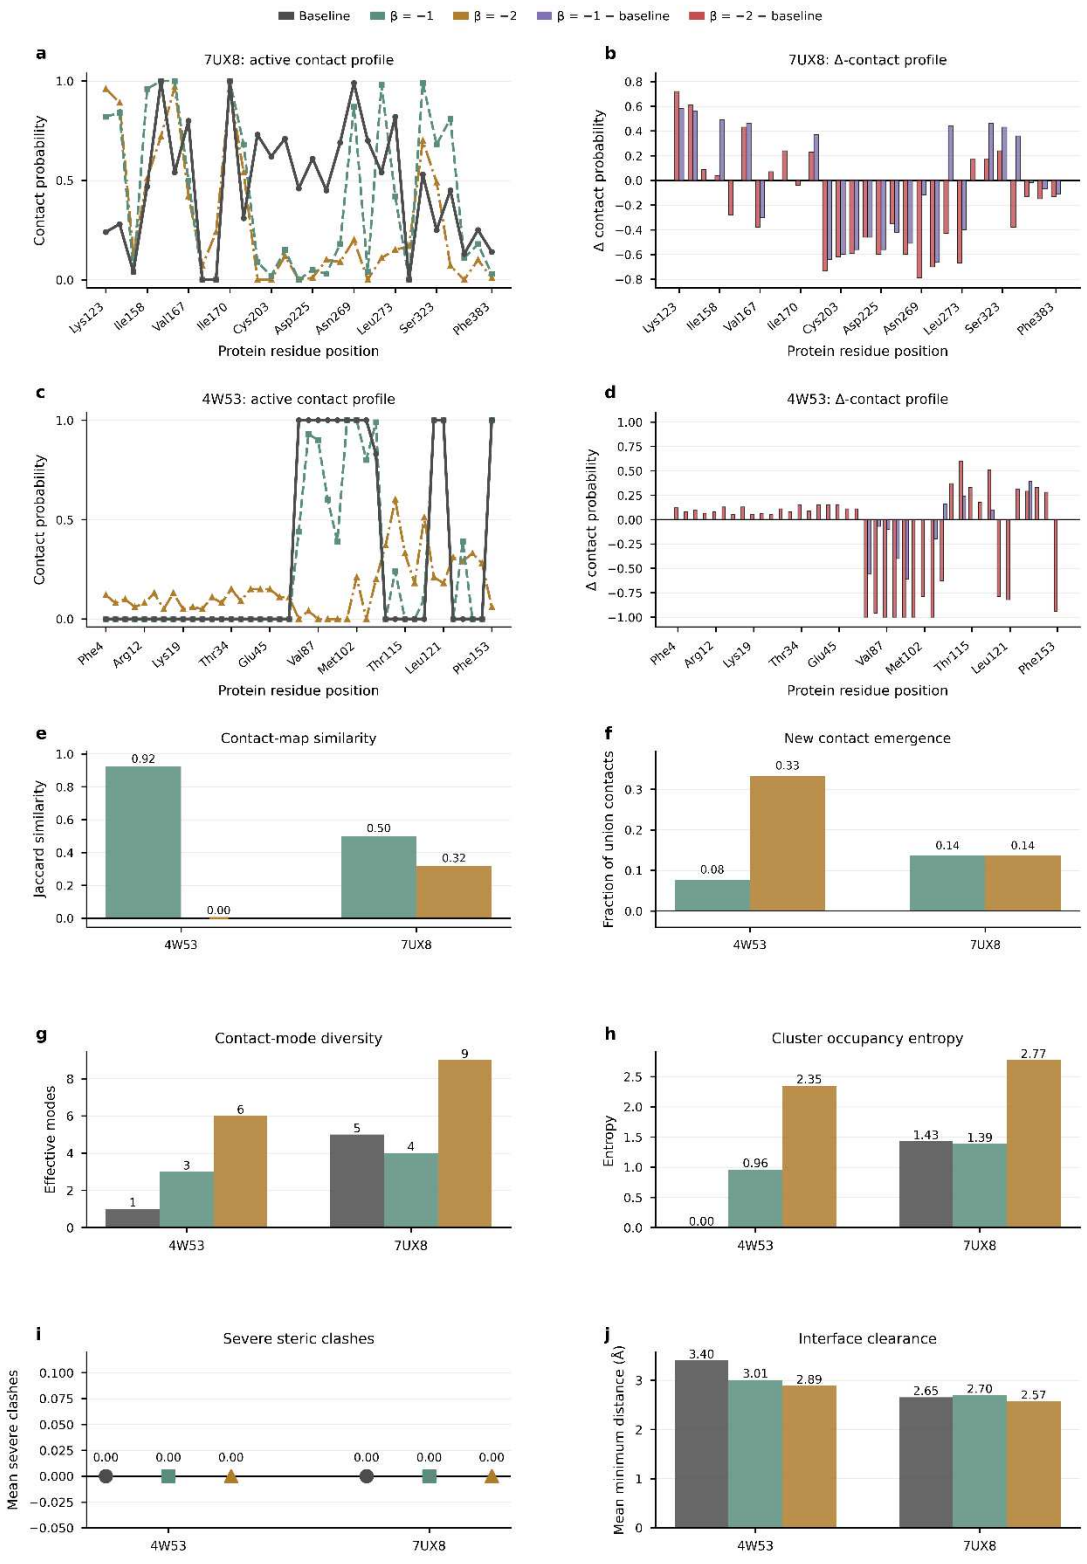


Figure 2. Contact-aware β perturbation expands protein-small-molecule contact-mode landscapes. a-d, Active residue-level protein-ligand contact profiles and Δ-contact profiles for 7UX8 and 4W53. For each system, 100 conformations generated with $\beta = -1$ or $\beta = -2$ were compared with 100 Boltz-2 baseline conformations. Δ-

contact profiles indicate perturbation-induced changes relative to the baseline. **e-f**, Contact-map Jaccard similarity and new-contact emergence rate computed at a contact-probability threshold of 0.3. **g-h**, Effective contact-mode counts and occupancy entropy derived from residue-level contact-fingerprint clustering. **i-j**, Severe steric clash counts and minimum protein-ligand interface distances used to assess geometric plausibility. Stronger perturbation, particularly $\beta = -2$, increased contact-mode diversity and induced alternative interfacial contact patterns without introducing severe steric clashes.

MCMC Sampling Resolves Fine-Grained Protein-Small-Molecule Interface Trajectories at β = −2

The fixed-β analysis established that stronger interface perturbation can broaden contact-mode landscapes, but it did not reveal how perturbed interface states are connected in feature space or whether contact changes occur through gradual remodeling or abrupt switching. We therefore focused on $\beta = -2$, where the strongest contact redistribution was observed, and used contact-guided MCMC sampling to organize perturbed pair-feature states into ordered interface trajectories.

Starting from the anchor perturbed pair feature Z_0, nine retained MCMC sampling states produced a sequence of perturbed pair features, $Z_1 - Z_9$. Each feature state was independently submitted to the Boltz-2 structure prediction module, yielding ten conformational ensemble blocks with 100 predicted conformations per state. A representative structure was selected from each block, and the ordered representatives were connected to form a contact-resolved interface trajectory. In this analysis, RMSD was calculated as the Kabsch-aligned matched heavy-atom RMSD of the ligand chain; therefore, it primarily measures ligand-internal conformer variation rather than protein-aligned ligand pose displacement.

For 7UX8, the trajectory showed contact redistribution under small ligand-internal RMSD. Adjacent representative RMSD ranged from 0.075 Å to 0.258 Å, indicating that neighboring MCMC states retained highly similar ligand internal geometry (**Fig. 3a**). Nevertheless, adjacent contact overlap and patch overlap varied along the trajectory, with minima of 0.545 and 0.583, respectively (**Fig. 3b,c**). Thus, local transitions between neighboring MCMC states did not involve large ligand conformer changes, but still produced measurable rearrangements in ligand-contacting residues and interface patches. This indicates that the 7UX8 trajectory explored a contact-

remodeling regime with partial interface continuity rather than a transition between unrelated binding modes.

The anchor-referenced metrics support this interpretation from a different perspective. Ligand RMSD relative to the anchor remained small, with a maximum of 0.222 Å at block 3 (**Fig. 3d**). Yet the minimum contact overlap and patch overlap relative to the anchor remained at 0.462 and 0.500, respectively (**Fig. 3e,f**). The sampled 7UX8 states retained similar ligand internal geometry while losing approximately half of the anchor-associated contact topology and interface footprint. This pattern is interpreted as partial contact redistribution within an anchor-related interface regime, rather than as large coordinate-level structural drift.

4W53 showed an even stronger decoupling between ligand-internal RMSD and contact topology. Adjacent representative RMSD remained extremely small across the trajectory, ranging from 0.003 Å at 5→6 to 0.008 Å at 7→8 (**Fig. 3g**). However, contact overlap and patch overlap showed broader variation, with contact overlap reaching 0.000 at several transitions and patch overlap decreasing to 0.067 at 2→3 (**Fig. 3h,i**). These results indicate that nearly indistinguishable ligand internal geometries can be associated with sharply different residue-level contact assignments. Thus, 4W53 represents a more discrete contact-switching regime, in which small coordinate differences near the contact cutoff are amplified into pronounced changes in contact topology.

The anchor-referenced view further emphasized this behavior. RMSD relative to the anchor remained below 0.008 Å across the sampled states (**Fig. 3j**), confirming that the representative structures stayed geometrically close to the anchor at the coordinate level. However, contact overlap relative to the anchor reached 0.000 and patch overlap decreased to 0.071 (**Fig. 3k,l**). Thus, the 4W53 trajectory was dominated not by ligand conformer displacement, but by redistribution of the protein-ligand contact topology. Compared with 7UX8, where approximately half of the anchor-associated contacts and patches were retained, 4W53 showed near-complete contact reassignment at selected trajectory stages.

We further evaluated intra-block conformational consistency and representative quality using within-block pairwise ligand RMSD, pairwise contact-overlap consistency, clash-outlier statistics and representative-level diagnostic metrics (**Supplementary Fig. S5, Supplementary Table S1**). PyMOL visualization of selected representative structures provided direct structural context for

the contact-retentive remodeling observed in 7UX8 and the sharper contact-switching behavior observed in 4W53 (**Supplementary Fig. S6**). These analyses provided a diagnostic layer for interpreting the representative-based inter-block trajectory metrics.

Thus, MCMC sampling under $\beta = -2$ provides a path-resolved view of interface remodeling that is not accessible from ensemble-level comparison alone. Both 7UX8 and 4W53 show that residue-level contact topology can change substantially despite small ligand-internal RMSD. However, the two systems differ in the extent of contact retention: 7UX8 follows a partial contact-redistribution trajectory with retained interface continuity, whereas 4W53 follows a sharper contact-switching trajectory with near-complete loss of anchor-associated contacts at selected stages. This comparison highlights the value of combining ligand RMSD with contact and patch overlap analyses, because RMSD alone would underestimate interface-level remodeling in both systems.

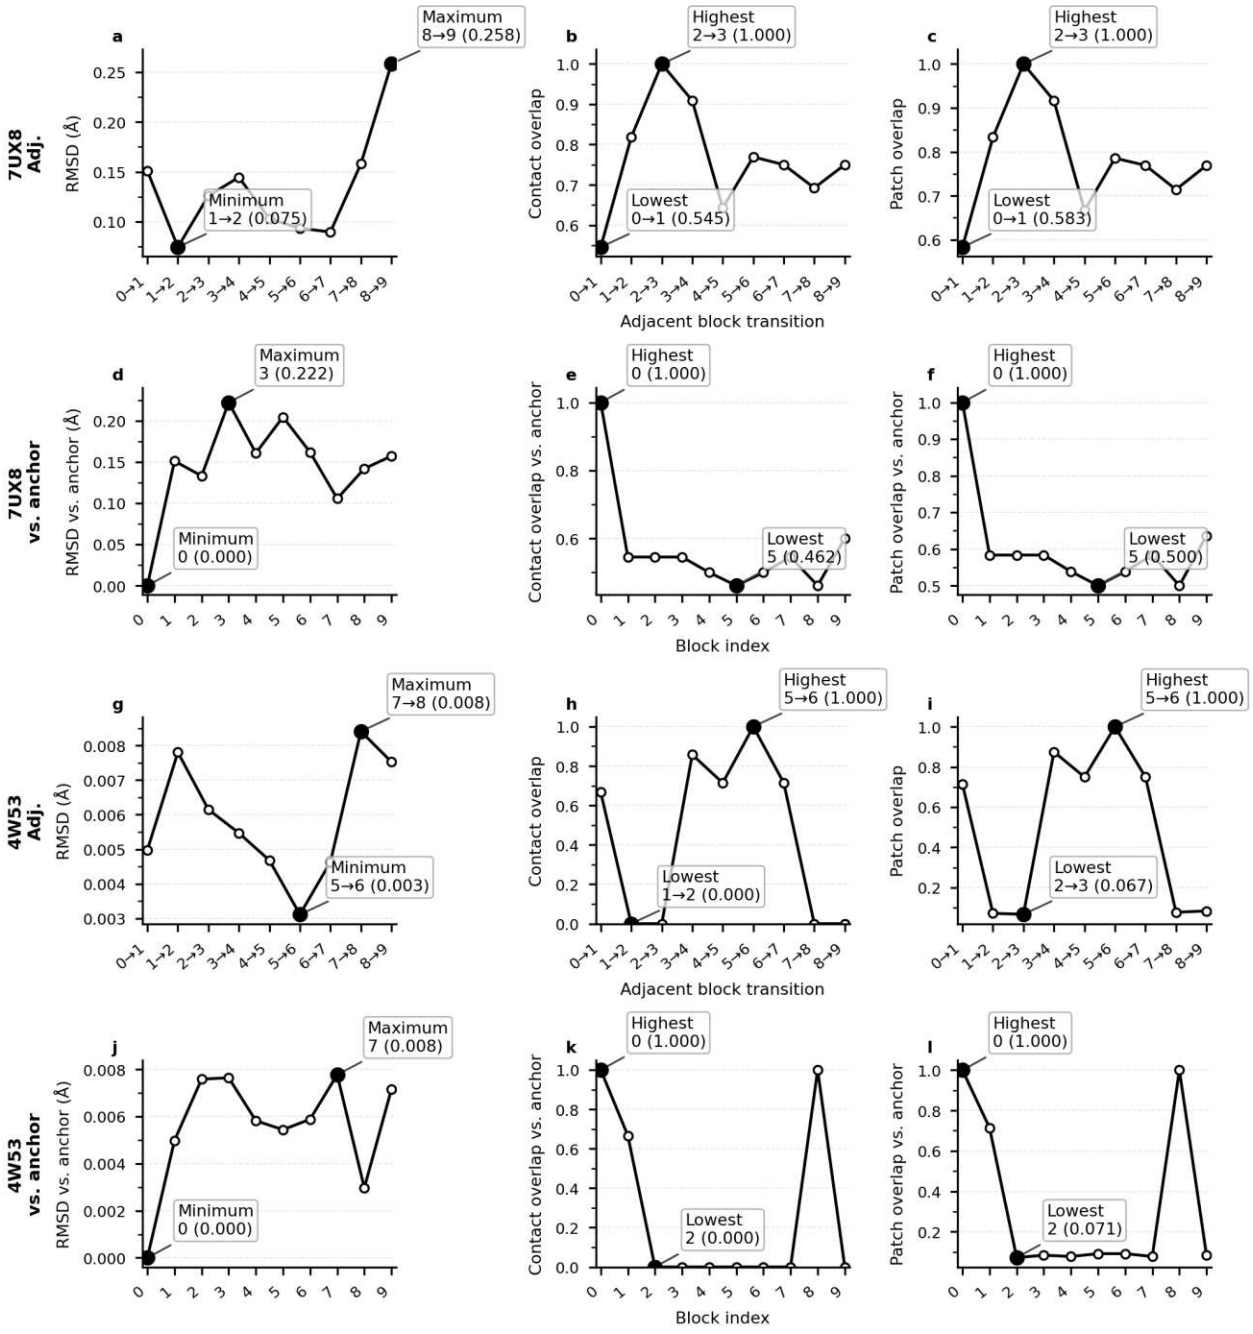


Figure 3. MCMC Sampling Resolves Fine-Grained Protein-Small-Molecule Interface Trajectories Under β = −2. **a-c**, Adjacent-block representative ligand RMSD, contact overlap and patch overlap for 7UX8. **d-f**, Corresponding anchor-referenced ligand RMSD, contact overlap and patch overlap for 7UX8. **g-i**, Adjacent-block representative ligand RMSD, contact overlap and patch overlap for 4W53. **j-l**, Corresponding anchor-referenced ligand RMSD, contact overlap and patch overlap for 4W53. Black curves indicate metric values along the MCMC-sampled trajectory. Filled black markers and linked annotation boxes denote extrema. Adjacent-block metrics quantify local changes between neighboring sampled blocks, whereas anchor-referenced metrics quantify deviation from the initial anchor representative structure.

Repulsion-Guided Boltz-2 Sampling Selectively Expands Ligand-Conformer and Binding-Environment Diversity

The MCMC analysis organized interface perturbation responses into contact-resolved paths, but it did not directly optimize diversity within a generation batch. We therefore next evaluated a complementary during-generation strategy: descriptor-specific repulsive guidance designed to expand selected diversity dimensions among simultaneously generated samples. Using 7UX8 and 4W53 as two independent test systems, we compared baseline generation with three guided sampling strategies: conformer-guided, pose-guided and environment-guided sampling. Diversity was assessed at three complementary levels: ligand conformer variation, protein-aligned pose displacement and binding-environment descriptor distance. For each system, guided-sampling gains were normalized to the corresponding baseline value within the same diversity dimension, enabling direct comparison of mode-specific expansion across guidance strategies. The summary statistics underlying these baseline-normalized gains are provided in **Supplementary Table S2**.

The baseline-normalized heatmaps showed that repulsion guidance produced clear and mode-dependent increases in diversity (**Fig. 4a**). In 7UX8, conformer-guided sampling increased the global conformer diversity by approximately 1.9-fold and the nearest-neighbor conformer diversity by approximately 2.3-fold, indicating both broader conformational coverage and reduced local redundancy among generated samples. Environment-guided sampling produced the most prominent change in binding-environment diversity, increasing the global environment descriptor distance by approximately 6.9-fold and the nearest-neighbor environment distance by approximately 3.1-fold. By contrast, pose-guided sampling produced a more modest global pose-diversity increase but a clearer nearest-neighbor-level improvement, suggesting that pose repulsion primarily enhanced local separation among ligand placements rather than inducing a large global redistribution.

The diversification effect was more pronounced in 4W53. Conformer-guided sampling strongly increased conformer diversity, with an approximately 82.7-fold increase in global pairwise conformer distance and an approximately 57.7-fold increase in nearest-neighbor conformer distance. Environment-guided sampling produced the largest global expansion of binding-environment diversity, reaching approximately 118.1-fold over baseline. Pose-guided sampling also increased environment-related and pose-related separation, including an approximately 3.6-

fold increase in nearest-neighbor pose diversity and an approximately 6.8-fold increase in nearest-neighbor environment diversity. These results indicate that different repulsive objectives can reshape the generated ensemble in distinct ways, with conformer guidance primarily expanding ligand-internal geometry and environment guidance preferentially diversifying the local binding context.

The raw pairwise distributions further supported these observations (**Fig. 4b**). Compared with baseline generation, guided sampling broadened the conformer-distance and environment-distance distributions, particularly in conditions matched to the intended guidance objective. The distributional shifts were not uniform across the two systems, suggesting that the magnitude of repulsion-induced diversification depends on the intrinsic flexibility of the ligand and the structural constraints imposed by the binding pocket. Nevertheless, both systems showed that guided sampling reduced baseline-like repetitive generations and expanded ligand conformational and binding-environment states.

Because **Fig. 4** primarily evaluates diversity rather than structural plausibility, we further examined whether the observed diversity gains were accompanied by severe ligand-protein overlaps. Minimum ligand-protein heavy-atom distances and clash statistics were calculated for all generated complexes（**Supplementary Table S3 and Supplementary Fig. S7**）. Across both 7UX8 and 4W53, baseline, conformer-guided and pose-guided generations showed no clash-containing complexes under the 2.0 Å cutoff. Environment-guided sampling introduced a small fraction of local clashes, with clash rates of 2% in 7UX8 and 9% in 4W53, and at most one clash pair per affected structure. Thus, the observed diversity expansion was not primarily driven by widespread sterically implausible ligand-protein overlaps, although environment-guided sampling can occasionally introduce local heavy-atom clashes.

Overall, descriptor-specific repulsion-guided Boltz-2 sampling provided controllable access to distinct diversity modes in protein-ligand complex generation. Conformer-guided sampling mainly expanded ligand-internal geometry, whereas environment-guided sampling strongly diversified the local binding context surrounding the ligand. This separation of conformer-, pose- and environment-level effects supports repulsion-guided sampling as a mode-specific strategy for reducing redundancy and broadening generated protein-ligand interface ensembles.

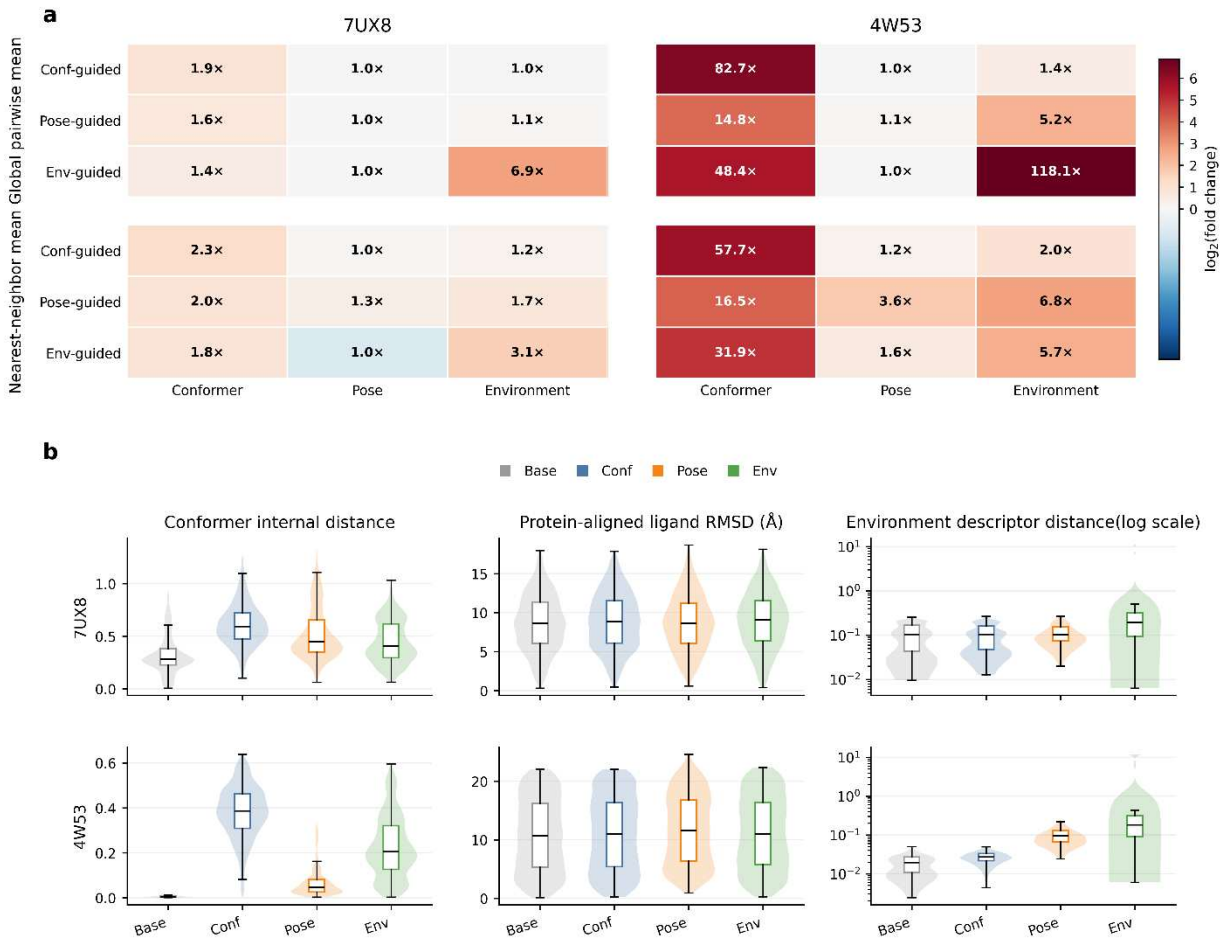


Figure 4. Repulsion-guided Boltz-2 sampling expands ligand and binding-environment diversity relative to baseline generation. a, Baseline-normalized diversity gains for repulsion-guided Boltz-2 sampling in two parallel protein-ligand systems, 7UX8 and 4W53. Heatmaps show fold changes relative to the corresponding baseline within each system. Columns indicate the three diversity dimensions, including conformer diversity, protein-aligned pose diversity and binding-environment diversity, whereas rows indicate the three guidance strategies: conformer-guided, pose-guided and environment-guided sampling. The upper row reports global pairwise mean diversity, and the lower row reports nearest-neighbor mean diversity, allowing simultaneous assessment of overall distributional spread and local de-duplication. Cell annotations indicate raw fold changes, and color intensity represents log2-transformed fold change. **b,** Within-system pairwise diversity distributions for baseline and guided generations. Violin plots with overlaid boxplots summarize conformer internal distance, protein-aligned ligand RMSD and environment descriptor distance for baseline, conformer-guided, pose-guided and environment-guided sampling. The environment descriptor distance is shown on a logarithmic scale. Together, these results show that repulsion-guided sampling can selectively reshape different diversity modes, with the strongest effects observed for conformer expansion and binding-environment diversification, while pose-level changes are more moderate and system-dependent.

Interface Perturbation Extends to Protein-Nucleic-Acid Complexes with Base- and Backbone-Resolved Responses

To test whether the receptor-partner formulation could be applied beyond protein-ligand complexes, we extended the perturbation and MCMC analyses to protein-nucleic-acid interfaces. We focused on $\beta = -2$ as a representative perturbation strength and applied the same baseline-versus-perturbed ensemble comparison and anchor-referenced MCMC trajectory analysis to one protein-DNA complex (1TGH) and one protein-RNA complex (2ERR). Because nucleic acids contain chemically distinct nucleobase and sugar-phosphate backbone components, protein-nucleic-acid contacts were further decomposed into total, base-associated and backbone-associated contact channels.

Ensemble-level $\beta = -2$ perturbation revealed distinct contact-redistribution patterns in the DNA and RNA systems. In the protein-DNA complex 1TGH, the total Δ contact map showed a relatively sparse but high-amplitude perturbation pattern, indicating localized contact weakening and strengthening rather than a uniform contact redistribution (**Fig. 5a**). Base- and backbone-resolved maps further showed that these changes were distributed across both nucleobase- and backbone-associated contact channels (**Fig. 5b,c**). In contrast, the protein-RNA complex 2ERR exhibited a broader perturbation response, with more extensive redistribution in the total and base-associated contact maps (**Fig. 5d,e**). The backbone-associated RNA map also showed detectable remodeling, although the response was less broadly distributed than that of the base-associated channel (**Fig. 5f**). These results indicate that the DNA and RNA cases respond differently to $\beta = -2$ perturbation, highlighting the ability of the framework to resolve case-specific base- and backbone-associated contact remodeling.

Anchor-referenced MCMC trajectory analysis further distinguished the remodeling behavior of the DNA and RNA interfaces. For 1TGH, the trajectory showed moderate structural deviation from the anchor while retaining a substantial fraction of both base- and backbone-associated contacts (**Fig. 5g-i**). By contrast, 2ERR displayed stronger anchor-relative remodeling, including larger structural displacement and near-complete loss of the original base- and backbone-associated contact patterns at specific trajectory stages (**Fig. 5j-l**). Thus, the RNA trajectory sampled more extensively remodeled interface states than the DNA trajectory under the same perturbation condition.

Detailed numerical summaries of the global Δ contact maps are provided in **Supplementary Table S4**, while adjacent-block MCMC transitions and ensemble-level quality-control analyses are shown in **Supplementary Figs. S8** and **S9**, respectively.

These analyses support the extension of the receptor-partner perturbation framework to protein-nucleic-acid complexes and show that perturbation responses can be resolved into base-associated and backbone-associated contact channels.

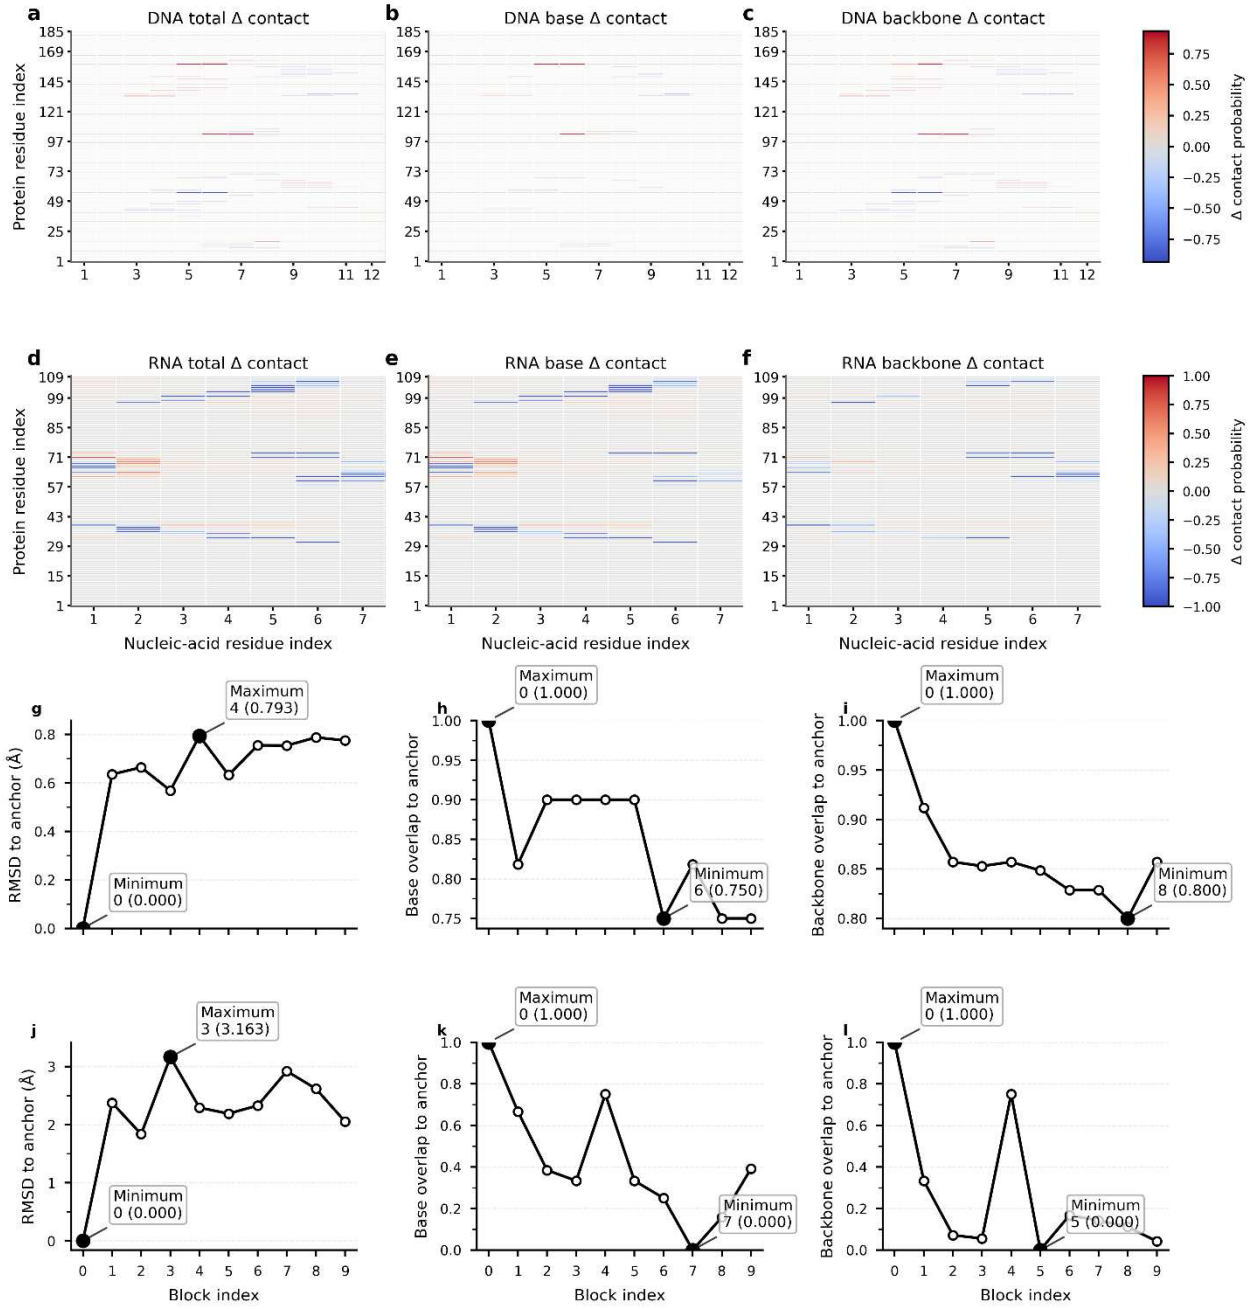


Figure 5. Interface Perturbation Extends to Protein-Nucleic-Acid Complexes with Base- and Backbone-Resolved Responses. a-c, $\beta = -2$ global perturbation analysis for the protein-DNA complex 1TGH. **a,** Total protein-DNA Δ contact probability relative to the baseline ensemble. **b,** Nucleobase-resolved Δ contact probability. **c,** Backbone-resolved Δ contact probability. **d-f,** Corresponding $\beta = -2$ global perturbation analysis for the protein-RNA complex 2ERR. **d,** Total protein-RNA Δ contact probability relative to the baseline ensemble. **e,** Nucleobase-resolved Δ contact probability. **f,** Backbone-resolved Δ contact probability. **g-i,** Anchor-referenced $\beta = -2$ MCMC trajectory metrics for 1TGH. **g,** RMSD to the anchor representative state. **h,** Nucleobase contact overlap to the anchor. **i,** Backbone contact overlap to the anchor. **j-l,** Corresponding anchor-referenced $\beta = -2$ MCMC trajectory metrics for 2ERR. **j,** RMSD to the anchor representative state. **k,** Nucleobase contact overlap to the anchor. **l,** Backbone contact overlap to the anchor. Δ contact probability denotes the difference between the $\beta = -2$ perturbed ensemble and the baseline ensemble. Base- and backbone-resolved metrics decompose the nucleic-acid partner into nucleobase-

associated and sugar-phosphate-backbone-associated contacts. Separate color scales are shown for the DNA and RNA Δ contact maps to preserve contrast within each system.

Interface Perturbation Extends to Protein-Peptide Complexes and Captures Peptide-Associated Interface Rearrangements

Protein-peptide interfaces provide a distinct test case in which the partner is a flexible polymeric chain rather than a small-molecule or a nucleic-acid. The peptide partner in 1BBZ was treated as a single interacting interface, and total contact probability, Δ contact probability, contact overlap and patch overlap were quantified.

The $\beta = -2$ perturbation substantially redistributed the 1BBZ protein-peptide contact landscape. The baseline ensemble showed a compact and highly conserved contact pattern, whereas the perturbed ensemble displayed a broader and more diffuse contact distribution (**Fig. 6a,b**). Consistently, the Δ contact probability map showed both loss of baseline-enriched contacts and emergence of alternative protein-peptide contacts under perturbation (**Fig. 6c**). Quantitatively, the number of non-zero contact-map entries increased from 29 in the baseline ensemble to 192 in the $\beta = -2$ ensemble, and 97 entries showed detectable redistribution using $|\Delta P| \geq 0.10$. Detailed global contact-map statistics are provided in **Supplementary Table S5**.

Anchor-referenced MCMC analysis further resolved the trajectory-level remodeling of the protein-peptide interface. RMSD to the anchor indicated moderate structural displacement of representative interface states, whereas contact overlap and patch overlap showed that detailed contact identities were more strongly remodeled than the broader interface-participating residue patch (**Fig. 6d-f**). In particular, contact overlap was completely lost at several blocks, while patch overlap remained partially retained along most of the trajectory, indicating that MCMC sampling captured contact-level rearrangement within a partially preserved interface region. Adjacent-block analyses further showed that the trajectory contained both locally smooth transitions and abrupt contact-level changes (**Supplementary Fig. S10**).

Ensemble-level quality control supported the interpretation that the observed contact redistribution was not driven by severe steric artifacts. Both baseline and $\beta = -2$ ensembles showed no detected clashes or severe clashes, and the minimum protein-peptide interface distance remained

comparable between the two ensembles. Meanwhile, the $\beta = -2$ ensemble occupied substantially more conformational clusters than the baseline ensemble, indicating increased ensemble diversity under perturbation (**Supplementary Fig. S11 and Supplementary Table S5**).

Thus, the 1BBZ analysis extends the framework to a polymeric peptide partner and shows that $\beta = -2$ perturbation can expand a compact baseline contact landscape into a broader set of peptide-associated contact states. The MCMC trajectory further indicates that detailed contact identities can be remodeled while part of the broader interface patch remains retained.

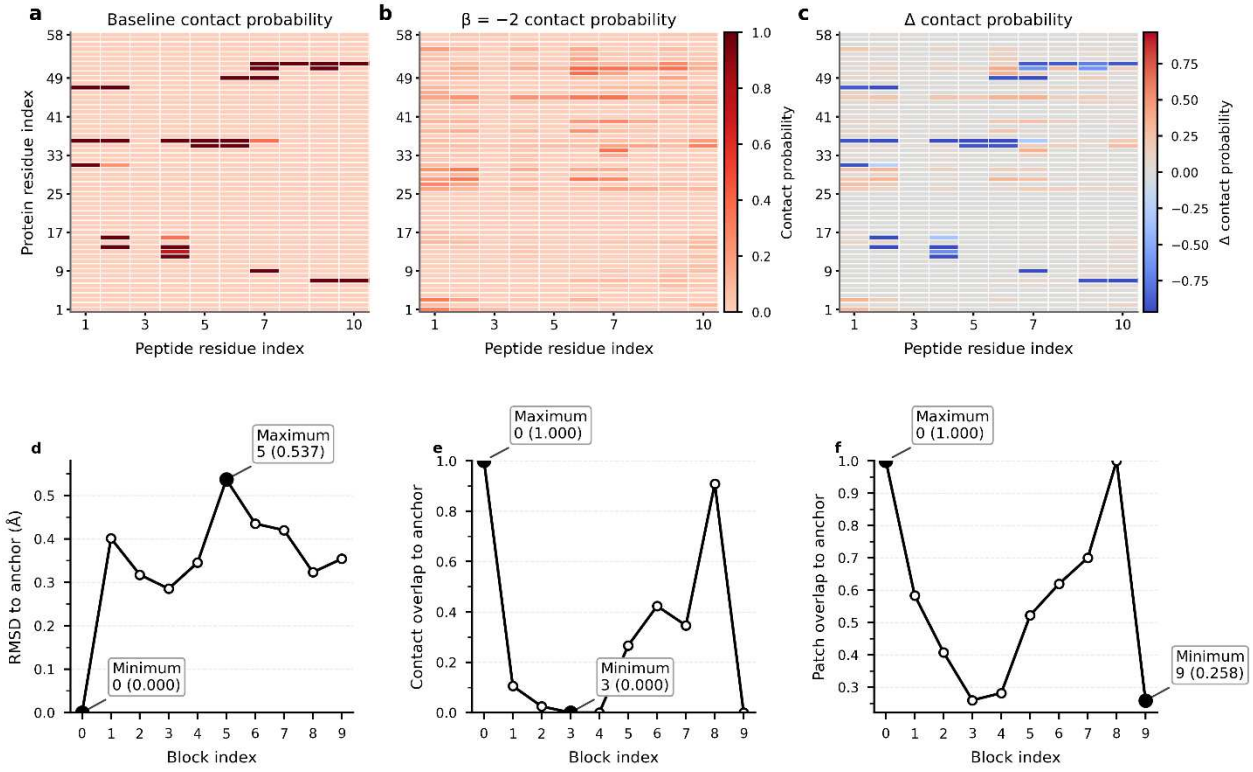


Figure 6. Interface Perturbation Extends to Protein-Peptide Complexes and Captures Peptide-Associated Interface Rearrangements. a-c, $\beta = -2$ global perturbation analysis for the protein-peptide complex 1BBZ using the total-interface protocol. **a,** Baseline protein-peptide contact probability map. **b,** $\beta = -2$ perturbed protein-peptide contact probability map. Panels a and b share the same contact-probability color scale. **c,** Δ contact probability map between the $\beta = -2$ perturbed ensemble and the baseline ensemble. **d-f,** Anchor-referenced $\beta = -2$ MCMC trajectory metrics for 1BBZ. **d,** RMSD to the anchor representative state. **e,** Contact overlap to the anchor. **f,** Patch overlap to the anchor. Δ contact probability denotes the difference between the $\beta = -2$ perturbed ensemble and the baseline ensemble. Contact overlap measures the retention of protein-peptide contact sets relative to the anchor state, whereas patch overlap measures the retention of interface-participating residue patches. The peptide was treated as a single interacting partner, following the same total-interface analysis used for protein-small-molecule complexes. Quantitative summaries are provided in Supplementary Table S5.

Discussion

In this study, ITRG shows that Boltz-2-generated biomolecular complexes can be steered and interpreted at the receptor-partner interface level by combining contact-aware pair-feature perturbation, MCMC-based path construction and descriptor-specific repulsive guidance. Across protein-small-molecule, protein-nucleic-acid and protein-peptide systems, interface perturbation redistributed contact patterns, revealed alternative interface states and generated interpretable conformational trajectories, whereas repulsion-guided sampling provided an additional mechanism for expanding selected diversity dimensions. Overall, these analyses suggest that interface-level control can broaden the exploration of generated biomolecular complexes while remaining broadly compatible with post-generation clash and minimum-distance analyses.

The protein-small-molecule analyses showed that the effect of interface perturbation is both perturbation-strength-dependent and system-dependent. In 7UX8, increasing the perturbation strength from $\beta = -1$ to $\beta = -2$ reduced contact-map Jaccard similarity from 0.50 to 0.32. Relative to the baseline ensemble, $\beta = -2$ also increased the effective number of contact modes from 5 to 9, with occupancy entropy rising from 1.43 to 2.77. In 4W53, the response was even more discrete: Jaccard similarity decreased from 0.92 at $\beta = -1$ to 0.00 at $\beta = -2$ while the effective mode count increased from 1 in the baseline ensemble to 6 under $\beta = -2$. These changes occurred without increasing severe clash counts in the evaluated ensembles, indicating that the expanded contact exploration was not simply driven by severe steric failure.

The MCMC trajectories further showed that different systems can respond to perturbation through distinct interface-remodeling modes. In 7UX8, MCMC sampling revealed partial contact redistribution under minimal ligand-internal RMSD. Adjacent representative RMSD ranged from 0.075 to 0.258 Å, and the maximum anchor-referenced RMSD reached only 0.222 Å, indicating that the sampled representatives retained highly similar ligand internal geometry. Nevertheless, contact and patch overlap relative to the anchor decreased to 0.462 and 0.500, respectively, showing that part of the original interface topology was remodeled despite limited coordinate-level variation. Thus, 7UX8 is better described as a contact-retentive remodeling trajectory, in which residue-level contacts can be partially redistributed while the broader interface footprint remains partly preserved. In contrast, 4W53 showed a sharper contact-switching trajectory: anchor-

referenced RMSD remained below 0.008 Å, but contact overlap could drop to 0.000 and patch overlap to 0.071 under the distance-based contact definition. This contrast is consistent with the distinct interface properties of the two systems. The larger and more polar L-tyrosine interface in 7UX8 may maintain a broader interaction footprint while allowing exchange of individual residue-level contacts, whereas the smaller hydrophobic toluene interface in 4W53 may be more sensitive to small coordinate changes near the contact cutoff. These observations emphasize that small RMSD should not be interpreted as unchanged interface topology; contact overlap and patch overlap provide complementary information about residue-level interface remodeling.

The MCMC-based path-construction module and the repulsion-guided sampling module serve complementary purposes. Contact-guided MCMC sampling does not directly impose a diversity objective during generation; instead, it samples β-perturbed pair-feature states and organizes the resulting Boltz-2 ensembles into contact-resolved interface trajectories. By contrast, repulsion-guided Boltz-2 sampling acts within the generation process itself, using descriptor-specific repulsive potentials to reduce local redundancy and expand targeted conformer-, pose- or environment-level diversity modes. In 7UX8, conformer-guided sampling increased global and nearest-neighbor conformer diversity by approximately 1.9-fold and 2.3-fold, respectively, whereas environment-guided sampling increased global and nearest-neighbor environment diversity by approximately 6.9-fold and 3.1-fold. In 4W53, the diversification effect was larger, with conformer-guided sampling increasing global and nearest-neighbor conformer diversity by approximately 82.7-fold and 57.7-fold, and environment-guided sampling increasing global environment diversity by approximately 118.1-fold. These results suggest that conformer guidance primarily expands ligand-internal geometry, whereas environment guidance preferentially diversifies the local binding context. The minimum ligand-protein distance analysis further indicated that these diversity gains were not mainly driven by widespread steric overlap, although environment-guided sampling occasionally introduced local clashes, with clash-containing structures observed in 2% of 7UX8 and 9% of 4W53 generations.

These cross-interface analyses support the extensibility of the receptor-partner formulation to protein-nucleic-acid and protein-peptide complexes. For nucleic-acid systems, decomposing contacts into nucleobase and sugar-phosphate backbone components enabled base- and backbone-resolved interpretation of perturbation responses. The DNA and RNA examples showed different

remodeling patterns, with 1TGH displaying more localized contact redistribution and partial trajectory-level contact retention, whereas 2ERR showed broader redistribution and stronger anchor-relative remodeling. For the protein-peptide complex 1BBZ, perturbation substantially broadened the contact landscape, increasing the number of non-zero contact-map entries from 29 in the baseline ensemble to 192 under $\beta = -2$, with 97 entries showing detectable redistribution using $|\Delta P| \geq 0.10$. These examples suggest that the framework is not restricted to small-molecule ligands when receptor and partner components can be explicitly defined.

Several limitations remain. First, the current study used representative examples rather than large-scale benchmarking, so the dependence of perturbation behavior on ligand size, pocket flexibility, interface polarity and partner type requires further evaluation. Second, contact-based metrics depend on distance thresholds, which can amplify small coordinate changes into discrete contact changes, especially in compact interfaces such as 4W53. Third, Boltz-2-generated complexes do not include explicit water molecules, ions or solvent-mediated interactions, which may be particularly important for nucleic-acid recognition. Finally, increased diversity does not necessarily imply improved physical stability, binding accuracy or functional relevance. Clash counts and minimum interface distances provide first-pass diagnostics for excluding widespread steric artifacts in the generated ensembles. Future applications that aim to assign thermodynamic weights, estimate binding stability or make experimentally testable mechanistic claims may further combine ITRG with relaxation, energy evaluation, molecular dynamics or experimental data. Despite these limitations, this work suggests that interface-level perturbation and diversity steering can make generative biomolecular complex modeling more interpretable, tunable and potentially extensible to broader classes of molecular interfaces.

Methods

β-Parameterized Interface Perturbation and Boltz-2 Ensemble Generation

To enable interface-centered conformational exploration across different biomolecular complex types, we constructed a β-parameterized perturbation framework on top of the pair features used by Boltz-2 (**Fig. 1A-C**). The framework was designed to perturb protein-partner pairwise interface representations before downstream structure generation, thereby allowing controlled expansion of the interface conformational landscape.

For each complex, the molecular system was first converted into the corresponding Boltz-2 input representation to obtain a pair-feature tensor Z (**Fig. 1A**). For protein-small-molecule complexes, the input consisted of protein residues and a small-molecule representation; for protein-peptide complexes, the input consisted of protein and peptide residues; and for protein-nucleic-acid complexes, the input consisted of protein residues and nucleic-acid residues. These inputs were processed through the Boltz-2 feature-extraction module to generate pairwise features describing protein-partner relationships.

An interfacial contact matrix M was derived from the corresponding 3D complex structure (**Fig. 1B**). Each matrix entry $M_{i,u}$indicates whether protein residue i and partner-side interaction unit u form an interface contact:

$$M_{i,u} = f(x) = \begin{cases} 1, & dmin(i,u) \leq 4.5 \\ 0, & otherswise \end{cases},$$

The partner-side interaction unit denotes ligand-heavy-atoms for protein-small-molecule complexes, a peptide residue for protein-peptide complexes and a nucleic-acid residue for protein-nucleic-acid complexes. A binary contact was assigned when the minimum distance between any heavy-atom pair, excluding hydrogen atoms, from protein residue i and partner-side interaction unit u was within the contact cutoff 4.5 Å. The contact matrix M provides a structure-derived interface mask that specifies which pair-feature entries are subject to perturbation.

The original pair features Z were updated according to the contact matrix M to generate β-dependent perturbed pair features Z_0 (**Fig. 1C**). Specifically, we applied an element-wise perturbation of the form

$$Z_0 = Z \odot (1.0 + \beta * M),$$

where $\odot$ denotes element-wise multiplication with broadcasting over feature channels, and the parameter $\beta \in \mathbb{R}$ was treated as a dimensionless interface-steering coefficient. It controls the magnitude and direction of the modulation applied to contact-associated pair-feature entries. It provides an inference-time handle for probing how Boltz-2-generated ensembles respond to controlled perturbation of receptor-partner interface representations. In this formulation, contact-associated pair features are selectively reweighted, whereas non-contact pair features remain unchanged. The baseline condition corresponds to Boltz-2 sampling without additional interface perturbation, whereas $\beta = -1$ and $\beta = -2$ were selected as two representative negative perturbation regimes. $\beta = -1$ was used to examine whether moderate contact-aware feature modulation could induce measurable contact redistribution while retaining partial baseline topology, whereas $\beta = -2$ was used to test whether stronger modulation could expose alternative contact modes and trajectory-level interface remodeling. These values were not intended to represent optimized or universally transferable settings. Instead, they were used to evaluate whether ITRG produces strength-dependent interface responses under moderate and stronger perturbation. A more exhaustive β sweep may be useful in future applications where β needs to be tuned for a specific system or design objective. In the subsequent MCMC analysis, the selected perturbed feature state was denoted as Z_0 and used as the anchor feature state.

For each β condition, the corresponding pair-feature representation was passed into the Boltz-2 structure generation module to produce a conformational ensemble. The baseline and perturbed ensembles were then compared using residue-level contact probability maps. For each generated structure s under perturbation β, a binary contact matrix $C_{\beta,s}$ was calculated using the same heavy-atom contact criterion:

$$C_{\beta,s}(i,u) = \begin{cases} 1, & dmin_{\beta,s}(i,u) \leq 4.5 \\ 0, & otherwise \end{cases},$$

where i denotes a protein residue, u denotes a partner-side interaction unit, $dmin_{\beta,s}(i,u)$ is the minimum heavy-atom distance between protein residue i and partner-side interaction unit u in structure s under perturbation β.

The ensemble-level contact probability under perturbation β condition was calculated by averaging binary contact matrices across all generated structures:

$$P_\beta(i,u) = \frac{1}{N_\beta}\sum_{s=1}^{N_\beta} C_{\beta,s}(i,u),$$

where N_β is the number of generated structures under perturbation condition β.

For the baseline ensemble, contact probability $P_{base}(i,u)$ was calculated analogously. Perturbation-induced contact redistribution was quantified as:

$$\Delta P_\beta(i,u) = P_\beta(i,u) - P_{base}(i,u),$$

Positive $\triangle P_\beta(i,u)$ values indicate contacts enriched in the β -perturbed ensemble, whereas negative values indicate contacts reduced relative to the baseline ensemble.

For each baseline or β-perturbed ensemble, binary contact matrices were flattened into contact fingerprints for contact-mode analysis. Contact fingerprints were compared using Jaccard distance, and greedy representative-based clustering was performed with a distance threshold of 0.30. Ensemble-level contact-mode diversity was summarized by cluster occupancy, occupancy entropy and effective cluster count, where the effective cluster count was defined as the number of clusters with occupancy of at least 5%. Steric and interface-distance diagnostics were calculated from protein-partner heavy-atom distances, including clash count using a 2.0 Å cutoff, severe clash count using a 1.5 Å cutoff and the minimum interface heavy-atom distance. These metrics were reported together with contact probability and Δ contact probability maps to compare baseline and β-perturbed ensembles in terms of contact redistribution, contact-mode occupancy and interface-distance quality. New and lost contacts were defined by threshold-based changes in contact activity. Thresholds of 0.1, 0.3 and 0.5 were used for multi-threshold contact-shift analysis. Detailed definitions and clustering procedures are provided in **Supplementary Method 1**.

MCMC-Based Pair-Feature Sampling and Conformational Trajectory Construction

To sample local variations around the β-perturbed anchor feature state while restricting proposals to interface-associated regions, we performed contact-guided Markov chain Monte Carlo (MCMC) sampling in the pair-feature space. Starting from the anchor pair features Z_0 , contact-guided MCMC sampling was performed in the pair-feature space to generate a sequence of interface feature states, denoted as $\{Z_1, Z_2, \dots, Z_r\}$. Here, Z_0 acts as the anchor feature state, and each MCMC step proposes a local perturbation of contact-associated pair-feature entries. Sampling was implemented as a masked block Gaussian Metropolis-Hastings procedure: The interface pair mask was first expanded and symmetrized to match the pair-feature tensor. At each step, a local block centred on a randomly selected active interface pair was perturbed with Gaussian noise. The proposed perturbation was then symmetrized and clipped to a predefined range before Metropolis-Hastings acceptance. The MCMC energy consisted of an anchor-restraint term that penalizes squared deviations from Z_0, together with an optional fast scoring term for additional guidance. Detailed definitions of contact-guided MCMC sampling procedures are provided in **Supplementary Method 2**.

Each retained pair-feature state Z_i was independently passed into the Boltz-2 structure prediction module to generate a conformational ensemble associated with that feature state. For each Z_i −associated ensemble, a representative conformation was selected according to a medoid-based criterion after excluding clash outliers when present. The selected representatives were then ordered according to the MCMC sampling sequence to assemble a representative conformational trajectory. This procedure converts a single perturbed feature state into a sequential series of structure ensembles and representative interface states, enabling trajectory-level exploration rather than independent ensemble sampling alone.

The sampled MCMC trajectory was converted into a block-wise representative trajectory before structural and contact-level evaluation. For each MCMC block, pairwise RMSD and contact-overlap matrices were calculated among all generated structures. Clash outliers were excluded when present, and the block representative was selected as the medoid structure with the lowest mean pairwise RMSD, using clash count as a secondary criterion. Mean pairwise RMSD was calculated as the average matched heavy-atom RMSD of ligand chain between all pairs of

predicted conformations within the same MCMC block. The representative conformation from block 0 was defined as the anchor state.

For a representative conformation k, the interface-contact set C_k was defined as the set of protein residue-partner interaction-unit pairs whose minimum inter-chain heavy-atom distance was within the contact cutoff (4.5 Å):

$$C_k = \{(r,u) \mid dmin(r,u) \leq 4.5\},$$

where r denotes a protein residue, u denotes a partner-side interaction unit. The partner interaction unit denotes the ligand-associated contact unit for protein-small-molecule complexes, the peptide residue for protein-peptide complexes and the nucleic-acid residue for protein-nucleic-acid complexes.

The corresponding interface patch P_k was then derived by projecting C_k onto individual interface-participating units:

$$P_k^{prot} = \{r \mid \exists u, (r,u) \in C_k\},$$

$$P_k^{part} = \{u \mid \exists r, (r,u) \in C_k\},$$

$$P_k = P_k^{prot} \cup P_k^{part},$$

Contact overlap evaluates the preservation of specific pairwise interface contacts, whereas patch overlap evaluates the preservation of the broader interface-participating region. Both contact overlap and patch overlap were calculated using the Jaccard index,

$$\mathrm{J}(A,B) = |A \cap B| / |A \cup B|.$$

Anchor-referenced metrics were then defined as $\mathrm{RMSD}(R_k, R_0)$, $\mathrm{J}(C_k, C_0)$ and $\mathrm{J}(P_k, P_0)$, where R_0, C_0 and P_0 denote the anchor representative structure, contact set and interface patch, respectively. Adjacent-block metrics were calculated analogously between consecutive representatives, including $\mathrm{RMSD}(R_k, R_{k-1})$, $\mathrm{J}(C_k, C_{k-1})$ and $\mathrm{J}(P_k, P_{k-1})$. Thus, anchor-referenced metrics quantify cumulative deviation from the initial interface state, whereas adjacent-block metrics quantify local stepwise transitions along the MCMC trajectory.

For protein-nucleic-acid complexes, the same definitions were additionally applied to nucleobase-associated and backbone-associated contact subsets.

Repulsion-Guided Boltz-2 Sampling for Conformer, Pose, and Environment Diversification

To promote structural diversity during Boltz-2 complex generation, we introduced a set of batch-level repulsion-guided sampling potentials. In contrast to conventional structure-level restraints that act independently on each generated complex, these potentials couple multiple samples generated in the same diffusion batch. For a batch of B complexes with atomic coordinates $\{x_i\}_{i=1}^{B}$, each complex was first mapped to a mode-specific descriptor $\phi(x_i)$, where m denotes conformer, pose or binding-environment diversity. A pairwise distance $d_m(x_i, x_j)$ was then computed between descriptors, and a flat-bottom repulsive energy was applied only when two samples were closer than a predefined diversity margin δ_m:

$$E_i^m = \frac{1}{B-1}\sum_{j\neq i}\left[max\left(0, \delta_m - d_m(x_i, x_j)\right)\right]^2,$$

This formulation penalizes near-duplicate samples while leaving sufficiently separated samples unaffected. The resulting energy gradients were computed with respect to atomic coordinates and used to guide the diffusion sampling trajectory.

Ligand conformer diversity was quantified using an internal-distance descriptor ϕ_{conf}. For each generated complex, matched ligand heavy atoms were used to construct the upper triangle of the intraligand distance matrix, and the resulting vector was normalized by the square root of its dimensionality. Pairwise Euclidean distances between these normalized internal-distance descriptors D_{conf} were used to quantify conformer diversity. Because this descriptor depends only on intraligand distances, it is invariant to global translation and rotation. Details of the ligand conformer-distance descriptor D_{conf} are provided in **Supplementary Method 3**.

For pose-guided sampling, diversity was defined in a protein-aligned coordinate frame. For each pair of generated complexes, interface protein atoms were aligned using a batched Kabsch[52,53] rigid transformation. The ligand coordinates of one complex were then transformed into the protein

frame of the other, and the ligand RMSD after protein alignment was used as D_{pose}. This design encourages different ligand binding poses relative to the protein binding site, rather than merely separating complexes by global rigid-body displacement.

For environment-guided sampling, we constructed a dynamic-shell environment descriptor ϕ_{env} to capture the local protein environment surrounding the ligand. Interface protein atoms were used as the candidate environment pool, and dynamic ligand-centred shells were generated from current protein-ligand distances using smooth radial gates with radii of 3.0, 4.5 and 6.0 Å. The descriptor summarized shell occupancy, shell-wise distance moments, protein atom-type composition within each shell, protein-ligand pair-type contact features and ligand-type-resolved nearest-environment features. The final descriptor was normalized by its dimensionality, and pairwise descriptor distances were used as D_{env}. This mode was designed to diversify the local binding environment sampled around the ligand, complementing conformer-level and pose-level repulsion. Details of the dynamic-shell environment descriptor are provided in **Supplementary Method 4**.

For each diversity mode, the repulsive potential was applied at regular intervals during diffusion sampling with a time-dependent guidance weight. The guidance schedule was designed to be stronger during early and intermediate sampling stages, where structural exploration is most effective, and weaker during late-stage refinement, where excessive repulsion may destabilize local geometry. Mode-specific diversity margins were used to activate repulsion only when two samples were closer than the desired separation in conformer, pose or environment descriptor space. The detailed margins, guidance intervals and weight schedules used for the three modes are summarized in **Supplementary Table S6**.

These three repulsion-guided sampling modes provide complementary control over Boltz-2 generation. Conformer guidance expands ligand-internal geometry, pose guidance separates ligand placements in a protein-aligned frame, and environment guidance diversifies the chemically informed binding context around the ligand. The generated ensembles were subsequently evaluated using matched conformer, pose and environment diversity metrics.

To examine whether each repulsive guidance strategy promoted the intended mode of diversity, we performed a post-generation pairwise analysis for each sampling strategy. Generated complexes were compared in a pairwise manner. Three complementary diversity metrics were

calculated. Ligand conformer diversity was measured using D_{conf}. Ligand pose diversity was quantified as D_{pose}. Protein-ligand environment diversity was evaluated by D_{env}, which summarize protein atom-type composition, distance distributions, and contact patterns within predefined radial shells.

For each metric, we reported the mean pairwise distance to quantify global diversity and the mean nearest-neighbor distance to assess sample de-duplication. Diversity gains were further computed relative to the baseline generation strategy using both absolute differences and multiplicative ratios.

Cross-Interface Extension to Protein-Nucleic-Acid and Protein-Peptide Complexes

To evaluate the generality of the interface-perturbation framework beyond protein-small-molecule complexes, we extended the same receptor-partner formulation to protein-nucleic-acid and protein-peptide systems. For each complex, receptor and partner chains were explicitly specified, and cross-interface contacts were defined as receptor-partner residue pairs containing at least one heavy-atom contact within the predefined contact cutoff (4.5 Å). Interface patches were defined as the receptor-side and partner-side residues involved in these contacts.

The same interface-steering and MCMC path generation procedure was applied to peptide and nucleic-acid complexes using a steering strength of $\beta = -2$. For protein-nucleic-acid complexes, the nucleic-acid partner was decomposed into nucleobase and sugar-phosphate backbone components, enabling total-interface, base-associated and backbone-associated contact analyses. For protein-peptide complexes, the peptide chain was treated as a single interacting partner and analyzed using the total-interface protocol.

The baseline and $\beta = -2$ perturbed ensembles were compared using contact-probability maps and Δ contact-probability maps. MCMC-generated trajectories were further analyzed by comparing block representative structures relative to the anchor state and between adjacent trajectory blocks. The trajectory-level metrics included representative RMSD, contact overlap and patch overlap. For nucleic-acid complexes, contact-overlap analyses were additionally performed using base- and

backbone-resolved contact definitions. Ensemble-level distance and clash diagnostics were calculated using the same cutoff definitions as in the protein-small-molecule analyses.

This unified design allowed the same perturbation, MCMC path construction and ensemble-level analysis to be applied consistently across protein-small-molecule, protein-nucleic-acid and protein-peptide complexes.

Representative Test Systems and Selection Rationale

To evaluate the algorithmic robustness of ITRG across distinct biophysical interaction modalities, we designed a targeted evaluation panel of five experimentally resolved receptor-partner complexes, rather than conducting blind large-scale benchmarking. This system panel was designed to cover distinct interface-perturbation scenarios, including contact-rich polar small-molecule binding, compact hydrophobic pocket binding, double-stranded protein-DNA recognition, single-stranded protein-RNA recognition and modular protein-peptide interaction.

The two protein-small-molecule systems were used as the main contrastive cases. MfnG O-methyltransferase bound to L-tyrosine (PDB 7UX8) was selected as the primary small-molecule case because the corresponding MfnG-ligand setting is distributed as a protein-ligand example in the Boltz repository, making it a natural starting point for evaluating an inference-time extension of Boltz-2. This system also represents a polar, contact-rich ligand-binding interface: L-tyrosine is an amino-acid-like ligand with multiple polar functional groups and can form several directional protein-ligand contacts. Thus, 7UX8 provides a chemically structured interface for evaluating contact-aware pair-feature perturbation in a multi-contact small-molecule binding site.

T4 lysozyme L99A bound to toluene (PDB 4W53) was selected as a complementary small-molecule case because it represents a different binding regime. Toluene is a small hydrophobic ligand buried in an engineered cavity, and the resulting protein-ligand interface is compact, contact-sparse and dominated by close-range hydrophobic packing rather than directional polar interactions. This makes 4W53 suitable for evaluating interface perturbation in a low-contact hydrophobic pocket, where residue-level contact assignment can be sensitive to small local geometric changes. Together, 7UX8 and 4W53 provide two complementary protein-small-

molecule test cases for protein-small-molecule interfaces, stressing the framework across two biophysically distinct extremes: a polar, contact-rich environment governed by directional interactions, and a compact, hydrophobic cavity dominated by cutoff-sensitive steric packing.

To examine whether the same receptor-partner formulation can be extended beyond small molecules, we selected one DNA complex, one RNA complex and one peptide complex with distinct partner organizations. Human TATA-binding protein bound to TATA-sequence DNA (PDB 1TGH) was selected as the protein-DNA case because it represents a canonical double-stranded DNA recognition interface. The DNA partner is an extended polymer composed of repeated nucleotide units, with chemically distinct nucleobase and sugar-phosphate backbone components. This makes 1TGH suitable for evaluating protein-DNA contact maps in which perturbation responses can be separated into base-associated and backbone-associated channels.

The Fox-1 RNA-binding domain bound to UGCAUGU RNA (PDB 2ERR) was selected as the protein-RNA case because it represents a compact sequence-specific RNA-recognition interface involving a short single-stranded RNA motif. Unlike the double-stranded DNA partner in 1TGH, the RNA partner in 2ERR is single-stranded and organized around a compact RNA-binding surface. This provides a distinct nucleic-acid case for evaluating RNA-specific contact organization and nucleobase- versus backbone-associated perturbation responses.

The Abl-SH3 domain bound to a designed high-affinity peptide ligand (PDB 1BBZ) was selected as the protein-peptide case because SH3-peptide recognition is a representative modular domain-peptide interaction. The peptide partner is a short peptide chain composed of amino-acid residues, with contacts distributed along a residue-level peptide partner rather than concentrated in a small-molecule pocket or partitioned into nucleic-acid base/backbone channels. This makes 1BBZ suitable for evaluating the total-interface protocol on a flexible polymeric peptide partner and for comparing detailed contact rearrangement with broader interface-patch retention.

Data availability

The generated structural ensembles supporting this study are available at Zenodo under https://doi.org/10.5281/zenodo.21126111. The deposited data include the generated complex ensembles (7UX8, 4W53, 1TGH, 2ERR, 1BBZ) and the processed outputs required to support the contact-map, trajectory, diversity, steric and interface-distance analyses reported in the manuscript. Additional analysis files are available from the corresponding authors upon reasonable request.

Code availability

The code for ITRG is available at https://github.com/lfs119/ITRG.git. The repository provides the inference implementation and scripts for interface perturbation, MCMC-based feature sampling, repulsion-guided Boltz-2 sampling and post-generation structural analysis. ITRG is an independent research implementation and does not represent an official release of the Boltz-2 developers.

References

1. Jumper, J. *et al.* Highly accurate protein structure prediction with AlphaFold. *Nature* **596**, 583–589 (2021).

2. Abramson, J. *et al.* Accurate structure prediction of biomolecular interactions with AlphaFold 3. *Nature* **630**, 493–500 (2024).

3. Krishna, R. et al. Generalized biomolecular modeling and design with RoseTTAFold All-Atom. *Science* **384**, eadl2528 (2024).

4. Chai Discovery *et al.* Chai-1: Decoding the molecular interactions of life. *bioRxiv* 2024.10.10.615955 (2024) doi:10.1101/2024.10.10.615955.

5. Wohlwend, J. *et al.* Boltz-1 Democratizing Biomolecular Interaction Modeling. *bioRxiv* 2024.11.19.624167 (2024) doi:10.1101/2024.11.19.624167.

6. Passaro, S. *et al.* Boltz-2: Towards Accurate and Efficient Binding Affinity Prediction. *bioRxiv* 2025.06.14.659707 (2025) doi:10.1101/2025.06.14.659707.

7. Corso, G., Stärk, H., Jing, B., Barzilay, R. & Jaakkola, T. DiffDock: Diffusion Steps, Twists, and Turns for Molecular Docking. In International Conference on Learning Representations (2023).

8. Qiao, Z., Nie, W., Vahdat, A., Miller, T. F. & Anandkumar, A. State-specific protein–ligand complex structure prediction with a multiscale deep generative model. *Nature Machine Intelligence* **6**, 195–208 (2024).

9. Lu, W. *et al.* DynamicBind: predicting ligand-specific protein-ligand complex structure with a deep equivariant generative model. *Nature Communications* **15**, 1071 (2024).

10. Wang, Y. *et al.* Protein Conformation Generation via Force-Guided SE(3) Diffusion Models. *Proc. Mach. Learn. Res.* **235**, 56835–56859 (2024).

11. Lewis, S. *et al.* Scalable emulation of protein equilibrium ensembles with generative deep learning. *Science* **389**, eadv9817 (2025).

12. Boehr, D. D., Nussinov, R. & Wright, P. E. The role of dynamic conformational ensembles in biomolecular recognition. *Nat. Chem. Biol.* **5**, 789–796 (2009).

13. Henzler-Wildman, K. & Kern, D. Dynamic personalities of proteins. *Nature* **450**, 964–972 (2007).

14. Hammes, G. G., Chang, Y. C. & Oas, T. G. Conformational selection or induced fit: A flux description of reaction mechanism. *Proc. Natl. Acad. Sci. U. S. A.* **106**, 13737–13741 (2009).

15. Changeux, J. P. & Edelstein, S. Conformational selection or induced fit? 50 years of debate resolved. *F1000 Biol. Rep.* **3**, 19 (2011).

16. Kalakoti, Y. & Wallner, B. AFsample3: Generating and selecting multiple conformational states with AlphaFold3. *bioRxiv* 2026.01.16.699904 (2026) doi:10.64898/2026.01.16.699904.

17. Bryant, P. & Noé, F. Structure prediction of alternative protein conformations. *Nature Communications* **15**, 7328 (2024).

18. Wei, H. & McCammon, J. A. Structure and dynamics in drug discovery. *npj Drug Discovery* **1**, 1 (2024).

19. Rohs, R. *et al.* The role of DNA shape in protein–DNA recognition. *Nature* **461**, 1248–1253 (2009).

20. London, N., Raveh, B. & Schueler-Furman, O. Peptide docking and structure-based characterization of peptide binding: from knowledge to know-how. *Curr. Opin. Struct. Biol.* **23**, 894–902 (2013).

21. Mitra, R. *et al.* Geometric deep learning of protein–DNA binding specificity. *Nature Methods* **21**, 1674–1683 (2024).

22. Gromiha, M. M. & Harini, K. Protein-nucleic acid complexes: Docking and binding affinity. *Curr. Opin. Struct. Biol.* **90**, 102955 (2025).

23. Dai, X., Wang, R. & Zhang, Y. Topological deep learning for enhancing peptide-protein complex prediction. *Communications Chemistry* **8**, 347 (2025).

24. Salentin, S., Schreiber, S., Haupt, V. J., Adasme, M. F. & Schroeder, M. PLIP: fully automated protein–ligand interaction profiler. *Nucleic Acids Res.* **43**, W443–W447 (2015).

25. Adasme, M. F. *et al.* PLIP 2021: expanding the scope of the protein–ligand interaction profiler to DNA and RNA. *Nucleic Acids Res.* **49**, W530–W534 (2021).

26. Bouysset, C. & Fiorucci, S. ProLIF: a library to encode molecular interactions as fingerprints. *Journal of Cheminformatics* **13**, 72 (2021).

27. Da, C. & Kireev, D. Structural Protein–Ligand Interaction Fingerprints (SPLIF) for Structure-Based Virtual Screening: Method and Benchmark Study. *J. Chem. Inf. Model.* **54**, 2555–2561 (2014).

28. Yin, S., Mi, X. & Shukla, D. Leveraging machine learning models for peptide–protein interaction prediction. *RSC Chem. Biol.* **5**, 401–417 (2024).

29. Buttenschoen, M., Morris, G. M. & Deane, C. M. PoseBusters: AI-based docking methods fail to generate physically valid poses or generalise to novel sequences. *Chem. Sci.* **15**, 3130–3139 (2024).

30. Škrinjar, P., Eberhardt, J., Tauriello, G., Schwede, T. & Durairaj, J. Have protein-ligand co-folding methods moved beyond memorisation? *bioRxiv* 2025.02.03.636309 (2025) doi:10.1101/2025.02.03.636309.

31. Xu, S. *et al.* Benchmarking all-atom biomolecular structure prediction with FoldBench. *Nature Communications* **17**, 442 (2026).

32. Masters, M. R., Mahmoud, A. H. & Lill, M. A. Investigating whether deep learning models for co-folding learn the physics of protein-ligand interactions. *Nature Communications* **16**, 8854 (2025).

33. Stein, R. A. & McHaourab, H. S. SPEACH_AF: Sampling protein ensembles and conformational heterogeneity with AlphaFold2. *PLoS Comput. Biol.* **18**, e1010483 (2022).

34. Wayment-Steele, H. K. *et al.* Predicting multiple conformations via sequence clustering and AlphaFold2. *Nature* **625**, 832–839 (2024).

35. Wallner, B. AFsample: improving multimer prediction with AlphaFold using massive sampling. *Bioinformatics* **39**, btad573 (2023).

36. Kalakoti, Y. & Wallner, B. AFsample2 predicts multiple conformations and ensembles with AlphaFold2. *Communications Biology* **8**, 373 (2025).

37. Monteiro da Silva, G., Cui, J. Y., Dalgarno, D. C., Lisi, G. P. & Rubenstein, B. M. High-throughput prediction of protein conformational distributions with subsampled AlphaFold2. *Nature Communications* **15**, 2464 (2024).

38. Del Alamo, D., Sala, D., McHaourab, H. S. & Meiler, J. Sampling alternative conformational states of transporters and receptors with AlphaFold2. *Elife* **11**, e75751 (2022).

39. Suzuki, S. & Amagasa, T. Steering Conformational Sampling in Boltz-2 via Pair Representation Scaling. *bioRxiv* 2026.01.23.701250 (2026) doi:10.64898/2026.01.23.701250.

40. Yuen, H. *et al.* Metadiffusion: inference-time meta-energy biasing of biomolecular diffusion models. *bioRxiv* 2026.02.10.704873 (2026) doi:10.64898/2026.02.10.704873.

41. Corso, G., Xu, Y., de Bortoli, V., Barzilay, R. & Jaakkola, T. Particle Guidance: non-I.I.D. Diverse Sampling with Diffusion Models. *12th International Conference on Learning Representations, ICLR 2024* https://arxiv.org/pdf/2310.13102 (2023).

42. Bansal, A. *et al.* Universal Guidance for Diffusion Models. *IEEE Computer Society Conference on Computer Vision and Pattern Recognition Workshops*, 843–852 (2023).

43. Nam, J. *et al.* Enhancing Diffusion-Based Sampling with Molecular Collective Variables. Preprint at https://arxiv.org/pdf/2510.11923 (2025).

44. Maddipatla, S. A. *et al.* Inverse Problems with Experiment-Guided AlphaFold. Preprint at https://mlanthology.org/icml/2025/maddipatla2025icml-inverse/ (2025).

45. Metropolis, N., Rosenbluth, A. W., Rosenbluth, M. N., Teller, A. H. & Teller, E. Equation of State Calculations by Fast Computing Machines. *J. Chem. Phys.* **21**, 1087–1092 (1953).

46. Hastings, W. K. Monte Carlo sampling methods using Markov chains and their applications. *Biometrika* **57**, 97–109 (1970).

47. Wu, K. L. *et al.* Expanding the eukaryotic genetic code with a biosynthesized 21st amino acid. *Protein Sci.* **31**, e4443 (2022).

48. Merski, M., Fischer, M., Balius, T. E., Eidam, O. & Shoichet, B. K. Homologous ligands accommodated by discrete conformations of a buried cavity. *Proc. Natl. Acad. Sci. U. S. A.* **112**, 5039–5044 (2015).

49. Juo, Z. S. *et al.* How proteins recognize the TATA box. *J. Mol. Biol.* **261**, 239–254 (1996).

50. Auweter, S. D. *et al.* Molecular basis of RNA recognition by the human alternative splicing factor Fox-1. *EMBO J.* **25**, 163–173 (2006).

51. Pisabarro, M. T., Serrano, L. & Wilmanns, M. Crystal structure of the Abl-SH3 domain complexed with a designed high-affinity peptide ligand: Implications for SH3-ligand interactions. *J. Mol. Biol.* **281**, 513–521 (1998).

52. Kabsch, W. A solution for the best rotation to relate two sets of vectors. *Acta Crystallographica Section A* **32**, 922–923 (1976).

53. Kabsch, W. A discussion of the solution for the best rotation to relate two sets of vectors. *Acta Crystallographica Section A* **34**, 827–828 (1978).

Acknowledgements

This work was supported by Macao Science and Technology Development Fund (0030/2024/RIA1) and Macao Polytechnic University (RP/FCA-15/2023).

Author contributions

K.X. conceived the study, designed the framework, performed the computational experiments, analysed and interpreted the data, prepared the figures, and wrote the manuscript. Y.T. and B.H. contributed to manuscript revision and editing after the initial draft. J.Z. contributed to the preparation of Fig. 1. H.L. and X.Y. supervised the study, provided guidance, and revised the manuscript. All authors reviewed and approved the final manuscript.

Competing interests

The authors declare no competing interests.

Supplementary Information

Supplementary Fig. S1

Ensemble-level interface-contact redistribution in 7UX8 under $\beta = -1$.

Supplementary Fig. S2

Ensemble-level interface-contact redistribution in 7UX8 under $\beta = -2$.

Supplementary Fig. S3

Ensemble-level interface-contact redistribution in 4W53 under $\beta = -1$.

Supplementary Fig. S4

Ensemble-level interface-contact redistribution in 4W53 under $\beta = -2$.

Supplementary Fig. S5

Intra-block ensemble consistency and representative quality assessment.

Supplementary Fig. S6

Structural visualization of selected representative protein-ligand interface states.

Supplementary Fig. S7

Minimum ligand-protein heavy-atom distance distributions for generated complexes.

Supplementary Fig. S8

Adjacent-block MCMC trajectory metrics for protein-nucleic-acid complexes.

Supplementary Fig. S9

Ensemble-level quality control for $\beta = -2$ protein-nucleic-acid perturbation analyses.

Supplementary Fig. S10

Adjacent-block MCMC trajectory metrics for the protein-peptide complex 1BBZ.

Supplementary Fig. S11

Ensemble-level quality control for $\beta = -2$ protein-peptide perturbation analysis.

Supplementary Table S1

intra-block ensemble consistency and representative diagnostics.

Supplementary Table S2

Summary statistics used for baseline-normalized diversity gains in Fig. 4a.

Supplementary Table S3

Summary of ligand-protein minimum heavy-atom distance and clash analysis.

Supplementary Table S4

Global Δ contact probability summary for $\beta = -2$ protein-nucleic-acid perturbation.

Supplementary Table S5

Summary statistics for $\beta = -2$ protein-peptide perturbation analysis of 1BBZ.

Supplementary Table S6

Parameter settings for repulsion-guided Boltz-2 sampling.

Supplementary Method 1

Contact-mode clustering, contact-shift analysis and steric/interface-distance diagnostics.

Supplementary Method 2

Contact-guided MCMC sampling of pair features.

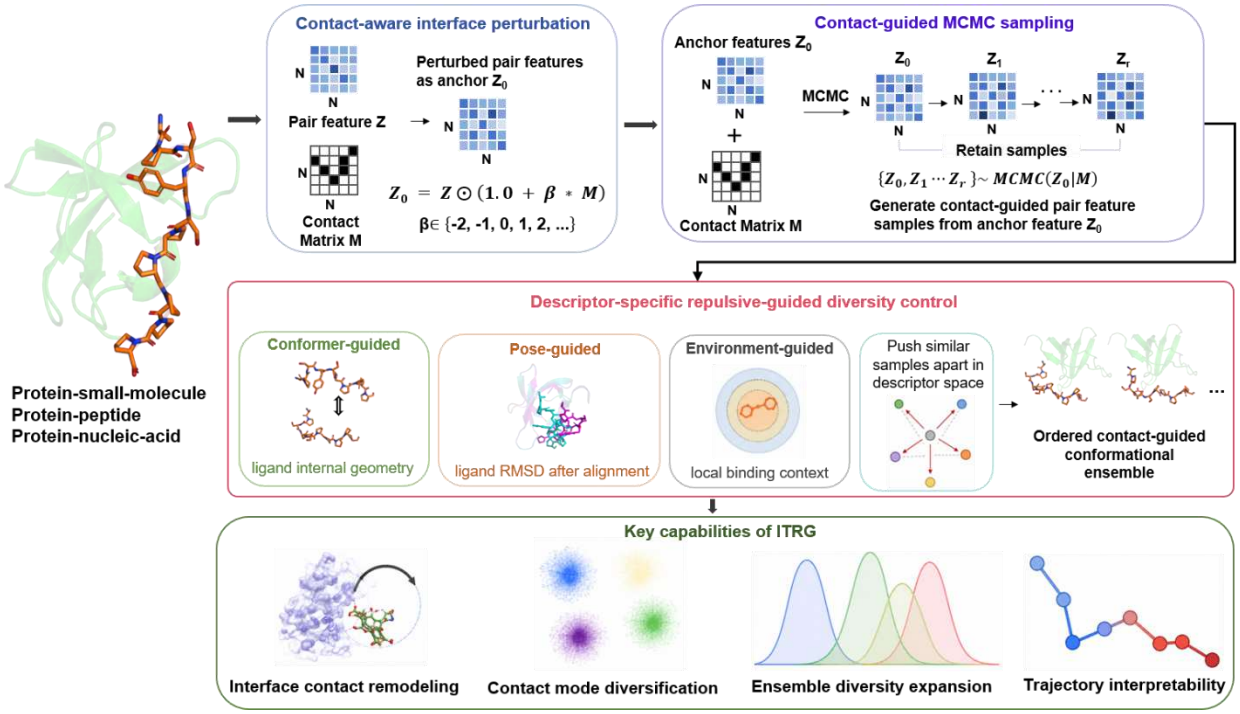
Supplementary Method 3

Ligand conformer-distance calculation.

Supplementary Method 4

Dynamic-shell environment descriptor for binding-environment diversity.

TOC



ITRG is a unified inference-time framework that enables controllable exploration of alternative biomolecular interface states with Boltz-2. ITRG integrates contact-aware interface perturbation, MCMC-based interface path construction, and descriptor-specific repulsive guidance to generate diverse and interpretable ensembles of protein–small-molecule, protein–peptide, and protein–nucleic-acid complexes. Comprehensive analyses show substantial interface remodeling, contact-mode diversification, ensemble expansion, and trajectory interpretability, offering a general strategy for studying and designing biomolecular interactions.

Supplementary Files

This is a list of supplementary files associated with this preprint. Click to download.

- ITRGsupplementaryinformation.docx