ALLOSHIFT: FINITE-DOSE POPULATION-DIRECTED BINDER DESIGN

Anonymous authors
Paper under double-blind review

ABSTRACT

Protein function often depends on the populations of different conformational states. Binders can alter this balance by preferentially binding one state, but state preference alone does not determine the effect at a feasible concentration: a selective binder may be too weak to engage the target, while a potent binder may bind both states. We introduce ALLOSHIFT, linking state-resolved affinities, baseline populations, and free binder concentration to an exact finite-dose rule for maximizing a desired state or matching a target population. Q_{bind} learns state-resolved affinities from sparse and censored measurements and supplies scores for binder screening and peptide design. Across three measured-affinity panels, affinity-only screening misses the optimum in 70.4% of separating one-binder tasks; the balance of potency and selectivity changes with dose. Q_{bind} improves receptor-cold affinity and state-direction prediction, and its scores improve model-evaluated peptide guidance and reranking. Together, these results provide a population-aware design objective and a learned route to candidate prioritization.

1 INTRODUCTION

Modern generative protein design can produce high-affinity proteins, peptides, and macrocycles from target structures or sequences (Watson et al., 2023; Pacesa et al., 2025; Rettie et al., 2025; Chen et al., 2026a; Bhat et al., 2025; Tang et al., 2025; Chen et al., 2025). For many therapeutic targets, the functional effect of binding depends on which conformational states a binder stabilizes. Kinases and receptors interconvert among states whose populations influence activity (Motlagh et al., 2014; Xie et al., 2020), and conformation-specific binders can alter signaling (Rasmussen et al., 2011; Adams et al., 2025). Recent methods account for structural diversity. For example, AlloGen uses differential structural compatibility to guide state-selective binder design (Cao et al., 2026b), while TD3B directs generation toward specified state transitions (Cao et al., 2026a). These advances motivate a practical question of how a binder's state preference translates into population shifts at achievable doses.

Consider two binders with the same affinity ratio between states but a 100-fold difference in absolute affinity. Under a two-state equilibrium model, they approach the same population bias at saturation, yet only one may appreciably engage the target at a feasible dose. Thermodynamic linkage connects binding to conformational populations (Monod et al., 1965; Wyman, 1964). A state preference specifies the direction of the shift; selectivity determines its limiting endpoint; binding strength and the available free concentration determine whether that endpoint can be approached. These are distinct requirements for selecting an intervention. Learning the relevant affinities is difficult because measurements span heterogeneous assays and chemistries, are often censored, and rarely cover the same binder in multiple states.

We introduce ALLOSHIFT as a finite-dose population-control formulation. Given state-resolved K_d values, a baseline population, and a dose ceiling, its equilibrium decision layer selects a binder to maximize a desired state or to approach a specified target population; in the latter task, it also determines reachability and the required dose. This objective also defines the prediction problem: learned affinities should improve which binder is tested under a limited budget. Q_{bind} uses binder chemistry, biochemical context, and state-specific receptor structures to predict K_d and auxiliary K_i measurements without docked complexes. Its scores inform candidate selection through the finite-dose rule and guide peptide generation. We make three contributions:

Reviewers: please read the Reviewer Guidelines (iclr.cc/Conferences/2027/ReviewerGuidelines) and the AI Policy for Reviewers (iclr.cc/Conferences/2027/AIPolicyForReviewers). If you used AI to expand, edit, or polish your review, please provide the input text to the LLM. Better still, consider skipping the LLM and submitting your original text: we, and the authors, are much more interested in your unedited thoughts than in what an LLM has to say. AI-assisted or not, you are putting your name and reputation behind your review: LLM-generated falsehoods, hallucinations or misrepresentations are subject to disciplinary action, which may include desk-rejecting all papers you have authored.

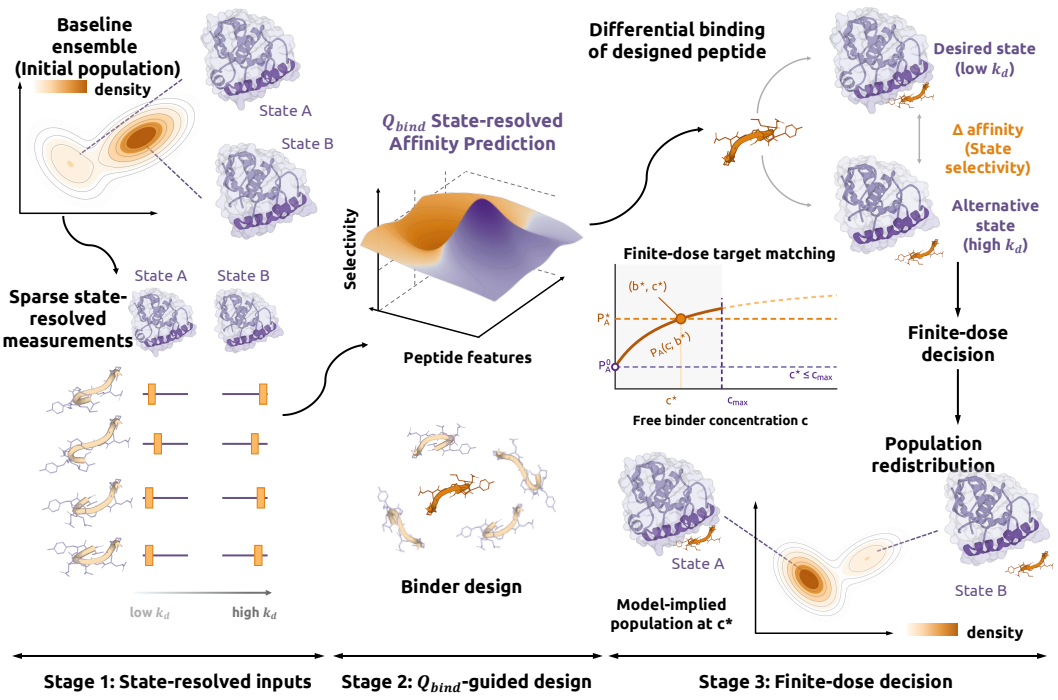

Figure 1: **Overview of AlloShift.** State-resolved affinity measurements train Q_{bind}, whose predictions guide candidate generation and ranking. A finite-dose decision layer then selects binder–dose interventions: selectivity sets the saturation population, while the absolute-affinity scale controls the concentration needed to approach it.

- **Population-aware formulation of binder design beyond affinity optimization.** We formulate finite-dose population control as a binder–concentration objective and characterize its exact reachable frontier. On measured affinity landscapes, it selects different interventions from affinity-only optimization.
- **Dose-dependent decisions and accuracy requirements.** We show when engagement and state selectivity change the preferred proxy, and use controlled errors to quantify when the population-aware advantage persists.
- **Learned signals for low-budget candidate prioritization.** Q_{bind} improves receptor-cold affinity and state-direction prediction and supplies scores for finite-dose candidate screening, guided peptide generation, and reranking.

2 Related Work

State-selective binder design. Protein and peptide generators optimize binding to chosen targets through diffusion, structure prediction, and language-model objectives (Watson et al., 2023; Pacesa et al., 2025; Stark et al., 2025; Rettie et al., 2025; Chen et al., 2026a; Bhat et al., 2025; Chen et al., 2026b; Tang et al., 2025). Multistate and negative-design methods extend this objective by comparing desired and competing conformations (Leaver-Fay et al., 2011; Davey & Chica, 2012). Recent approaches further connect conformational preference with function. Designed GPCR binders stabilize distinct receptor states (Cheng et al., 2026; Muratspahić et al., 2026), TD3B optimizes directional transitions (Cao et al., 2026a), and AlloGen learns conformation-selective scores for reranking and guided generation (Cao et al., 2026b). Building on this emphasis on state preference, AlloShift learns state-resolved affinities from experiments and accounts for the dose required to shift conformational populations.

Experimental affinity learning. Affinity predictors use sequence, molecular graphs, and structural information to estimate binding strength, as demonstrated by GraphDTA, PRODIGY, ProAffinity-GNN, Boltz-2, and PeptiVerse (Nguyen et al., 2021; Xue et al., 2016; Zhou et al., 2024; Passaro et al., 2025; Zhang et al., 2026). Predicting state preference additionally requires comparing affinities across conformations, often with censored observations and few matched measurements for the same binder. These comparisons place different demands on a model because accurate absolute affinities can still yield incorrect differences between states. Accordingly, Q_{bind} separates shared affinity contributions from state-specific effects, and we evaluate absolute-affinity prediction and state direction independently.

Guidance and population objectives. Inference-time guidance steers sampling toward specified conditions or properties (Dhariwal & Nichol, 2021; Ho & Salimans, 2022), with extensions to discrete generators and lookahead decoding (Nisonoff et al., 2025; Li et al., 2025). ProteinGuide supports protein property guidance (Xiong et al., 2026), while AlloGen and TD3B guide binder generation using state-aware objectives (Cao et al., 2026b;a). Connecting these objectives to population changes requires accounting for conformational equilibrium. Ensemble generators and dynamical models describe conformational distributions and transitions (Lewis et al., 2025; Mardt et al., 2018; Wu & Noé, 2020), while classical allostery and linkage models relate state populations to ligand binding (Monod et al., 1965; Wyman, 1964). ALLOSHIFT draws on this equilibrium description to connect learned state-specific affinities with both directional generation and dose-dependent binder selection.

3 PROBLEM FORMULATION

3.1 STATE-RESOLVED BINDING AFFINITY

A binder's affinity for each receptor state determines both its preference between states and the concentration required for appreciable binding. For binder b and state $s \in \mathcal{S}$, express the dissociation constant on a common logarithmic scale:

$$q_s(b) = \log_{10}\left(\frac{K_{d,s}(b)}{1\,\mathrm{M}}\right), \tag{1}$$

so smaller q_s indicates tighter binding. For states A and B, write

$$\Delta(b; A, B) = q_A(b) - q_B(b), \qquad \mu(b; A, B) = \frac{q_A(b) + q_B(b)}{2}. \tag{2}$$

Here $\Delta > 0$ favors B, with $10^{\Delta} = K_{d,A}/K_{d,B}$, while μ captures the shared affinity scale. Binders with identical state preference can therefore differ in the dose required to act.

3.2 FINITE-DOSE POPULATION RESPONSE

Thermodynamic linkage connects binding to conformational populations (Monod et al., 1965; Wyman, 1964). Consider interconverting states A and B with ligand-free population P_B^0 in state B. With reversible 1:1 binding at free binder concentration c, the total population in state B, including bound and unbound receptors, is

$$P_B(c; b) = \frac{P_B^0\,(1 + c/K_{d,B}(b))}{(1 - P_B^0)\,(1 + c/K_{d,A}(b)) + P_B^0\,(1 + c/K_{d,B}(b))}. \tag{3}$$

The numerator weights the baseline state-B population by its binding contribution; the denominator normalizes over both states. At saturation, binding multiplies the baseline population odds by 10^{Δ}. Selectivity sets the limiting population, while the common affinity scale controls the dose needed to approach it.

3.3 POPULATION-DIRECTED DESIGN

We study two interventions using the same equilibrium response. Desired-state maximization seeks the highest attainable population of state s; target matching seeks a specified population, which can require a dose below the ceiling. Both select a binder and a feasible free concentration, unlike

affinity-only selection (Leaver-Fay et al., 2011; Davey & Chica, 2012; Cao et al., 2026b). For library $\mathcal{B}$, target population vector $\mathbf{P}^\star$, and ceiling $c_{\max}$, target matching is

$$(b^\star, c^\star) = \arg \min_{b \in \mathcal{B},\ 0 \le c \le c_{\max}} D(\mathbf{P}(c; b), \mathbf{P}^\star)\,, \tag{4}$$

Here $\mathbf{P}(c; b)$ is the population vector and D measures discrepancy; for two states, $D = |P_B(c; b) - P_B^\star|$. Desired-state maximization instead uses $-P_s(c; b)$ with the same binder–dose constraint. Evaluating the attainable response for each candidate determines whether the target is reachable and, if not, identifies the closest intervention.

4 METHODS

4.1 Q_{bind} STATE-RESOLVED AFFINITY MODEL

The affinity measurements connect relatively few binders to multiple receptor states, so binder–state interactions are weakly identified. We use a restricted additive readout to combine binder and biochemical-context information with state-specific structural features, and examine its identifiability in Appendix A.

4.1.1 REPRESENTATIONS AND READOUT

Each binder is encoded from monomer-level physicochemical, stereochemical, and modification descriptors, including descriptors for noncanonical monomers. A shared encoder produces a pooled representation $\bar{\mathbf{p}}_b$, and $\mathbf{d}_b$ summarizes its composition. Each receptor state combines ESM-2 features (Lin et al., 2023), residue identity, local structural descriptors, and a C_α spatial graph into summaries $\bar{\mathbf{r}}_s$ and $\boldsymbol{\sigma}_s$. The biochemical measurement context is encoded separately as $\mathbf{c}_s$.

For measurement type $a \in \{K_d, K_i\}$, the predicted log affinity adds a binder-and-context contribution to a structural-state contribution:

$$\hat{q}_{s,a}(b) = f_a^{\text{ctx}}(\bar{\mathbf{p}}_b, \mathbf{c}_s, \mathbf{d}_b) + g_a^{\text{state}}(\bar{\mathbf{r}}_s, \boldsymbol{\sigma}_s)\,. \tag{5}$$

The outputs use $\log_{10}(K/1\,\text{M})$ for their respective measurement types. We call the context-only variant Q_{bind}-*Affinity* and the additive variant Q_{bind}-*State*. Within either K_d or K_i, state preference is

$$\hat{\Delta}_a(b; A, B) = \hat{q}_{A,a}(b) - \hat{q}_{B,a}(b), \tag{6}$$

with positive values favoring B.

Binder information enters absolute-affinity prediction through f_a^{ctx}, while the structural contribution is shared by binders in the same state. When measurement contexts match across states, the context contribution cancels in Eq. 6, so predicted Δ cannot rank binders within that pair (Section 5.6). The K_d and K_i heads remain separate: K_i provides auxiliary supervision, whereas the population analyses use measured K_d. The predicted K_d values can also be passed to the finite-dose decision layer for candidate screening.

4.1.2 LEARNING OBJECTIVE

The heterogeneous observations provide different levels of supervision. Exact affinities use Huber regression in log space (Huber, 1964); censored measurements define admissible intervals, including one-sided bounds at detection limits. Matched measurements of both states constrain $\hat{\Delta}_a$, while direction-only comparisons constrain its sign. The combined objective is

$$\mathcal{L}_Q = \mathcal{L}_{K_d} + \mathcal{L}_{K_i} + \mathcal{L}_\Delta + 0.25\mathcal{L}_{\text{dir}} + 10^{-3}\mathcal{L}_{\text{state}}. \tag{7}$$

$\mathcal{L}_{K_d}$ and $\mathcal{L}_{K_i}$ average their respective absolute-measurement streams: exact rows use a one-log-unit Huber threshold, and censored rows incur a half-squared penalty outside the admissible bound. $\mathcal{L}_\Delta$ averages squared violations of measured difference intervals; $\mathcal{L}_{\text{dir}}$ averages $\text{softplus}(-y\hat{\Delta}/\tau)$ over determined, non-near-zero comparisons, with $\tau = 1$. $\mathcal{L}_{\text{state}}$ regularizes the structural outputs of the K_d and K_i heads. No pairwise ranking term is active; empty eligible streams contribute zero. Appendix B specifies exact reductions and training settings.

4.2 FINITE-DOSE POPULATION CONTROL

Unless specified otherwise, population calculations use measured state-resolved K_d and declared baseline populations. For $0 < P_B^0 < 1$ and positive, finite dissociation constants, Eq. 3 separates into a limiting population and a concentration scale:

$$P_B(c; b) = P_B^0 + \left(P_B^\infty(b) - P_B^0\right) \frac{c}{c + c_{50}(b)}, \tag{8}$$

where

$$\begin{aligned} P_B^\infty(b) &= \frac{P_B^0\, 10^{\Delta(b;A,B)}}{1 - P_B^0 + P_B^0\, 10^{\Delta(b;A,B)}}, \\ c_{50}(b) &= \frac{K_{d,A}(b) K_{d,B}(b)}{(1 - P_B^0) K_{d,B}(b) + P_B^0 K_{d,A}(b)}. \end{aligned} \tag{9}$$

P_B^∞ sets the selectivity-dependent endpoint, whereas c_{50} is the free concentration at half of the binder-induced population change. Adding a common offset η to both log affinities multiplies c_{50} by 10^η without changing P_B^∞. Thus binders can have the same saturation effect but differ at a feasible concentration.

For a target strictly between baseline and saturation, the required free concentration is

$$c_{\mathrm{req}}(b) = c_{50}(b) \frac{P_B^\star - P_B^0}{P_B^\infty(b) - P_B^\star}. \tag{10}$$

The required dose depends on both the response scale and the remaining distance from saturation. The baseline is attained at zero concentration; a nonconstant response only approaches saturation as $c \to \infty$. Targets beyond the limiting population or opposite the response direction are unreachable. Equal state affinities produce no population shift at any dose.

A finite dose ceiling restricts each binder's reachable populations to

$$\mathcal{I}_b = \left[\min\{P_B^0, P_B(c_{\max}; b)\}, \max\{P_B^0, P_B(c_{\max}; b)\}\right]. \tag{11}$$

Monotonicity means this interval contains every population attainable between zero concentration and the ceiling. For a fixed binder, projecting $P_B^\star$ onto $\mathcal{I}_b$ minimizes absolute target error: the projection uses zero dose, the ceiling, or the interior dose obtained by inverting Eq. 8. Comparing the resulting errors across binders solves Eq. 4 exactly.

The target is reachable by the library precisely when it lies in $\bigcup_{b \in \mathcal{B}} \mathcal{I}_b$. Every candidate interval contains P_B^0, so the union is the interval between the smallest and largest ceiling responses. Outside that interval, the nearest endpoint determines the best attainable intervention. This construction uses the stated baseline, free-dose ceiling, reversible equilibrium, and state-resolved affinities; it does not by itself measure a post-binding population.

For learned one-binder screening, we rank candidates by the same finite-dose objective after substituting Q_{bind}'s predicted state-specific K_d values. Regret is evaluated at the selected binder's best dose under *measured* affinities, isolating binder selection from dose-prediction error.

4.3 Q_{bind}-GUIDED GENERATION AND PRIORITIZATION

We use frozen Q_{bind} scores both during pretrained PepTune sampling and for post-hoc selection of completed candidates. For requested direction $y \in \{-1, +1\}$, the guidance score is

$$J_\Delta(b) = y\, \hat{\Delta}_a(b; A, B), \tag{12}$$

where $y = +1$ favors state B. We compare this score with native sampling and desired-state affinity guidance. Eight-step lookahead favors high-scoring continuations. Separate guide and reporting ensembles limit direct score reuse, but the reporting evaluator informed development screening and guidance-strength selection. Its direction scores therefore assess candidate prioritization under learned evaluators rather than physical validation. Appendix B specifies the budgets, lookahead, and scoring protocol.

5 EXPERIMENTS

5.1 EVALUATION OVERVIEW

Datasets. The affinity corpus supplies K_d and K_i observations for training and receptor-cold affinity testing; paired comparisons test state-preference direction. Three measured-affinity panels support finite-dose selection under declared baseline populations, and 48 target–direction cells test peptide guidance and reranking. Table 1 summarizes their sizes and roles; Appendix Table A1 details the paired data.

Table 1: Evaluation datasets, sizes, and experimental tasks.

Dataset	Size	Task
Affinity corpus	2,919 measurements; 719 receptors; 1,992 binders	Affinity prediction
Paired state comparisons	117 comparisons	State-direction prediction
Measured-affinity panels	3 panels; 28 binders	Finite-dose selection
Generation targets	48 target–direction cells; 3 seeds	Guidance and reranking

Metrics. Affinity prediction uses $\log_{10} K_d$ mean absolute error; paired comparisons use balanced direction accuracy. Finite-dose screens report Regret@K; desired-state tasks also report expected population and gain recovered, while target matching reports population error and reachability. Generation reports evaluator-assigned direction success, validity, and diversity. Appendix C.1 defines aggregation and ties. Population values derive from measured affinities and declared baselines.

Baselines. Decision baselines include random, measured-affinity potency and selectivity, their combination, and the exact objective. Frozen-predictor screens compare Q_{bind} with learned affinity and selectivity rankings; selected binders receive their best measured-affinity dose. Generation compares native and guided PepTune with post-hoc selection and RFpeptides/PepMLM pools (Appendix B).

5.2 POPULATION CONTROL REQUIRES OPTIMIZATION BEYOND AFFINITY

At a one-binder screening budget, affinity-only selection loses achievable desired-state population relative to selectivity-aware selection (Figure 2a). It selects a non-optimal binder in 361 of 513 separating tasks (70.4%; Figure 2b) and reaches expected population 0.713 rather than the measured-affinity optimum of 0.752. This is a decision gap conditional on the measured affinity landscapes and declared baselines, not an independent population measurement.

The gap depends on feasible dose: affinity-only recovers 99.3% of the available gain in the lowest band but 75.6% in the highest (Figure 2c). The selectivity-only control shows why affinity cannot serve as a universal selection proxy. Additional baselines, larger screening budgets, and panel-level uncertainty are reported in Appendix G and F.

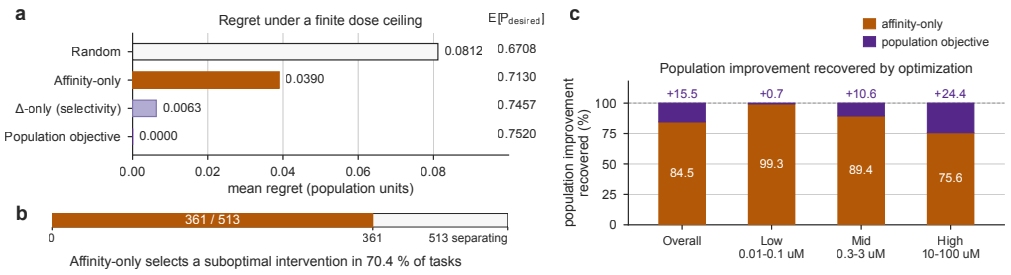

Figure 2: Desired-state selection on 513 separating finite-dose tasks: (a) Regret@1 and expected population; (b) suboptimal affinity-only choices (361/513); (c) gain recovered by affinity-only and added by the population objective, overall and across dose bands.

5.3 THE FINITE-DOSE OBJECTIVE DEFINES THE ATTAINABLE FRONTIER

For the distinct task of matching a requested population, projecting that target onto each binder's reachable interval gives the exact finite-dose decision (Section 4.2). The rule attains the optimum in all 3,956 separating target-matching tasks (Figure 3a); reachability follows from the same interval construction (Appendix E). These are implications of measured affinities and declared baselines, not independent population observations.

At limited screening budgets, selectivity-only is a better one-coordinate proxy than affinity-only in this target-matching grid, while both approach the measured-affinity optimum as the screened set grows (Figure 3a). Exact budget values and the small-library frontier appear in Appendix Tables A2 and A3.

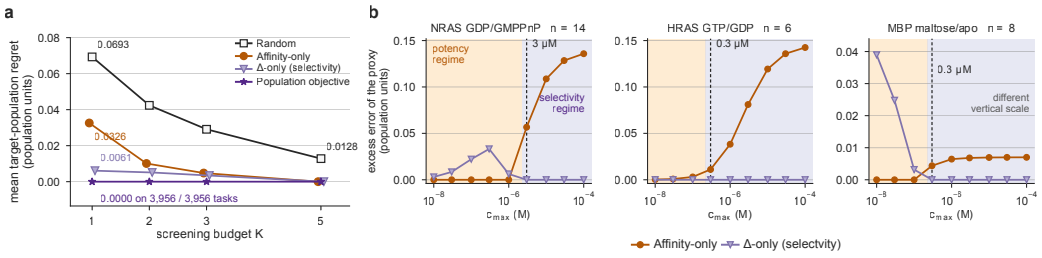

Figure 3: Target-matching decisions and dose-dependent proxy errors on measured-affinity landscapes. Panel (a) shows regret across screening budgets on 3,956 separating tasks; panel (b) shows affinity-only and selectivity-only proxy errors across dose ceilings in three measured panels.

5.4 DOSE DETERMINES THE RELEVANT THERMODYNAMIC COORDINATE

At low ceilings, engagement limits the response even for selective binders; at higher ceilings, selectivity matters more (Section 4.2). Across the three measured panels, the error of affinity-only and selectivity-only proxies changes with the dose ceiling (Figure 3b). The crossovers occur near 3 μM for NRAS and 0.3 μM for HRAS and MBP, showing that no fixed concentration defines the useful proxy across systems. The declared baseline populations and panel-specific reference scales are recorded in Appendix Table A7.

5.5 ACCURACY AND ROBUSTNESS DELIMIT USEFUL PREDICTION

Controlled perturbations show that predictive accuracy matters through the downstream decision and dose regime. The population-aware advantage over affinity-only selection persists to joint-noise scales of 0.456 log units at $K = 1$ and 0.328 at $K = 3$ on the target-matching grid (Appendix Table A4). On the separate desired-state grid, adding noise to measured selectivity progressively increases low-budget regret (Appendix Figure A1). These sensitivity analyses define decision requirements rather than a calibrated model of Q_{bind} residuals.

5.6 Q_{bind} SIGNALS FOR CANDIDATE SELECTION

On receptor-cold splits, Q_{bind} improves absolute-affinity prediction and state-direction accuracy over the corresponding controls (Table 2). Its affinity head supplies a potency-aware signal; its state head reaches 0.6195 balanced accuracy, compared with 0.5639 for a state-blind control. A state-only model reaches 0.6113, indicating that much of this direction information is associated with the receptor-state representation.

Table 2: Receptor-cold affinity prediction and paired state-direction benchmarks. Each metric is evaluated on its corresponding held-out cohort.

Affinity predictor	K_d MAE ↓
Cold constant	1.4075
Q_{bind}-Affinity	**1.3244**

Direction predictor	Balanced accuracy ↑
State-blind	0.5639
State-only control	0.6113
Q_{bind}-State	**0.6195**

Composing the frozen predictions with the finite-dose objective gives the lowest numerical one-candidate regret among the learned rankings on the 513-task desired-state benchmark (Table 3). It is 0.0397 versus 0.0529 for ranking by the strongest predicted state affinity; a separate measured desired-state affinity baseline scores 0.0390 in Figure 2. The gain here is a binder-screening result, with the selected binder's dose optimized using measured affinities. The additive readout (Section 4.1.1) does not establish binder-specific Δ ranking in matched contexts.

Table 3: One-binder screening on 513 separating desired-state tasks. Gain recovered and regret use the measured-affinity landscape; predicted rankings use frozen scores.

Selection rule	Gain recovered ↑	Regret@1 ↓
Random screen	0.594	0.0812
Predicted best-state affinity	0.673	0.0529
Predicted $\lvert\hat{\Delta}\rvert$	0.697	0.0511
Q_{bind} plug-in finite dose	**0.704**	**0.0397**
Measured $\lvert\Delta\rvert$ selector	0.924	0.0063
Exact measured-affinity objective	1.000	0.0000

5.7 Q_{bind}-GUIDED GENERATION AND RERANKING

Score-guided PepTune sampling and post-hoc Q_{bind} selection contribute complementary model-evaluated gains (Figure 4a). Native proposals score 0.500, reranking scores 0.569, and guidance followed by reranking reaches 0.613 under the original evaluator. The combined strategy also improves over native proposals under a second, architecturally disjoint evaluator. Selectivity-guided sampling raises pool validity from 0.357 to 0.407 with similar within-pool diversity (Figure 4c).

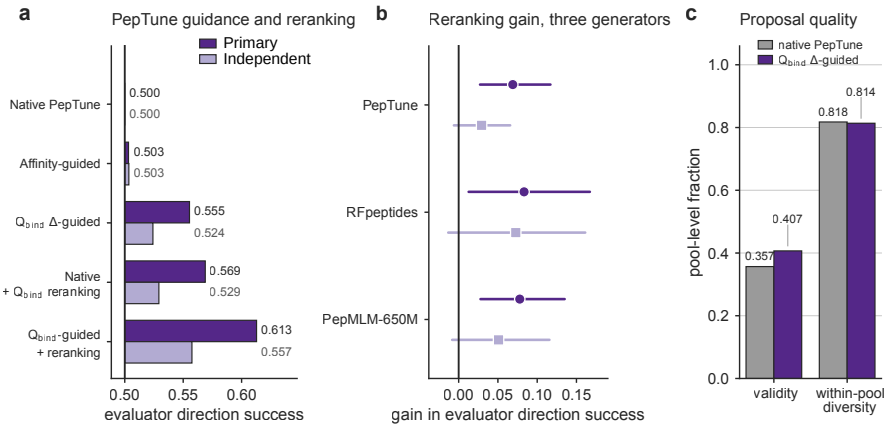

Figure 4: Q_{bind}-guided generation and candidate prioritization. Panels show PepTune direction success under primary and independent evaluators, post-hoc reranking gain across three generators with 95% target-bootstrap intervals, and PepTune proposal quality.

Reranking shows positive mean gains for RFpeptides and PepMLM as well (Figure 4b), although their proposal budgets and target coverage differ. The second-evaluator intervals for individual generator

families include zero; the combined PepTune intervention retains a positive target-bootstrap interval (Appendix B). These are evaluator-assigned outcomes rather than measured binding.

5.8 Structural illustration of candidate prioritization

We selected one 16-residue PepTune peptide, KNDSGRDGTGIDCTWG, with the frozen Q_{bind} affinity score for α-bungarotoxin and modeled the same sequence against free and peptide-engaged toxin templates (Figure 5). The toxin blocks nicotinic acetylcholine receptors; experimentally, toxin-binding antibodies can prevent receptor engagement (Chiappinelli & Zigmond, 1978; Pachner & Ricalton, 1989). A peptide binding near its receptor-recognition surface could therefore motivate competition and neutralization assays. Boltz-2 models retain each imposed binding-site geometry, but these templates do not establish a solution-state population shift. The figure's equal-budget percentiles summarize a separate 24-target cohort: Q_{bind} selections average 0.32/0.40 under two learned evaluators, versus 0.65/0.67 for generator order (lower is better). This cohort gain reverses in the rendered seed-1 pool (Appendix I). Binding, receptor competition, and post-binding population remain unmeasured for this peptide.

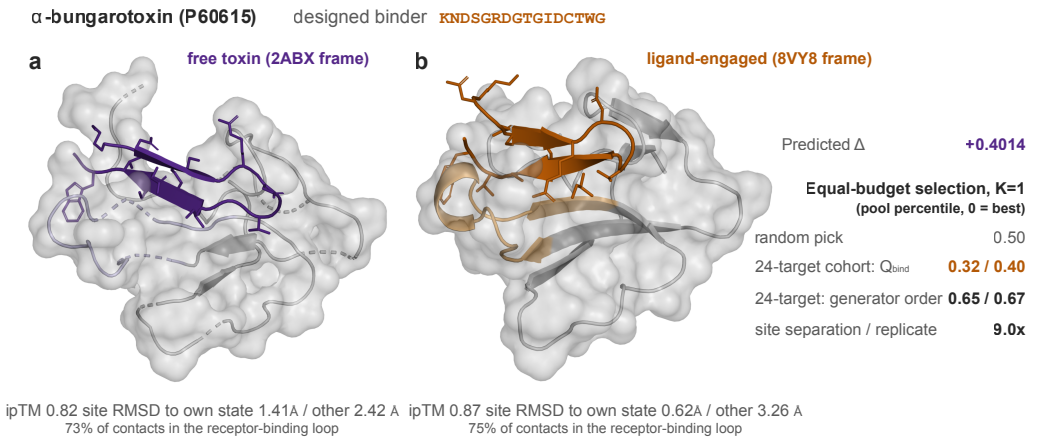

Figure 5: Computational illustration of a Q_{bind}-prioritized α-bungarotoxin peptide. The same sequence is modeled with free (2ABX) and HAP-peptide-engaged (8VY8) toxin templates; the right-hand ranking metrics refer to a separate 24-target cohort under two learned evaluators.

6 Conclusion

Population control requires a finite-dose objective beyond affinity alone. On measured-affinity landscapes, affinity-only misses the optimal binder in 70.4% of separating one-binder decisions, while the exact rule solves all 3,956 separating target-matching tasks on a distinct grid. Q_{bind} provides receptor-cold affinity and state-direction signals; its plug-in ranking and generation guidance improve numerical or model-evaluated candidate prioritization in the tested low-budget settings. The single-peptide structural illustration shows how that shortlist can be taken toward a binding assay. Resolving binder-specific state selectivity and establishing an actual population shift now require matched-state affinity measurements and independent baseline and post-binding population readouts under known free concentrations.

AI Use Statement

Generative AI tools assisted with literature discovery, hypothesis framing and critique, experiment planning, code and figure workflows, interpretation of computational diagnostics, and manuscript drafting and editing. The authors verified source records and quantitative analyses and take responsibility for the experimental design and all final scientific claims, including AI-assisted material.

Ethics Statement

This study analyzes previously reported binding measurements and computationally generated candidate sequences. It includes no human- or animal-participant experiments and makes no claim of therapeutic efficacy for the proposed binders. The reported population responses are thermodynamic calculations under stated assumptions, rather than independently observed treatment effects.

Reproducibility Statement

The population equations and decision conditions appear in Sections 3 and 4.2. Appendix B specifies the model, training, and scoring protocols; Appendix C records dataset sizes and evaluation cohorts; Appendices E and F give the finite-dose grids, baseline conventions, and sensitivity analyses. All conclusions about population response distinguish measured affinities from derived or model-scored quantities.

References

Chloe S. Adams, Hyojin Kim, Abigail E. Burtner, Dong Sun Lee, Craig Dobbins, Cameron Criswell, Brian Coventry, Adri Tran-Pearson, Ho Min Kim, and Neil P. King. De novo design of protein minibinder agonists of TLR3. *Nature Communications*, 16(1):1234, 2025. doi: 10.1038/s41467-025-56369-w. URL https://doi.org/10.1038/s41467-025-56369-w.

Suhaas Bhat, Kalyan Palepu, Lauren Hong, Joey Mao, Tianzheng Ye, Rema Iyer, Lin Zhao, Tianlai Chen, Sophia Vincoff, Rio Watson, Tian Z. Wang, Divya Srijay, Venkata Srikar Kavirayuni, Kseniia Kholina, Shrey Goel, Pranay Vure, Aniruddha J. Deshpande, Scott H. Soderling, Matthew P. DeLisa, and Pranam Chatterjee. De novo design of peptide binders to conformationally diverse targets with contrastive language modeling. *Science Advances*, 11(4):eadr8638, 2025. doi: 10.1126/sciadv.adr8638. URL https://doi.org/10.1126/sciadv.adr8638.

Hanqun Cao, Aastha Pal, Sophia Tang, Yinuo Zhang, Jingjie Zhang, Pheng Ann Heng, and Pranam Chatterjee. TD3B: Transition-directed discrete diffusion for allosteric binder generation. In *Forty-third International Conference on Machine Learning*, 2026a. URL https://openreview.net/forum?id=RNuC8Nj6rD.

Hanqun Cao, Zachary Quinn, Aastha Pal, Sumi Kimura, Jingjie Zhang, Pheng Ann Heng, and Pranam Chatterjee. Allogen: Conformation-selective binder generation with differential state scoring. *arXiv preprint arXiv:2606.05474*, 2026b. doi: 10.48550/arXiv.2606.05474. URL https://arxiv.org/abs/2606.05474.

Leo Tianlai Chen, Zachary Quinn, Madeleine Dumas, Christina Peng, Lauren Hong, Moises Lopez-Gonzalez, Alexander Mestre, Rio Watson, Sophia Vincoff, Lin Zhao, Jianli Wu, Audrey Stavrand, Mayumi Schaepers-Cheu, Tian Zi Wang, Divya Srijay, Connor Monticello, Pranay Vure, Rishab Pulugurta, Sarah Pertsemlidis, Kseniia Kholina, Shrey Goel, Matthew P. DeLisa, Jen-Tsan Ashley Chi, Ray Truant, Hector C. Aguilar, and Pranam Chatterjee. Target sequence-conditioned design of peptide binders using masked language modeling. *Nature Biotechnology*, 44(6):1002–1010, 2026a. doi: 10.1038/s41587-025-02761-2. URL https://doi.org/10.1038/s41587-025-02761-2.

Tong Chen, Yinuo Zhang, Sophia Tang, and Pranam Chatterjee. Multi-objective-guided discrete flow matching for controllable biological sequence design. *arXiv preprint arXiv:2505.07086*, 2025.

Tong Chen, Zachary Quinn, Kunal Mishra, Erin C O'Connor, Sophia E Silver, Yinuo Zhang, Mary Jo Valencia, Ying Mei, Jacques Behmoaras, Leonardo MR Ferreira, and Pranam Chatterjee. moPPIt: De Novo Generation of Motif-Specific and Functionally Active Peptide Binders via Discrete Flow Matching. *bioRxiv*, pp. 2026–01, 2026b.

Shizhuo Cheng, Jia Guo, Yun-li Zhou, Xumei Luo, Gufang Zhang, Ya-zhi Zhang, Yixin Yang, Jiannan Xie, Ping Xu, Dan-dan Shen, Shaokun Zang, Huicui Yang, Xuechu Zhen, Min Zhang, and Yan Zhang. De novo design of GPCR exoframe modulators. *Nature*, 651(8104):242–250, 2026. doi: 10.1038/s41586-025-09957-1. URL https://doi.org/10.1038/s41586-025-09957-1.

VyA Chiappinelli and RE Zigmond. alpha-bungarotoxin blocks nicotinic transmission in the avian ciliary ganglion. *Proceedings of the National Academy of Sciences*, 75(6):2999–3003, 1978.

James A. Davey and Roberto A. Chica. Multistate approaches in computational protein design. *Protein Science*, 21(9):1241–1252, 2012. doi: 10.1002/pro.2128. URL https://doi.org/10.1002/pro.2128.

Prafulla Dhariwal and Alexander Nichol. Diffusion Models Beat GANs on Image Synthesis. In *Advances in Neural Information Processing Systems*, volume 34, pp. 8780–8794, 2021. URL https://proceedings.neurips.cc/paper/2021/hash/49ad23d1ec9fa4bd8d77d02681df5cfa-Abstract.html.

Jonathan Ho and Tim Salimans. Classifier-Free Diffusion Guidance. *arXiv preprint arXiv:2207.12598*, 2022. doi: 10.48550/arXiv.2207.12598. URL https://arxiv.org/abs/2207.12598.

Peter J. Huber. Robust Estimation of a Location Parameter. *The Annals of Mathematical Statistics*, 35(1):73–101, 1964. doi: 10.1214/aoms/1177703732. URL https://doi.org/10.1214/aoms/1177703732.

Andrew Leaver-Fay, Ron Jacak, P. Benjamin Stranges, and Brian Kuhlman. A Generic Program for Multistate Protein Design. *PLoS ONE*, 6(7):e20937, 2011. doi: 10.1371/journal.pone.0020937. URL https://doi.org/10.1371/journal.pone.0020937.

Sarah Lewis, Tim Hempel, José Jiménez-Luna, Michael Gastegger, Yu Xie, Andrew Y. K. Foong, Victor García Satorras, Osama Abdin, Bastiaan S. Veeling, Iryna Zaporozhets, Yaoyi Chen, Soojung Yang, Adam E. Foster, Arne Schneuing, Jigyasa Nigam, Federico Barbero, Vincent Stimper, Andrew Campbell, Jason Yim, Marten Lienen, Yu Shi, Shuxin Zheng, Hannes Schulz, Usman Munir, Roberto Sordillo, Ryota Tomioka, Cecilia Clementi, and Frank Noé. Scalable emulation of protein equilibrium ensembles with generative deep learning. *Science*, 389(6761):eadv9817, 2025. doi: 10.1126/science.adv9817. URL https://doi.org/10.1126/science.adv9817.

Xiner Li, Yulai Zhao, Chenyu Wang, Gabriele Scalia, Gokcen Eraslan, Surag Nair, Tommaso Biancalani, Shuiwang Ji, Aviv Regev, Sergey Levine, and Masatoshi Uehara. Derivative-Free Guidance in Continuous and Discrete Diffusion Models with Soft Value-based Decoding. In *Advances in Neural Information Processing Systems*, volume 38, pp. 95507–95545, 2025. doi: 10.52202/085713-3194. URL https://papers.nips.cc/paper_files/paper/2025/hash/899af0d66d8850318a20781484416152-Abstract-Conference.html.

Zeming Lin, Halil Akin, Roshan Rao, Brian Hie, Zhongkai Zhu, Wenting Lu, Nikita Smetanin, Robert Verkuil, Ori Kabeli, Yaniv Shmueli, Allan dos Santos Costa, Maryam Fazel-Zarandi, Tom Sercu, Salvatore Candido, and Alexander Rives. Evolutionary-scale prediction of atomic-level protein structure with a language model. *Science*, 379(6637):1123–1130, 2023. doi: 10.1126/science.ade2574. URL https://doi.org/10.1126/science.ade2574.

Andreas Mardt, Luca Pasquali, Hao Wu, and Frank Noé. VAMPnets for deep learning of molecular kinetics. *Nature Communications*, 9(1):5, 2018. doi: 10.1038/s41467-017-02388-1. URL https://doi.org/10.1038/s41467-017-02388-1.

Jacque Monod, Jeffries Wyman, and Jean-Pierre Changeux. On the nature of allosteric transitions: A plausible model. *Journal of Molecular Biology*, 12(1):88–118, 1965. doi: 10.1016/s0022-2836(65)80285-6. URL https://doi.org/10.1016/s0022-2836(65)80285-6.

Hesam N. Motlagh, James O. Wrabl, Jing Li, and Vincent J. Hilser. The ensemble nature of allostery. *Nature*, 508(7496):331–339, 2014. doi: 10.1038/nature13001. URL https://doi.org/10.1038/nature13001.

Edin Muratspahić, David Feldman, David E. Kim, Xiangli Qu, Ana-Maria Bratovianu, Paula Rivera-Sánchez, Jan Hendrik Voss, Emil P. T. Hertz, Mads Jeppesen, Federica Dimitri, Kensuke Sakamoto, Amrita Nallathambi, Pia Peceli, Jianjun Cao, Brian P. Cary, Matthew J. Belousoff, Peter Keov, Phuc N. H. Trinh, Qingchao Chen, Yue Ren, Justyn Fine, Sudha Mishra, Annu Dalal, Shachie Sinha, Ramanuj Banerjee, Manisankar Ganguly, Karthik Varappalayam Karuppusamy, Isaac Sappington, Thomas Schlichthaerle, Jason Z. Zhang, Arvind Pillai, Brian Coventry, Ljubica

Mihaljević, Magnus S. Bauer, Susana Vázquez Torres, Amir Motmaen, Gyu Rie Lee, Long Tran, Xinru Wang, Inna Goreshnik, Dionne K. Vafeados, Justin E. Svendsen, Parisa Hosseinzadeh, Nicolai Lindegaard, Matthäus Brandt, Yann Waltenspühl, Kristine Deibler, Lukas Deweid, Anja Bennett, Jendrik Schöppe, Tiantang Dong, Xiaoli Yan, Luke Oostdyk, William Cao, Lakshmi Anantharaman, Johan J. Weisser, Jesper Frank Bastlund, Christoffer Bundgaard, Ayodeji A. Asuni, Justin G. English, Lance J. Stewart, Lauren Halloran, Jamie B. Spangler, André Lieber, Arun K. Shukla, Patrick M. Sexton, Bryan L. Roth, Brian E. Krumm, Denise Wootten, Christopher G. Tate, Christoffer Norn, and David Baker. De novo design of miniproteins targeting GPCRs. *Nature*, 656 (8129):1044–1053, 2026. doi: 10.1038/s41586-026-10656-8. URL https://doi.org/10.1038/s41586-026-10656-8.

Thin Nguyen, Hang Le, Thomas P Quinn, Tri Nguyen, Thuc Duy Le, and Svetha Venkatesh. GraphDTA: predicting drug–target binding affinity with graph neural networks. *Bioinformatics*, 37(8):1140–1147, 2021. doi: 10.1093/bioinformatics/btaa921. URL https://doi.org/10.1093/bioinformatics/btaa921.

Hunter Nisonoff, Junhao Xiong, Stephan Allenspach, and Jennifer Listgarten. Unlocking Guidance for Discrete State-Space Diffusion and Flow Models. In *The Thirteenth International Conference on Learning Representations*, 2025. URL https://openreview.net/forum?id=XsgHl54yO7.

Martin Pacesa, Lennart Nickel, Christian Schellhaas, Joseph Schmidt, Ekaterina Pyatova, Lucas Kissling, Patrick Barendse, Jagrity Choudhury, Srajan Kapoor, Ana Alcaraz-Serna, Yehlin Cho, Kourosh H. Ghamary, Laura Vinué, Brahm J. Yachnin, Andrew M. Wollacott, Stephen Buckley, Adrie H. Westphal, Simon Lindhoud, Sandrine Georgeon, Casper A. Goverde, Georgios N. Hatzopoulos, Pierre Gönczy, Yannick D. Muller, Gerald Schwank, Daan C. Swarts, Alex J. Vecchio, Bernard L. Schneider, Sergey Ovchinnikov, and Bruno E. Correia. One-shot design of functional protein binders with BindCraft. *Nature*, 646(8084):483–492, 2025. doi: 10.1038/s41586-025-09429-6. URL https://doi.org/10.1038/s41586-025-09429-6.

Andrew R Pachner and Nancy Ricalton. In vitro neutralization by monoclonal antibodies of α-bungarotoxin binding to acetylcholine receptor. *Toxicon*, 27(12):1263–1268, 1989.

Saro Passaro, Gabriele Corso, Jeremy Wohlwend, Mateo Reveiz, Stephan Thaler, Vignesh Ram Somnath, Noah Getz, Tally Portnoi, Julien Roy, Hannes Stark, David Kwabi-Addo, Dominique Beaini, Tommi Jaakkola, and Regina Barzilay. Boltz-2: Towards Accurate and Efficient Binding Affinity Prediction. *bioRxiv*, 2025. doi: 10.1101/2025.06.14.659707. URL https://doi.org/10.1101/2025.06.14.659707.

Søren G. F. Rasmussen, Hee-Jung Choi, Juan Jose Fung, Els Pardon, Paola Casarosa, Pil Seok Chae, Brian T. DeVree, Daniel M. Rosenbaum, Foon Sun Thian, Tong Sun Kobilka, Andreas Schnapp, Ingo Konetzki, Roger K. Sunahara, Samuel H. Gellman, Alexander Pautsch, Jan Steyaert, William I. Weis, and Brian K. Kobilka. Structure of a nanobody-stabilized active state of the β2 adrenoceptor. *Nature*, 469(7329):175–180, 2011. doi: 10.1038/nature09648. URL https://doi.org/10.1038/nature09648.

Stephen A. Rettie, David Juergens, Victor Adebomi, Yensi Flores Bueso, Qinqin Zhao, Alexandria N. Leveille, Andi Liu, Asim K. Bera, Joana A. Wilms, Alina Üffing, Alex Kang, Evans Brackenbrough, Mila Lamb, Stacey R. Gerben, Analisa Murray, Paul M. Levine, Maika Schneider, Vibha Vasireddy, Sergey Ovchinnikov, Oliver H. Weiergräber, Dieter Willbold, Joshua A. Kritzer, Joseph D. Mougous, David Baker, Frank DiMaio, and Gaurav Bhardwaj. Accurate de novo design of high-affinity protein-binding macrocycles using deep learning. *Nature Chemical Biology*, 21(12):1948–1956, 2025. doi: 10.1038/s41589-025-01929-w. URL https://doi.org/10.1038/s41589-025-01929-w.

Hannes Stark, Felix Faltings, MinGyu Choi, Yuxin Xie, Eunsu Hur, Timothy O'Donnell, Anton Bushuiev, Talip Uçar, Saro Passaro, Weian Mao, Mateo Reveiz, Roman Bushuiev, Tally Portnoi, Tomáš Pluskal, Josef Sivic, Karsten Kreis, Arash Vahdat, Shamayeeta Ray, Jonathan T. Goldstein, Andrew Savinov, Jacob A. Hambalek, Anshika Gupta, Diego A. Taquiri-Diaz, Yaotian Zhang, Samuel J. Snyder, A. Katherine Hatstat, Angelika Arada, Nam Hyeong Kim, Haoyu Fan, Ethel Tackie-Yarboi, Dylan Boselli, Lee Schnaider, Chang C. Liu, Gene-Wei Li, Denes Hnisz, David M.

Sabatini, William F. DeGrado, Jeremy Wohlwend, Gabriele Corso, Regina Barzilay, and Tommi Jaakkola. BoltzGen: Toward Universal Binder Design. *bioRxiv*, 2025. doi: 10.1101/2025.11.20.689494. URL https://doi.org/10.1101/2025.11.20.689494.

Sophia Tang, Yinuo Zhang, and Pranam Chatterjee. PepTune: De Novo Generation of Therapeutic Peptides with Multi-Objective-Guided Discrete Diffusion. In *Proceedings of the 42nd International Conference on Machine Learning*, volume 267 of *Proceedings of Machine Learning Research*, pp. 59017–59065. PMLR, 2025. URL https://proceedings.mlr.press/v267/tang25n.html.

Joseph L. Watson, David Juergens, Nathaniel R. Bennett, Brian L. Trippe, Jason Yim, Helen E. Eisenach, Woody Ahern, Andrew J. Borst, Robert J. Ragotte, Lukas F. Milles, Basile I. M. Wicky, Nikita Hanikel, Samuel J. Pellock, Alexis Courbet, William Sheffler, Jue Wang, Preetham Venkatesh, Isaac Sappington, Susana Vázquez Torres, Anna Lauko, Valentin De Bortoli, Emile Mathieu, Sergey Ovchinnikov, Regina Barzilay, Tommi S. Jaakkola, Frank DiMaio, Minkyung Baek, and David Baker. De novo design of protein structure and function with RFdiffusion. *Nature*, 620(7976):1089–1100, 2023. doi: 10.1038/s41586-023-06415-8. URL https://doi.org/10.1038/s41586-023-06415-8.

Hao Wu and Frank Noé. Variational Approach for Learning Markov Processes from Time Series Data. *Journal of Nonlinear Science*, 30(1):23–66, 2020. doi: 10.1007/s00332-019-09567-y. URL https://doi.org/10.1007/s00332-019-09567-y.

Jeffries Wyman, Jr. Linked Functions and Reciprocal Effects in Hemoglobin: A Second Look. In *Advances in Protein Chemistry*, volume 19, pp. 223–286. Academic Press, 1964. doi: 10.1016/s0065-3233(08)60190-4. URL https://doi.org/10.1016/s0065-3233(08)60190-4.

Tao Xie, Tamjeed Saleh, Paolo Rossi, and Charalampos G. Kalodimos. Conformational states dynamically populated by a kinase determine its function. *Science*, 370(6513):eabc2754, 2020. doi: 10.1126/science.abc2754. URL https://doi.org/10.1126/science.abc2754.

Junhao Xiong, Ishan Gaur, Maria Lukarska, Hunter Nisonoff, Luke M. Oltrogge, David F. Savage, and Jennifer Listgarten. Property guidance for protein sequence generative models with ProteinGuide. *Nature Biotechnology*, 2026. doi: 10.1038/s41587-026-03207-z. URL https://doi.org/10.1038/s41587-026-03207-z.

Li C. Xue, João Pglm Rodrigues, Panagiotis L. Kastritis, Alexandre Mjj Bonvin, and Anna Vangone. PRODIGY: a web server for predicting the binding affinity of protein–protein complexes. *Bioinformatics*, 32(23):3676–3678, 2016. doi: 10.1093/bioinformatics/btw514. URL https://doi.org/10.1093/bioinformatics/btw514.

Yinuo Zhang, Sophia Tang, Tong Chen, Elizabeth Mahood, Sophia Vincoff, and Pranam Chatterjee. PeptiVerse: A unified platform for therapeutic peptide property prediction. *Nature Communications*, 17(1):6819, 2026. doi: 10.1038/s41467-026-74167-w. URL https://doi.org/10.1038/s41467-026-74167-w.

Zhiyuan Zhou, Yueming Yin, Hao Han, Yiping Jia, Jun Hong Koh, Adams Wai-Kin Kong, and Yuguang Mu. ProAffinity-GNN: A Novel Approach to Structure-Based Protein–Protein Binding Affinity Prediction via a Curated Data Set and Graph Neural Networks. *Journal of Chemical Information and Modeling*, 64(23):8796–8808, 2024. doi: 10.1021/acs.jcim.4c01850. URL https://doi.org/10.1021/acs.jcim.4c01850.

Appendix

A Additional identifiability statements

Let $G = (\mathcal{P}, \mathcal{T}, \mathcal{E})$ denote the binder–receptor measurement graph. The following statements summarize the admissibility checks used in the main paper.

Free cell interactions. If every measured binder–receptor cell receives an unconstrained interaction parameter in addition to a receptor offset, the model has more free parameters than observations. Interaction values are not identified without additional structure.

Free binder effects. If receptor-specific likelihood terms contain disjoint binder sets, the likelihood separates into independent receptor blocks. A held-out receptor has no likelihood connection to binder effects learned on the remaining receptors. More rows within the disconnected components cannot repair this problem.

Featurized interactions. For $f(p,t) = z(p)^\top W x(t)$, the identifiable subspace of W is determined by the sum of receptor-specific tensor-product spans. Zero binder overlap can be compatible with transfer when multiple receptors have sufficient binder depth and the features connect their interaction spaces. A single deep receptor cannot identify a shared receptor interaction map regardless of its sample count.

Matched contrasts. A matched contrast between two binders measured in the same receptor context eliminates a shared receptor-level offset exactly. Ranking chemistry within a fold still requires at least two test contrasts. A test fold containing only one pair makes chemistry-only ranking metrics undefined rather than zero.

B Reproduction and measurement semantics

The frozen affinity model has separate K_d and K_i heads. Its receptor path combines frozen ESM-2 (35M) residue features (480 dimensions) with 33 geometry/identity features, projects these to 48 dimensions per residue, and applies one residual message-passing layer. The C_α graph uses 24 spatial nearest neighbours, continuous chain bonds, 16 sequence-spaced landmarks, and self edges, with 19 edge channels. Binder monomers use frozen 14-channel descriptors projected to 64 dimensions and pooled; a separate 14-channel aggregate descriptor and two 26-dimensional biochemical-context vectors enter the base readout. In the selected pooled configuration, the receptor-conditioned binder feature is identically zero: the structural contribution is additive in the binder score. The structural branch's final layer starts at zero. The reported model does not use an inverse-folding or peptide language-model branch; it ensembles 25 final-weight models from five receptor-cluster-cold folds and five seeds per fold.

Every update draws 32 absolute-affinity rows and eight complete state pairs. Each of the two absolute-head losses is the mean over its own eligible rows: exact $\log_{10}$ affinities contribute Huber loss (threshold 1); $>$ and $<$ bounds contribute half the squared one-sided violation. The pair loss is the mean squared distance of predicted Δ outside the observed interval. The direction loss averages $\mathrm{softplus}(-s\hat{\Delta})$ over sign-determined, non-near-zero pairs; an empty eligible set contributes zero. The structural penalty is the mean squared output of the two-head structural branch. The complete weights are 1 for each absolute and pair loss, 0.25 for direction, 10^{-3} for the structural penalty, and 0 for within-transition ranking (no ranking comparisons were constructed). AdamW uses learning rate 3×10^{-4} throughout, weight decay 10^{-2} for non-normalization multi-dimensional parameters and zero otherwise, dropout 0.1, and global gradient clipping at 1.0. Training runs 20 passes without early stopping and retains final weights; the sampling-stream seed is $1000s + 7$ for seed s.

The unmodified PepTune model is sampled over 128 reverse steps with sequences of length 100. In the last 96 steps, each eight-step block evaluates four independent continuations with the guide, standardizes their rewards across siblings, and resamples by $\mathrm{softmax}(2z)$. At zero guidance strength the sampler reduces exactly to its native transition. Guidance strength was selected from $\{1, 2, 4, 8\}$

on three excluded development targets. PepTune uses 200 proposals per cell across 48 target–direction cells and three seeds. Validity requires successful chemistry adaptation and finite binder features; valid candidates are reranked by the guide's signed Δ, retaining the top 25. The guide has 25 members and each reporting evaluator has five; the latter enters development screening and tuning, but neither in-loop scoring nor reranking. Native PepTune samples one unconditional pool per seed and reuses it for all targets and both directions. PepMLM uses 300 proposals per run over 48 cells and three seeds; RFpeptides supplies one run in 31 cells, mean 216 proposals. For evaluator-assigned direction success, 95% intervals use a target bootstrap. The combined guided-and-reranked PepTune strategy gains 0.0575 [0.0011, 0.1158] under the independent evaluator; the individual generator-family reranking intervals include zero.

K_i is auxiliary supervision and never enters the equilibrium equations. All three finite-dose panels use measured state-resolved K_d and declared, unmeasured P_B^0 to imply $P_B(c)$; these outputs are not independent population readouts. Each affinity record retains its construct, binder, assay condition, censoring status, and publication provenance.

C Dataset census

Table A1: Core state-selectivity datasets. Exact values and interval-determined values are listed separately because censored observations can retain ordering information.

Axis	Receptors	Binders	Cells	Exact	Determined	Shared	Primary role
RAS GDP/GMPPnP	4	31	46	24	40	5	receptor and chemistry transfer
Hsp70 ADP/ATP	9	9	39	39	39	7	transfer ceiling and external models
Hsp70 ADP/NN	6	9	17	17	17	7	supporting geometry
Hsp70 ATP/NN	6	9	15	15	15	5	supporting geometry
Calmodulin apo/holo	2	2	4	4	4	2	reserved family
10 further axes	–	–	60	26	37	–	geometry and boundary cases
Total	19	71	181	125	154	–	

C.1 Evaluation metrics and aggregation

Absolute-affinity error is the mean absolute difference between predicted and observed $\log_{10}(K_d/1\,\mathrm{M})$ on exact held-out measurements. Balanced state-direction accuracy is one half the sum of sensitivity and specificity for the sign of the paired state-affinity difference; the receptor-cold evaluation holds out receptor clusters.

For a desired-state task t, define the measured-affinity utility $U_t(b) = \max_{0\le c\le c_{\max,t}} P_{s_t}(c;b)$ and the full-pool optimum $U_t^\star = \max_{b\in\mathcal{B}_t} U_t(b)$. For a requested-population task, define the attainable error $E_t(b) = \min_{0\le c\le c_{\max,t}} |P_B(c;b) - P_{B,t}^\star|$ and $E_t^\star = \min_{b\in\mathcal{B}_t} E_t(b)$. If a ranking screens its top K binders $S_{t,K}$, the two forms of regret are

$$\mathrm{Regret}_t\,@K = \begin{cases} U_t^\star - \max_{b\in S_{t,K}} U_t(b), & \text{desired-state maximization,} \\ \min_{b\in S_{t,K}} E_t(b) - E_t^\star, & \text{target matching.} \end{cases} \tag{13}$$

For learned rankings, predictions order the shortlist; measured affinities determine the best dose and the reported regret. Exact score ties are averaged over uniform permutations before the top-K shortlist is evaluated. The two task families have separate declared grids: desired-state maximization uses both requested states, 19 baseline populations, and nine dose ceilings; target matching uses 19 baselines, 19 target populations, and nine ceilings. We retain only tasks with unequal candidate outcomes when reporting separating-task regret, yielding 513 and 3,956 tasks, respectively; these are repeated decisions on three biological panels, not independent systems.

For desired-state maximization, expected population is the mean shortlisted utility. The fraction of achievable gain recovered averages $(\max_{b\in S_{t,K}} U_t(b) - P_{s_t}^0)/(U_t^\star - P_{s_t}^0)$ over tasks with positive

attainable gain. Both metrics are averaged within panel before averaging the three panels; the fraction is computed per task rather than as a ratio of aggregate gains. For target matching, Hit@0.05 denotes the fraction of tasks with shortlisted absolute population error at most 0.05; a reachable-task Hit conditions on exact full-pool reachability and is reported separately. Reachability requires the request to lie in the union of the candidate intervals in Eq. 11. The critical noise scale σ^* marks where perturbed population-based selection loses its mean-regret advantage over the affinity-only control under the specified perturbation experiment (Appendix E).

Generation direction success is the fraction of selected valid proposals whose state-preference sign agrees with the requested direction under a separately trained reporting evaluator, averaged across target–direction pools. Validity is the fraction of proposed sequences with successful chemistry adaptation and finite binder features. Within-pool diversity averages one minus Jaccard overlap between candidate monomer sets, then averages across pools. The case-study percentile ranks a selected sequence within its own evaluator-scored proposal pool (zero is best); neither evaluator score nor structural confidence is a measured binding or population endpoint.

D External-model applicability

External tools require different inputs and produce different affinity endpoints: a sequence-only model cannot resolve the GDP/GMPPnP state change from a common receptor sequence, an IC_{50} estimate cannot be identified with K_d without an assay model, and a pose-conditioned state scorer requires bound complexes in both states. We therefore do not pool these external outputs with the measured-K_d selection benchmarks or infer a performance ranking across incompatible endpoints.

E Finite-dose decision evidence

The target-matching grid crosses three measured panels with 19 baseline populations, 19 requested target populations, and nine dose ceilings. Its exactness follows from interval projection in Section 4.2. A response-normalizing binder with $A_b \neq 0$ reaches a target exactly when $0 \leq (P_B^\star - P_B^0)/A_b \leq (1 + 10^{-x})^{-1}$, $x = \log_{10}(c_{\max}/c_{50,b})$; for $A_b = 0$ only the baseline is reachable. A ceiling-specific reach leader gives the exact pool reachability classification on the declared grid. Replacing it with a ceiling-independent saturation-amplitude leader agrees on 9,688/9,747 (0.9939); other fixed normalizers based on the tightest c_{50} and pool-median c_{50} agree on 0.9598 and 0.9238 of tasks, respectively. These are identities and comparisons computed from the input affinities, not independent biological population outcomes.

Among the declared requests, 47.4% seek a shift opposite to every candidate, 12.0% are attainable by every candidate, and for the remaining 40.6% the candidate choice affects attainability. These fractions are properties of the declared task grid, not a distribution of observed population responses.

Measured-affinity selection baselines and their exact screening-budget values are in Table A2. At $K = 1$, selectivity-only attains the full-pool optimum on 3,029 of 3,956 separating tasks, versus 2,143 for affinity-only.

Table A2: Measured-affinity target-matching regret by screening budget on 3,956 separating tasks.

Selection rule	$K = 1$	$K = 2$	$K = 3$	$K = 5$
Random	0.0693	0.0424	0.0291	0.0128
Affinity-only	0.0326	0.0101	0.0048	0.0000
Δ-only	0.0061	0.0051	0.0034	0.0000
Population objective	0.0000	0.0000	0.0000	0.0000

Table A3: Small-library coverage on the declared finite-dose task grid. Full-pool coverage is the fraction of tasks with a candidate within 0.05 population units of the target; minimum size is found by exhaustive subset enumeration. Selectivity and response-diversity columns show the sizes at which each heuristic first equals full-pool coverage. These single-signed panels have substantial wrong-side target mass.

Panel	Pool	Full coverage	Minimum	Selectivity	Diversity
NRAS	14	0.320	2	4	5
HRAS	6	0.249	1	1	2
MBP	8	0.506	2	3	5

The coverage endpoints persist under a separate dense-grid check; these three small libraries do not establish a general library-size law.

Perturbation protocol and empirical residuals. We perturb $\mu = (q_A + q_B)/2$ and $\Delta = q_A - q_B$ independently over a 7×7 scale grid with 2,000 draws per cell and evaluate decisions on the unperturbed measured affinities. The low-regret cells in the main text correspond to $(\sigma_\mu, \sigma_\Delta) \in \{(0,0), (0,0.1), (0.1,0), (0.1,0.1)\}$. A 0.90 binder-choice agreement threshold passes only $(0,0)$ and $(0.1,0)$ on the same grid; reachability accuracy ≥ 0.95 permits $\sigma_\mu = 0.25$ when $\sigma_\Delta \leq 0.1$ among tested cells. The oracle Hit@0.05 across all separating requests is 0.5111 when task-weighted and 0.548 when first averaged within each of the three panels. Neither value is the Hit@0.05 conditioned on reachable requests.

Table A4: Perturbation scale at which the measured-affinity population selection rule loses its advantage over affinity-only selection.

Perturbation regime	Critical scale σ^* ($\log_{10}$ units)
Joint noise, $K = 1$	0.456
Joint noise, $K = 3$	0.328
Joint noise, low dose	0.049
Joint noise, high dose	0.777
Selectivity-only noise, $K = 1$	0.508

Out-of-fold predictions on the 28 measured candidates have mean μ bias $+0.8176$. These empirical errors are neither zero-mean nor independent draws from the synthetic perturbation model. Replaying the actual residuals gives task-weighted regret 0.0581 and reachability accuracy 0.818 on all 9,747 tasks, versus reachability accuracy 0.633 on the separating 3,956. The two cohorts and the absolute-regret and relative-advantage criteria must remain distinct.

Reachable-target success. Hit@0.05 restricted to the 2,197 oracle-reachable tasks at $K = 1$ is 0.6632 for random selection, 0.837 for affinity-only, 0.975 for Δ-only, and 1.000 for the measured-affinity population objective. These task-weighted values have a different denominator from the all-separating oracle Hit@0.05 above.

F Desired-state selection: protocol and scope

The desired-state benchmark is indexed by panel, requested state, 19 declared baseline populations, and nine dose ceilings. Affinity-only ranks the requested state's measured K_d; each selection rule screens the same pool and can choose any free dose within its ceiling, including zero. Ties are scored as an expectation over uniform permutations of the tied candidate set. The affinity-plus-selectivity coefficient is chosen leave-one-panel-out. Because the three measured candidate pools are each single-signed, half the declared tasks do not distinguish binder choices.

Effective units and uncertainty. The same 19 of 27 panel–state–ceiling cells give an affinity-only/optimal choice inversion at every baseline population, so the baseline sweep does not create 513

independent biological observations. The regret varies with baseline despite the invariant inversion count. Resampling the three panels gives a $K = 1$ population-minus-affinity regret contrast of -0.0390 with interval $[-0.0670, -0.0078]$. The arithmetic mean no-treatment population 0.500 is fixed by the symmetric baseline grid and is not an independently measured chance level. All population outcomes in this benchmark derive from measured state-specific affinities and declared baseline populations.

At $K = 1$, the affinity-plus-selectivity heuristic and selectivity-only rule have mean regrets 0.0066 and 0.0063, versus 0.0390 for affinity-only. Potency-only ties affinity-only under this task definition. Table A5 gives the corresponding expected populations, and Table A6 records exact larger-budget entries.

G Additional measured-affinity selection results

These measured-affinity comparisons complement the main finite-dose results with selectivity baselines, larger screening budgets, and panel-specific dose references.

Table A5: Selectivity-based rules at $K = 1$ on the 513 separating desired-state tasks.

Selection rule	Mean regret ↓	Expected population ↑
Affinity $+\lambda\Delta$	0.0066	0.7454
Δ-only	0.0063	0.7457

Table A6: Desired-state mean regret at larger screening budgets.

Selection rule	$K = 3$	$K = 5$
Affinity-only	0.0072	0.0000
Population objective	0.0000	0.0000

Table A7: Declared baselines and reference-normalized dose crossovers for the measured panels.

Panel	Baseline P_B	Crossover/K_{ref}
NRAS	0.10	6.82
HRAS	0.90	0.31
MBP	0.90	75.00

H Controlled selectivity-noise diagnostic

On the desired-state benchmark, Gaussian noise is added to measured Δ before the one-binder screen; final utility uses unperturbed measured affinities. Figure A1 shows 20-repeat means for $K = 1$ and $K = 3$, alongside the numerical frozen Q_{bind} and predicted-affinity $K = 1$ regrets. This is an oracle perturbation, not an estimate of the frozen model's residual distribution, and comparison at an equivalent noise scale does not validate the learned selectivity ranking.

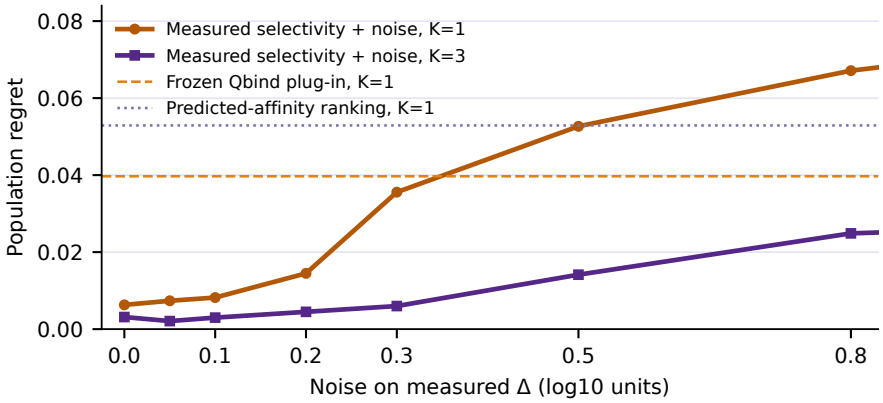

Figure A1: Desired-state candidate-selection regret after controlled noise on measured state differences at screening budgets $K = 1$ and $K = 3$. Horizontal lines identify frozen-model ranking benchmarks.

I STRUCTURAL CANDIDATE ILLUSTRATION

The main-text Figure 5 visualizes one candidate selected from a native PepTune seed-1 pool by predicted desired-state affinity. The toxin templates originate from free PDB 2ABX and HAP-peptide-bound PDB 8VY8; the HAP peptide was removed for prediction. Both Boltz-2 runs use the same mature 74-residue toxin sequence and the same peptide, with separate enforced state templates. This geometry comparison does not establish an interconverting ensemble under one solution condition. Peptide pLDDT is 53.6/58.9 for the depicted complexes; the free-frame model contains sub-van-der-Waals contacts. Neither ipTM nor contact count is an experimental affinity.

The ranking numbers printed in Figure 5 are *cohort* means over 24 targets and 72 target–seed pools under two learned evaluators (Table A8). They are not the performance of the depicted seed-1 binder: on that seed, the Q_{bind} choice scores 0.881/0.925, versus 0.522/0.582 for generator order. The model's predicted state difference is identical for all 214 candidates on this target at four-decimal resolution, so the choice is driven by the absolute-affinity channel and coincides with the desired-state affinity control. The case is an illustration of a computational shortlist and its structural context, not physical binding or population redistribution.

Table A8: Percentile of the $K = 1$ candidate selected from α-bungarotoxin pools and from the separate 24-target cohort (lower is better).

Pool	Q_{bind} primary	Q_{bind} independent	Generator primary	Generator independent
Target seed 0	0.000	0.014	0.671	0.686
Target seed 1 (depicted)	0.881	0.925	0.522	0.582
Target seed 2	0.195	0.299	0.818	0.857
24-target cohort	0.320	0.404	0.652	0.667

No paired K_d, baseline population, dose, or independent post-binding population is available for this peptide. Such measurements under matched solution conditions are necessary to test whether it is selective and shifts population.

J POPULATION-CONTROL EXTENSIONS

Affinity uncertainty. For a predictive affinity distribution $p(\mathcal{K})$ or interval uncertainty set $\mathcal{U}$, the decision layer can replace the point objective by

$$\begin{aligned} J_{\mathrm{exp}}(b,c) &= \mathbb{E}_{\mathcal{K}\sim p(\mathcal{K})}[D(\mathbf{P}(\mathcal{K},b,c),\mathbf{P}^\star)]\,, \\ J_{\mathrm{rob}}(b,c) &= \max_{\mathcal{K}\in\mathcal{U}} D(\mathbf{P}(\mathcal{K},b,c),\mathbf{P}^\star)\,. \end{aligned} \tag{14}$$

These are possible decision criteria under predictive uncertainty; the main perturbation experiment evaluates point-estimate decisions and does not benchmark these extensions.

Ligand mass balance. When receptor depletion is non-negligible, the experimentally controlled total concentration c_{tot} and free concentration c satisfy

$$c_{\mathrm{tot}} = c + R_{\mathrm{tot}}\frac{c}{c+c_{50}}. \tag{15}$$

We solve the positive root for c before evaluating the population response. The free-ligand approximation is recovered when R_{tot}/c_{50} is small.

Multiple states. For m states, the equilibrium population of state i is

$$P_i(c;b) = \frac{P_i^0\,(1+c/K_{d,i}(b))}{\sum_{j=1}^{m} P_j^0\,(1+c/K_{d,j}(b))}. \tag{16}$$

The same binder–dose objective applies with a vector-valued population discrepancy. A single target-state occupancy can sometimes be reduced exactly to an effective two-state problem, whereas controlling the full distribution requires the complete state-resolved landscape.

K LIMITATIONS

The decision analysis derives population responses from measured state-resolved K_d and declared baseline populations under a reversible two-state, 1:1 equilibrium model. It does not establish that any tested or generated binder changes a measured conformational population. Testing that claim requires matched state affinities and independent baseline and post-binding population measurements at known free concentrations. These will be addressed by further wet-lab experiments.

Paired-state affinity data remain sparse. Q_{bind} provides affinity and state-direction signals, but its current additive readout has limited ability to distinguish binders by state preference within a matched context, and its one-candidate selection gain is modest. Generation and structural examples are evaluated computationally. More matched-state measurements, better binder–state models, and direct binding and population assays are needed to move from candidate prioritization to validated population control.