SOAPIA: Specificity-Guided Generation of Off-Target-Avoiding Protein Interactions with High Target Affinity

Anonymous Authors[1]

Abstract

Therapeutic molecules must selectively interact with a target protein while avoiding structurally or functionally similar off-targets, such as alternate isoforms, point mutants, and homologous family members. As an example, the development of safe therapeutics for fusion-driven cancers requires interacting with the fusion oncoprotein, which arises from a chromosomal translocation, while avoiding its homologous parental head and tail proteins. However, no existing generative strategy explicitly optimizes both target affinity and off-target avoidance. To address this, we introduce **SOAPIA**, a framework for the **S**pecificity-guided generation of **O**ff-target-**A**voiding **P**rotein **I**nteractions with high target **A**ffinity. SOAPIA generates *de novo* peptide binders by steering the generative process of a discrete flow matching model, enabling Pareto-efficient exploration of discrete sequence space without gradient access. Affinity is optimized via a pre-trained predictor, while specificity is enforced using SiameseCat, a novel Siamese Neural Network. SiameseCat identifies isoform-specific, mutant-specific, and fusion oncoprotein parent-specific interactions at hit rates of up to 43.2%, 28.5%, and 30.7%, respectively, encompassing both protein-protein and peptide-protein interactions. We benchmark SOAPIA on 17 clinically relevant fusion oncoproteins, where it produces binders that preferentially engage the full fusion over its wild-type head and tail counterparts with predicted nanomolar-range affinity. SOAPIA thus offers a general framework for designing selective biologics across isoform-, mutant-, and fusion-specific targets, with promise for currently untreatable diseases.

[1]Anonymous Institution, Anonymous City, Anonymous Region, Anonymous Country. Correspondence to: Anonymous Author <anon.email@domain.com>.

Submitted to the 2026 Workshop on Generative and Agentic AI for Biology (ICML 2026). Do not distribute.

1. Introduction

Selective modulation of pathogenic proteins is essential for drug design (Nada et al., 2024). Off-target interactions can reduce efficacy or induce toxicity, a challenge shared across small molecules, PROTACs, and biologics (Garon et al., 2017; Chen et al., 2023). Since large-scale *in vitro* screening is costly and impractical, computational methods for designing drug-target interactions (DTIs) and protein-protein interactions (PPIs) are increasingly important. Structure-based approaches offer atomistic resolution but fail on intrinsically disordered residues, which comprise an estimated 37-50% of the human proteome (Oates et al., 2012; Tunyasuvunakool et al., 2021; Chen et al., 2023). Sequence-based models for DTIs (Singh et al., 2023; McNutt et al., 2024; Gao et al., 2024), PPIs (Sledzieski et al., 2021; Singh et al., 2022), and peptide design (Bhat et al., 2025a; Tang et al., 2025a) address this limitation by operating directly on primary sequences.

Yet, most generative approaches optimize only for target binding, without explicitly avoiding off-target interactions. This is particularly problematic when highly homologous off-targets are present, such as mutants of the target, alternate isoforms spliced from the same mRNA transcript, or different members of the same protein family. Selectivity is especially critical when targeting fusion oncoproteins, key drivers of many pediatric cancers. These rare proteins arise from chromosomal translocations that join a“head” and “tail” gene, retaining high sequence identity with one or both parental counterparts (Vincoff et al., 2025). Designing binders for such targets requires a multi-objective formulation that simultaneously maximizes affinity and enforces specificity against a set of known off-targets. Including multiple explicit decoys during training and generation better reflects real-world therapeutic constraints, where safe and selective binding is essential (Chen et al., 2023).

Discrete flow matching (DFM) (Gat et al., 2024; Davis et al., 2024; Stark et al., 2024; Dunn & Koes, 2024; Tang et al., 2025b) has recently emerged as a powerful framework for sequence design. Like its predecessor, discrete diffusion (Campbell et al., 2024; Shi et al., 2024; Sahoo et al., 2024; Wang et al., 2024; Alamdari et al., 2023), DFM enables high-quality, controllable generation at the token level without

requiring 3D structures or continuous embeddings (Hong et al., 2026). Peptide generators built from DFM (Chen et al., 2025d) and rectified DFM (Chen et al., 2025c) paradigms can deliver naturalistic, foldable peptide sequences using an amino acid or SMILES alphabet. To enable sequence design that satisfies particular design objectives, several guidance strategies have been proposed, both gradient-based (Tang et al., 2025b) and gradient-free (Nisonoff et al., 2024; Wan et al., 2026). Recently, MOG-DFM (Chen et al., 2025d) introduced a multi-objective strategy which performs gradient-free, guided token transitions to enforce progression towards the optimal Pareto front, where no single objective can be improved without compromising another. MOG-DFM produces target-binding peptides with any number of optimized properties, given that each property can be predicted with some reward function (Chen et al., 2025d).

We build on these advances with **SOAPIA**: a framework for the **S**pecificity-guided generation of **O**ff-target-**A**voiding **P**rotein **I**nteractions with high target **A**ffinity. SOAPIA introduces a novel specificity predictor called SiameseCat, whose scores reflect confidence that a binder interacts with a particular target protein, while avoiding multiple predefined off-targets. SiameseCat was trained on a newly curated dataset, SNOOPPIXL, which expands the sequence-based PPI dataset SNOOPPI (Vincoff & Chatterjee, 2026) with additional PDB-derived positive and literature-derived negative interactions. To produce target-specific binders, SOAPIA leverages MOG-DFM and the pretrained binding affinity predictor from PeptiVerse (Zhang et al., 2026) (hereafter referred to as PVaffinity). MOG-DFM was selected for its unique level of control over the generative process; additionally, MOG-DFM with a PVaffinity objective has been demonstrated to produce target-binding peptides *in vitro* (Chen et al., 2025b). We show that SOAPIA outperforms unguided sampling baselines on both affinity and specificity, and generates peptide-like binders that preferentially dock to fusion proteins while avoiding their head and tail domains. *In silico* docking with AlphaFold-Multimer (Evans et al., 2021) confirms SOAPIA's ability to design safe and selective binders for undruggable and isoform-sensitive targets, such as fusion oncoproteins.

2. Methods

2.1. Data curation and handling

A new sequence-based PPI dataset called SNOOPPIXL was curated to train SiameseCat (Fig. 1A). Although SNOOPPIXL contains over six thousand verified negatives, this dataset alone was not large enough to form a sufficient number of training samples. Therefore, SNOOPPIXL was supplemented with random negatives, which were carefully selected to maximize difficulty while avoiding erroneous false negatives (Fig. 1B). The full set of SNOOPPIXL interactions and random negatives was then organized into candidate quadruplets containing one binder, one target, and two off-targets each. A greedy algorithm selected a set of quadruplets with minimal biases and high difficulty for the training, validation, and test sets, leaving aside the vast majority of extremely difficult cases (*e.g.*, one or two off-targets that are identical to the target aside from a single mutation) for the gold-standard benchmarking set (Fig. 1C). Detailed information on the attributes of each split – including isoform-specific, mutant-specific, high-homology, and particularly high-confidence interactions – can be found in Tables S1 and S2. The final quadruplet counts were: 35,156 train, 4,308 validation, 4,344 test, and 17,593 gold standard. These splits are cluster-disjoint on the binder sequence, preventing data leakage. Comprehensive data curation, splitting, and clustering details can be found in the Supplementary Methods.

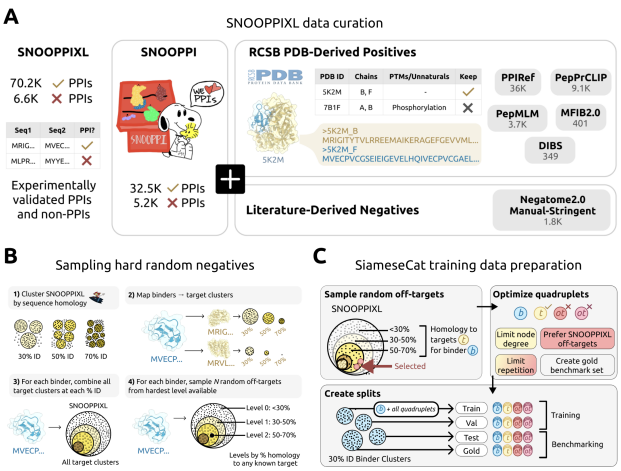

Figure 1. **A** SNOOPPIXL combines SNOOPPI (Vincoff & Chatterjee, 2026) with Negatome2.0 (Blohm et al., 2014) and various RCSB PDB-derived resources (Berman et al., 2000): PPIRef (Bushuiev et al., 2023), PepPrCLIP (Brixi et al., 2023), PepMLM (Chen et al., 2025a), MFIB2.0 (Fichó et al., 2025), and DIBS (Schad et al., 2018). **B** Hard random negative sampling. Each binder's known target proteins are clustered at 30%, 50%, and 70% identity. Random negatives are selected from these clusters, with greater homology preferred. **C** Training data preparation: supplement SNOOPPIXL with random negatives, greedily select optimal quadruplets, and perform a cluster-based split on binder sequence.

2.2. SiameseCat specificity predictor

Protein encoding SiameseCat is a two-track Siamese Neural Network, meaning that the same protein will produce a unique output embedding depending on its role (Fig. 2A). Only the binder sequence is submitted to the binder track, while the target track consumes all proteins that could potentially interact with the binder: namely, the target and all off-targets. In each track, the protein is encoded by a

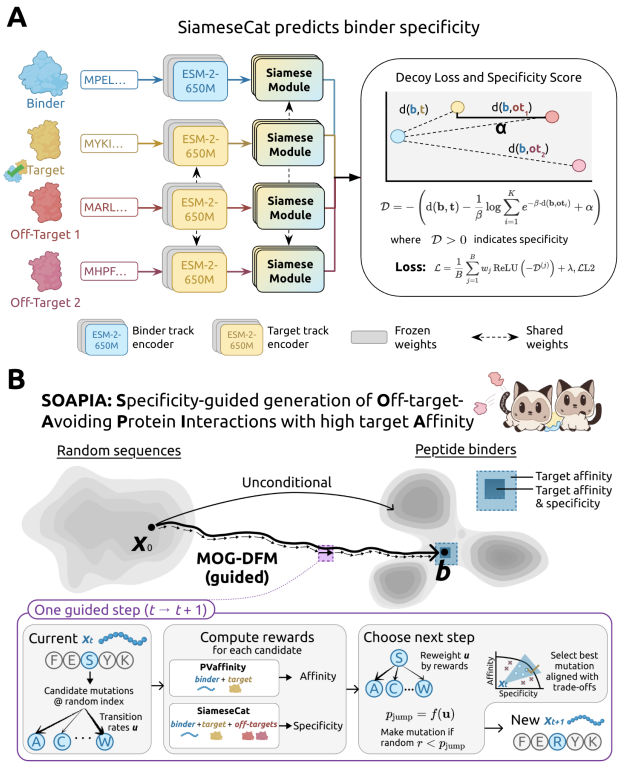

Figure 2. **SOAPIA algorithm. A** Architecture of SiameseCat, which predicts whether a binder interacts with one target while avoiding multiple off-targets. Loss includes specificity score and L2-norm to control embedding dimension. **B** SOAPIA employs MOG-DFM's multi-objective generative process (Chen et al., 2025d) with PVaffinity (Zhang et al., 2026) and SiameseCat as rewards. Generates high-affinity, specific binders using only the target and off-target sequences.

33-layer ESM-2-650M model (Lin et al., 2023). Before pretraining, these are identical; however, we unfreeze the last two layers for training, resulting in different encoders. This setup guarantees that a single protein sequence can be evaluated for homomeric interaction with distinct binder versus target representations. After encoding, all proteins pass through the shared Siamese Module (Fig. S1). A positional multi-head attention module (`n_heads` = 10) with rotary positional embeddings (RoPE) (Su et al., 2024) captures sequence-order information. Outputs are passed through two SiLU-activated linear layers with skip connections, and attention pooling produces fixed-length embeddings, which are passed directly to the loss function (Fig. 2A, Fig. S1).

Specificity loss To enforce off-target avoidance, we utilize a Log-Sum-Exp Decoy Loss. Let $\mathbf{b}$, $\mathbf{t}$, and $\mathbf{ot}_i \in \mathbb{R}^{1280}$ represent the embeddings of the binder, target, and ith off-targets out of K total off-targets. During training, K=2; at inference time, K can be any number of off-targets and is only limited by compute. We calculate specificity score

$$\mathcal{D} = -\left(\mathrm{d}(\mathbf{b},\mathbf{t}) - \frac{1}{\beta}\log\sum_{i=1}^{K} e^{-\beta\cdot\mathrm{d}(\mathbf{b},\mathbf{ot}_i)} + \alpha\right), \quad (1)$$

where d represents Euclidean distance, β controls sharpness, and α is a margin. The log-sum-exp term computes a differentiable soft minimum of all off-target distances (Fig. S2), effectively comparing the target distance to the closest (most difficult) off-target. The specificity score becomes positive when the binder is closer to the target than to its nearest off-target by at least margin α, indicating a specific interaction. Euclidean distance was selected based on the finding that cosine similarity can treat high-dimensional embeddings as random vectors, while Euclidean distance does not suffer from the same collapse (Tessari et al., 2024).

We define the specificity term of the loss as follows:

$$\mathcal{L}_{\text{spec}} = \mathrm{ReLU}\left(-\mathcal{D}\right)$$

such that the loss is greater than zero (punishing the model) only when a quadruplet is deemed non-specific. The margin term α prevents loss from easily plateauing at zero. An additional L2-norm loss term penalizes large embedding magnitudes, focusing the training process on rearranging embeddings rather than amplifying them:

$$\mathcal{L}_{\text{L2}} = \sum_{z\in\mathcal{E}} \|z\|_2^2$$

Here, $\mathcal{E}$ denotes the set of binder, target, and off-target embeddings. The loss combines $\mathcal{L}_{\text{spec}}$ and $\mathcal{L}_{\text{L2}}$ with weighting factors w_j and λ, respectively:

$$\mathcal{L} = \frac{1}{B}\sum_{j=1}^{B} w_j \mathcal{L}_{\text{spec}}^{(j)} + \lambda\mathcal{L}_{\text{L2}}, \quad (2)$$
$$\text{where} \quad w_j = 1 + \gamma k_j^{\text{real}}.$$

During training, examples containing experimentally validated ("real") off-targets from SNOOPPIXL are upweighted relative to examples containing random decoys through the weighting term w_j. γ is a tunable hyperparameter and k_j^{real} denotes the number of real off-targets in quadruplet j within a batch of size B. The λ term is also a tunable hyperparameter controlling the influence of $\mathcal{L}_{\text{L2}}$.

Implementation details SiameseCat was implemented using PyTorch Lightning (Falcon, 2019) and trained on 4xA6000 NVIDIA GPUs with an effective batch size of 64. Learning rate was initialized at 1×10^{-4} and decayed using cosine annealing with 200 warmup steps. See Table S3 for full hyperparameters. Each generation task for SOAPIA (Fig. 2B) was conducted on one A6000 NVIDIA GPU.

Embedding separation visualization Before the start of training and after each epoch, we computed Euclidean distances between binder-target, binder-off-target 1, and binder-off-target 2 embeddings across a consistent set of 576 randomly sampled quadruplets in the training and validation sets. These distances were visualized using `matplotlib v3.8.2` (Figure 3A, S3).

Model selection SiameseCat was trained until the validation specificity decreased for three continuous epochs. Epoch 3, which reached the greatest validation specificity of 8.49, was selected as the model checkpoint (Fig. 3). All downstream benchmarking was performed with this checkpoint only.

SiameseCat benchmarking SiameseCat was benchmarked on various subsets of the gold-standard split (details in Tables S1 and S2). For the structured vs. disordered comparison, PPIs from the Disordered Binding Site (DIBS) database (Schad et al., 2018), which do originate from the RCSB PDB, were selected as the "disordered" set. PPIs that satisfied the following criteria were selected for the "structured" set: (1) present in the RCSB PDB, (2) neither partner is a member of the DisProt database (Nugnes et al., 2026), (3) both partners have high foldability (>= 90 pLDDT) in AlphaFoldDB (Varadi et al., 2024). For the protein-protein interaction (PPI) vs. peptide-protein interaction (PepPI) comparison, all PPIs in the gold-standard split were selected for the protein-protein set and all PepPIs were selected for the peptide-protein set. Peptides are defined as being shorter than 50 amino acids. It should be noted that a high proportion of these interactions contain other difficult attributes such as mutant, isoform, and binding-site-rearrangement relationships. The isoform and mutant subsets were collected by identifying groupings where at least one off-target is an isoform or mutant of the target, respectively. In each of these tasks, SiameseCat scores the specificity of the designated quadruplets with margin α set to 0. Because there are no known examples of specific binders to fusion oncoproteins, the fusions themselves act as off-targets in the fusion task. A dataset of proteins predicted to bind one of the wild-type parents (either head or tail), but *not* the fusion or the opposite parent, was assembled by combining Table S9 of ChiPPI (Frenkel-Morgenstern et al., 2017) with FusOn-DB (Vincoff et al., 2025). ChiPPI maps fusion constituents to their binding partners through BioGRID (Oughtred et al., 2021) and predicts missing interactors in the fusion based on domain loss caused by the chromosomal translocation event. It should be noted that the head-partner and tail-partner interactions from ChiPPI/BioGRID were not also present in SNOOPPIXL, a very strict PPI dataset. This is another way that the fusion task represents an out-of-distribution assessment. To compare robustness across specificity tasks with different baseline difficulties, we computed a relative performance retention metric (Fig. 3B). For each task category c, performance at a given specificity margin α was normalized by the corresponding baseline performance at $\alpha = 0$:

$$\text{Retention}_c(\alpha) = \frac{\text{Performance}_c(\alpha)}{\text{Performance}_c(0)}$$

This normalization isolates how rapidly performance deteriorates as specificity constraints become increasingly strict.

SOAPIA generation tasks SOAPIA was employed to generate 200 18-mer peptide binders each for a set of 17 fusion oncoproteins, selected from FusOn-DB (Vincoff et al., 2025) based on suitability for AlphaFold-Multimer validation (Fig. 4). All 17 fusion oncoproteins have short, foldable sequences (<= 400 amino acids, pLDDT > 70 in AlphaFold2), and the same is true of their heads and tails. These criteria ensure quick folding and maximize the chances of the predictions being reliable. The actifpTM (Varga et al., 2025) metric was selected over ipTM for evaluation due to its focus on the true binding interface between two chains, and its strong performance for peptide-protein interactions.

Discrete Flow Matching Binder generation is based on a discrete flow matching (DFM) (Gat et al., 2024) model. For a peptide sequence $x = (x_1, \ldots, x_n) \in \mathcal{T}^n$ where $\mathcal{T}$ is the 20-letter amino acid vocabulary, DFM models a continuous-time Markov chain (CTMC) $\{X_t\}_{t\in[0,1]}$ with time-dependent transition rates $u_t(y, x)$ that transport probability mass from an initial distribution p_0 to the target distribution of peptide binders p_1. The marginal probability distribution evolves according to the Kolmogorov forward equation

$$\frac{d}{dt}p_t(y) = \sum_{x\in\mathcal{S}} u_t(y, x)p_t(x),$$

where $u_t(y, x)$ denotes the transition rate from sequence x to y at time t. In practice, PepDFM factorizes the velocity field over sequence positions as $u_t(y, x) = \sum_i \delta(y_{\bar{i}}, x_{\bar{i}})\, u_t^i(y_i, x)$, enabling efficient token-level updates during generation. The model is trained via conditional flow matching:

$$\mathcal{L}_{\text{CDFM}}(\theta) = \mathbb{E}_{t,\, Z,\, X_t\sim p_{t|Z}} \sum_i D^i_{X_t}\big(u_t^i(\cdot, X_t \mid Z),\ u^i_{\theta,t}(\cdot, X_t)\big), \quad (3)$$

where $D^i_{X_t}$ is a Bregman divergence. Sampling proceeds by iteratively selecting candidate amino acid substitutions according to the learned CTMC dynamics, gradually transporting random initial sequences toward the distribution of realistic peptide binders.

Sampling and generation strategy (SOAPIA) SOAPIA produces peptide binders by utilizing the generator and multi-objective strategy introduced in MOG-DFM (Chen et al., 2025d) (Fig. 2B). MOG-DFM steers pre-trained discrete flow models toward Pareto-efficient solutions across multiple objectives $\{s_n\}_{n=1}^{N}$ by combining rank-based improvement with directional alignment toward a weight vector ω. In SOAPIA, the generative process of PepDFM is guided by two soft-value reward functions that are maximized for the optimal binder: the specificity score ($r(x) = \mathcal{D}$) (Eq. 1) of SiameseCat, and the output of a pre-trained peptide-protein binding affinity predictor introduced in PeptiVerse (Zhang et al., 2026), which we call PVaffinity. Predictions from PVaffinity correspond to experimenal K_d, K_i, and IC50 measurements. Values from 6-7 indicate weak micromolar binding, and values from 7-9 indicate stronger nanomolar binding. PepDFM, an unconditional generator, was initially trained to transport random starting sequences to the target distribution of generic peptide binders. With MOG-DFM guidance, PepDFM flows toward a much lower-density region of that target space: where binders have high affinity to the target and avoid all specified off-targets (Fig. 2B).

Generation is initialized with a random peptide of fixed length L, sampled uniformly from the 20 natural amino acids. At each step from $t = 0$ to τ, one random index i is selected for mutation. All 19 possible amino acid substitutions are considered, with CTMC transition rates captured by vector $\mathbf{u}_t$. For each candidate mutation y_i, SOAPIA evaluates improvements in affinity and specificity relative to the current sequence x, producing a multi-objective guidance score $\Delta S(y_i, x, \omega)$ that combines normalized rank scores across both objectives with a directional term $\Delta s(y_i, x) \cdot \omega$. Guided transition rates for $y_i \neq x_i$ are then computed as

$$\tilde{u}_t^i(y_i, x \mid \omega) = \beta\, u_t^i(y_i, x) \exp\big(\Delta S(y_i, x, \omega)\big),$$

with the diagonal entry $\tilde{u}_t^i(x_i, x \mid \omega) = -\sum_{y_i \neq x_i} \tilde{u}_t^i(y_i, x \mid \omega)$ set to conserve probability, thereby biasing generation toward mutations aligned with the desired affinity-specificity tradeoff. MOG-DFM then performs adaptive hypercone filtering, restricting transitions to lie within an angle Φ around ω, before probabilistically applying a mutation step. Repeating this process iteratively transports the sequence trajectory from random peptides toward binders optimized for both target affinity and off-target avoidance.

3. Results

3.1. SNOOPPIXL adds thousands of PPIs to the sequence-indexed interactome.

We significantly expand SNOOPPI (Vincoff & Chatterjee, 2026), the first sequence-indexed PPI and non-PPI database with PTM, mutation, and isoform-level granularity, by incorporating structure-resolved positive PPIs in the RCSB PDB (Berman et al., 2000; Bushuiev et al., 2023; Bhat et al., 2025b; Chen et al., 2025a; Schad et al., 2018; Fichó et al., 2025) and literature-derived negatives from Negatome2.0 (Blohm et al., 2014). SNOOPPIXL delivers 70.2K positive and 6.6K negative PPIs which are experimentally validated and, to our knowledge, represent the true sequences expressed in the original experiment (Fig. 1A).

3.2. SiameseCat learns specificity-aware representations across difficult biological tasks.

After training on SNOOPPIXL and sampled hard negatives (Fig. 1), we first evaluated whether SiameseCat learns embedding representations capable of separating target interactions from closely related off-targets. During training, the model was optimized to minimize binder-target distance while maximizing distance from the hardest off-target under the Log-Sum-Exp decoy objective (Eqn. 2). Across training epochs, we observed progressive separation between binder-target and binder-off-target distances in both the training and validation sets (Figs. 3A, S3). Importantly, this separation emerged despite the inclusion of highly challenging off-target relationships, including mutant proteins, alternate isoforms, and proteins sharing high sequence identity with the target (Tables S1, S2).

To quantify how well SiameseCat retains performance on increasingly difficult specificity tasks, we evaluated the model on multiple biologically relevant subsets of the gold-standard benchmark set (Fig. 3B). Structured protein-protein interactions exhibited the strongest performance, while disordered PPIs, peptide-protein interactions, mutant-specific interactions, isoform-specific interactions, and fusion parent-specific interactions exhibited progressively lower retention. Despite this increase in difficulty, SiameseCat maintained discriminative ability across all tasks, including mutant-specific interactions where as little as one amino acid may be different between a target and off-target. The 57% hit rate on peptide-protein interactions supports the notion that SiameseCat can learn both PPI and PepPI behavior even from a dataset where the vast majority of interactions are PPIs. The protein-protein interaction subset had the lowest hit rate, which is most likely a reflection of other conditions (*e.g.* many mutant-specific examples within this group) rather than SiameseCat struggling on PPI relative to PepPI prediction.

To ensure that SiameseCat is not biased towards calling random protein groups specific, a true positive rate was calculated on a subset of 290 gold-standard quadruplets where both "off-targets" are actually interactors. SiameseCat achieved a 79% hit rate in these cases, meaning it did not call the interaction specific at threshold $\alpha = 0$. The hit

rate increased to 97.2% at $\alpha = 5$, indicating that high specificity scores from SiameseCat are quite trustworthy. Overall, these results suggest that SiameseCat does not merely memorize broad interaction patterns, but instead learns representations sensitive to subtle sequence-dependent determinants of specificity.

3.3. SOAPIA generates binders optimized for both affinity and specificity

Next, we investigated whether SiameseCat specificity scores could be used as a generative reward within the MOG-DFM framework. SOAPIA combines SiameseCat specificity guidance with PVaffinity, a pretrained peptide-protein affinity predictor, to steer the discrete flow matching process toward Pareto-efficient regions of sequence space (Fig. 2B).

We applied SOAPIA to a panel of 17 clinically relevant fusion oncoproteins selected from FusOn-DB (Vincoff et al., 2025) (breakpoint and head-tail composition shown in Fig. S4). For each fusion, the corresponding wild-type head and tail proteins were treated as explicit off-targets during generation. SOAPIA generated 200 candidate binders per target fusion using gradient-free multi-objective guidance. Across all 17 targets, SOAPIA produced a set of peptides simultaneously satisfying both affinity and specificity criteria ("green-zone" binders; Fig. 4A). However, noticeable variation existed between fusion proteins, indicating that some fusion-specific interfaces are intrinsically more designable than others. The differences in performance were not correlated with fusion composition; a fusion closer to a 50/50 split between head and tail was not intrinsically more designable (Figs. S4, 3A), indicating that the challenge of engineering off-target-avoiding PPIs goes beyond sequence homology. Importantly, fusion targets with lower overall hit rates still frequently yielded several high-confidence candidate binders, suggesting that SOAPIA can effectively navigate difficult specificity landscapes. While SiameseCat and PVaffinity operate purely at the sequence level, many green-zone binders exhibited favorable structural behavior after AlphaFold-Multimer co-folding. In numerous complexes, the target actifpTM exceeded the corresponding scores for both off-target proteins ("structurally selective" binders; Fig. 4A), with hit rates of up to 63.7%.

We examined several representative AlphaFold-Multimer complexes for SOAPIA binders (Fig. 4B). In these examples, peptides preferentially engaged the fusion oncoprotein while exhibiting weaker predicted interactions with the wild-type head and tail proteins. Interestingly, several binders appeared to contact residues spanning both sides of the fusion breakpoint, suggesting that SOAPIA may exploit structural or contextual features unique to the fusion itself, rather than simply recognizing one parental domain in isolation.

Finally, we validated the usefulness of our strategy by comparing SOAPIA's performance with a non-guided baseline. Here, MOG-DFM was run with a dummy reward function, which returned 0 at each step regardless of the peptide sequence (Fig. 2B). SOAPIA binders heavily populated the high-affinity, high-specificity region; meanwhile, the control method failed to produce many candidates with both strong affinity and specificity.

Together, these results demonstrate that SOAPIA can generate candidate binders that jointly optimize affinity and specificity at the sequence level while also producing structurally selective interactions *in silico*.

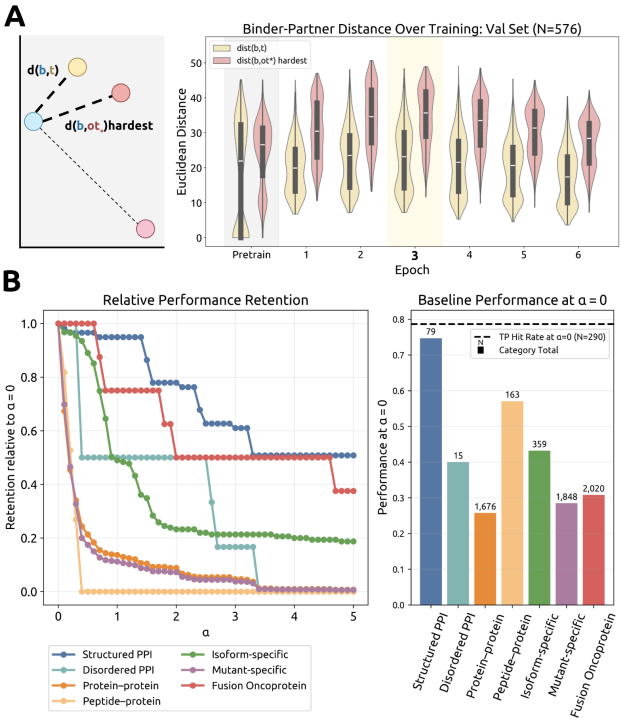

Figure 3. **SiameseCat training and performance on difficult specificity tasks. A** Comparison of Euclidean distance between binder and target, versus binder and hardest (closest) off-target, for a random sample of the validation set (N=576) before (grey highlight) and throughout training. Epoch 3 (yellow highlight) was selected as the SiameseCat checkpoint due to its high validation specificity. **B** SiameseCat performance on challenging specificity discrimination tasks. (Left) Relative performance retention across increasingly difficult specificity thresholds (α), normalized to each task's baseline performance at $\alpha = 0$. (Right) Specificity hits for each task, compared with the true positive hit rate of 79%.

4. Discussion

SOAPIA is a new framework for multi-objective binder design that jointly optimizes target affinity and off-target avoidance through gradient-free, guided discrete flow matching. Building on MOG-DFM (Chen et al., 2025d), SOAPIA combines soft-value signals from a Siamese contrastive

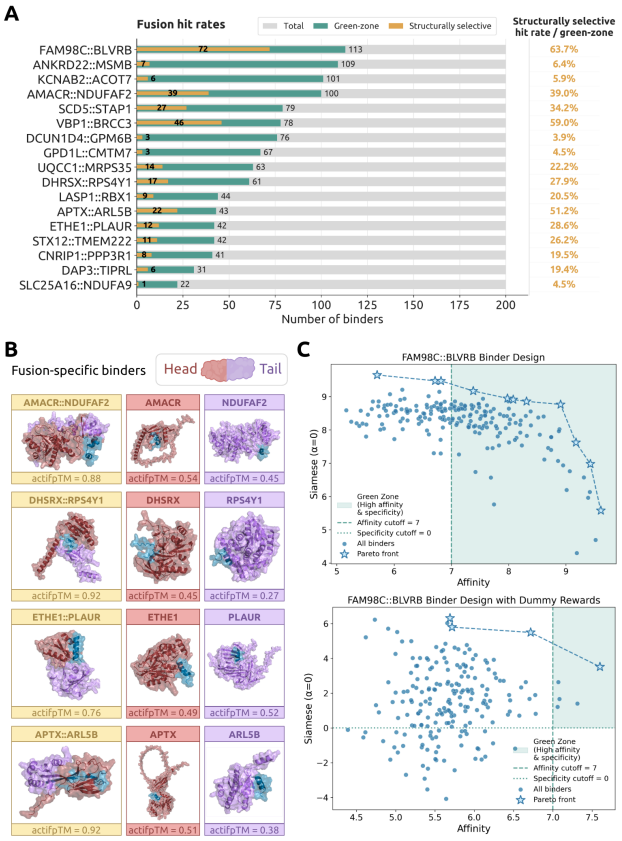

Figure 4. **SOAPIA generates fusion oncoprotein-specific peptide binders. A** SOAPIA generated 200 total 18-mer peptides to each of 17 short, highly foldable fusion oncoproteins. All green-zone binders (affinity $>= 7.0$, the predicted nanomolar-range, and specificity > 0) were folded in AlphaFold-Multimer (Evans et al., 2021); cases where binder-target actifpTM (Varga et al., 2025) exceeded all binder-off-target actifpTMs are deemed structurally selective. **B** Visualization of structurally selective peptide binders to four fusion oncoproteins. **C** Comparison of MOG-DFM scores when using SOAPIA (top) versus dummy rewards (bottom), with green-zone binders in the shaded region.

specificity predictor (SiameseCat) and a pretrained peptide-protein affinity predictor (PVaffinity) (Zhang et al., 2026) to steer the generative trajectory of PepDFM toward Pareto-efficient regions of sequence space. This dual-guidance mechanism allows SOAPIA to outperform unguided sampling baselines across both objectives, generating peptide sequences that exhibit strong predicted binding while avoiding closely related off-targets, including alternate isoforms, single-residue mutants, and homologous family members. Because SiameseCat treats every protein in the target track symmetrically, SOAPIA naturally generalizes beyond the one-target, many-off-targets formulation: the same framework can be inverted to design binders that engage a panel of pathogenic variants (e.g., multiple oncogenic mutants) while sparing a single wild-type counterpart, or to enforce specificity simultaneously across multiple desirable and undesirable interactions. As a concrete demonstration, SOAPIA generates candidate binders to fusion oncoproteins that show preferential AlphaFold-Multimer (Evans et al., 2021) docking to the full fusion but not to the individual head or tail proteins, with structurally selective hit rates of up to 63.7%. This level of selectivity is critical for therapeutic applications, particularly in pediatric fusion-driven cancers, where targeting the aberrant fusion while sparing wild-type counterparts is essential for minimizing toxicity (Vincoff et al., 2025).

Looking ahead, we will refine SOAPIA's guidance weights and adaptive hypercone schedules to further improve sample efficiency, and conduct experimental validation in cellular systems spanning isoform-specific, mutant-specific, and PTM-specific design tasks. For PTM-aware design, we plan to incorporate PTM-Mamba embeddings (Peng et al., 2025) so that SiameseCat can distinguish modified and unmodified residues at the sequence level, and for fusion oncoprotein targets we plan to integrate FusOn-pLM embeddings (Vincoff et al., 2025) to better capture breakpoint-localized context. SOAPIA's robustness could be further extended by incorporating phylogenetically informed objectives, broadening its applicability to rapidly mutating viral and bacterial targets where future escape variants must also be considered. When paired with experimental platforms such as ubiquibodies (uAbs) and deubiquibodies (duAbs) for targeted protein degradation (Brixi et al., 2023; Bhat et al., 2025b; Chen et al., 2025a) and programmable stabilization (Hong et al., 2025), SOAPIA provides a generalizable, sequence-only approach to modulating previously undruggable proteins, with potential impact across oncology, rare disease, infectious disease, and immunotherapy.

5. Data Availability

All code and data related to SOAPIA will be publicly released upon publication of the manuscript.

Impact Statement

SOAPIA supports the design of selective peptide binders for challenging therapeutic targets, including fusion oncoproteins that drive many pediatric cancers and other rare, currently undruggable diseases, with the potential to reduce off-target toxicity and broaden the therapeutic window of peptide interventions. As with other generative models for biomolecular design, outputs require experimental validation before any translational use. Beyond these considerations, the broader societal consequences of this work are consistent with those of machine learning research advancing computational biology, and none require further specific highlighting here.

References

Abdin, O., Nim, S., Wen, H., and Kim, P. M. Pepnn: a deep attention model for the identification of peptide binding sites. *Communications biology*, 5(1):503, 2022.

Alamdari, S., Thakkar, N., van den Berg, R., Tenenholtz, N., Strome, B., Moses, A., Lu, A. X., Fusi, N., Amini, A. P., and Yang, K. K. Protein generation with evolutionary diffusion: sequence is all you need. *BioRxiv*, pp. 2023–09, 2023.

Berman, H. M., Westbrook, J., Feng, Z., Gilliland, G., Bhat, T. N., Weissig, H., Shindyalov, I. N., and Bourne, P. E. The protein data bank. *Nucleic acids research*, 28(1): 235–242, 2000.

Bernett, J., Blumenthal, D. B., and List, M. Cracking the black box of deep sequence-based protein–protein interaction prediction. *Briefings in Bioinformatics*, 25(2): bbae076, 2024.

Bhat, S., Palepu, K., Hong, L., Mao, J., Ye, T., Iyer, R., Zhao, L., Chen, T., Vincoff, S., Watson, R., Wang, T. Z., Srijay, D., Kavirayuni, V. S., Kholina, K., Goel, S., Vure, P., Deshpande, A. J., Soderling, S. H., DeLisa, M. P., and Chatterjee, P. De novo design of peptide binders to conformationally diverse targets with contrastive language modeling. *Science Advances*, 11(4), January 2025a. ISSN 2375-2548. doi: 10.1126/sciadv.adr8638. URL http://dx.doi.org/10.1126/sciadv.adr8638.

Bhat, S., Palepu, K., Hong, L., Mao, J., Ye, T., Iyer, R., Zhao, L., Chen, T., Vincoff, S., Watson, R., et al. De novo design of peptide binders to conformationally diverse targets with contrastive language modeling. *Science Advances*, 11(4):eadr8638, 2025b.

Blohm, P., Frishman, G., Smialowski, P., Goebels, F., Wachinger, B., Ruepp, A., and Frishman, D. Negatome 2.0: a database of non-interacting proteins derived by literature mining, manual annotation and protein structure analysis. *Nucleic acids research*, 42(D1):D396–D400, 2014.

Brixi, G., Ye, T., Hong, L., Wang, T., Monticello, C., Lopez-Barbosa, N., Vincoff, S., Yudistyra, V., Zhao, L., Haarer, E., et al. Salt&peppr is an interface-predicting language model for designing peptide-guided protein degraders. *Communications Biology*, 6(1):1081, 2023.

Bushuiev, A., Bushuiev, R., Filkin, A., Kouba, P., Gabrielova, M., Gabriel, M., Sedlar, J., Pluskal, T., Damborsky, J., Mazurenko, S., et al. Learning to design protein-protein interactions with enhanced generalization. *arXiv preprint arXiv:2310.18515*, 2023.

Campbell, A., Yim, J., Barzilay, R., Rainforth, T., and Jaakkola, T. Generative flows on discrete state-spaces: enabling multimodal flows with applications to protein co-design. In *Proceedings of the 41st International Conference on Machine Learning*, ICML'24. JMLR.org, 2024.

Chen, L. T., Quinn, Z., Dumas, M., Peng, C., Hong, L., Lopez-Gonzalez, M., Mestre, A., Watson, R., Vincoff, S., Zhao, L., et al. Target sequence-conditioned design of peptide binders using masked language modeling. *Nature Biotechnology*, pp. 1–9, 2025a.

Chen, T., Hong, L., Yudistyra, V., Vincoff, S., and Chatterjee, P. Generative design of therapeutics that bind and modulate protein states. *Current Opinion in Biomedical Engineering*, 28:100496, December 2023. ISSN 2468-4511. doi: 10.1016/j.cobme.2023.100496. URL http://dx.doi.org/10.1016/j.cobme.2023.100496.

Chen, T., Quinn, Z., Zhang, Y., and Chatterjee, P. moppit-v3: Motif-specific peptides generated via multi-objective-guided discrete flow matching. In *NeurIPS 2025 Workshop on Structured Probabilistic Inference {\&} Generative Modeling*, 2025b.

Chen, T., Zhang, Y., and Chatterjee, P. Areuredi: Annealed rectified updates for refining discrete flows with multi-objective guidance. *arXiv preprint arXiv:2510.00352*, 2025c.

Chen, T., Zhang, Y., Tang, S., and Chatterjee, P. Multi-objective-guided discrete flow matching for controllable biological sequence design. In *ICML 2025 Generative AI and Biology (GenBio) Workshop*, 2025d. URL https://openreview.net/forum?id=8YIMLoHP9J.

Davis, O., Kessler, S., Petrache, M., Ceylan, İ. İ., Bronstein, M., and Bose, A. J. Fisher flow matching for generative modeling over discrete data. *Advances in Neural Information Processing Systems*, 37:139054–139084, 2024.

Del Toro, N., Shrivastava, A., Ragueneau, E., Meldal, B., Combe, C., Barrera, E., Perfetto, L., How, K., Ratan, P., Shirodkar, G., et al. The intact database: efficient access to fine-grained molecular interaction data. *Nucleic acids research*, 50(D1):D648–D653, 2022.

Dunn, I. and Koes, D. R. Exploring discrete flow matching for 3d de novo molecule generation. *ArXiv*, pp. arXiv–2411, 2024.

Evans, R., O'Neill, M., Pritzel, A., Antropova, N., Senior, A., Green, T., Žídek, A., Bates, R., Blackwell, S., Yim, J., Ronneberger, O., Bodenstein, S., Zielinski, M., Bridgland, A., Potapenko, A., Cowie, A., Tunyasuvunakool, K., Jain, R., Clancy, E., Kohli, P., Jumper,

J., and Hassabis, D. Protein complex prediction with alphafold-multimer. October 2021. doi: 10.1101/2021.10.04.463034. URL http://dx.doi.org/10.1101/2021.10.04.463034.

Falcon, W. A. Pytorch lightning. *GitHub*, 3, 2019.

Fichó, E., Pancsa, R., Magyar, C., Kalman, Z. E., Schád, É., Németh, B. Z., Simon, I., Dobson, L., and Tusnády, G. E. Mfib 2.0: a major update of the database of protein complexes formed by mutual folding of the constituting protein chains. *Nucleic acids research*, 53(D1):D487–D494, 2025.

Frenkel-Morgenstern, M., Gorohovski, A., Tagore, S., Sekar, V., Vazquez, M., and Valencia, A. Chippi: a novel method for mapping chimeric protein–protein interactions uncovers selection principles of protein fusion events in cancer. *Nucleic acids research*, 45(12):7094–7105, 2017.

Gao, B., Qiang, B., Tan, H., Jia, Y., Ren, M., Lu, M., Liu, J., Ma, W.-Y., and Lan, Y. Drugclip: Contrasive protein-molecule representation learning for virtual screening. *Advances in Neural Information Processing Systems*, 36, 2024.

Garon, S. L., Pavlos, R. K., White, K. D., Brown, N. J., Stone, C. A., and Phillips, E. J. Pharmacogenomics of off-target adverse drug reactions. *British Journal of Clinical Pharmacology*, 83(9):1896–1911, April 2017. ISSN 1365-2125. doi: 10.1111/bcp.13294. URL http://dx.doi.org/10.1111/bcp.13294.

Gat, I., Remez, T., Shaul, N., Kreuk, F., Chen, R. T., Synnaeve, G., Adi, Y., and Lipman, Y. Discrete flow matching. *Advances in Neural Information Processing Systems*, 37:133345–133385, 2024.

Hong, L., Ye, T., Wang, T. Z., Srijay, D., Liu, H., Zhao, L., Watson, R., Vincoff, S., Chen, T., Kholina, K., et al. Programmable protein stabilization with language model-derived peptide guides. *Nature Communications*, 16(1): 3555, 2025.

Hong, L., Vincoff, S., and Chatterjee, P. Ai-designed peptides as tools for biochemistry. *Biochemistry*, 2026.

Lin, Z., Akin, H., Rao, R., Hie, B., Zhu, Z., Lu, W., Smetanin, N., Verkuil, R., Kabeli, O., Shmueli, Y., et al. Evolutionary-scale prediction of atomic-level protein structure with a language model. *Science*, 379(6637): 1123–1130, 2023.

Liu, D., Young, F., Lamb, K. D., Claudio Quiros, A., Pancheva, A., Miller, C. J., Macdonald, C., Robertson, D. L., and Yuan, K. Plm-interact: extending protein language models to predict protein-protein interactions. *Nature Communications*, 16(1):9012, 2025.

Martins, P., Mariano, D., Carvalho, F. C., Bastos, L. L., Moraes, L., Paixão, V., and Cardoso de Melo-Minardi, R. Propedia v2. 3: A novel representation approach for the peptide-protein interaction database using graph-based structural signatures. *Frontiers in Bioinformatics*, 3:1103103, 2023.

McNutt, A. T., Adduri, A. K., Ellington, C. N., Dayao, M. T., Xing, E. P., Mohimani, H., and Koes, D. R. Sprint enables interpretable and ultra-fast virtual screening against thousands of proteomes. *arXiv preprint arXiv:2411.15418*, 2024.

Nada, H., Choi, Y., Kim, S., Jeong, K. S., Meanwell, N. A., and Lee, K. New insights into protein–protein interaction modulators in drug discovery and therapeutic advance. *Signal Transduction and Targeted Therapy*, 9(1), December 2024. ISSN 2059-3635. doi: 10.1038/s41392-024-02036-3. URL http://dx.doi.org/10.1038/s41392-024-02036-3.

Nisonoff, H., Xiong, J., Allenspach, S., and Listgarten, J. Unlocking guidance for discrete state-space diffusion and flow models. *arXiv preprint arXiv:2406.01572*, 2024.

Nugnes, M. V., Bouhraoua, K. E. A., Zoubiri, M., Pancsa, R., Fichó, E., Tompa, P., Piovesan, D., Tosatto, S. C., and Aspromonte, M. C. Disprot in 2026: enhancing intrinsically disordered proteins accessibility, deposition, and annotation. *Nucleic Acids Research*, 54(D1):D383–D392, 2026.

Oates, M. E., Romero, P., Ishida, T., Ghalwash, M., Mizianty, M. J., Xue, B., Dosztányi, Z., Uversky, V. N., Obradovic, Z., Kurgan, L., et al. D2p2: database of disordered protein predictions. *Nucleic acids research*, 41 (D1):D508–D516, 2012.

Oughtred, R., Rust, J., Chang, C., Breitkreutz, B.-J., Stark, C., Willems, A., Boucher, L., Leung, G., Kolas, N., Zhang, F., et al. The biogrid database: A comprehensive biomedical resource of curated protein, genetic, and chemical interactions. *Protein Science*, 30(1):187–200, 2021.

Peng, F. Z., Wang, C., Chen, T., Schussheim, B., Vincoff, S., and Chatterjee, P. Ptm-mamba: a ptm-aware protein language model with bidirectional gated mamba blocks. *Nature Methods*, pp. 1–5, 2025.

Sahoo, S. S., Arriola, M., Schiff, Y., Gokaslan, A., Marroquin, E., Chiu, J. T., Rush, A., and Kuleshov, V. Simple and effective masked diffusion language models. *arXiv preprint arXiv:2406.07524*, 2024.

Schad, E., Ficho, E., Pancsa, R., Simon, I., Dosztányi, Z., and Mészáros, B. Dibs: a repository of disordered binding sites mediating interactions with ordered proteins. *Bioinformatics*, 34(3):535–537, 2018.

Shi, J., Han, K., Wang, Z., Doucet, A., and Titsias, M. K. Simplified and generalized masked diffusion for discrete data. In *Advances in Neural Information Processing Systems*, 2024.

Singh, R., Devkota, K., Sledzieski, S., Berger, B., and Cowen, L. Topsy-turvy: integrating a global view into sequence-based ppi prediction. *Bioinformatics*, 38 (Supplement_1):i264–i272, 2022.

Singh, R., Sledzieski, S., Bryson, B., Cowen, L., and Berger, B. Contrastive learning in protein language space predicts interactions between drugs and protein targets. *Proceedings of the National Academy of Sciences*, 120(24): e2220778120, 2023.

Sledzieski, S., Singh, R., Cowen, L., and Berger, B. D-script translates genome to phenome with sequence-based, structure-aware, genome-scale predictions of protein-protein interactions. *Cell Systems*, 12(10):969–982, 2021.

Stark, H., Jing, B., Wang, C., Corso, G., Berger, B., Barzilay, R., and Jaakkola, T. Dirichlet flow matching with applications to dna sequence design. In *Forty-first International Conference on Machine Learning*, 2024.

Steinegger, M. and Söding, J. Mmseqs2 enables sensitive protein sequence searching for the analysis of massive data sets. *Nature biotechnology*, 35(11):1026–1028, 2017.

Su, J., Ahmed, M., Lu, Y., Pan, S., Bo, W., and Liu, Y. Roformer: Enhanced transformer with rotary position embedding. *Neurocomputing*, 568:127063, 2024.

Tang, S., Zhang, Y., and Chatterjee, P. Peptune: De novo generation of therapeutic peptides with multi-objective-guided discrete diffusion. *Proceedings of the 42nd International Conference on Machine Learning (ICML)*, 2025a.

Tang, S., Zhang, Y., Tong, A., and Chatterjee, P. Gumbel-softmax score and flow matching for discrete biological sequence generation. In *ICLR 2025 Workshop on AI for Nucleic Acids*, 2025b. URL https://openreview.net/forum?id=ITpCmDhSfu.

Tessari, F., Yao, K., and Hogan, N. Surpassing cosine similarity for multidimensional comparisons: Dimension insensitive euclidean metric. *arXiv preprint arXiv:2407.08623*, 2024.

Tunyasuvunakool, K., Adler, J., Wu, Z., Green, T., Zielinski, M., Žídek, A., Bridgland, A., Cowie, A., Meyer, C., Laydon, A., et al. Highly accurate protein structure prediction for the human proteome. *Nature*, 596(7873):590–596, 2021.

Varadi, M., Bertoni, D., Magana, P., Paramval, U., Pidruchna, I., Radhakrishnan, M., Tsenkov, M., Nair, S., Mirdita, M., Yeo, J., et al. Alphafold protein structure database in 2024: providing structure coverage for over 214 million protein sequences. *Nucleic acids research*, 52(D1):D368–D375, 2024.

Varga, J. K., Ovchinnikov, S., and Schueler-Furman, O. actifptm: a refined confidence metric of alphafold2 predictions involving flexible regions. *Bioinformatics*, 41(3): btaf107, 2025.

Vincoff, S. and Chatterjee, P. SNOOPPI: Sequence-normalized database of on- and off-target protein-protein interactions. In *ICLR 2026 Workshop on Foundation Models for Science: Real-World Impact and Science-First Design*, 2026. URL https://openreview.net/forum?id=270Ej8S67w.

Vincoff, S., Goel, S., Kholina, K., Pulugurta, R., Vure, P., and Chatterjee, P. Fuson-plm: a fusion oncoprotein-specific language model via adjusted rate masking. *Nature Communications*, 16(1):1436, 2025.

Wan, Z., Ouyang, Y., Xie, L., Fang, F., Zha, H., and Cheng, G. Discrete guidance matching: Exact guidance for discrete flow matching. In *The Fourteenth International Conference on Learning Representations*, 2026. URL https://openreview.net/forum?id=N1RYhOg6ib.

Wang, X., Zheng, Z., Ye, F., Xue, D., Huang, S., and Gu, Q. Diffusion language models are versatile protein learners. *arXiv preprint arXiv:2402.18567*, 2024.

Zhang, Y., Tang, S., Chen, T., Mahood, E., Vincoff, S., and Chatterjee, P. Peptiverse: A unified platform for therapeutic peptide property prediction. *bioRxiv*, pp. 2025–12, 2026.

A. Supplementary Methods

True positives and negatives Each sample in SiameseCat's training data is a protein quadruplet consisting of a binder, target, and two off-targets. The binder and target are known to form a protein-protein interaction (PPI), while the binder and off-targets are known (or assumed, if random) not to interact. We chose a quadruplet construction to maintain a relatively high ratio of negative to positive PPIs, which reflects the sparsity of PPIs in nature (Liu et al., 2025), while avoiding the drawbacks of using more than two off-targets: (1) over-reliance on random negatives, and (2) infeasible training due to time and compute requirements.

An initial set of 32,578 positive and 5,339 negative PPIs was downloaded from SNOOPPI (Vincoff & Chatterjee, 2026) (Fig. 1A). This recent dataset, curated from IntAct (Del Toro et al., 2022), contains the precise amino acid sequences of interacting and non-interacting pairs, including distinctions for post-translational modifications (PTMs), mutations, and different isoforms. Given SNOOPPI's relatively small size, producing under 2.2K training quadruplets at most, we chose to expand the dataset. First, additional positive PPIs were collected from a set of RCSB PDB-derived (Berman et al., 2000) resources: PPIRef (Bushuiev et al., 2023), PepPrCLIP (Bhat et al., 2025b), PepMLM (Chen et al., 2025a) (combines PepNN (Abdin et al., 2022) and Propedia (Martins et al., 2023)), MFIB2.0 (Fichó et al., 2025), and DIBS (Schad et al., 2018). All PPIs whose provided sequences could be verified in the RCSB PDB were added. Often, these sequences were simply a concatenation of resolved residues, skipping unresolved positions. For the training dataset, we chose to utilize the full experimental sequences of each chain instead; these sequences most accurately represent the molecules which were synthesized or expressed in the experiment, ultimately producing the solved structure. Data from the RCSB PDB was also utilized to identify and remove any interactions containing unnatural residues or PTMs. This left a set of interactions which could be precisely defined by the natural amino acid sequences of each chain. The final counts were as follows, including overlaps: PPIRef - 35,950 PPIs, PepPrCLIP - 9,081 PepPIs, PepMLM - 3,680 PepPIs, MFIB2.0 - 401 PPIs, DIBS - 349 PPIs. The exclusion of PTMs and unnatural residues also reduced SNOOPPI positives to 26,299 and negatives to 4,781. Next, a total of 1,806 literature-derived negative interactions were collected from Negatome2.0 (Blohm et al., 2014), a manually curated database of proteins with experimental evidence indicating the absence of physical interaction. Negatome2.0 also contains structure-derived negatives, but these were excluded because the absence of interaction within a crystal structure does not confirm that protein chains will fail to interact in a biological environment. Any positive or negative pairs that overlapped each other or SNOOPPI-Unknown were removed. The resulting dataset, containing 70,161 positive and 6,587 negative interactions, is called SNOOPPIXL (Fig. 1A). Finally, both positive and negative heteromeric PPIs were flipped to approximately double dataset size, a valid data augmentation due to SiameseCat's two-track architecture. This produced a training dataset of 116,816 positive PPIs and 12,949 negative interactions.

Random negatives Although SNOOPPIXL represents a significant collection of PPIs, it still produces a very limited number of training samples. Concerningly, when quadruplets are constructed from SNOOPPIXL alone, any protein with fewer than two true negatives must be excluded. To address this issue, we sample two additional negatives per binder from SNOOPPIXL (Fig. 1B-C). There are many established criteria for random negatives, including lack of known interaction, low sequence homology to known interactors, and different cellular localization than the binder. Adopting the first criterion alone can make the task too easy, while the latter two strategies have been shown to induce biases and limit generalizability for PPI classifiers (Bernett et al., 2024). Here, we overcome these limitations through hard negative mining. First, SNOOPPIXL is clustered at three percent identity levels - 30%, 50%, and 70% - using MMSeqs2 (Steinegger & Söding, 2017). For each binder, we consider proteins that are under 70% homologous to all known targets, avoiding false negatives. Under this 70% threshold, high homology to a known target becomes preferable, to increase the hardness of the task. When sampling random negatives, the 50-70% target homology bucket is checked first, followed by 30-50%, and finally $<$30% (Fig. 1B-C).

Quadruplet selection SiameseCat was trained and evaluated with four splits: training, validation, testing, and gold standard. To ensure that SiameseCat learned relationships rather than predefined roles, four constraints were imposed when selecting quadruplets: (1) each binder appears at most ten times (controlling for PPI node degree bias), (2) each target-off-target grouping appears only once, (3) each target appears once per binder, and (4) each off-target appears once per binder. Constraint 1 was enforced preemptively by subsampling positive PPIs. When making choices to satisfy the other constraints, quadruplets with smaller Euclidean distances between targets and off-targets were prioritized. Random negatives were assigned a very high distance to ensure they were deprioritized in comparison with real negatives. Solving this optimization problem exactly through linear programming was computationally infeasible, leading to the implementation of a greedy algorithm.

Once the quadruplets were selected, we determined which groups contained special relationships - for example, an off-target which is a mutant or isoform of the target, or an interaction (either positive or negative) that was assigned a high confidence score in IntAct (Del Toro et al., 2022). We sequestered 25% of these difficult, high-quality examples for the training, validation, and test sets; the other 75% were reserved for a gold-standard benchmarking set. To prevent data leakage, all binders that clustered at 30% ID with binders in the gold-standard dataset were set aside, along with all of their quadruplets. Finally, all binders whose quadruplets remained in consideration for training were clustered at 30% identity, and these clusters were randomly divided at 80/10/10 ratio into train, validation, and test.

B. Supplementary Tables

Table S1. **Types of off-target pairings in SiameseCat training data.**

Feature	Train	Val	Test	Gold Standard
Random-Random	34079 (96.94%)	4216 (97.86%)	4143 (95.37%)	387 (2.20%)
Random-Real	411 (1.17%)	44 (1.02%)	75 (1.73%)	722 (4.10%)
Real-Real	666 (1.89%)	48 (1.11%)	126 (2.90%)	16484 (93.70%)
Total quadruplets	35156	4308	4344	17593

Table S2. **Special feature annotations across dataset splits.**

Feature	Description	Who satisfies it	Train	Val	Test	Gold Standard
Isoform	Off-target = isoform of target	1 off-target	3 (0.01%)	0 (0.00%)	0 (0.00%)	162 (0.92%)
		Both off-targets	1 (0.00%)	0 (0.00%)	0 (0.00%)	203 (1.15%)
		At least 1 off-target	4 (0.01%)	0 (0.00%)	0 (0.00%)	365 (2.07%)
Mutant	Off-target = mutant of target	1 off-target	65 (0.18%)	6 (0.14%)	16 (0.37%)	1174 (6.67%)
		Both off-targets	18 (0.05%)	2 (0.05%)	1 (0.02%)	679 (3.86%)
		At least 1 off-target	83 (0.24%)	8 (0.19%)	17 (0.39%)	1853 (10.53%)
Binding site change	Off-target = target with binding-site change	1 off-target	6 (0.02%)	1 (0.02%)	1 (0.02%)	433 (2.46%)
		Both off-targets	21 (0.06%)	0 (0.00%)	1 (0.02%)	527 (3.00%)
		At least 1 off-target	27 (0.08%)	1 (0.02%)	2 (0.05%)	960 (5.46%)
High target homology	Off-target clusters with target at ID80/Cov80	1 off-target	253 (0.72%)	26 (0.60%)	25 (0.58%)	3894 (22.13%)
		Both off-targets	106 (0.30%)	7 (0.16%)	8 (0.18%)	5732 (32.58%)
		At least 1 off-target	359 (1.02%)	33 (0.77%)	33 (0.76%)	9626 (54.71%)
High MI	PPI/non-PPI with binder has high MI score (¿0.8)	Target	180 (0.51%)	33 (0.77%)	35 (0.81%)	4309 (24.49%)
		1 off-target	105 (0.30%)	24 (0.56%)	30 (0.69%)	6784 (38.56%)
		Both off-targets	33 (0.09%)	3 (0.07%)	17 (0.39%)	3531 (20.07%)
		At least 1 off-target	117 (0.33%)	24 (0.56%)	33 (0.76%)	7924 (45.04%)
		At least 1 interaction	256 (0.73%)	46 (1.07%)	64 (1.47%)	9945 (56.53%)
		Target + both off-targets	12 (0.03%)	0 (0.00%)	3 (0.07%)	1140 (6.48%)
None			34580 (98.36%)	4240 (98.42%)	4255 (97.95%)	0 (0.00%)
Total			35156	4308	4344	17593

Table S3. **SiameseCat Model Architecture and Training Hyperparameters.**

Hyperparameter	Value
Model Architecture	
ESM Model Base	ESM2_t33_650M_UR50D
Embedding Dimension	1280
ESM Unfrozen Layers	2
Linear Layers	2
Positional Attention Head: `n_heads`	10
Log-Sum Decoy Loss	
α	5
β	3
γ	1.25
ϵ_{weight}	1e-6
λ_{reg}	1e-4
Training	
Max Sequence Length	1800
Batch Size / Device	8
Gradient accumulation steps	2
Effective Batch Size	64
Dataloader `num_workers`	30
Learning Rate (LR)	1e-4
LR Scheduler: Warmup Steps	200
LR Scheduler: Total Steps	9000
LR Scheduler: Min/Max LR Ratio	0.1
Gradient Clipping	0.5
Embedding saving during training	
Embedding subset pairs	576

C. Supplementary Figures

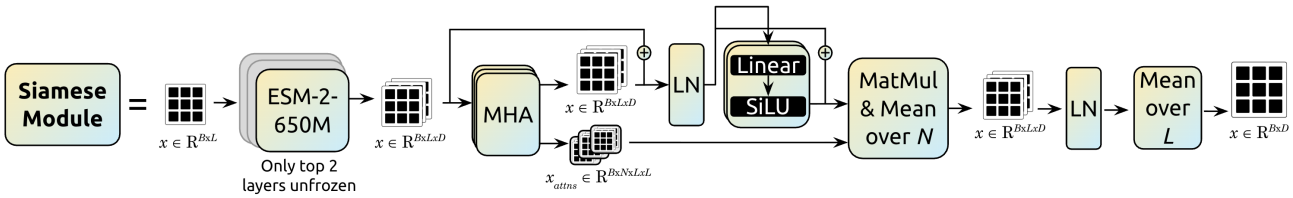

Figure S1. **Architecture of the Siamese Module within SiameseCat.**

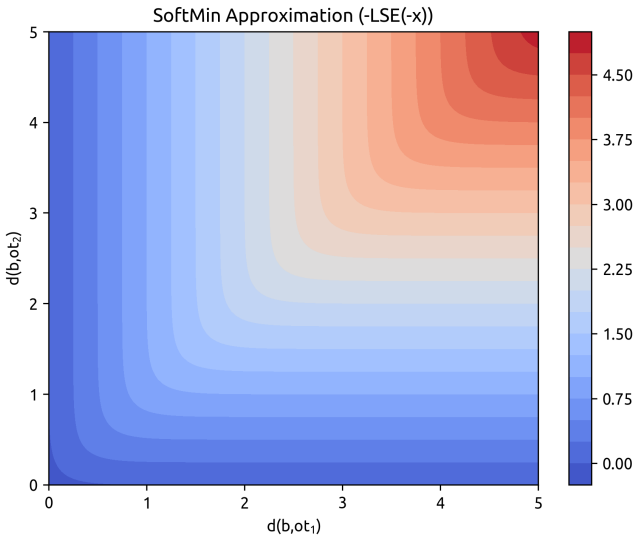

Figure S2. **Log-Sum-Exp computes a differentiable soft minimum distance of multiple off-targets.**

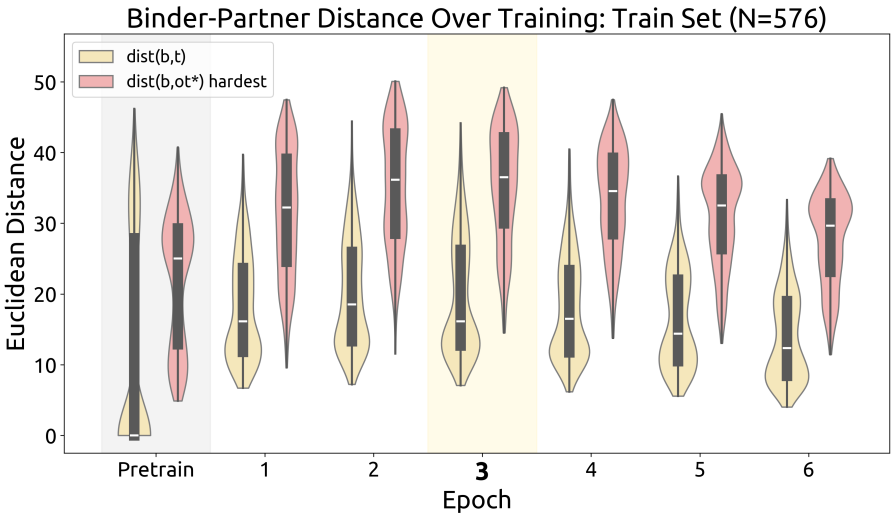

Figure S3. **Euclidean distance between binder and target, versus binder and hardest off-target, throughout training.**

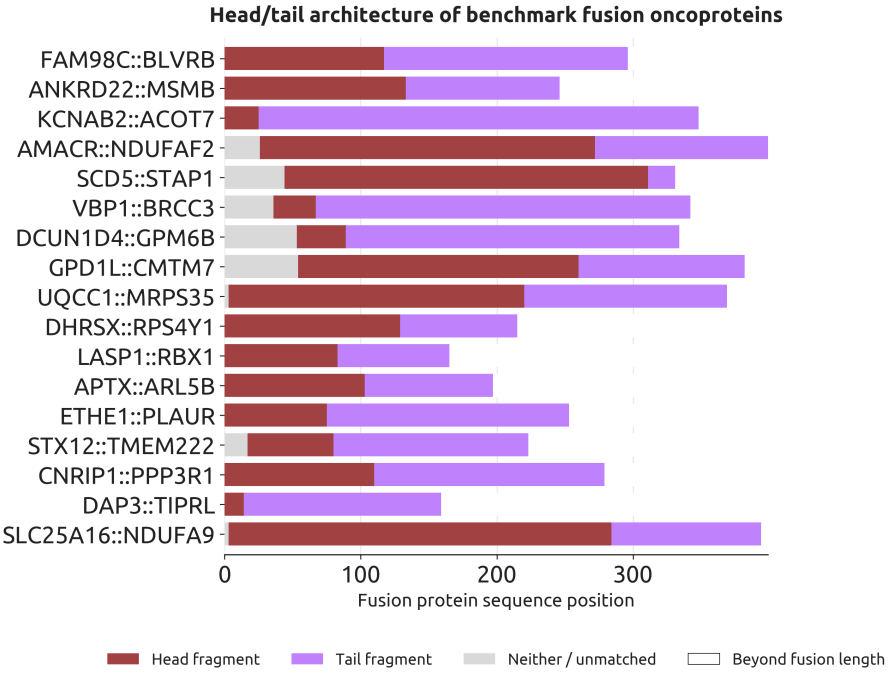

Figure S4. **Fusion oncoprotein composition for the structured fusion benchmark.**