

SOAPI: SIAMESE-GUIDED GENERATION OF OFF-TARGET-AVOIDING PROTEIN INTERACTIONS

Sophia Vincoff,¹ Oscar Davis,² Alexander Tong,^{3,4} Joey Bose,² Pranam Chatterjee^{1,5,†}

¹Department of Biomedical Engineering, Duke University, Durham, NC

²University of Oxford, Oxford, UK

³Mila – Quebec AI Institute, Montréal, Canada

⁴Université de Montréal, Montréal, Canada

⁵Department of Computer Science, Duke University, Durham, NC, USA

†Corresponding author: pranam.chatterjee@duke.edu

ABSTRACT

Therapeutics that modulate pathogenic proteins while avoiding off-target interactions are essential for effective drug design. However, designing binders that selectively engage a target protein while minimizing interactions with structurally or functionally similar proteins remains a major challenge. To address this, we introduce a Siamese-guided strategy for the generation of **Off** target-**A**voiding **P**rotein **I**nteractions, termed **SOAPI**. SOAPI leverages a Siamese protein language model with an adaptive Log-Sum-Exp Decoy Loss to enforce specificity by embedding fusion-specific binders close to their target while maintaining separation from off-targets. These optimized embeddings then guide a diffusion protein language model (DPLM), which generates binders using soft-value-based decoding (SVDD) and Sequential Monte Carlo resampling to iteratively refine candidates. *In silico* validation demonstrates significant off-target avoidance, highlighting SOAPI’s potential for generating precise and selective protein interactions.

1 INTRODUCTION

Selective modulation of pathogenic proteins is essential for drug design (Nada et al., 2024). Nonspecific drugs risk reduced efficacy due to off-target interference or, worse, disruption of cellular homeostasis (Garon et al., 2017). This challenge spans small molecules, PROTACs, and biologics. Since exhaustive *in vitro* screening is impractical, computational methods for predicting drug-target (DTIs) and protein-protein interactions (PPIs) are crucial. Structure-based methods like molecular docking can model stable proteins but are too slow for high-throughput analysis and ineffective for intrinsically disordered proteins (IDPs) (Chen et al., 2023b). To address this, state-of-the-art DTI (Singh et al., 2023; McNutt et al., 2024; Gao et al., 2024), PPI (Sledzieski et al., 2021; Singh et al., 2022), and peptide-protein predictors (Bhat et al., 2025; Tang et al., 2024) use contrastive learning to infer interactions directly from protein sequences and molecular representations.

More than 80% of protein targets lack structured binding pockets, making them “undruggable” by small molecules (Xie et al., 2023). Antibodies and peptides offer alternatives, but their efficacy depends on precise off-target avoidance (Chen et al., 2023b). Fusion oncoproteins, key drivers of pediatric and adult cancers, exemplify this challenge (Vincoff et al., 2025). These chimeric proteins, formed by chromosomal translocations, share high sequence homology with their wild-type counterparts, complicating selective targeting (Vincoff et al., 2025).

Recent protein language models (pLMs) have enabled peptide binder design for disordered targets (Bhat et al., 2025; Brixi et al., 2023; Chen et al., 2023a). However, existing approaches lack an explicit off-target exclusion mechanism. Most one-to-one PPI predictors select negatives randomly rather than based on experimental evidence (Bhat et al., 2025), and no existing models account for two off-targets simultaneously.

To address this, we introduce a Siamese-guided framework that integrates an adaptive Log-Sum-Exp Decoy Loss to enforce specificity in binder design. The Siamese model structures an embedding space where target-specific binders cluster near their target while maximizing separation from homologous off-targets. These embeddings guide a diffusion protein language model (DPLM) with soft-value-based decoding (SVDD) and Sequential Monte Carlo (SMC) resampling for iterative refinement (Wang et al., 2024), ensuring high-affinity binders while filtering out off-target interactions. *In silico* validation with AlphaFold-Multimer (Evans et al., 2021) confirms strong target engagement, improved interface quality, and minimal off-target binding. By combining the Siamese model’s discriminative power with DPLM’s generative flexibility, SOAPI provides a principled solution for designing highly selective protein interactions.

2 METHODS

Data curation and handling. All data curation, splitting, and clustering details can be found in the Supplementary Methods.

SOAPI architecture. SOAPI is a Siamese Neural Network that encodes four protein sequences—a binder, target, and two off-targets—into a shared latent space using pLM embeddings. Each sequence passes through the 33-layer ESM-2-650M encoder, with the final two layers unfrozen for training. A positional multi-head attention (MHA) module ($n_heads = 10$) with rotary positional embeddings (RoPE) (Su et al., 2024) enhances encoding of relative positional information across varying sequence lengths. Skip connections retain key embedding features, while two SiLU-activated linear layers model complex nonlinear relationships. Finally, attention pooling computes a weighted sum of token embeddings, producing fixed-length sequence representations. The architecture is presented in Figure 1.

Loss function. To handle two off-targets, we propose an Adaptive Log-Sum-Exp Decoy Loss (eq. (1)), where β controls the sharpness of the approximation, $\alpha > 0$ is the margin, and regularization term $\epsilon = 1e - 8$ prevents the logarithm from becoming too small. This formulation provides a differentiable loss that considers both off-target distances. Euclidean distances between proteins are represented by $\text{dist}(\mathbf{b}, \mathbf{t})$ (binder and target), $\text{dist}(\mathbf{b}, \mathbf{ot}_1)$ (binder and off-target 1), and $\text{dist}(\mathbf{b}, \mathbf{ot}_2)$ (binder and off-target 2).

$$\mathcal{L} = \sum_b \max(0, \text{dist}(\mathbf{b}, \mathbf{t}) - \frac{1}{\beta} \log(e^{\beta \cdot \text{dist}(\mathbf{b}, \mathbf{ot}_1)} + e^{\beta \cdot \text{dist}(\mathbf{b}, \mathbf{ot}_2)} + \epsilon) + \alpha) \quad (1)$$

Implementation details. SOAPI was implemented using PyTorch Lightning (Falcon, 2019) and trained on 4xA100 NVIDIA GPUs with an effective batch size of 128. The learning rate was initially set to $1e-4$ and modulated using a cosine annealing scheduler with 200 warmup steps. Full model hyperparameters can be found in Table A2. The model was trained until training loss and validation loss had stabilized for at least three epochs, resulting in 10 total training epochs.

Embedding separation visualization. At the conclusion of each epoch, the entire training set was passed through the model and SOAPI embeddings ($\mathbf{x} \in \mathbb{R}^{1280}$) were stored. For each quadruplet datum, the distance $d = \sqrt{\sum (x_i - y_i)^2}$ was calculated between the embeddings of the binder and target, binder and off-target 1, and binder and off-target 2. The resulting measurements $\text{dist}(\mathbf{b}, \mathbf{t})$, $\text{dist}(\mathbf{b}, \mathbf{ot}_1)$, and $\text{dist}(\mathbf{b}, \mathbf{ot}_2)$ were visualized every three epochs throughout training using `matplotlib v3.8.2`.

Binder recovery screen. A binder recovery test was performed by screening 50,000 candidate binders through SOAPI for each target in the test set. The `random` package in Python was used to generate 50,000 random combinations of the 20 natural amino acids, with minimum and maximum lengths corresponding to the shortest and longest binders in the test set (56 and 856 amino acids), respectively. SOAPI was tasked to rank the true binder among the top N for all $N \in \{1, 10, 100, 500, 5000, 25000\}$. Rankings were based on an adapted version of the loss function (Equation (2)), where rank increased as D decreased:

$$D = \text{dist}(\mathbf{b}, \mathbf{t}) - \frac{1}{\beta} \log(e^{\beta \cdot \text{dist}(\mathbf{b}, \mathbf{ot}_1)} + e^{\beta \cdot \text{dist}(\mathbf{b}, \mathbf{ot}_2)} + \epsilon) + \alpha \quad (2)$$

Results for each threshold are plotted in Figure 3.

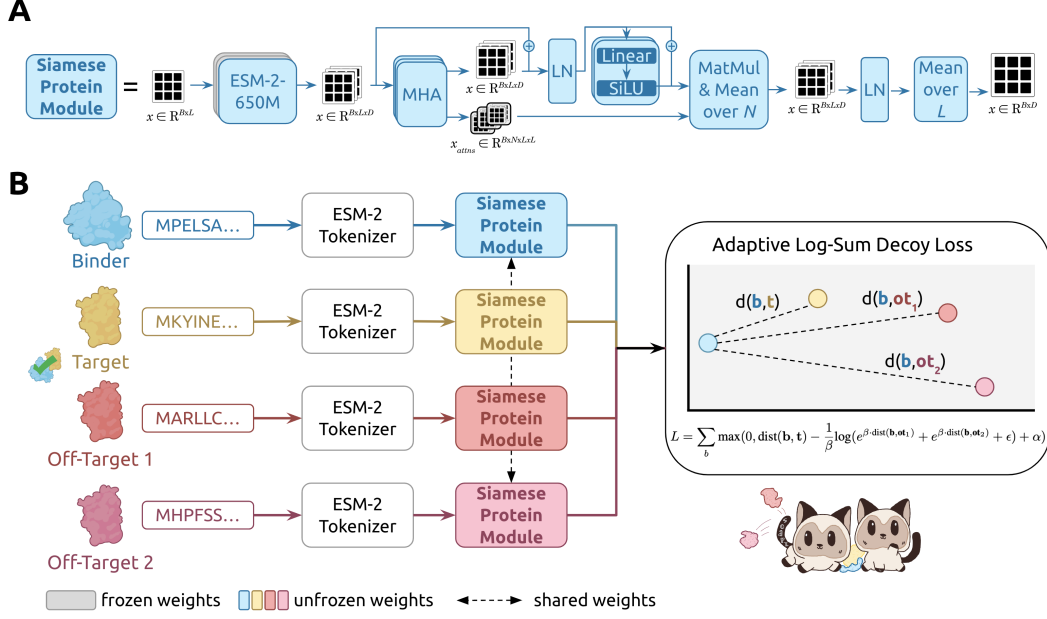


Figure 1: **SOAPI model architecture.** **A** Siamese Protein Module, the core unit of the quadruplet network. ESM-2-650M encodes sequences into $[B, 1280]$ embeddings, refined via positional attention (10 heads) with rotary embeddings, skip-connected linear layers, and attention pooling. **B** Full SOAPI pipeline. Binder, target, and off-target sequences pass through the Siamese module with shared weights. Euclidean distances between embeddings define a loss that pulls the binder toward the target while pushing it away from off-targets.

Masked discrete diffusion. SOAPI employs a masked discrete diffusion model (MDM) (Sahoo et al., 2024; Shi et al., 2024; Gat et al., 2024; Zhao et al., 2024) to iteratively refine protein binder sequences. The forward process corrupts a sequence \mathbf{x}_0 by progressively masking tokens with transition probability:

$$p_t(\mathbf{x}_t | \mathbf{x}_0) = \prod_{i=1}^n \text{Cat}(x_t^i; \alpha_t \delta(x_0^i) + (1 - \alpha_t) \delta(m)), \quad (3)$$

where α_t is a noise schedule and m is a masked token. The reverse process reconstructs the sequence by sampling:

$$p_t(x_{t-1}^i | x_t^i, x_0^i) \propto \alpha_{t-1} \delta(x_0^i) + (1 - \alpha_{t-1}) \delta(m). \quad (4)$$

The model is trained using a weighted cross-entropy loss:

$$\mathcal{L} = -\mathbb{E}_{\mathbf{x}_t \sim p_t} \sum_{i=1}^n x_0^i \log \mu_\theta(x_t^i, t). \quad (5)$$

Siamese-guided sampling. Binder sequences are generated using a diffusion protein language model (DPLM) (Wang et al., 2024), guided by the Siamese model’s specificity loss. Soft-value-based decoding (SVDD) (Li et al., 2024) assigns resampling weights:

$$w_t = \exp(\mathcal{L}_t - \hat{\mathcal{L}}_{t-1}), \quad (6)$$

steering diffusion toward high-specificity binders via Sequential Monte Carlo (SMC) resampling.

A full description of MDMs and sampling can be found in the Appendix.

Binder generation. For binder generation, we use SVDD (Li et al., 2024), which performs particle filtering using Sequential Monte Carlo (SMC) (Doucet et al., 2001) on intermediate trajectories at inference using a soft-value function. For binder generation using SVDD, we use the difference

in loss under the pre-trained Siamese model as the soft-value function. In essence, this soft-value function captures how close a proposed binder is to a target protein and simultaneously how far it is from the two off-targets. Using this, we steer an unconditional DPLM protein language model (Wang et al., 2024) to produce generated protein sequences with higher binding efficacy by up-weighting promising (according to the soft-value) intermediate trajectories by performing SMC, which resamples points using the soft-value function. Precisely, the soft-value function we use is the difference in the SOAPI loss before and after the denoising step of the DPLM protein model:

$$w_{t-1}^{(m)} = \exp \left(\mathcal{L}_t^{(m)} - \hat{\mathcal{L}}_{t-1}^{(m)} \right), \quad m \in [M], \quad (7)$$

following the notation in Li et al. (2024). Here $\mathcal{L}_t^{(m)}$ is the loss of the m -th sample at step t , and $\hat{\mathcal{L}}_{t-1}^{(m)}$ is the loss following the denoising step at time $t - 1$ (the next step in the reverse process), while M is the total number of particles.

3 RESULTS

SOAPI effectively learned to distinguish binders from off-targets, embedding true binders closer to their targets while maintaining separation from non-interacting proteins in its latent space (Figure 2). The model achieved consistently low loss across training, validation, and test sets, with no signs of overfitting (Table A1). Visualization of SOAPI’s learned representations across training epochs further confirmed a progressive improvement in the embedding space, with a clear separation between binder-target pairs and off-target interactions (Figure 2). This structured organization of latent space ensures that SOAPI-generated binders maintain high specificity.

To assess SOAPI’s ability to rank binders based on target specificity, we evaluated its performance against 50,000 randomly generated candidates. In the top-1 ranking test, SOAPI correctly identified the true binder for 19.2% of targets, with performance improving to 24.0% in the top-10 and reaching 56.8% in the top-25% evaluation (Figure 3).

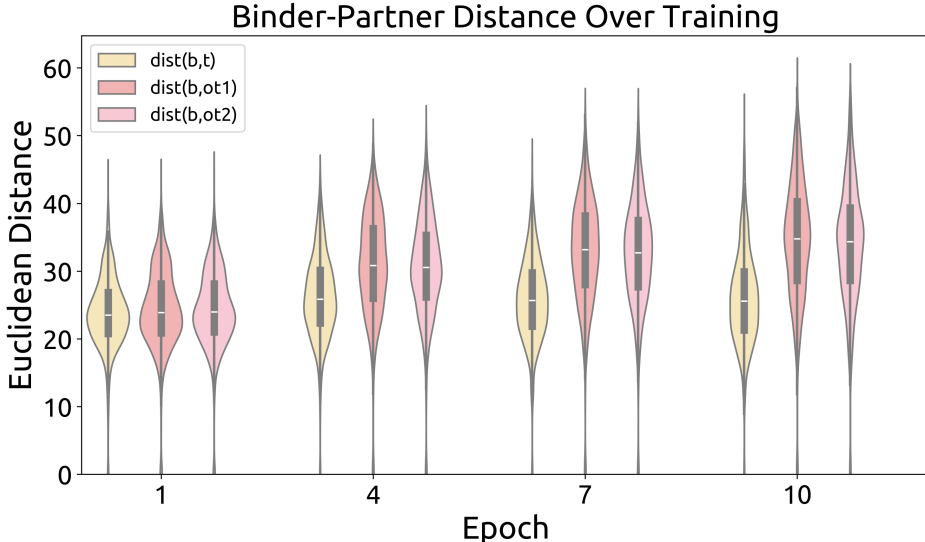


Figure 2: **Embedding separation throughout model training.** Euclidean distances between SOAPI embeddings of the proteins in each training quadruplet: $\text{dist}(\mathbf{b}, \mathbf{t})$ (binder and target), $\text{dist}(\mathbf{b}, \mathbf{ot}_1)$ (binder and off-target 1), and $\text{dist}(\mathbf{b}, \mathbf{ot}_2)$ (binder and off-target 2). The inner box of each violin plot indicates the median (white) and inter-quartile range, representing the middle 50% of distances (grey box). Distances are plotted every three epochs throughout training, starting at the end of epoch 1.

We report our results for binder generation using SVDD in Table 1. Due to computational constraints, only a small subset of generated binders was evaluated, as each requires *in silico* co-folding to assess structural compatibility and binding specificity.

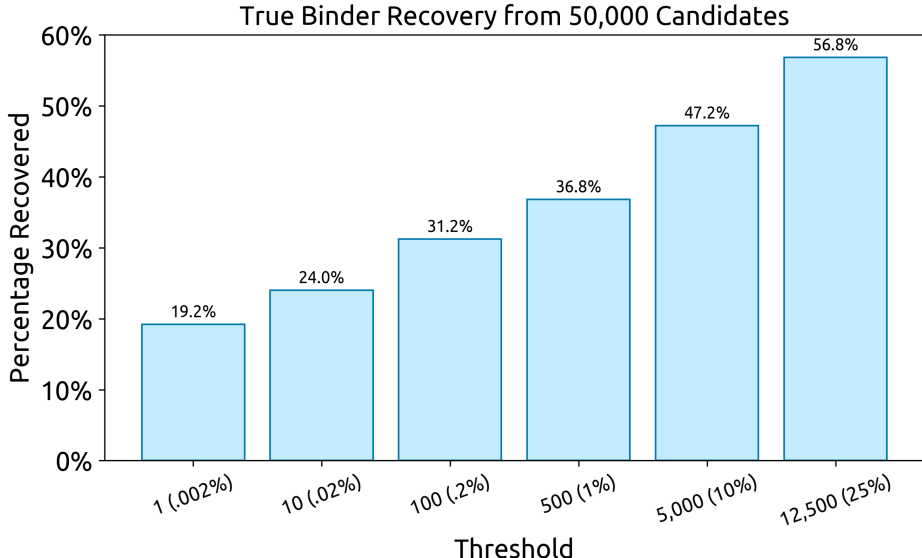


Figure 3: **SOAPI ranks true binders against 50,000 randomly generated candidates.** SOAPI produced a specificity-based ranking of 50,000 randomly generated binders and one true binder to 125 targets from the test set. Relative distance metrics D were calculated using Equation (2). Any target where SOAPI ranked the true binder among the top- N or top- $N\%$ was considered a hit. Top-1, top-10, top-100, top-1%, top-10%, and top-25% evaluations were conducted.

Relative interface pTM scores from AlphaFold-Multimer show that designed binders preferentially engage their target over off-targets, confirming effective specificity enforcement. The corresponding sequences for each target are listed in Table A3, demonstrating a range of high-specificity binders designed using our approach.

Target	Target pLDDT (\uparrow)	Relative ipTM 1 (\uparrow)	Relative ipTM 2 (\uparrow)
APTX : : ARL5B	83.09	0.43	0.21
DAP3 : : TIPRL	86.80	0.11	0.14
VBP1 : : BRCC3	71.72	0.34	0.14
VBP1 : : BRCC3	71.72	0.21	0.21

Table 1: Results on conditional binder generation. The relative ipTMs are the differences in ipTMs between the target and the off-targets 1 and 2, respectively.

4 DISCUSSION

SOAPI introduces a principled framework for generating protein binders with high specificity while minimizing off-target interactions. By integrating a Siamese protein language model with an adaptive Log-Sum-Exp Decoy Loss, we enforce specificity constraints that separate target-specific binders from homologous off-targets. This guided embedding space is then leveraged by DPLM with SVDD to iteratively refine binder candidates. Our results demonstrate strong target engagement and significant off-target avoidance in *in silico* validation.

Future work will expand SOAPI’s benchmarking with more extensive *in silico* analysis and *in vitro* validation against diverse targets, including fusion oncoproteins. To enhance fusion oncoprotein selectivity, we will incorporate FusOn-pLM embeddings (Vincoff et al., 2025), which capture fusion-specific sequence and structural properties. Additionally, we will evaluate alternative sampling strategies, such as Path Planning (P2) (Peng et al., 2025), to improve the efficiency and precision of masked discrete diffusion in binder generation. These advancements will further establish SOAPI as a robust tool for engineering highly selective therapeutics.

REFERENCES

- 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), January 2025. ISSN 2375-2548. doi: 10.1126/sciadv.adr8638. URL <http://dx.doi.org/10.1126/sciadv.adr8638>. (Cited on page 1)
- Philipp Blohm, Goar Frishman, Pawel Smialowski, Florian Goebels, Benedikt Wachinger, Andreas Ruepp, and Dmitrij Frishman. 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. (Cited on page 10)
- Garyk Brixi, Tianzheng Ye, Lauren Hong, Tian Wang, Connor Monticello, Natalia Lopez-Barbosa, Sophia Vincoff, Vivian Yudistyra, Lin Zhao, Elena Haarer, Tianlai Chen, Sarah Pertsemlidis, Kalyan Palepu, Suhaas Bhat, Jayani Christopher, Xinning Li, Tong Liu, Sue Zhang, Lillian Petersen, Matthew P. DeLisa, and Pranam Chatterjee. Saltamp;peppr is an interface-predicting language model for designing peptide-guided protein degraders. *Communications Biology*, 6(1), October 2023. ISSN 2399-3642. doi: 10.1038/s42003-023-05464-z. URL <http://dx.doi.org/10.1038/s42003-023-05464-z>. (Cited on page 1)
- Anton Bushuiev, Roman Bushuiev, Anatolii Filkin, Petr Kouba, Marketa Gabrielova, Michal Gabriel, Jiri Sedlar, Tomas Pluskal, Jiri Damborsky, Stanislav Mazurenko, et al. Learning to design protein-protein interactions with enhanced generalization. *arXiv preprint arXiv:2310.18515*, 2023. (Cited on page 10)
- Tianlai Chen, Madeleine Dumas, Rio Watson, Sophia Vincoff, Christina Peng, Lin Zhao, Lauren Hong, Sarah Pertsemlidis, Mayumi Shaepers-Cheu, Tian Zi Wang, Divya Srijay, Connor Monticello, Pranay Vure, Rishab Pulugurta, Kseniia Kholina, Shrey Goel, Matthew P. DeLisa, Ray Truant, Hector C. Aguilar, and Pranam Chatterjee. Pepmlm: Target sequence-conditioned generation of therapeutic peptide binders via span masked language modeling. *arXiv*, 2023a. doi: 10.48550/ARXIV.2310.03842. URL <https://arxiv.org/abs/2310.03842>. (Cited on page 1)
- Tianlai Chen, Lauren Hong, Vivian Yudistyra, Sophia Vincoff, and Pranam Chatterjee. Generative design of therapeutics that bind and modulate protein states. *Current Opinion in Biomedical Engineering*, 28:100496, December 2023b. ISSN 2468-4511. doi: 10.1016/j.cobme.2023.100496. URL <http://dx.doi.org/10.1016/j.cobme.2023.100496>. (Cited on page 1)
- Noemi Del Toro, Anjali Shrivastava, Eliot Ragueneau, Birgit Meldal, Colin Combe, Elisabet Barrera, Livia Perfetto, Karyn How, Prashansa Ratan, Gautam Shirodkar, et al. The intact database: efficient access to fine-grained molecular interaction data. *Nucleic acids research*, 50(D1):D648–D653, 2022. (Cited on page 10)
- Arnaud Doucet, Nando De Freitas, Neil James Gordon, et al. *Sequential Monte Carlo methods in practice*, volume 1. Springer, 2001. (Cited on page 3)
- Richard Evans, Michael O’Neill, Alexander Pritzel, Natasha Antropova, Andrew Senior, Tim Green, Augustin Židek, Russ Bates, Sam Blackwell, Jason Yim, Olaf Ronneberger, Sebastian Bodenstein, Michal Zielinski, Alex Bridgland, Anna Potapenko, Andrew Cowie, Kathryn Tunyasuvunakool, Rishub Jain, Ellen Clancy, Pushmeet Kohli, John Jumper, and Demis Hassabis. 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>. (Cited on page 2)
- William A Falcon. Pytorch lightning. *GitHub*, 3, 2019. (Cited on page 2)
- Bowen Gao, Bo Qiang, Haichuan Tan, Yinjun Jia, Minsi Ren, Minsi Lu, Jingjing Liu, Wei-Ying Ma, and Yanyan Lan. Drugclip: Contrastive protein-molecule representation learning for virtual screening. *Advances in Neural Information Processing Systems*, 36, 2024. (Cited on page 1)

- Sarah L. Garon, Rebecca K. Pavlos, Katie D. White, Nancy J. Brown, Cosby A. Stone, and Elizabeth J. Phillips. 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>. (Cited on page 1)
- Itai Gat, Tal Remez, Neta Shaul, Felix Kreuk, Ricky TQ Chen, Gabriel Synnaeve, Yossi Adi, and Yaron Lipman. Discrete flow matching. *arXiv preprint arXiv:2407.15595*, 2024. (Cited on pages 3 and 9)
- 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, 2024. URL <https://arxiv.org/abs/2408.08252>. (Cited on pages 3 and 4)
- Zeming Lin, Halil Akin, Roshan Rao, Brian Hie, Zhongkai Zhu, Wenting Lu, Nikita Smetanin, Robert Verkuil, Ori Kabeli, Yaniv Shmueli, et al. Evolutionary-scale prediction of atomic-level protein structure with a language model. *Science*, 379(6637):1123–1130, 2023. (Cited on page 10)
- Andrew T McNutt, Abhinav K Adduri, Caleb N Ellington, Monica T Dayao, Eric P Xing, Hosein Mohimani, and David R Koes. Sprint enables interpretable and ultra-fast virtual screening against thousands of proteomes. *arXiv preprint arXiv:2411.15418*, 2024. (Cited on page 1)
- Hossam Nada, Yongseok Choi, Sungdo Kim, Kwon Su Jeong, Nicholas A. Meanwell, and Kyeong Lee. 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>. (Cited on page 1)
- Rose Oughtred, Jennifer Rust, Christie Chang, Bobby-Joe Breitkreutz, Chris Stark, Andrew Willems, Lorrie Boucher, Genie Leung, Nadine Kolas, Frederick Zhang, et al. The biogrid database: A comprehensive biomedical resource of curated protein, genetic, and chemical interactions. *Protein Science*, 30(1):187–200, 2021. (Cited on page 10)
- Fred Zhangzhi Peng, Zachary Bezemek, Sawan Patel, Jarrod Rector-Brooks, Sherwood Yao, Alexander Tong, and Pranam Chatterjee. Path planning for masked diffusion model sampling. *arXiv*, 2025. doi: 10.48550/ARXIV.2502.03540. URL <https://arxiv.org/abs/2502.03540>. (Cited on page 5)
- Subham Sekhar Sahoo, Marianne Arriola, Yair Schiff, Aaron Gokaslan, Edgar Marroquin, Justin T Chiu, Alexander Rush, and Volodymyr Kuleshov. Simple and effective masked diffusion language models. *arXiv preprint arXiv:2406.07524*, 2024. (Cited on pages 3 and 9)
- Jiaxin Shi, Kehang Han, Zhe Wang, Arnaud Doucet, and Michalis K Titsias. Simplified and generalized masked diffusion for discrete data. *arXiv preprint arXiv:2406.04329*, 2024. (Cited on pages 3 and 9)
- Rohit Singh, Kapil Devkota, Samuel Sledzieski, Bonnie Berger, and Lenore Cowen. Topsy-turvy: integrating a global view into sequence-based ppi prediction. *Bioinformatics*, 38(Supplement_1): i264–i272, 2022. (Cited on page 1)
- Rohit Singh, Samuel Sledzieski, Bryan Bryson, Lenore Cowen, and Bonnie Berger. Contrastive learning in protein language space predicts interactions between drugs and protein targets. *Proceedings of the National Academy of Sciences*, 120(24):e2220778120, 2023. (Cited on page 1)
- Samuel Sledzieski, Rohit Singh, Lenore Cowen, and Bonnie Berger. 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. (Cited on page 1)
- Martin Steinegger and Johannes Söding. Mmseqs2 enables sensitive protein sequence searching for the analysis of massive data sets. *Nature biotechnology*, 35(11):1026–1028, 2017. (Cited on page 10)
- Jianlin Su, Murtadha Ahmed, Yu Lu, Shengfeng Pan, Wen Bo, and Yunfeng Liu. Roformer: Enhanced transformer with rotary position embedding. *Neurocomputing*, 568:127063, 2024. (Cited on page 2)

- Sophia Tang, Yinuo Zhang, and Pranam Chatterjee. Peptune: De novo generation of therapeutic peptides with multi-objective-guided discrete diffusion. *arXiv*, 2024. doi: 10.48550/ARXIV.2412.17780. URL <https://arxiv.org/abs/2412.17780>. (Cited on page 1)
- Sophia Vincoff, Shrey Goel, Kseniia Kholina, Rishab Pulugurta, Pranay Vure, and Pranam Chatterjee. Fuson-plm: a fusion oncoprotein-specific language model via adjusted rate masking. *Nature Communications*, 16(1), February 2025. ISSN 2041-1723. doi: 10.1038/s41467-025-56745-6. URL <http://dx.doi.org/10.1038/s41467-025-56745-6>. (Cited on pages 1 and 5)
- Xinyou Wang, Zaixiang Zheng, Fei Ye, Dongyu Xue, Shujian Huang, and Quanquan Gu. Diffusion language models are versatile protein learners. *International Conference on Machine Learning*, 2024. doi: 10.48550/ARXIV.2402.18567. URL <https://arxiv.org/abs/2402.18567>. (Cited on pages 2, 3, and 4)
- Xin Xie, Tingting Yu, Xiang Li, Nan Zhang, Leonard J. Foster, Cheng Peng, Wei Huang, and Gu He. Recent advances in targeting the “undruggable” proteins: from drug discovery to clinical trials. *Signal Transduction and Targeted Therapy*, 8(1), September 2023. ISSN 2059-3635. doi: 10.1038/s41392-023-01589-z. URL <http://dx.doi.org/10.1038/s41392-023-01589-z>. (Cited on page 1)
- Lingxiao Zhao, Xueying Ding, Lijun Yu, and Leman Akoglu. Unified discrete diffusion for categorical data. *arXiv preprint arXiv:2402.03701*, 2024. (Cited on pages 3 and 9)

A APPENDIX

A.1 PRELIMINARIES

A.1.1 MASKED DISCRETE DIFFUSION MODELS

We now outline the model class known as masked discrete diffusion models (MDMs) (Sahoo et al., 2024; Shi et al., 2024; Gat et al., 2024; Zhao et al., 2024) which operates on a vocabulary of $\mathcal{X} = \{1, \dots, d-1\}$ and an additional masked token \mathbf{m} . The forward masking process is designed to take a clean sequence x_0 of size n at time $t = 0$ to a sequence of fully masked tokens $p_1 = [\delta(\mathbf{m})]^n$ at time $t = 1$. Formally, this forward process can be written analytically given a transition kernel $p_t(x_t^i|x_0^i)$:

$$p_t(x_t|x_0) = \prod_{i=1}^n p_t(x_t^i|x_0^i) = \prod_{i=1}^n \text{Cat}(x_t^i; \alpha_t \delta(x_0^i) + (1 - \alpha_t) \delta(\mathbf{m})), \quad (8)$$

where we used a superscript x^i to denote the i -th index of a sequence of length n , while the subscript t denotes a discretized timestep $t \in [0, 1]$. Furthermore, α_t is an invertible reparameterization of time such that $\alpha_0 = 1$ and $\alpha_1 = 0$ and can be understood as the noise schedule. Indeed, once a token is masked in the forward process it remains masked for the remainder of the process—i.e. until $t = 1$.

The reverse process associated with the masked forward process iteratively denoises a sample such that at $t = 0$ all tokens in the sequence are unmasked. The reverse transition probability $p_t(x_{t-1}^i|x_t^i, x_0^i)$ conditioned on x_0^i and x_t^i is given by the following expression which is the Bayesian posterior:

$$p_t(x_{t-1}^i|x_t^i, x_0^i) = \begin{cases} \text{Cat}(x_{t-1}^i; \delta(x_t^i)) & x_t^i \neq \mathbf{m} \\ \text{Cat}\left(x_{t-1}^i; \frac{(1-\alpha_{t-1})\delta(\mathbf{m}) + (\alpha_{t-1}-\alpha_t)\delta(x_0^i)}{1-\alpha_t}\right) & x_t^i = \mathbf{m}. \end{cases} \quad (9)$$

In an analogous manner to the forward process, once a token transitions out of the masked state for a time $t > 0$ it remains in this state for the remainder of the trajectory. The reverse process suggests a parametrization for an MDM that predicts the mean of this denoising posterior in Equation (9). More precisely, $\mu_\theta : \mathcal{X}^n \times \mathbb{R} \rightarrow (\Delta^d)^n$, which predicts the clean sample at $t = 0$ by denoising a noisy x_t^i ,

$$q_{t,\theta}(x_{t-1}^i|x_t^i, \mu_\theta(x_t^i, t)) = \begin{cases} \text{Cat}(x_{t-1}^i; \delta(x_t^i)) & x_t^i \neq \mathbf{m} \\ \text{Cat}\left(x_{t-1}^i; \frac{(1-\alpha_{t-1})\delta(\mathbf{m}) + (\alpha_{t-1}-\alpha_t)\mu_\theta(x_t^i, t)}{1-\alpha_t}\right) & x_t^i = \mathbf{m}. \end{cases} \quad (10)$$

Using this mean-parametrization we can form a simple objective for learning that forms a valid *evidence lower bound* on the log marginal of the clean data distribution $\log p(x_0)$. Specifically, for discrete diffusion models, this operationally amounts to a weighted cross-entropy loss:

$$\log p(x_0) \geq - \int_0^1 \frac{d\alpha_t}{dt} \cdot \frac{1}{1 - \alpha_t} \mathbb{E}_{x_t \sim p_t(x_t|x_0)} \left[\sum_{i=1}^n (x_0^i)^T \log \mu_\theta(x_t^i, t) \right] dt. \quad (11)$$

We will use Equation (11) as an objective to pre-train MDMs before controlling them at inference, e.g. classifier-based guidance, with task-specific information.

A.2 SUPPLEMENTARY METHODS

Data collection. Each sample in the training data is a protein quadruplet consisting of a binder, target, off-target 1, and off-target 2. Positive protein-protein interactions (PPIs) were collected from BioGRID (October 2022) (Oughtred et al., 2021), IntAct (October 2022) (Del Toro et al., 2022), and PPIRef (January 2025) (Bushuiev et al., 2023). Negative interactions were collected from Negatome2.0 (Blohm et al., 2014), a manually curated database of proteins with experimental evidence indicating the absence of direct interaction. Positive PPIs were filtered by cleaning (e.g. dropping sequences with non-natural amino acids), swapping targets and binders to double the dataset size, removing duplicate homomer interactions, applying a length limit of 1022 amino acids for both target and binder sequences, and only retaining rows where one partner is included in the Negatome. The Negatome was filtered by removing any listed target and off-target pairs that are included as binder and target pairs in the PPI database. This produced a dataset of 245,587 positive PPIs and 7,198 negative interactions.

Quadruplet selection. The training, validation, and testing datasets were designed to (1) enhance generalizability by including diverse binders, targets, and off-targets, (2) prevent rigid role assignments by allowing any sequence to act in any interaction context, and (3) maximize difficulty to improve learning. Quadruplet selection was formulated as a linear programming problem using length-averaged ESM-2-650M (Lin et al., 2023) embeddings. With PuLP v2.9.0, quadruplets were optimized for difficulty while minimizing role repetition (Supplementary Algorithm 1). Difficulty increased with higher cosine similarity between target and off-target embeddings, ensuring the Siamese model learned to distinguish subtle differences. Euclidean distance was used in the loss function, but cosine similarity was preferred for selection due to its bounded range (-1 to 1). Selecting closely related targets and off-targets better reflects SOAPI’s real-world applications, such as designing binders that avoid wild-type protein interactions. To ensure the model learned relationships rather than predefined roles, four constraints were imposed: (1) each binder appears at most ten times, (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 pre-enforced by subsampling positive PPIs, and Constraint 2 required a tiebreaker term ($\lambda = 0.001$) based on Euclidean binder-target distance. A total of 1,352 quadruplets were selected, consisting of 565 unique binders, 1,054 targets, and 969 off-targets.

Clustering and splitting. Quadruplets were clustered on binder sequence using MMSeqs2 easy clustering module (Steinegger & Söding, 2017) with a minimum sequence identity of 30% and a coverage threshold of 70%. The resulting clusters were randomly split at 80-10-10 ratio using sklearn v1.2.0 into a training set (1,111 quadruplets, 82.2%), validation set (116 quadruplets, 8.6%), and test set (125 quadruplets, 9.2%). The distribution of roles (binder, target, off-target) played by each sequence in the full dataset and individual splits can be found in Figure A1.

A.3 SUPPLEMENTARY TABLES

Table A1: SOAPI loss on training data.

Split	Size	Loss
Train	1111	0.04
Validation	116	1.43
Test	125	1.57

Table A2: SOAPI 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: <code>n_heads</code>	10
Adaptive Log-Sum Decoy Loss	
α	5
β	0.5
ϵ	1e-8
Training	
Max Sequence Length	1022
Batch Size / Device	16
Effective Batch Size	128
Dataloader <code>num_workers</code>	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

Target	Proposed sequence
APT _X : : ARL5B	FNLKNKVLITGSRGIGKGIAELAKEGVNLIARTLE KKEKSKVLKDVTDLINIFEKAKQFNNKIDILVNNA
DAP ₃ : : TIPRL	MAASPSHLSSSSSLHLSPRRLGPTAAEAATAL RFRPRNVPLTPFRPSRRSSLSPASPSSAPPPQTT
VBP ₁ : : BRCC3	MGCQFNPTPQRTALWLCGTDPLPLPLPLQL PGAVAAVAADSVYVREVQLGRGPGLLEGEAQ
VBP ₁ : : BRCC3	YLDQRRQGAQEAGSESVLKKCWTAERLNAP ASRRAAQSVGGDVPRVVAAGMSLLCLLHW

Table A3: Best generated sequences for the conditional generation experiment.

A.4 SUPPLEMENTARY FIGURES

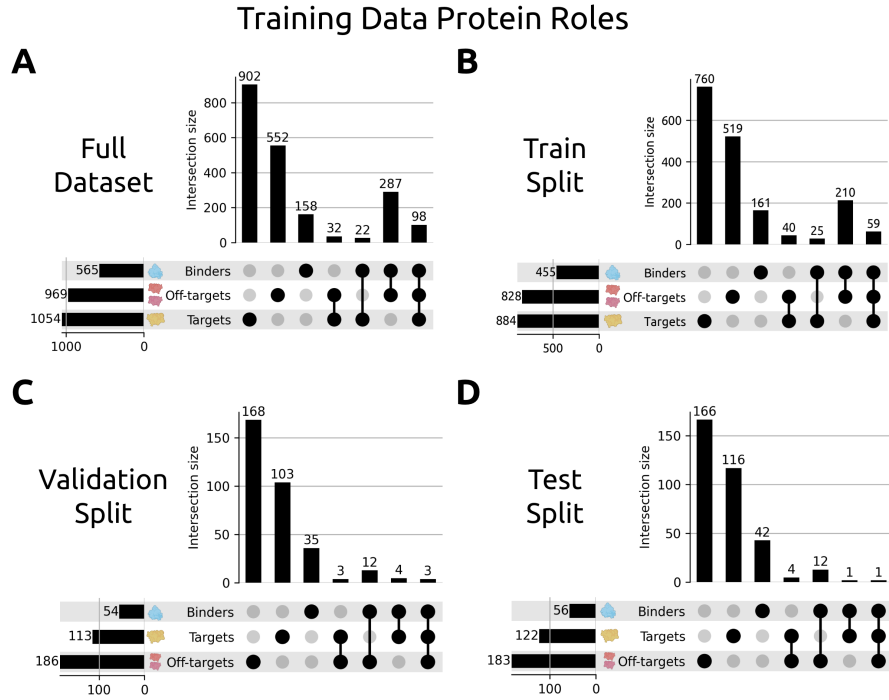


Figure A1: **Training Data Protein Roles.** Distribution of roles (binder, target, off-target) played by proteins in the quadruplets comprising the **A** full training dataset (1,352 quadruplets), **B** train split (1,111 quadruplets), **C** validation split (116 quadruplets), and **D** test split (125 quadruplets).

Algorithm 1 Quadruplet Selection via Linear Programming

-
- 1: **Input:** Set of positive and negative PPIs P_{pos} and P_{neg} , cosine similarity matrix of target t and off-target ot embeddings C , Euclidean distance matrix of binder b and target t embeddings E , small weighting factor λ , and constraints: m_{group} (max occurrences of a specific (t, ot_1, ot_2) group), m_b (max binder occurrences), m_{tpb} (max repeats of same target per binder), m_{otpb} (max repeats of same off-target per binder).
 - 2: **Output:** Q^* , optimal quadruplets (b, t, ot_1, ot_2) satisfying the constraints.
 - 3: $P_{\text{pos}}^* \leftarrow \text{RandomSample}(P_{\text{pos}}, m_b)$
 - 4: $Q \leftarrow (b, t, ot_1, ot_2) \forall (b, ot_1, ot_2) \in P_{\text{neg}}, (b, t) \in P_{\text{pos}}^*$
 - 5: **for** each quadruplet (b, t, ot_1, ot_2) **do**
 - 6: $d(b, t) \leftarrow E_{b,t}$
 - 7: $s(t, ot_1) \leftarrow C_{t,ot_1}$
 - 8: $s(t, ot_2) \leftarrow C_{t,ot_2}$
 - 9: Compute adjusted similarity:

$$S(b, t, ot_1, ot_2) = \frac{s(t, ot_1) + s(t, ot_2)}{2} + \lambda \cdot d(b, t).$$

- 10: **end for**
- 11: Define decision variables $x(b, t, ot_1, ot_2) \in \{0, 1\}$ for all candidate quadruplets.
- 12: Maximize:

$$\sum_{(b,t,ot_1,ot_2)} x(b, t, ot_1, ot_2) \cdot S(b, t, ot_1, ot_2).$$

- 13: Subject to:
 - $\sum_{(ot_1, ot_2)} x(b, t, ot_1, ot_2) \leq m_{\text{tpb}} \quad \forall b, t.$
 - $\sum_{t, ot_2} x(b, t, ot_1, ot_2) + \sum_{t, ot_1} x(b, t, ot_2, ot_1) \leq m_{\text{otpb}} \quad \forall b, ot.$
 - $\sum_b x(b, t, ot_1, ot_2) + x(b, t, ot_2, ot_1) \leq m_{\text{group}} \quad \forall t, ot_1, ot_2.$
 - 14: Solve the linear programming problem.
 - 15: Extract and return Q^* , selected quadruplets where $x(b, t, ot_1, ot_2) = 1$.
-