CupOFMoCA: Coupled Objective-Guided Discrete Flows for Molecular Conjugate Assembly

Ruoxi Rosie Zhang[1] **Pranam Chatterjee**[1,2,†]
[1]Department of Computer and Information Science, University of Pennsylvania
[2]Department of Bioengineering, University of Pennsylvania
[†]Correspondence to: pranam@seas.upenn.edu

Abstract

Molecular conjugates, including PROTACs and peptide-drug conjugates (PDCs), derive their function from the joint behavior of multiple coupled components, yet most generative approaches design these components independently and combine them only after generation. Such staged pipelines ignore cross-component dependencies and often produce conjugates that are chemically invalid or lose function upon assembly. We introduce **Co**u**p**led **O**bjective-Guided Discrete **F**lows for **Mo**lecular **C**onjugate **A**ssembly (**CupOFMoCA**), a discrete generative framework that formulates conjugate design as a constrained, coupled generation problem. CupOFMoCA operates by restricting discrete generative trajectories to a chemically feasible conjugate manifold and biasing local transitions using a target-specific binding predictor, ensuring that peptide, linker, and drug components remain compatible throughout generation. Focusing on PDCs, we show that coupling constraints and objective guidance enable anticipatory design that preserves post-conjugation predicted binding and produces structurally realistic conjugates, outperforming staged baselines that attach drugs to independently generated peptides. These results demonstrate that explicit coupling and constraint enforcement are sufficient to recover functional molecular conjugates and provide a principled foundation for generative modeling in settings where function emerges only at the level of the assembled system.

1 Introduction

Molecular conjugates play a central role in modern therapeutics and chemical biology by integrating complementary functional components into a single molecular entity Cooper et al. (2021); Hua et al. (2022). Conjugation strategies are widely used to improve specificity, modulate pharmacokinetics, expand chemical reach, and access mechanisms unavailable to individual components alone (Marei et al., 2022; Eras et al., 2022; Drago et al., 2021). Representative examples include peptide-drug conjugates (PDCs), which combine sequence-based target recognition with small-molecule potency (Rizvi et al., 2024; Heh et al., 2023; Dean et al., 2024), proteolysis-targeting chimeras (PROTACs) that induce proximity between a target protein and the ubiquitin-proteasome machinery (Vikal et al., 2025; Wang et al., 2024a), and other bifunctional or modular conjugates that integrate sensing, delivery, and actuation (Hua et al., 2022; Conway, 2020). In all of these settings, function arises from the coupled behavior of the assembled conjugate rather than from its components in isolation, making conjugate design an inherently joint problem.

Structure-based drug discovery using three-dimensional protein-ligand complexes is a mature and widely adopted paradigm Wang et al. (2018); Albanese et al. (2025); however, it depends on high-quality structural information, is less effective for flexible or disordered systems, and primarily captures local geometric interactions rather than end-to-end functional behavior (Isert et al., 2023; Wang et al., 2018). Generative sequence modeling has therefore emerged as a complementary approach for molecular design, enabling data-driven discovery of functional peptides, proteins, and small molecules. Autoregressive models such as ProGen (Madani et al., 2020) and ProtGPT2 (Ferruz et al., 2022), masked and diffusion-based language models (MDLMs) including DPLM (Wang

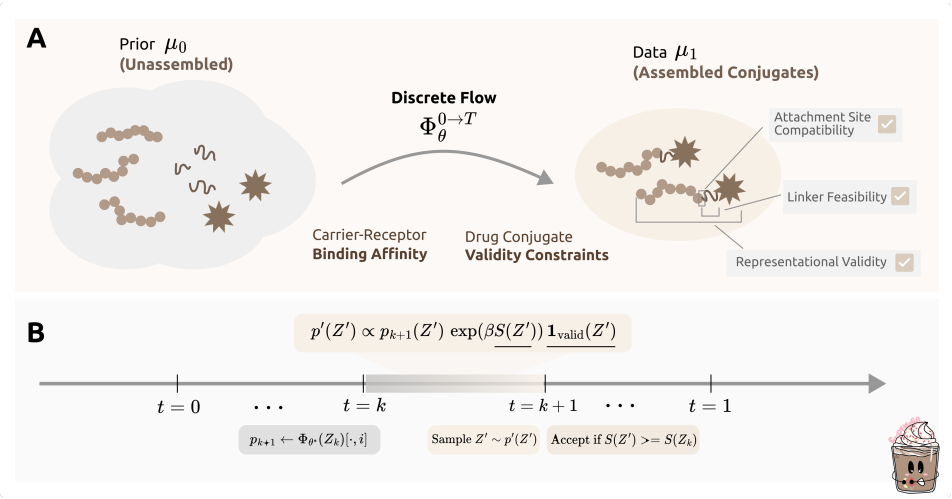

Figure 1: **Overview of CupOFMoCA for peptide-drug conjugate (PDC) generation. (A)** CupOfMoCA guides a pre-trained discrete flow model from unassembled peptide, linker, and drug components toward valid assembled conjugates, incorporating carrier-receptor binding affinity optimization and validity constraints (attachment site compatibility, linker feasibility, and representational validity). **(B)** Guided sampling with validity-constrained acceptance. At each timestep, candidate sequences are sampled from a reweighted distribution incorporating binding affinity and validity. Samples are accepted if they meet or exceed the current score, ensuring monotonic improvement toward high-affinity, valid conjugates.

et al., 2024b) and EvoDiff (Alamdari et al., 2023), and peptide-focused MDLMs such as PepTune (Tang et al., 2025a) and TR2-D2 (Tang et al., 2025c) have demonstrated strong performance within individual molecular modalities. Related MDLM formulations have also been applied to small-molecule generation, for example in GenMol (Lee et al., 2025). Discrete flow matching (DFM)-based methods (Campbell et al., 2024; Stark et al., 2024a; Davis et al., 2024; Tang et al., 2025b; Chen et al., 2025b;a) have furthered extended these ideas through transport-based updates. Despite this progress, existing approaches typically generate molecular components independently and combine them only after generation, relying on fixed corruption or interpolation schemes that traverse chemically implausible or unstable intermediate states. This limitation is particularly acute for molecular conjugates, where both validity and function emerge only at the level of the assembled system.

This gap motivates a different formulation of conjugate generation. Rather than designing peptides, linkers, and payloads independently and attaching them post hoc, we ask whether conjugates can be generated directly as coupled systems, with chemical feasibility and functional compatibility enforced throughout generation. We introduce **C**o**up**led **O**bjective-Guided **M**olecular **C**onjugate **A**ssembly (**CupOfMoCA**), a discrete generative framework that constructs molecular conjugates by jointly sampling peptide, linker, and drug components under explicit conjugation constraints and target-specific functional guidance. By operating on discrete flows rather than post-hoc attachment, CupOfMoCA enables anticipatory assembly, in which intermediate states remain chemically meaningful and design decisions account for downstream conjugation constraints.

2 CupOfMoCA

2.1 Conjugate State Space and Data Distribution

Let $\mathcal{V}$ denote a finite vocabulary of SMILES tokens and fix a length $L \in \mathbb{N}$. We write a single molecular-string space as

$$\mathcal{X} := \mathcal{V}^L. \tag{1}$$

In CupOfMoCA, each component of a conjugate is represented as a SMILES token sequence. In particular, peptide components are represented as *peptide SMILES strings* (rather than residue-level

amino-acid sequences), consistent with prior peptide generative frameworks. A molecular conjugate is represented as a structured state composed of multiple coupled components:

$$z = \left(x^{\text{pep}}, x^{\text{link}}, x^{\text{drug}}\right) \in \mathcal{Z} \quad \text{with} \quad \mathcal{Z} := \mathcal{X}_{\text{pep}} \times \mathcal{X}_{\text{link}} \times \mathcal{X}_{\text{drug}}, \tag{2}$$

where each factor is a SMILES-token string space, potentially with a different tokenizer and length bound. Although the formulation applies broadly to molecular conjugates, we focus on PDCs where the state explicitly includes peptide SMILES, a conjugating linker, and a drug scaffold.

We assume access to an unknown data distribution μ_1 over $\mathcal{Z}$ that describes the empirical law of valid assembled conjugates. In practice, we observe an i.i.d. dataset

$$\mathcal{D}_{\text{data}} = \{z^{(i)}\}_{i=1}^{N}, \qquad z^{(i)} \sim \mu_1. \tag{3}$$

We also specify a simple prior distribution μ_0 on $\mathcal{Z}$, such as a factorized initialization over components with masked or generic tokens. This prior serves as the starting distribution for the generative process.

The learning objective is to construct a transport mechanism from μ_0 to μ_1:

$$Z_0 \sim \mu_0, \qquad Z_T \sim \mu_1, \tag{4}$$

where intermediate states remain chemically valid and preserve cross-component compatibility throughout the assembly trajectory.

2.2 Reference Edit Structure and Neighborhoods

As in Section B, we use local edit neighborhoods to define admissible discrete moves. For a component string $x \in \mathcal{X}$, let $\mathcal{N}(x)$ denote the set of states obtainable by a single token substitution. For a conjugate state $z \in \mathcal{Z}$, we define a structured neighborhood

$$\mathcal{N}(z) := \{z' \in \mathcal{Z} \mid d(z, z') = 1\}, \tag{5}$$

where $z' = \left(x'^{\text{pep}}, x'^{\text{link}}, x'^{\text{drug}}\right)$ differs from z by one admissible token edit in exactly one component. This neighborhood provides a neutral reference notion of "small" changes in conjugate space and is used both for defining feasibility constraints and for constructing training signals. In practice, admissible edits are restricted to those that preserve syntactic validity of each component representation (e.g., SMILES well-formedness under the chosen tokenization).

2.3 Discrete Flows and Transport

Rather than learning stepwise stochastic dynamics, we learn a *discrete normalizing flow* that transports samples from a simple prior distribution to the data distribution through a sequence of deterministic maps. Let

$$0 = t_0 < t_1 < \cdots < t_K = T$$

denote a discrete time grid. We parameterize a family of invertible transport maps

$$\Phi_\theta^{t_k} : \mathcal{Z} \to \mathcal{Z}, \qquad k = 0, \ldots, K-1, \tag{6}$$

where each map advances a state from time t_k to t_{k+1}.

Starting from $z_0 \sim \mu_0$, the discrete flow evolves by composition,

$$z_{k+1} = \Phi_\theta^{t_k}(z_k), \tag{7}$$

so that the state at time t_k is given by

$$z_k = \Phi_\theta^{t_{k-1}} \circ \cdots \circ \Phi_\theta^{t_0}(z_0).$$

The distribution induced at time t_k is the pushforward of the prior,

$$\mu_{t_k}^\theta := \left(\Phi_\theta^{t_{k-1}} \circ \cdots \circ \Phi_\theta^{t_0}\right)_{\#} \mu_0, \tag{8}$$

where $\#$ denotes the pushforward of measures. The terminal distribution $\mu_T^\theta = \mu_{t_K}^\theta$ defines the model distribution.

This formulation mirrors classical normalizing flows, where a complex target distribution is obtained by composing simple, learnable transport maps. From a flow matching perspective, each map $\Phi_\theta^{t_k}$ approximates the finite-time transport induced by an underlying continuous probability flow, enabling few-step generation by discretizing and composing these learned transports.

2.4 CONSTRAINTS AND OBJECTIVE GUIDANCE

Conjugate assembly imposes hard validity constraints that couple components (e.g., attachment-site compatibility, linker feasibility, and representational validity). We encode these via a constraint indicator

$$\mathbf{1}_{\text{valid}}(z) \in \{0, 1\}, \tag{9}$$

which returns 1 if and only if z satisfies the prescribed chemical and representation constraints.

In addition, we assume access to one or more objective functions that score conjugates or their components. In this work we focus on a pre-trained peptide-target predictor and define an objective

$$S(z) \in \mathbb{R}, \tag{10}$$

which may depend on the assembled state through its peptide SMILES component and conjugation context. These objectives guide transport toward functionally relevant regions of $\mathcal{Z}$ but do not define a likelihood model for μ_1.

2.5 ENDPOINT COUPLINGS AND INTERPOLATION

Learning discrete flows benefits from relating prior samples to data samples through an endpoint coupling. We assume a coupling γ between μ_0 and μ_1:

$$\gamma \in \Pi(\mu_0, \mu_1), \tag{11}$$

where $\Pi(\mu_0, \mu_1)$ denotes the set of joint distributions on $\mathcal{Z} \times \mathcal{Z}$ with marginals μ_0 and μ_1. In practice, γ may be approximated by independent pairing or by a heuristic cost-based matching between conjugates.

For each sampled pair $(z_0, z_1) \sim \gamma$ and time $t \in [0, T]$, we construct an intermediate state z_t by applying a bounded number of admissible token edits that move from z_0 toward z_1 while maintaining validity. This interpolation provides training triples (z_0, z_t, z_1) that expose intermediate conjugate configurations and support learning of maps for arbitrary time intervals.

2.6 GENERATIVE SAMPLING

Once the discrete flow $\{\Phi_\theta^{t_k}\}_{k=0}^{K-1}$ has been learned, sampling proceeds by drawing an initial state

$$Z_0 \sim \mu_0 \tag{12}$$

and transporting it forward through the flow to the terminal time T.

In the multi-step regime, we choose a discrete time grid

$$0 = t_0 < t_1 < \cdots < t_K = T$$

and generate samples by composing the corresponding transport maps:

$$Z_{t_{k+1}} = \Phi_\theta^{t_k}(Z_{t_k}), \qquad k = 0, \ldots, K-1. \tag{13}$$

The final sample $Z_T = Z_{t_K}$ is obtained after K compositions.

This procedure produces samples from the learned model distribution μ_T^θ using a small number of deterministic transport steps, enabling a flexible tradeoff between computational cost and sample quality through the choice of discretization.

3 RESULTS

For evaluation of CupOFMoCA, we focus on three key questions: (1) whether coupled generation guarantees conjugate validity while avoiding functional penalties that could be associated with post-hoc staged conjugation; (2) whether discrete flows with exponential guidance enable high-quality generation; and (3) whether the model exhibits competitive design against therapeutic baselines under in-silico structural validation.

3.1 COUPLED VERSUS STAGED MOLECULAR CONJUGATE ASSEMBLY

3.1.1 MOLECULAR PROPERTIES AGAINST PHARMA REFERENCE

Current standard practices for PDC design rely on staged pipelines: a peptide is optimized for a target in isolation and subsequently chemically conjugated to a linker-payload. We hypothesize that this decoupling forces trajectories through invalid intermediates and ignores steric or electrostatic clashes introduced by the payload. To test this, we compared CupOFMoCA against a strong staged baseline on two therapeutically-relevant targets: ERBB2 and GNRHR. Specifically, ERBB2 (also known as HER2) is a well-established oncogenic driver in breast and gastric cancers (Harari & Yarden, 2000) and a validated target for PDCs Armstrong et al. (2025), while GNRHR is a clinically important endocrine receptor targeted in reproductive disorders and hormone-driven cancers (Maggi et al., 2016).

For the baseline, we generated high-affinity peptide SMILES strings using a pre-trained discrete flow matching model (Chen et al., 2025a) and attached fixed linker-payload post-generation. For CupOFMoCA, during generation process, the model produces peptides that can be attached to given linker-payload SMILES under hard chemical validity constraints including attachment site compatibility, linker feasibility, and representational validity. We evaluated both methods using an independent oracle predictor for peptide-target binding affinity. And we compared our results with pharmaceutical reference PDCs with the same linker, payload, and conjugation method from PDCdb (Sun et al., 2025).

Table 1: Mean molecular properties for top-20 *de novo* PDCs targeting receptor ERBB2 (peptide length: 80 SMILES tokens). All conjugates comprise a peptide linked to Camptothecin via Disulfide bond linker. Binding affinity: predicted peptide-receptor binding score. MW = molecular weight (Da); LogP = lipophilicity; TPSA = topological polar surface area (Å^2); HBD/HBA = hydrogen bond donors/acceptors.

Property	Pharma	Staged	Coupled
Binding affinity	6.80	6.67	**7.39**
MW (Da)	3000.20	2127.72	2073.41
LogP	-8.95	-2.38	-1.95
TPSA (Å^2)	1220.87	758.03	720.88
HBD	36.71	23.50	22.80
HBA	42.14	44.20	41.20
Rotatable bonds	85.57	68.50	68.55
Heavy atoms	213.71	143.50	137.00

Table 2: Mean molecular properties for top-20 *de novo* PDCs targeting receptor GNRHR (peptide length: 80 SMILES tokens). All conjugates comprise a peptide linked to Daunorubicin via an Aminooxyacetic acid linker.

Property	Pharma	Staged	Coupled
Binding affinity	8.02	**10.65**	**10.65**
MW (Da)	1980.00	2133.17	2130.94
LogP	-0.94	-3.31	-4.12
TPSA (Å^2)	705.84	795.02	810.57
HBD	21.86	27.40	28.30
HBA	29.38	44.60	47.00
Rotatable bonds	45.10	55.45	71.50
Heavy atoms	142.90	144.60	143.60

The Coupled approach demonstrated improved binding affinity compared to pharmaceutical reference compounds across both targets. For GNRHR, the Coupled method achieved a 33% improvement in predicted binding affinity (10.65 vs 8.02), matching the performance of the Staged approach (10.65). For ERBB2, the Coupled method showed a 9% binding improvement (7.39 vs 6.80), while the Staged approach underperformed relative to the baseline (6.67).

Molecular properties remained within acceptable ranges for both conjugate design methods. For the peptide-Aminooxyacetic acid-Daunorubicin conjugates targeting receptor GNRHR, *de novo* generated samples resulted in moderate increases in molecular weight ($\sim$150 Da), polar surface area ($\sim$90-105 Å^2), and hydrogen bonding capacity, while maintaining similar heavy atom counts. The Coupled and Staged approaches produced comparable molecular properties, with the Coupled method showing slightly more favorable TPSA (811 vs 795 Å^2) and LogP (-4.12 vs -3.31).

For the peptide-disulfide bond-Camptothecin conjugates targeting ERBB2, both Staged and Coupled methods generated molecules with more favorable drug-like properties than the pharmaceutical reference, including substantially reduced molecular weight (Coupled: 2073 Da vs Pharma: 3000 Da) and lower polar surface area (Coupled: 721 vs Pharma: 1221 Å^2). Notably, only the Coupled approach maintained improved binding affinity while achieving these favorable physicochemical properties.

3.1.2 Conjugate Validity Guarantee

In the Staged baseline, though applying peptide SMILES validity checks ensure the generation quality of peptides, many of the generated candidates are lost due to not satisfying conjugation constraints, which most of the case would be lacking a Lysine presence.

Table 3: Validity rate across receptor targets and generation methods. Pep. and Conj. stand for peptide validity and peptide drug conjugate validity respectively. The validity rates are reported over 100 samples.

Receptor	Length	Staged		**Coupled**
		Pep.	Conj.	Conj.
GNRHR	80	93%	27%	**94%**
	100	98%	55%	**96%**
ERBB2	80	97%	33%	**92%**
	100	98%	56%	**95%**

Table 3 demonstrates the substantial validity gains achieved by the Coupled generation approach compared to the Staged baseline. While the Staged method produces highly valid peptide sequences (93-98% validity), a significant portion of these candidates fail to satisfy conjugation constraints during the post-generation attachment phase, resulting in conjugate validity rates of only 27-56%. This attrition is primarily due to the absence of appropriate attachment sites, particularly the lack of Lysine residues required for linker conjugation.

In contrast, the Coupled approach directly generates valid PDCs with consistently high success rates (92-96%), effectively eliminating the inefficiency of the two-stage process. By incorporating conjugation constraints during the generation process rather than as a post-hoc filter, the Coupled method ensures that generated peptides inherently possess the necessary chemical features for successful linker attachment. This results in a 2-3$\times$ improvement in overall generation efficiency, with the Coupled method producing nearly as many valid conjugates as the Staged method produces valid peptides, while simultaneously achieving superior binding affinity as shown in Tables 1 and 2.

3.2 Generation Quality of Discrete Flows with Exponential Tilting Guidance

We benchmarked CupOFMoCA against two SMILES-based peptide design baselines to evaluate the effectiveness of discrete flow matching with exponential tilting guidance. PepTune (Tang et al., 2025a) represents the discrete diffusion baseline with staged generation, i.e. the peptide is generated independently and conjugated post-hoc. AReURedi (Chen et al., 2025a) shares our discrete flow maching backbone but employs an alternative annealed rectified updates guidance approach. For completeness and isolation of the impact of guidance strategy on the same generative backbone, we include both AReUReDi's original staged results and a CupOFMoCA-style coupled generation variant where we integrate PDC validity constraints into AReURedi's guidance framework.

Tables 4 and 5 present the molecular properties of valid *de novo* PDCs from 100 generation runs for ERBB2 and GNRHR receptors, respectively. The discrete diffusion baseline (PepTune) achieved

Table 4: Mean molecular properties for valid PDCs in 100 runs of *de novo* generation targeting receptor ERBB2 (peptide length: 80 SMILES tokens). All conjugates comprise a peptide linked to Camptothecin via Disulfide bond linker. **Conj. Validity** represents the proportion of runs yielding valid peptides capable of forming valid conjugates. The rest of the molecular properties are calculated using only the valid conjugates from the 100 runs.

Property	PepTune	AReUReDi		**Ours**
	Staged	Staged	Coupled	Coupled
Conj. Validity	39.28%	30.1%	**100%**	**100%**
Binding affinity	5.03	**6.26**	6.08	<u>6.18</u>
MW (Da)	2135.35	2114.37	2062.85	2069.22
LogP	0.92	-2.32	-3.25	-3.07
TPSA (Å^2)	674.01	762.60	750.90	750.40
HBD	20.80	23.49	25.01	24.35
HBA	32.80	42.14	42.76	42.59
Rotatable bonds	59.20	63.26	71.77	71.07
Heavy atoms	147.60	142.49	138.96	138.72

Table 5: Mean molecular properties for valid PDCs in 100 runs of *de novo* generation targeting receptor GNRHR (peptide length: 80 SMILES tokens). All conjugates comprise a peptide linked to Daunorubicin via an Aminooxyacetic acid linker.

Property	PepTune	AReUReDi		**Ours**
	Staged	Staged	Coupled	Coupled
Conj. Validity	27.8%	26.8%	**100%**	**100%**
Binding affinity	8.29	<u>10.00</u>	<u>10.00</u>	**10.12**
MW (Da)	2180.29	2139.73	2136.08	2128.04
LogP	-0.98	-3.26	-5.03	-3.84
TPSA (Å^2)	750.86	790.62	821.02	795.98
HBD	24.80	28.56	29.63	28.76
HBA	34.80	45.26	45.27	45.70
Rotatable bonds	57.20	61.91	69.04	70.62
Heavy atoms	152.60	143.91	143.71	142.67

binding affinities of 5.03 and 8.29 for ERBB2 and GNRHR, lower than the two flow-matching model baselines, highlighting the advantages of the flow-based generation framework. Comparing the two discrete flow matching approaches under coupled generation, CupOFMoCA demonstrated modest but consistent improvements in binding affinity across both targets. For ERBB2, CupOFMoCA achieved a binding score of 6.18 compared to 6.08 for AReUReDi (1.6% improvement). For GNRHR, CupOFMoCA reached 10.12 versus 10.00 for AReUReDi (1.2% improvement).

Importantly, both flow matching guidance methods produced conjugates with highly similar physicochemical properties, indicating that the improved validity rate and binding affinity achieved by CupOFMoCA does not come at the cost of degraded molecular characteristics. Molecular weight, lipophilicity (LogP), polar surface area (TPSA), hydrogen bonding capacity, and structural complexity metrics remained comparable between the two coupled generation approaches. This suggests that CupOFMoCA's exponential tilting guidance effectively steers generation toward higher-affinity candidates while maintaining appropriate drug-like properties and structural diversity within the chemical space of valid PDCs.

3.3 Structural Validation of Generated Designs

To evaluate the structural quality of designs generated by CupOFMoCA, we performed structure prediction using Boltz-2, the current state-of-the-art molecule-protein complex predictor (Passaro et al., 2025), on receptor-ligand complexes for ERBB2 and GNRHR. We compared our generated designs (Coupled) against reference pharmaceutical compounds (Pharma) across four metrics: overall confidence score, interface predicted TM-score (ipTM), complex pLDDT, and interface pLDDT (ipLDDT).

Table 6: Structure prediction quality for top 10 ERBB2-targeting designs (length 80). Comparison of peptides and PDCs binding to ERBB2 receptor, with metrics measured on the receptor-ligand complex using Boltz-2. Pharma represents reference pharmaceutical designs, while Coupled represents our generation method. conf. = confidence score, ipTM = interface predicted TM-score, pLDDT = complex pLDDT, ipLDDT = complex ipLDDT.

Metric	Peptide		**Conjugate**	
	Pharma	**Ours**	Pharma	**Ours**
Conf.	0.6653	0.6878	0.6415	0.6971
ipTM	0.7217	0.7571	0.6281	0.8043
pLDDT	0.6511	0.6705	0.6448	0.6703
ipLDDT	0.4344	0.4976	0.4021	0.5023

Table 7: Structure prediction quality for top 10 GNRHR-targeting designs (length 80). Comparison of peptides and PDCs binding to GNRHR receptor, with metrics measured on the receptor-ligand complex using Boltz-2. Pharma represents reference pharmaceutical designs, while Coupled represents our generation method.

Metric	Peptide		**Conjugate**	
	Pharma	**Ours**	Pharma	**Ours**
Conf.	0.8037	0.7764	0.7111	0.7251
ipTM	0.9217	0.8826	0.8581	0.8695
pLDDT	0.7742	0.7498	0.6744	0.6890
ipLDDT	0.6522	0.5827	0.4639	0.4961

As shown in Tables 6 and 7, CupOFMoCA demonstrates strong performance in generating structurally valid receptor-binding designs. For ERBB2-targeting molecules, our method outperforms pharmaceutical references across all metrics for both peptides and PDCs, with particularly substantial improvements observed for conjugates: notably a 28% relative improvement in ipTM (0.8043 vs. 0.6281) and 25% improvement in ipLDDT (0.5023 vs. 0.4021). For GNRHR-targeting designs, while reference peptides achieve higher confidence scores, CupOFMoCA-generated conjugates consistently surpass pharmaceutical baselines across all structural quality metrics. These results suggest that our coupled generation approach is particularly effective for the more challenging task of designing PDCs, where the model must jointly optimize peptide sequence and conjugation chemistry while maintaining favorable binding geometry.

4 Discussion

We present **CupOFMoCA**, a framework for objective-guided discrete generation in which molecular conjugates are generated as coupled systems. During generation, peptide, linker, and drug components are constrained to remain chemically compatible, so that intermediate states correspond to valid partial conjugates. Applied to peptide-drug conjugates (PDCs), this coupled formulation produces fully assembled molecules that maintain predicted binding after conjugation and achieves higher validity than staged baselines where payloads are attached after peptide design.

Several limitations remain. Our current evaluations do not fully capture steric effects, linker conformational dynamics, or cellular context. Experimental validation will therefore be required to assess the stability, target engagement, and functional efficacy of generated conjugates. Future work includes extending the framework to other conjugation modalities such as PROTACs and multi-component degraders, incorporating objectives related to pharmacokinetics and linker properties, and relaxing the current assumption of fixed linker-payload chemistry to enable adaptive component selection. As conjugate-based therapeutics continue to mature, CupOFMoCA provides a practical step toward system-level molecular generation for therapeutic design.

REFERENCES

Takuya Akiba, Shotaro Sano, Toshihiko Yanase, Takeru Ohta, and Masanori Koyama. Optuna: A next-generation hyperparameter optimization framework, 2019. URL https://arxiv.org/abs/1907.10902.

S Alamdari, N Thakkar, R van den Berg, A.X Lu, N Fusi, A.P Amini, and K.K Yang. Protein generation with evolutionary diffusion: Sequence is all you need. *bioRxiv*, 2023. doi: 10.1101/2023.09.11.556673. URL https://www.biorxiv.org/content/10.1101/2023.09.11.556673v2.full-text. Preprint.

Katherine I Albanese, Sophie Barbe, Shunsuke Tagami, Derek N Woolfson, and Thomas Schiex. Computational protein design. *Nature Reviews Methods Primers*, 5(1):13, 2025.

Amy Armstrong, Fleur Coburn, Yanyamba Nsereko, and Othman Al Musaimi. Peptide-drug conjugates: A new hope for cancer. *Journal of Peptide Science*, 31(8):e70040, 2025.

Andrew Campbell, Jason Yim, Regina Barzilay, Tom Rainforth, and Tommi Jaakkola. Generative flows on discrete state-spaces: Enabling multimodal flows with applications to protein co-design. In *International Conference on Machine Learning*, pp. 5453–5512. PMLR, 2024.

Michael Cardei, Jacob K Christopher, Bhavya Kailkhura, Thomas Hartvigsen, and Ferdinando Fioretto. Constrained molecular generation with discrete diffusion for drug discovery. In *NeurIPS 2025 Workshop on AI Virtual Cells and Instruments: A New Era in Drug Discovery and Development*, 2025.

Tong Chen, Yinuo Zhang, and Pranam Chatterjee. Areuredi: Annealed rectified updates for refining discrete flows with multi-objective guidance, 2025a. URL https://arxiv.org/abs/2510.00352.

Tong Chen, Yinuo Zhang, Sophia Tang, and Pranam Chatterjee. Multi-objective-guided discrete flow matching for controllable biological sequence design, 2025b. URL https://arxiv.org/abs/2505.07086.

Jacob K Christopher, Stephen Baek, and Nando Fioretto. Constrained synthesis with projected diffusion models. *Advances in Neural Information Processing Systems*, 37:89307–89333, 2024.

Stuart J Conway. Bifunctional molecules beyond protacs. *Journal of Medicinal Chemistry*, 63(6): 2802–2806, 2020.

Bethany M Cooper, Jessica Iegre, Daniel H O'Donovan, Maria Ölwegård Halvarsson, and David R Spring. Peptides as a platform for targeted therapeutics for cancer: peptide–drug conjugates (pdcs). *Chemical society reviews*, 50(3):1480–1494, 2021.

Oscar Davis, Samuel Kessler, Mircea Petrache, İsmail İ Ceylan, Michael Bronstein, and Avishek J Bose. Fisher flow matching for generative modeling over discrete data. *Advances in Neural Information Processing Systems*, 37:139054–139084, 2024.

Trevor T Dean, Juliet Jelu-Reyes, A'Lester C Allen, and Terry W Moore. Peptide–drug conjugates: An emerging direction for the next generation of peptide therapeutics. *Journal of medicinal chemistry*, 67(3):1641–1661, 2024.

Joshua Z Drago, Shanu Modi, and Sarat Chandarlapaty. Unlocking the potential of antibody–drug conjugates for cancer therapy. *Nature Reviews Clinical Oncology*, 18(6):327–344, 2021.

Alexis Eras, Danna Castillo, Margarita Suárez, Nelson Santiago Vispo, Fernando Albericio, and Hortensia Rodriguez. Chemical conjugation in drug delivery systems. *Frontiers in chemistry*, 10: 889083, 2022.

Aaron L Feller and Claus O Wilke. Peptide-aware chemical language model successfully predicts membrane diffusion of cyclic peptides. *Journal of Chemical Information and Modeling*, 65(2): 571–579, 2025.

N Ferruz, S Schmidt, and B Höcker. Protgpt2 is a deep unsupervised language model for protein design. *Nature Communications*, 13, 2022. doi: 10.1038/s41467-022-32007-7. URL https://www.nature.com/articles/s41467-022-32007-7.

Itai Gat, Tal Remez, Neta Shaul, Felix Kreuk, Ricky TQ Chen, Gabriel Synnaeve, Yossi Adi, and Yaron Lipman. Discrete flow matching. *Advances in Neural Information Processing Systems*, 37: 133345–133385, 2024.

Nate Gruver, Samuel Stanton, Nathan C. Frey, Tim G. J. Rudner, Isidro Hotzel, Julien Lafrance-Vanasse, Arvind Rajpal, Kyunghyun Cho, and Andrew Gordon Wilson. Protein design with guided discrete diffusion, 2023. URL https://arxiv.org/abs/2305.20009.

Daniel Harari and Yosef Yarden. Molecular mechanisms underlying erbb2/her2 action in breast cancer. *Oncogene*, 19(53):6102–6114, 2000.

Ethan Heh, Jesse Allen, Fabiola Ramirez, Daniel Lovasz, Lorena Fernandez, Tanis Hogg, Hannah Riva, Nathan Holland, and Jessica Chacon. Peptide drug conjugates and their role in cancer therapy. *International Journal of Molecular Sciences*, 24(1):829, 2023.

Liwen Hua, Qiuyue Zhang, Xinyue Zhu, Ruoning Wang, Qidong You, and Lei Wang. Beyond proteolysis-targeting chimeric molecules: designing heterobifunctional molecules based on functional effectors. *Journal of Medicinal Chemistry*, 65(12):8091–8112, 2022.

Clemens Isert, Kenneth Atz, and Gisbert Schneider. Structure-based drug design with geometric deep learning. *Current Opinion in Structural Biology*, 79:102548, 2023.

Wengong Jin, Regina Barzilay, and Tommi Jaakkola. Junction tree variational autoencoder for molecular graph generation. In *International conference on machine learning*, pp. 2323–2332. PMLR, 2018.

Sachin Kumar, Biswajit Paria, and Yulia Tsvetkov. Gradient-based constrained sampling from language models. *arXiv preprint arXiv:2205.12558*, 2022.

Greg Landrum, Paolo Tosco, Brian Kelley, Ricardo Rodriguez, David Cosgrove, Riccardo Vianello, Peter Gedeck, Gareth Jones, Eisuke Kawashima, Dan Nealschneider, et al. rdkit/rdkit: 2025_03_1 (q1 2025) release. *Zenodo*, 2025.

Seul Lee, Karsten Kreis, Srimukh Prasad Veccham, Meng Liu, Danny Reidenbach, Yuxing Peng, Saee Gopal Paliwal, Weili Nie, and Arash Vahdat. Genmol: A drug discovery generalist with discrete diffusion. In *Forty-second International Conference on Machine Learning*, 2025. URL https://openreview.net/forum?id=KM7pXWG1xj.

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.

Ali Madani, Bryan McCann, Nikhil Naik, Nitish Shirish Keskar, Namrata Anand, Raphael R. Eguchi, Po-Ssu Huang, and Richard Socher. Progen: Language modeling for protein generation. *bioRxiv*, 2020. doi: 10.1101/2020.03.07.982272. URL https://www.biorxiv.org/content/10.1101/2020.03.07.982272v2.full-text. Preprint.

Roberto Maggi, Anna Maria Cariboni, Marina Montagnani Marelli, Roberta Manuela Moretti, Valentina Andre, Monica Marzagalli, and Patrizia Limonta. Gnrh and gnrh receptors in the pathophysiology of the human female reproductive system. *Human reproduction update*, 22(3): 358–381, 2016.

Hany E Marei, Carlo Cenciarelli, and Anwarul Hasan. Potential of antibody–drug conjugates (adcs) for cancer therapy. *Cancer Cell International*, 22(1):255, 2022.

Hunter Nisonoff, Junhao Xiong, Stephan Allenspach, and Jennifer Listgarten. Unlocking guidance for discrete state-space diffusion and flow models. *arXiv preprint arXiv:2406.01572*, 2024.

Saro Passaro, Gabriele Corso, Jeremy Wohlwend, Mateo Reveiz, Stephan Thaler, Vignesh Ram Somnath, Noah Getz, Tally Portnoi, Julien Roy, Hannes Stark, et al. Boltz-2: Towards accurate and efficient binding affinity prediction. *BioRxiv*, 2025.

Syed Faheem Askari Rizvi, Linjie Zhang, Haixia Zhang, and Quan Fang. Peptide-drug conjugates: design, chemistry, and drug delivery system as a novel cancer theranostic. *ACS Pharmacology & Translational Science*, 7(2):309–334, 2024.

Hannes Stark, Bowen Jing, Chenyu Wang, Gabriele Corso, Bonnie Berger, Regina Barzilay, and Tommi Jaakkola. Dirichlet flow matching with applications to dna sequence design. *arXiv preprint arXiv:2402.05841*, 2024a.

Hannes Stark, Bowen Jing, Chenyu Wang, Gabriele Corso, Bonnie Berger, Regina Barzilay, and Tommi Jaakkola. Dirichlet flow matching with applications to dna sequence design. In *Forty-first International Conference on Machine Learning*, 2024b.

Xiuna Sun, Hanyang Li, Zhen Chen, Yang Zhang, Zhangle Wei, Hangwei Xu, Yang Liao, Wanghao Jiang, Yichao Ge, Lingyan Zheng, et al. Pdcdb: the biological activity and pharmaceutical information of peptide–drug conjugate (pdc). *Nucleic Acids Research*, 53(D1):D1476–D1485, 2025.

Sophia Tang, Yinuo Zhang, and Pranam Chatterjee. Peptune: De novo generation of therapeutic peptides with multi-objective-guided discrete diffusion, 2025a. URL https://arxiv.org/abs/2412.17780.

Sophia Tang, Yinuo Zhang, Alexander Tong, and Pranam Chatterjee. 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.

Sophia Tang, Yuchen Zhu, Molei Tao, and Pranam Chatterjee. Tr2-d2: Tree search guided trajectory-aware fine-tuning for discrete diffusion, 2025c. URL https://arxiv.org/abs/2509.25171.

Akash Vikal, Rashmi Maurya, Brij Bihari Patel, Rajeev Sharma, Preeti Patel, Umesh K Patil, and Balak Das Kurmi. Protacs in cancer therapy: mechanisms, design, clinical trials, and future directions. *Drug delivery and translational research*, 15(6):1801–1827, 2025.

Chao Wang, Yujing Zhang, Wujun Chen, Yudong Wu, and Dongming Xing. New-generation advanced protacs as potential therapeutic agents in cancer therapy. *Molecular cancer*, 23(1):110, 2024a.

Xin Wang, Ke Song, Li Li, and Lijiang Chen. Structure-based drug design strategies and challenges. *Current Topics in Medicinal Chemistry*, 18(12):998–1006, 2018.

Xinyou Wang, Zaixiang Zheng, Fei Ye, Dongyu Xue, Shujian Huang, and Quanquan Gu. Diffusion language models are versatile protein learners. In Ruslan Salakhutdinov, Zico Kolter, Katherine Heller, Adrian Weller, Nuria Oliver, Jonathan Scarlett, and Felix Berkenkamp (eds.), *Proceedings of the 41st International Conference on Machine Learning*, volume 235 of *Proceedings of Machine Learning Research*, pp. 52309–52333. PMLR, 2024b. URL https://proceedings.mlr.press/v235/wang24ct.html.

Jaehoon Yoo, Wonjung Kim, and Seunghoon Hong. Redi: Rectified discrete flow. In *The Thirty-ninth Annual Conference on Neural Information Processing Systems*.

Yimeng Zeng, Hunter Elliott, Phillip Maffettone, Peyton Greenside, Osbert Bastani, and Jacob R Gardner. Antibody design with constrained bayesian optimization. In *GEM workshop, ICLR 2024*, 2024.

Appendix

A Related Works

Discrete Flow Matching. Discrete flow matching has recently emerged as a principled approach for modeling and sampling from complex discrete spaces. One line of work focuses on jump-process models, which learn time-dependent transition rates for token-level stochastic updates and induce continuous-time Markov chains with prescribed marginals (Gat et al., 2024; Campbell et al., 2024). In contrast, continuous-time simplex methods diffuse discrete data through continuous embeddings over the probability simplex, enabling training via flow-matching or optimal transport objectives defined on relaxed categorical representations (Tang et al., 2025b; Davis et al., 2024; Stark et al., 2024b). More recently, rectified formulations such as Rectified Discrete Flows (ReDi) (Yoo et al.) have been proposed to simplify discrete probability paths and improve sampling efficiency by iteratively straightening transport couplings, and AReUReDi (Chen et al., 2025a) extends this paradigm to multi-objective settings with Pareto-guided sampling while preserving distributional invariance.

Constrained Generation for Biomolecule Design. Generative models for biomolecule design have primarily focused on optimizing linear objectives, where higher (or lower) scores are always preferred, through techniques such as reward-weighted sampling, classifier guidance, or guided diffusion (Nisonoff et al., 2024; Gruver et al., 2023). However, many important design requirements are more naturally expressed as binary feasibility constraints rather than continuous objectives. Examples include representation validity, where generated graphs or SMILES strings must parse into chemically valid molecules (Jin et al., 2018; Tang et al., 2025a; Chen et al., 2025a); or property such as thermostability thresholds, where antibody candidates must exceed a minimum melting temperature to be developable (Zeng et al., 2024)

For diffusion-based generation, gradient-based constrained sampling enables constraint satisfaction by guiding iterates toward feasible regions (Kumar et al., 2022), while projected diffusion models recast the sampling process as constrained optimization to certify compliance with physical and geometric constraints (Christopher et al., 2024). Constrained Discrete Diffusion (CDD) (Cardei et al., 2025) integrates differentiable constraint optimization directly into sampling, enabling training-free enforcement.

While these methods can bias output distributions toward feasibility, most rely on soft guidance and do not guarantee constraint satisfaction. Notably, analogous constraint-enforcement mechanisms remain underdeveloped for discrete flow matching models, which have focused on controllable generation via reward reweighting rather than hard constraint satisfaction.

B Preliminaries

We introduce the mathematical objects required to define coupled discrete flows on structured molecular string spaces and to formulate objective-guided transport for molecular conjugate assembly. Throughout, $\mathcal{V}$ denotes a finite vocabulary of molecular string tokens (e.g., SMILES tokens), and $\mathcal{V}^L$ denotes the set of molecular strings of length L. In this work, peptides are represented directly as peptide SMILES strings, consistent with prior peptide generative frameworks such as PepTune, TR2-D2, and AReUReDi (Tang et al., 2025a; Chen et al., 2025a; Tang et al., 2025c). While we present notation for a single molecular string, subsequent sections consider structured states composed of multiple coupled molecular strings corresponding to different components of a conjugate.

B.1 Molecular Strings as States in a Discrete Space

A molecular object is represented as a discrete string state

$$x = (x_1, \ldots, x_L) \in \mathcal{V}^L, \tag{14}$$

where each token $x_i \in \mathcal{V}$ corresponds to a valid SMILES symbol. For peptide components, x encodes a peptide SMILES string rather than a residue-level amino-acid sequence.

Transitions between molecular strings occur through local token-level edit operations. We define the edit neighborhood

$$\mathcal{N}(x) := \left\{ x' \in \mathcal{V}^L \;:\; \text{single token edit} \right\}, \tag{15}$$

where a local edit corresponds to a token substitution. Extensions to insertions or deletions are straightforward and compatible with SMILES-based representations, but are omitted here for clarity.

B.2 Reference Dynamics

To define transport in discrete molecular string space, we introduce a simple reference process that captures uninformed local exploration. We consider a continuous-time Markov chain (CTMC) with generator R_0 defined by

$$R_0(x, x') = \begin{cases} \frac{1}{|\mathcal{N}(x)|} & x' \in \mathcal{N}(x), \\ -\sum_{y \neq x} R_0(x, y) & x' = x. \end{cases} \tag{16}$$

This reference dynamics defines a random walk over molecular strings that performs small, unbiased token edits and serves as a neutral baseline relative to which structured transport is learned.

B.3 Discrete Flow

Rather than modeling infinitesimal stochastic dynamics, we focus on learning *discrete flows* that directly transport molecular string states between times. For $0 < t \leq T$, a discrete flow is a function

$$\Phi_\theta^{t \to T} : \mathcal{V}^L \to \mathcal{V}^L, \tag{17}$$

parameterized by θ, which maps a molecular string at time t to a molecular string at time T. These maps are intended to approximate the solution operator of an underlying probability flow, enabling one- or few-step generation without explicit simulation of a continuous-time process.

B.4 Objective Functions

To guide transport toward functionally relevant regions of molecular string space, we assume access to one or more objective functions

$$S : \mathcal{V}^L \to \mathbb{R}, \tag{18}$$

which score molecular strings according to downstream criteria such as peptide-target binding, compatibility with conjugation chemistry, or chemical feasibility. Importantly, these objectives are evaluated on the *current transported molecular state* and do not define a probability density. Instead, they provide directional guidance that biases transport toward desirable configurations.

B.5 Generative Sampling

After training, sampling proceeds by drawing an initial state $X_0 \sim \mu_0$ from a simple molecular-string prior and applying the learned flow:

$$X_t = \Phi_\theta^t(X_0), \qquad 0 < t \leq T. \tag{19}$$

In the multistep regime, we choose a time grid $0 = t_0 < t_1 < \cdots < t_K = T$ and compose maps sequentially. Iterating this procedure from 0 to T yields samples from the learned objective-guided transport, producing valid molecular strings corresponding to peptides or assembled conjugates in a small number of steps.

C Theoretical Proofs

We provide a minimal theoretical justification for the objective-guided sampling mechanism used in CupOFMoCA. In particular, we show that the exponential tilting rule employed during discrete flow sampling arises as the unique solution to a KL-regularized objective improvement problem under hard validity constraints.

Algorithm 1 Exponential Tilting with Acceptance Criterion

Input: trained discrete flow Φ_{θ^*}, prior μ_0 on $\mathcal{Z}$
time horizon T, number of steps K
guidance score $S(z)$, strength β
Define time grid $0 = t_0 < t_1 < \cdots < t_K = T$
Sample $Z_0 \sim \mu_0$ s.t. $\mathbf{1}_{\text{valid}}(Z_0) = 1$; $Z^* \leftarrow Z_0$
for $k = 0$ **to** $K - 1$ **do**
  Select random position i
  $p_{k+1} \leftarrow \Phi_{\theta^*}(Z_k)[\cdot, i]$ (predicted distribution at i)
  for each $z' \in \text{TopK}(p_{k+1})$ **do**
    $p'(z') \propto p_{k+1}(z') \exp(\beta S(z'))\, \mathbf{1}_{\text{valid}}(z')$
    (exponential tilting)
  end for
  Sample $z^\dagger \sim p'$
  $Z' \leftarrow Z_k$ with position i set to $z^\dagger$
  if $S(Z') \geq S(Z_k)$ **then**
    (acceptance criterion)
    $Z_{k+1} \leftarrow Z'$
    if $S(Z') > S(Z^*)$ **then**
      $Z^* \leftarrow Z'$
    end if
  else
    $Z_{k+1} \leftarrow Z_k$ (reject, keep current)
  end if
end for
return Z^* (best molecular conjugate found)

C.1 Notation and Setup

We adopt the notation from the main text. Let $\mathcal{Z}$ denote the finite discrete state space of molecular conjugates. A discrete flow step is represented by a Markov kernel $K(\cdot \mid z)$ on $\mathcal{Z}$, where $K(z' \mid z)$ denotes the probability of proposing state z' given current state z. Deterministic updates are recovered as the special case where $K(\cdot \mid z)$ is a point mass.

Validity is encoded by an indicator function $\mathbf{1}_{\text{valid}} : \mathcal{Z} \to \{0, 1\}$, which specifies whether a candidate state satisfies all chemical and representational constraints. Objective guidance is defined by a score function $S : \mathcal{Z} \to \mathbb{R}$ and a guidance strength $\beta > 0$.

C.2 Objective Guidance as KL-Regularized Exponential Tilting

We show that exponential tilting is the unique optimal way to incorporate objective information into a discrete flow step while remaining close to a base kernel and enforcing hard validity constraints.

[Exponential tilting solves KL-regularized objective improvement] Fix a base kernel $K(\cdot \mid z)$ on $\mathcal{Z}$ and a score function $S : \mathcal{Z} \to \mathbb{R}$. For a fixed input state z, consider the optimization problem

$$\max_Q \left\{ \mathbb{E}_{Z' \sim Q}[S(Z')] - \frac{1}{\beta} \text{KL}(Q \,\|\, K(\cdot \mid z)) \right\} \quad \text{subject to} \quad Q(z') = 0 \;\text{ if }\; \mathbf{1}_{\text{valid}}(z') = 0, \tag{20}$$

where $\beta > 0$. Then the unique maximizer is

$$Q^\star_\beta(z' \mid z) = \frac{K(z' \mid z)\, \mathbf{1}_{\text{valid}}(z') \, \exp\big(\beta S(z')\big)}{\sum_{y \in \mathcal{Z}} K(y \mid z)\, \mathbf{1}_{\text{valid}}(y) \, \exp\big(\beta S(y)\big)}. \tag{21}$$

Let $\mathcal{Z}_{\text{ok}} := \{z' \in \mathcal{Z} : \mathbf{1}_{\text{valid}}(z') = 1\}$ and restrict Q to be supported on $\mathcal{Z}_{\text{ok}}$. The objective in equation 20 can be written as

$$\sum_{z' \in \mathcal{Z}_{\text{ok}}} Q(z')\, S(z') - \frac{1}{\beta} \sum_{z' \in \mathcal{Z}_{\text{ok}}} Q(z') \log \frac{Q(z')}{K(z' \mid z)}. \tag{22}$$

Introduce a Lagrange multiplier λ enforcing $\sum_{z' \in \mathcal{Z}_{\text{ok}}} Q(z') = 1$. Differentiating with respect to $Q(z')$ and setting the derivative to zero yields

$$S(z') - \frac{1}{\beta}(\log Q(z') - \log K(z' \mid z) + 1) - \lambda = 0. \tag{23}$$

Rearranging gives

$$\log Q(z') = \log K(z' \mid z) + \beta S(z') + c, \tag{24}$$

for a constant c, implying

$$Q(z') \propto K(z' \mid z) \exp\big(\beta S(z')\big) \quad \text{on } \mathcal{Z}_{\text{ok}}. \tag{25}$$

Normalizing over $\mathcal{Z}_{\text{ok}}$ yields equation 21. The objective is strictly concave in Q, so the maximizer is unique.

[Limits of guided kernels] Assume $K(\cdot \mid z)$ assigns positive mass to at least one valid state. Then: (i) as $\beta \to 0$, $Q^\star_\beta(\cdot \mid z)$ converges to $K(\cdot \mid z)$ restricted and renormalized to valid states; (ii) as $\beta \to \infty$, $Q^\star_\beta(\cdot \mid z)$ concentrates on valid maximizers of S within the support of $K(\cdot \mid z)$.

Both statements follow directly from standard properties of the log-sum-exp normalization in equation 21.

D Model Architectures and Training Details

D.1 SMILESReDi Base Model

We utilized the pre-trained SMILESReDi model from (Chen et al., 2025a) as our peptide SMILES generation model. SMILESReDi is a *rectified discrete flow* (ReDi) model (Yoo et al.) built on top of discrete flow matching (DFM) Gat et al. (2024), using a Diffusion Transformer (DiT) architecture as its backbone. The model operates within a masked discrete generation framework and learns probability paths between source and target distributions rather than simulating a forward diffusion process.

SMILES sequences are tokenized using the PeptideCLM-23M tokenizer (Feller & Wilke, 2025) and mapped to discrete tokens. Token embeddings and time-step embeddings are projected into continuous representations and processed through stacked DiT transformer blocks with Rotary Positional Embeddings (RoPE) and multi-head self-attention to capture long-range dependencies. The model outputs token-level transition probabilities via a final layer normalization followed by a linear projection.

D.1.1 Time-Dependent Bond-Aware Probability Path

A key feature of SMILESReDi is its bond-aware discrete probability path designed for peptide SMILES. Peptides contain a conserved backbone of alternating carbonyl and amide groups connected by peptide bonds, with variable side chains. To preserve this structure, SMILESReDi employs a time-dependent bond-aware probability path that transitions backbone bond tokens more slowly than side-chain tokens.

For each position j with bond indicator $b_j \in \{0, 1\}$, the time-t marginal of the probability path is defined as

$$p_t(x_t^{(j)} \mid x_0^{(j)}, x_1^{(j)}) = [b_j t^\gamma + (1 - b_j)t]\delta_{x_1^{(j)}} + [1 - b_j t^\gamma - (1 - b_j)t]\delta_{x_0^{(j)}}, \quad t \in [0, 1],\ \gamma > 1, \tag{26}$$

ensuring that backbone tokens ($b_j = 1$) transition more slowly along the discrete flow path than non-bond tokens.

D.1.2 Training Configuration

The SMILESReDi model was trained and rectified following the configuration reported in (Chen et al., 2025a), and we use the released pre-trained weights without further fine-tuning. The configuration is detailed in Table 8.

Table 8: SMILESReDi training configuration.

Parameter	Value
Hardware	4×A6000 NVIDIA RTX 6000 Ada GPU (48 GB VRAM)
Training epochs	5
Optimizer	Adam
Learning rate	1×10^{-4}
Learning rate schedule	Cosine decay with 10% warmup
Minimum learning rate	1×10^{-5}
Model dimension	768
Transformer layers	8
Attention heads	8
Gradient clipping	1.0
Noising parameter γ	2.0

D.2 Binding Affinity Score Model

We employed the binding affinity prediction model from (Chen et al., 2025a) to score protein-peptide binding affinity. This model uses a reciprocal attention transformer architecture that utilizes unpooled token-level representations from the ESM-2 650M protein language model (Lin et al., 2023). The architecture processes unpooled token embeddings from ESM-2 through convolutional layers followed by cross-attention layers, allowing the model to capture interaction patterns between protein and peptide sequences.

D.2.1 Model Architecture

The binding affinity predictor architecture is detailed in Table 9.

Table 9: Binding affinity model architecture.

Component	Configuration
Input	Unpooled token-level embeddings from ESM-2 650M
Convolutional layers	3 initial CNN kernel layers (dimension 384)
Cross-attention layers	4 layers (dimension 2048)
Prediction head	Shared prediction head (dimension 1024)

D.2.2 Training Details

As reported in the original work, the binding affinity data were split into a 0.8/0.2 ratio for training and validation, maintaining similar affinity score distributions across splits. The authors employed OPTUNA (Akiba et al., 2019) for hyperparameter optimization, tracking validation correlation scores as the objective metric. The training configuration is presented in Table 10.

Table 10: Binding affinity model training configuration.

Parameter	Value
Training epochs	50
Learning rate	3.84×10^{-5}
Dropout rate	0.15
CNN layers	3 initial layers (dimension 384)
Cross-attention layers	4 layers (dimension 2048)
Prediction head dimension	1024
Train/validation split	0.8/0.2
Validation performance	0.64 Spearman correlation

E CONJUGATION VALIDITY CONSTRAINTS

During the generative workflow, a candidate PDC is first validated by the `PDCAnalyzer` for general structural feasibility, then subjected to conjugation-specific checks by the builder's validation methods, and finally scored for binding affinity and other therapeutic properties by downstream prediction models.

To ensure chemically valid PDCs from de-novo designed peptide components, we developed a specialized conjugation and validation pipeline using RDKit (Landrum et al., 2025). Both the Aoa-Dau and SSCpt conjugation builders implement multi-layered validation comprising SMARTS-based pattern recognition, structural integrity checks, and valency analysis.

E.1 AMINOOXYACETIC ACID - DAUNORUBICIN CONJUGATES (AOA-DAU)

E.1.1 CONJUGATION REACTION

Among the 21 PDC entries in PDCdb sharing the same peptide scaffold, Aminooxyacetic acid (Aoa) linker, and Daunorubicin (Dau) payload, four distinct conjugation strategies are represented, including direct conjugation via Lys side chain ε-amine, spacer-mediated conjugation, ranched conjugates and bis-conjugates at two Lys sites (additional Lys introduced via peptide modification).

We adopted the direct conjugation approach via Lys side chain NH_2, which is the most prevalent strategy among the database entries. Our implementation follows the architecture of PDC_00208 (GnRH-III-[8Lys(Dau=Aoa)]) as a reference template. To ensure site-selectivity, conjugation was restricted to the ε-amine of Lys side chains; the α-amine at the N-terminus was excluded as a potential reaction site.

Reaction Scheme. Conjugation proceeds via amide bond formation between the Lysine ε-amine (side chain $-NH_2$) on valid peptide, and the carboxyl group on Aoa (pre-conjugated to Dau via oxime bond):

$$\underbrace{\text{Peptide-Lys}}_{\text{targeting moiety}}\text{-NH}_2 \;+\; \text{HOOC-}\underbrace{\text{Aoa-Dau}}_{\text{linker-payload}} \xrightarrow{\text{amide coupling}} \text{Peptide-Lys-NH}{-}\text{CO-Aoa-Dau} \;+\; \text{H}_2\text{O} \tag{27}$$

Molecular Architecture. The complete PDC structure is represented schematically as:

$$\underbrace{\text{Peptide-Lys}}_{\text{De-novo Designed}}\text{-}\underbrace{\text{NH}{-}\text{C}(=\text{O})}_{\substack{\text{Amide}\\(\text{formed})}}\text{-}\underbrace{\text{CH}_2{-}\text{O}{-}\text{N}{=}\text{C}(\text{CH}_3)}_{\substack{\text{Aoa oxime}\\(\text{pre-formed})}}\text{-}\underbrace{\text{Dau}}_{\text{Payload}} \tag{28}$$

E.1.2 ATTACHMENT SITE COMPATIBILITY CHECKS

The resulting PDC candidates need to pass all the checks in Table 11 for compatibility on the attachment site. For valency requirements, the amide nitrogen at the conjugation site must satisfy one of the following: a secondary amide, where N is bonded to 2 heavy atoms (one carbonyl C) and 1 H, or a tertiary amide, where N is bonded to 3 heavy atoms (one carbonyl C) and 0 H.

Method	SMARTS Pattern	Purpose
`has_aoa_amide_conjugation()`	`[NX3][CX3](=[OX1])[CH2][OX2][NX2]=[C]`	Confirms amide-CH_2-O-oxime linkage
`has_oxime_bond()`	`[C]-[O]-[N]=[C]`	Verifies Aoa-Dau linkage intact
`has_unreacted_aoa_carboxylic_acid()`	`[CX3](=[OX1])[OX2H1][CH2][OX2][NX2]=[C]`	Failure indicator: Free HOOC—CH_2—O—N=C
`find_amide_conjugation_sites()`	`[NX3:1][CX3:2](=[OX1])[CH2][OX2][NX2]=[C]`	Locates conjugation sites
`check_valency_at_conjugation_site()`	Structural analysis	Validates N valency

Table 11: AoaDau attachment site compatibility checks.

E.2 DISULFIDE-CAMPTOTHECIN CONJUGATES (SSCPT)

E.2.1 CONJUGATION REACTION

7 PDC entries exist in PDCdb sharing the same ERBB2-targeting peptide scaffold, disulfide linker, and Camptothecin (CPT) payload, multiple conjugation Their conjugation architectures include dimer architectures with internal attachment, monomer architectures with N-terminal region attachment, and monomer architectures with Lys side chain conjugation. Three linker-to-drug chemistries are observed: ester, carbonate, and carbamate.

We adopted the monomer architecture with Lys side chain NH_2 conjugation, following PDC_00294 (Z7, ester) and PDC_00295 (Z8, carbonate) as reference templates. To ensure site-selectivity, conjugation was restricted to the ε-amine of Lys side chains; the α-amine at the N-terminus was excluded as a potential reaction site.

Reaction Scheme. Conjugation proceeds via amide bond formation between the Lysine ε-amine (side chain $-NH_2$) on valid peptide, and the carboxyl group on Cys (pre-conjugated to CPT via disulfide-linked spacer):

$$\underbrace{\text{Peptide-Lys}}_{\text{targeting moiety}}\text{-NH}_2 + \text{HOOC-}\underbrace{\text{Cys-SS-Spacer-CPT}}_{\text{linker-payload}} \xrightarrow{\text{amide coupling}} \text{Peptide-Lys-NH}{-}\text{CO-Cys-SS-Spacer-CPT} + \text{H}_2\text{O} \tag{29}$$

Molecular Architecture. The complete PDC structure is represented schematically as:

$$\underbrace{\text{Peptide-Lys}}_{\text{De-novo Designed}}\text{-}\underbrace{\text{NH}{-}\text{C(=O)}}_{\substack{\text{Amide}\\\text{(formed)}}}\text{-Cys-}\underbrace{\text{S}{-}\text{S}}_{\substack{\text{Disulfide}\\\text{(cleavable)}}}\text{-Spacer-}\underbrace{\text{CPT}}_{\text{Payload}} \tag{30}$$

Two linker-to-drug chemistries are supported, differing in the bond connecting the spacer to Camptothecin:

- Ester (Z7): $\text{Cys}{-}\text{S}{-}\text{S}{-}\text{CH}_2{-}\text{CH}_2{-}\text{C(=O)}{-}\text{O}{-}\text{CPT}$
- Carbonate (Z8): $\text{Cys}{-}\text{S}{-}\text{S}{-}\text{CH}_2{-}\text{CH}_2{-}\text{O}{-}\text{C(=O)}{-}\text{O}{-}\text{CPT}$

E.2.2 ATTACHMENT SITE COMPATIBILITY CHECKS

The resulting PDC candidates need to pass all the checks in Table 12 for compatibility on the attachment site. For valency requirements, the conjugation site must satisfy the following connectivity: Lys nitrogen with 3 bonds (NH connected to amide), carbonyl carbon with 3 bonds (C=O and two single bonds), and Cys α-carbon with 4 bonds (tetrahedral geometry).

Method	SMARTS Pattern	Purpose
`has_lys_cys_amide_conjugation()`	`[NX3][CX3](=[OX1])[CH]([CH2][SX2][SX2])[NX3]`	Confirms Lys-ε-N$-$C(=O)-C$_\alpha$(Cys) pattern
`has_disulfide_bond()`	`[SX2]-[SX2]`	Verifies S-S linker intact
`has_unreacted_cys_carboxylic_acid()`	`[NH2][CH]([CH2][SX2][SX2])[CX3](=[OX1])[OX2H1]`	Failure indicator: Free Cys-COOH
`find_amide_conjugation_sites()`	`[NX3][CX3](=[OX1])[CH]([CH2][SX2][SX2])[NX3]`	Locates conjugation sites
`check_valency_at_conjugation_site()`	Structural analysis	Validates atom connectivity

Table 12: SSCpt attachment site compatibility checks.

E.3 PDC ANALYZER

While the builder classes encapsulate linker-payload and conjugation reaction specific validation logic, the `PDCAnalyzer` remains conjugation-agnostic and can be applied to novel PDC designs under different conjugation reaction and / or linker-payload components

E.3.1 HARD CONSTRAINT VALIDATION

The analyzer enforces two categories of hard constraints, both of which must be satisfied for a PDC to be considered structurally valid. Violation of any hard constraint results in immediate rejection of the candidate molecule.

Linker Feasibility Constraints. Linker regions in PDCs are particularly susceptible to structural infeasibility. The analyzer addresses two primary failure modes:

Ring-closure constraints are evaluated by identifying rings that are likely part of the linker rather than the peptide backbone or drug payload. Aromatic rings are excluded from linker classification as they typically belong to the drug core. The analyzer specifically searches for succinimide rings (five-membered rings containing nitrogen and two carbonyl groups), which commonly form from maleimide-based conjugation after thiol addition. For non-aromatic rings, size-based filtering is applied: three-membered rings are checked for sp^2 carbons participating in C=C double bonds (which introduce severe angle strain), four-membered rings are rejected if they contain more than two sp^2 centers, and rings larger than twelve atoms are flagged as synthetically implausible linker components.

Long carbon chain detection identifies excessively long aliphatic chains that would be synthetically challenging and prone to conformational heterogeneity. The analyzer uses a SMARTS pattern matching sixteen consecutive CH_2 groups, rejecting any PDC containing such chains.

Representational Validity Constraints. These constraints ensure that the generated SMILES string represents a chemically valid and processable molecule. Four checks are performed sequentially: (1) RDKit sanitization; (2) round-trip SMILES conversion, where the molecule is converted to SMILES and re-parsed to check if reconstruction could be conducted; (3) chemistry problem detection using RDKit's diagnostic tools, specifically checking for `AtomValenceException` and `KekulizeException` errors that indicate fundamental chemical inconsistencies; and (4) 2D coordinate generation, which serves as a final structural coherence test. Failure at any stage indicates that the molecule, while potentially representable as a string, does not correspond to a valid chemical structure.

E.3.2 PHYSICOCHEMICAL PROPERTIES

Beyond binary validity assessment, the analyzer can compute eight molecular descriptors that are relevant for drug-likeness and pharmacokinetic profiling. These properties are toggled via boolean flags and include: molecular formula, topological polar surface area (TPSA), molecular complexity (BertzCT index), lipophilicity (XLogP via Wildman-Crippen method), heavy atom count, hydrogen bond donors and acceptors, and rotatable bond count.

F SAMPLING IMPLEMENTATION

F.1 NOTATION

Table 13: Notation used in the algorithm description.

Symbol	Description
Φ_{θ^*}	Pretrained discrete flow model with parameters θ^*
$\mathcal{Z}$	Discrete token space (SMILES vocabulary)
μ_0	Prior distribution (uniform over $\mathcal{Z}^L$)
N	Number of flow steps for initial generation
K	Number of guided optimization steps
$S(z)$	Guidance score function
β	Guidance strength (temperature parameter)
$\mathbf{1}_{\text{valid}}(z)$	Validity indicator function
Z^*	Best sequence found during optimization
L	Sequence length

F.2 TWO-PHASE GENERATION FRAMEWORK

Following Chen et al. (2025a), our approach separates generation into two phases, each serving a specific purpose. This design decouples *generation* (producing valid molecules) from *optimization* (improving target properties), allowing each phase to be tuned independently.

F.2.1 PHASE 1: INITIAL SEQUENCE GENERATION

The first phase rapidly generates an initial sequence Z_0 using unguided discrete flow matching. Starting from random noise, we iteratively denoise through the learned flow to obtain a molecular sequence satisfying the validity constraints. Running $N = 16$ steps provides sufficient denoising while remaining computationally efficient, and ensures the initial sequence lies on the data manifold. When the `require_valid_x0` flag is set, we regenerate Z_0 until the final sequence in this phase satisfies $\mathbf{1}_{\text{valid}}(Z_0) = 1$, enforcing that the optimization phase starts from a feasible point in the constraint manifold. This was adapted from the sampling process in Chen et al. (2025a), which was originally used to guarantee SMILES peptide validity.

Table 14: Phase 1 (initial generation) parameters.

Parameter	Value	Description
Flow Steps (N)	16	Number of denoising steps
Temperature	1.0	Softmax temperature for predictions
`require_valid_x0`	Enabled	Regenerate until $\mathbf{1}_{\text{valid}}(Z_0) = 1$
`max_x0_retries`	10	Maximum regeneration attempts

F.2.2 PHASE 2: GUIDED OPTIMIZATION

The second phase refines Z_0 through $K = 128$ steps of guided optimization using exponential tilting with an acceptance criterion, as in Algorithm 1. At each step, we select a random position, compute a tilted distribution over candidate tokens, and accept the proposal only if it improves (or maintains) the guidance score. This allows fine-grained control through local modifications, while the acceptance criterion prevents the optimization from degrading sequence quality.

Log-Space Computation. For numerical stability, we compute the tilted distribution in log-space:

$$\tilde{p}(z') = \log p_k(z') + \beta \cdot S(z') + \log \mathbf{1}_{\text{valid}}(z') \tag{31}$$

$$\tilde{p}_{\text{stable}}(z') = \tilde{p}(z') - \max_{z''} \tilde{p}(z'') \tag{32}$$

$$p'(z') = \frac{\exp(\tilde{p}_{\text{stable}}(z'))}{\sum_{z''} \exp(\tilde{p}_{\text{stable}}(z''))} \tag{33}$$

Fallback Behavior. If exponential tilting produces invalid probabilities (due to all candidates being invalid or numerical instability), we fall back to sampling from the base distribution p_k.

Acceptance Criterion. After sampling $z^\dagger \sim p'$ and constructing the candidate sequence Z', we apply an acceptance criterion:

$$Z_{k+1} = \begin{cases} Z' & \text{if } S(Z') \geq S(Z_k) \\ Z_k & \text{otherwise} \end{cases} \tag{34}$$

This ensures the score never decreases, providing monotonic improvement guarantees. Combined with tracking Z^*, this guarantees we return the best sequence encountered during optimization.

Token Filtering. Invalid tokens are masked during both phases. The filtered indices $\{0, 1, 2, 3, 4, 585, 586\}$ correspond to special tokens: `[PAD]`, `[UNK]`, `[CLS]`, `[SEP]`, `[MASK]`, and reserved indices.

Table 15: Phase 2 (guided optimization) parameters.

Parameter	**Value**	**Description**
Optimization Steps (K)	128	Number of guided refinement steps
Guidance Strength (β)	50.0	Exponential tilting temperature
Sequence Length (L)	80, 100	Generated sequence length (tokens)
Top-p	1.0	Nucleus sampling threshold
Top-K	200	Maximum candidates per position
Position Selection	Random	Uniform random position each step
`use_acceptance`	Enabled	Reject proposals with $S(Z') < S(Z_k)$

F.3 Computational Requirements

Sampling experiments were each conducted on a single A6000 NVIDIA GPU with 48GB of VRAM. For each sample, Phase 1 (16 flow steps) completes in under 1 minute; Phase 2 (128 optimization steps) takes 5-10 minutes depending on the sequence length and candidate size.

G Additional Results

Table 16: Mean molecular properties for top-20 *de novo* PDCs targeting receptor ERBB2 (peptide length: 100 SMILES tokens). All conjugates comprise a peptide linked to Daunorubicin via an Aminooxyacetic acid linker. Binding affinity: predicted peptide-receptor binding score. MW = molecular weight (Da); LogP = lipophilicity; TPSA = topological polar surface area (Å^2); HBD/HBA = hydrogen bond donors/acceptors.

Property	AReURedi		CupOFMoCA	
	Staged	Coupled	Staged	Coupled
Conj. Validity	62.5%	100%	57.14%	100%
Binding affinity	7.67	7.69	8.01	8.41
MW (Da)	2505.77	2471.69	2515.23	2474.41
LogP	-3.33	-3.16	-1.32	-0.15
TPSA (Å^2)	934.17	874.73	893.45	830.93
HBD	26.70	27.25	26.15	23.05
HBA	50.0	48.35	49.45	47.65
Rotatable bonds	72.15	83.8	69.5	80.25
Heavy atoms	169.30	165.80	168.80	164.50

Table 17: Mean molecular properties for top-20 *de novo* PDCs targeting receptor GNRHR (peptide length: 100 SMILES tokens). All conjugates comprise a peptide linked to Daunorubicin via an Aminooxyacetic acid linker. Binding affinity: predicted peptide-receptor binding score. MW = molecular weight (Da); LogP = lipophilicity; TPSA = topological polar surface area (Å^2); HBD/HBA = hydrogen bond donors/acceptors.

Property	AReURedi		CupOFMoCA	
	Staged	Coupled	Staged	Coupled
Conj. Validity	57.5%	100%	56.12%	100%
Binding affinity	11.0	10.09	10.09	10.91
MW (Da)	2536.03	2289.89	2578.66	2604.29
LogP	-3.55	-6.07	-3.58	-3.81
TPSA (Å^2)	947.59	909.87	971.61	957.04
HBD	30.95	31.05	31.35	30.20
HBA	48.00	47.80	52.15	51.85
Rotatable bonds	67.25	74.15	61.75	82.40
Heavy atoms	171.50	154.75	175.35	173.7