PROPHET: Phylogenetically Robust Antiviral Peptide Design Against Heterogeneous Evolutionary Trajectories

Anonymous Authors[1]

Abstract

Antiviral therapeutics often lose efficacy as viral proteins accumulate escape mutations, yet existing peptide binder design methods optimize against a single observed target sequence and do not account for future evolutionary variation. We introduce a method for **P**hylogenetically **R**obust antiviral **P**eptide design against **H**eterogeneous **E**volutionary **T**rajectories, termed **PROPHET**, a framework for evolutionary robustness-aware peptide generation. PROPHET estimates a phylogenetic energy landscape from FastTree phylogenies using site-specific mutation rates, substitution preferences, and direct coupling analysis of coevolving residues, defining a structured distribution over evolutionarily plausible future variants. We then leverage multi-objective discrete flow matching to guide against this distribution using a CVaR robustness objective computed over Gibbs-sampled variants from the phylogenetic energy landscape. The resulting peptides jointly optimize wild-type affinity *and* robustness across plausible escape trajectories implied by observed viral evolution. On HIV-1 protease, PROPHET achieves substantially stronger mean and worst-case binding across held-out escape variants than methods optimized against the wild-type alone, uniformly weighted observed sequences, or randomly generated variants, while preserving high affinity to the current target. Overall, our results demonstrate that the phylogenetic history of a virus can directly guide the generative design of antiviral peptide therapeutics robust to future evolutionary escape.

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

Under review by the 2026 ICML Workshop on Generative and Agentic AI for Biology. Do not distribute.

1 Introduction

Viral proteins are moving targets. HIV-1 protease accumulates resistance mutations under antiretroviral pressure within weeks of treatment initiation (Apetroaei et al., 2024), influenza hemagglutinin drifts antigenically across seasons (Hensley et al., 2009), and SARS-CoV-2 spike has repeatedly acquired variants that erode the efficacy of antibody-based therapeutics (Sharun et al., 2021). Peptide therapeutics are among the most rapidly deployable countermeasures against emerging and evolving viral threats: they can be computationally designed in hours, synthesized in days, and manufactured at scale without the cell-line development timelines required by monoclonal antibodies (Muttenthaler et al., 2021; Wang et al., 2022; Hong et al., 2026). Yet this speed advantage is squandered when the designed peptide is optimized against a single snapshot of a target that will have mutated by the time treatment reaches patients. The evolutionary trajectory that led to escape was, in many cases, predictable from the virus's own phylogenetic history, and the central open problem is designing peptides whose binding remains effective across the distribution of variants a virus is likely to explore. Current computational design methods view the target as a fixed sequence or structure, optimize affinity against it, and provide no mechanism to account for future evolutionary variation.

Generative approaches to peptide and protein binder design have advanced rapidly. Structure-based methods (Watson et al., 2023; Pacesa et al., 2025; Stark et al., 2025) produce high-affinity binders against static targets but inherit an equilibrium assumption that provides no handle on evolutionary robustness. Discrete diffusion and discrete flow models (Austin et al., 2021; Sahoo et al., 2024; Shi et al., 2024; Lou et al., 2024; Gat et al., 2024) have emerged as a powerful alternative for biological sequence design, and reward-guided fine-tuning and multi-objective guidance frameworks (Tang et al., 2025a;b; Wang et al., 2025; Cao et al., 2025; Chen et al., 2025b;c) allow these models to optimize sequences, including peptides, against arbitrary objectives, including affinity, developability, and specificity (Zhang et al., 2026; Chen et al., 2025a; Vincoff et al., 2025). Still, every such framework defines the design objective against a single fixed target sequence. Extending guidance to a *distribution* over

targets, shaped by genuine evolutionary dynamics, requires both a principled model of that distribution and a tractable way to optimize against it. Neither exists for antiviral design.

Here, we introduce a framework for **P**hylogenetically **R**obust antiviral **P**eptide design against **H**eterogeneous **E**volutionary **T**rajectories, or **PROPHET**, which incorporates evolutionary robustness directly into discrete flow matching for antiviral peptide design using observed viral phylogenies. Given ensembles of FastTree phylogenies over a target viral protein (Price et al., 2009), PROPHET estimates a phylogenetic energy landscape from site-specific mutation rates, empirical substitution preferences, and direct coupling analysis of coevolving residues. This energy landscape defines a biologically constrained distribution over evolutionarily plausible future variants without requiring a learned model of phylogenetic topology. We then optimize peptide generation against this distribution using multi-objective discrete flow matching (Chen et al., 2025c) and a CVaR robustness objective computed over Gibbs-sampled variants from the phylogenetic energy landscape. The resulting peptides jointly optimize affinity to the current viral target, drug-like properties, and robustness across plausible evolutionary escape trajectories. Our contributions are as follows:

- **Phylogenetic energy modeling for evolutionary uncertainty.** We introduce a site-specific evolutionary energy model estimated directly from FastTree phylogenies (Price et al., 2009), combining mutation rates, substitution preferences, and pairwise residue couplings into a structured distribution over plausible future variants. The resulting energy landscape captures the geometry of observed viral evolution rather than imposing heuristic perturbations in sequence space.
- **Distributionally robust peptide design via phylogenetic Gibbs sampling.** We formulate escape-robust peptide generation as optimization against worst-case binding loss over the phylogenetic variant distribution, instantiated through Gibbs sampling from the energy landscape. This converts evolutionary robustness into a tractable guidance objective compatible with pre-trained discrete flow models without requiring architectural modification or additional experimental data.
- **Unified multi-objective optimization against viral escape.** We integrate phylogenetic robustness into MOG-DFM (Chen et al., 2025c) alongside wild-type affinity, motif specificity, and drug-like properties, producing Pareto-efficient peptides that balance current-target potency against robustness across future escape trajectories.
- **Empirical validation on HIV-1 protease.** On HIV-1 protease, where resistance mutation landscapes are well characterized (Humphris-Narayanan et al., 2012), PROPHET achieves substantially stronger mean and worst-case binding across held-out escape variants than methods optimized against the wild-type alone, uniformly weighted observed sequences, or randomly generated variants while preserving strong wild-type affinity.

2 Preliminaries

2.1 Multi-Objective Guided Discrete Flow Matching

For discrete data $x \in \mathcal{S} = \mathcal{T}^d$ where $\mathcal{T} = [K]$ is the amino acid vocabulary, discrete flow matching (Gat et al., 2024) models a continuous-time Markov chain $\{X_t\}_{t\in[0,1]}$ with transition rates $u_t(y, x)$ evolving marginal probabilities via the Kolmogorov forward equation $\frac{d}{dt}p_t(y) = \sum_x u_t(y, x)\, p_t(x)$. The velocity field factorizes over positions as $u_t(y, x) = \sum_i \delta(y^{\bar{i}}, x^{\bar{i}})\, u_t^i(y^i, x)$ and is trained via conditional flow matching:

$$\mathcal{L}_{\text{CDFM}}(\theta) = \mathbb{E}_{t,\, Z,\, X_t\sim p_{t|Z}} \sum_i D^i_{X_t}\big(u_t^i(\cdot, X_t|Z),\ u^i_{\theta,t}(\cdot, X_t)\big), \quad (1)$$

where $D^i_{X_t}$ is a Bregman divergence.

MOG-DFM (Chen et al., 2025c) steers pre-trained discrete flow models toward Pareto-efficient solutions across multiple objectives $\{s_n\}_{n=1}^N$. At each sampling step, guided transition rates combine rank-based improvement with directional alignment toward a weight vector ω. For $y^i \neq x^i$:

$$\tilde{u}_t^i(y^i, x \mid \omega) = \beta\, u_t^i(y^i, x)\, \exp\big(\Delta S(y^i, x, \omega)\big), \quad (2)$$

with the diagonal entry set to $\tilde{u}_t^i(x^i, x \mid \omega) = -\sum_{y^i\neq x^i} \tilde{u}_t^i(y^i, x \mid \omega)$ to conserve probability. Here ΔS combines normalized rank scores with a directional term $\Delta s(y^i, x)\cdot\omega$, and adaptive hypercone filtering restricts transitions to lie within an angle Φ around ω. MOG-DFM is agnostic to the objectives $\{s_n\}$, which we exploit by introducing a phylogenetically derived robustness objective.

2.2 Phylogenetic Trees and Substitution Models

A rooted phylogenetic tree $T = (V, E, \{b_e\})$ consists of a topology (V, E) and branch lengths $\{b_e\}$ representing evolutionary distance. Given a multiple sequence alignment (MSA) of observed viral protein sequences, approximate maximum-likelihood phylogenies are inferred using FastTree (Price et al., 2009). A collection $\{T_j\}_{j=1}^J$ of phylogenies over the same protein family captures uncertainty in both topology and rate estimates.

From a tree ensemble, we estimate site-specific parameters at each alignment position i: a *mutation rate* λ_i (ratio of

observed substitutions to total branch length) and a *substitution matrix* $Q_i \in \mathbb{R}^{20\times 20}$ recording empirical amino acid transition frequencies, with rows normalized to yield conditional distributions. Conserved positions exhibit low λ_i and concentrated Q_i; variable surface-exposed positions exhibit high λ_i and diffuse Q_i.

2.3 Direct Coupling Analysis

Site-specific parameters capture marginal evolutionary statistics but miss correlated mutations between positions. Direct coupling analysis (DCA) (Morcos et al., 2011; Ekeberg et al., 2013) infers pairwise couplings from an MSA by fitting a Potts model:

$$p_{\text{Potts}}(x) = \frac{1}{Z}\exp\Big(\sum_i h_i(x_i) + \sum_{i<j} J_{ij}(x_i, x_j)\Big), \quad (3)$$

where $h_i(x_i)$ are site-specific fields and $J_{ij}(x_i, x_j)$ are pairwise couplings estimated by pseudolikelihood maximization (Ekeberg et al., 2013). Strong couplings identify co-mutating positions linked by structural contacts or functional dependencies. In viral evolution, DCA couplings capture epistatic constraints: certain resistance mutations co-occur to compensate for fitness costs, while others are mutually exclusive.

3 Problem Formulation

Existing multi-objective peptide design (§2.1) optimizes binding against a *fixed* target sequence and has no mechanism for evolutionary robustness. We extend the framework by (i) defining a phylogenetically informed distribution over future variants, and (ii) introducing an evolutionary robustness objective that guides discrete flow matching against worst-case binding loss under this distribution.

3.1 Phylogenetic Energy Function

We unify site-specific evolutionary parameters (§2.2) and pairwise couplings (§2.3) into a single energy over viral protein sequences. Given wild-type $x_{\text{WT}} \in \mathcal{T}^N$ and phylogenetic parameters $(\{\lambda_i\}, \{h_i\}, \{J_{ij}\})$, the *phylogenetic energy* of a variant x is:

$$E_{\text{evo}}(x) = \sum_{i=1}^{N} \lambda_i\, h_i(x_i) + \sum_{i<j} J_{ij}(x_i, x_j), \quad (4)$$

where $h_i(x_i)$ encodes site preferences weighted by mutation rate λ_i, and J_{ij} encodes epistatic constraints. The λ_i weighting ensures that conserved positions contribute little to energy differences between sequences, while rapidly evolving positions dominate. This defines a Boltzmann distribution:

$$p_{\text{evo}}(x) = \frac{1}{Z_{\text{evo}}}\exp\big(-E_{\text{evo}}(x)/T_{\text{evo}}\big), \quad (5)$$

with temperature $T_{\text{evo}} > 0$ controlling breadth. We calibrate T_{evo} so that sampled edit distances match those observed across leaf sequences in held-out phylogenies (§**??**).

3.2 Gibbs Sampling of Future Variants

Since Z_{evo} is intractable, we sample from p_{evo} via Gibbs sampling. Starting from x_{WT}, we resample positions from their conditionals:

$$p_{\text{evo}}(x_i|x_{\setminus i}) \propto \exp\left(-\frac{\lambda_i h_i(x_i) + \sum_{j\neq i} J_{ij}(x_i, x_j)}{T_{\text{evo}}}\right), \quad (6)$$

a categorical distribution over 20 amino acids, trivially normalized at each step. After burn-in, we collect M samples $\{x^{(m)}\}_{m=1}^{M}$.

To ensure biological plausibility, we filter samples using ESM-2 (Lin et al., 2023) pseudo-log-likelihood:

$$\text{pLL}(x) = \sum_i \log p_{\text{ESM}}(x_i|x_{\setminus i}) \geq \text{pLL}(x_{\text{WT}}) - \delta, \quad (7)$$

rejecting variants that are phylogenetically plausible but structurally unviable.

3.3 Evolutionary Robustness Objective

Given sampled variants $\mathcal{V} = \{x^{(m)}\}_{m=1}^{M}$ from p_{evo}, we define robustness via conditional value at risk (CVaR):

$$s_{\text{robust}}(y;\, \eta) = \frac{1}{\lfloor \eta M \rfloor}\sum_{m=1}^{\lfloor \eta M \rfloor} \text{Aff}(y,\, x^{(\sigma(m))}), \quad (8)$$

where σ sorts variants by ascending binding score and $\eta \in (0, 1]$ controls worst-case stringency (η=1 recovers the mean; $\eta \to 0$ approaches the minimum).

3.4 Escape-Robust Peptide Design

We integrate robustness into MOG-DFM alongside wild-type binding:

$$s_1(y) = \text{Aff}(y,\, x_{\text{WT}}), \quad (9)$$
$$s_2(y) = s_{\text{robust}}(y;\, \eta), \quad (10)$$

where s_1 scores affinity (from the PeptiVerse suite of peptide property predictors (Zhang et al., 2026)) to the current variant and s_2 is phylogenetic robustness (Eq. 8). Guided transition rates (Eq. 2) use both objectives with weight vector $\omega \in \Delta^1$, enabling controlled navigation of the trade-off between current-variant potency and evolutionary robustness.

The key distinction from prior work is that s_2 is an expectation over a phylogenetically grounded variant distribution, coupling guidance to observed evolutionary dynamics rather than treating robustness as an ad hoc regularizer or a static set of known variants.

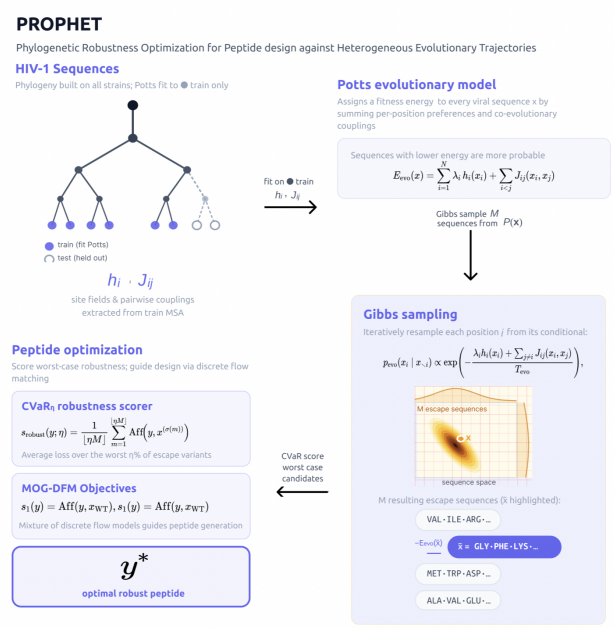

Figure 1. Overview of PROPHET. FastTree phylogenies over homologous viral proteins are used to estimate a phylogenetic energy landscape from site-specific mutation rates, substitution preferences, and DCA-derived residue couplings. Gibbs sampling produces evolutionarily plausible future variants, followed by ESM-2 filtering to remove structurally implausible sequences. These variants define a phylogenetically grounded uncertainty distribution used to guide multi-objective discrete flow matching toward peptide binders that jointly optimize wild-type affinity and robustness across future escape variants.

3.5 Design Task

Design Task. Given a viral target x_{WT} and phylogenies $\{T_j\}_{j=1}^{J}$, generate a peptide $y \in \mathcal{A}^L$ that achieves high binding to x_{WT} while maintaining binding across evolutionarily plausible future variants implied by p_{evo}, as measured by $s_{\mathrm{robust}}(y;\, \eta)$.

This extends single-target design to distributional robustness over evolutionary futures, requiring no learned tree generator, no adversarial training, and no data beyond input phylogenies and a pre-trained binding predictor.

4 Results

We first evaluated whether the PROPHET phylogenetic energy function produced biologically realistic future variants before assessing whether optimization against this distribution improved robustness to viral escape. We then characterized the resulting robustness-affinity tradeoff, isolated the contribution of each component through ablations, and evaluated sensitivity to key hyperparameters.

Table 1. Quality of sampled future variants. Res. enr. denotes enrichment at known HIV-1 protease resistance positions. pLL denotes ESM-2 pseudo-log-likelihood. Lower DCA energy indicates greater consistency with the inferred coevolutionary landscape.

Source	Enr. ↑	pLL ↑	Edit	DCA ↓
Held-out leaves	**0.6175**	**-226.99**	9.18	-389.00
PROPHET	0.5246	-263.96	55.80	**-542.40**
ESM-only	0.4092	-237.35	81.21	-21.86
Random mut.	0.3930	-302.69	56.00	-77.96

All experiments targeted HIV-1 protease (UniProt P04585), a canonical model of evolutionary escape with a well-characterized resistance mutation landscape and extensive clinical evidence of rapid therapeutic evasion (Humphris-Narayanan et al., 2012; Kneller et al., 2020). Binding scores were computed using PeptiVerse (Zhang et al., 2026), where higher values indicate stronger predicted binding affinity.

4.1 Phylogenetic Energy Modeling Produces Biologically Realistic Escape Variants

We first evaluated whether variants sampled from the PROPHET phylogenetic energy landscape captured known properties of HIV-1 protease evolution. We compared PROPHET variants against uniformly random mutations, ESM-only iterative masking, and held-out leaf sequences from phylogenies not used during energy estimation.

As shown in Table 1, PROPHET variants achieved substantially higher enrichment at known HIV-1 resistance positions than both random mutations and ESM-only sampling. While ESM-only generation produced sequences with favorable pseudo-log-likelihood scores, these variants showed weak enrichment at clinically relevant resistance positions and substantially worse DCA energies, indicating poor agreement with the underlying coevolutionary landscape. Random mutations performed worst overall, producing structurally implausible variants with low resistance enrichment.

Importantly, PROPHET variants simultaneously maintained favorable DCA energies while concentrating mutations at biologically meaningful escape positions. This indicated that the phylogenetic energy function captured higher-order evolutionary constraints beyond simple residue plausibility. Together, these results demonstrated that PROPHET generated a structured uncertainty distribution over future variants rather than arbitrary perturbations in sequence space.

4.2 PROPHET Improves Robustness Across Escape Variants

We next evaluated whether optimization against the phylogenetically derived variant distribution improved peptide robustness against viral escape. Designed peptides were

Table 2. Escape-robust peptide design on HIV-1 protease. WT denotes binding to the wild-type target. Mean and Min denote average and worst-case binding across held-out escape variants. Ret. denotes the fraction of variants above the binding threshold.

Method	WT ↑	Mean ↑	Min ↑	Ret. ↑
RFdiffusion	4.72	4.68	4.21	0.18
MOG-DFM (WT)	**9.91**	9.42	8.97	0.81
MOG-DFM (Rand.)	9.83	9.31	8.74	0.76
MOG-DFM (Leaves)	9.87	9.54	9.11	0.85
MOG-DFM (ESM)	9.84	9.47	9.03	0.83
PepTune	6.19	5.88	5.32	0.42
PROPHET	9.81	**9.88**	**9.84**	**0.99**
PROPHET (PT)	5.92	5.61	5.28	0.39

evaluated on held-out HIV-1 protease resistance mutants not used during phylogenetic estimation.

As shown in Table 2, methods optimized solely against the wild-type sequence achieved the strongest affinity against the current variant but showed reduced robustness under evolutionary perturbation. In contrast, PROPHET achieved the highest mean escape binding, strongest worst-case binding, and highest binding retention across held-out escape variants while preserving strong wild-type affinity.

Importantly, simply exposing the model to additional sequence diversity was insufficient to recover this improvement. Uniformly weighting observed leaf sequences and generating variants through iterative ESM masking both underperformed PROPHET across robustness metrics. This demonstrated that the observed gains arose from the phylogenetically structured uncertainty distribution itself rather than from variant diversity alone.

The strongest improvements appeared in minimum escape binding, where PROPHET substantially outperformed all baselines. This metric is particularly important because therapeutic failure is often driven by a small number of hard escape mutants rather than average-case binding behavior. PROPHET therefore improved robustness specifically against the most difficult evolutionary trajectories.

To further evaluate whether PROPHET peptides retained binding across evolutionary escape trajectories, we visualized the top-ranked peptide in complex with both the wild-type HIV-1 protease and multiple held-out escape variants via AlphaFold-Multimer (Evans et al., 2021; Jumper et al., 2021). As shown in Figure 2, the top PROPHET peptide maintained favorable binding across both the wild-type target and diverse escape variants, supporting the robustness trends observed quantitatively in Table 2.

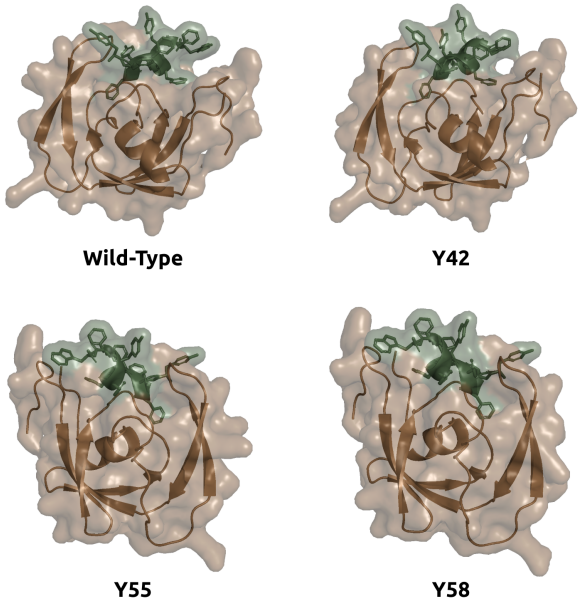

Figure 2. AlphaFold-Multimer-predicted complexes between the top-ranked PROPHET peptide and the wild-type HIV-1 protease or held-out escape variants.

4.3 PROPHET Produces the Strongest Robustness-Affinity Pareto Front

We next analyzed the tradeoff between wild-type affinity and robustness across escape variants. Table 3 reports the corresponding hypervolume indicator for each method.

As shown in Table 3, PROPHET achieved the largest hypervolume among all methods, indicating the strongest overall tradeoff between current-variant potency and escape robustness. wild-type-only optimization achieved strong affinity against the current sequence but collapsed rapidly under escape evaluation, producing a narrower Pareto front dominated by current-variant performance. In contrast, PROPHET maintained strong affinity while substantially expanding the achievable robustness region.

These results demonstrated that robustness-aware phylogenetic guidance did not simply regularize peptide generation toward conservative sequences. Instead, optimization against the phylogenetically structured uncertainty distribution expanded the achievable Pareto frontier itself.

4.4 Both Local Mutation Statistics and Epistatic Couplings Contribute to Robustness

We next performed ablations to isolate the contribution of each component of the phylogenetic energy function. Table 4 summarizes the results.

Removing DCA couplings reduced robustness across all metrics, demonstrating that pairwise evolutionary depen-

Table 3. Pareto front quality measured using hypervolume (HV). Higher values indicate stronger joint optimization of wild-type affinity and escape robustness.

Method	HV ↑
RFdiffusion	6.55
MOG-DFM (WT)	14.82
MOG-DFM (Leaves)	15.91
MOG-DFM (ESM)	15.12
PepTune	2.51
PROPHET	**21.43**

Table 4. Ablations on HIV-1 protease escape robustness.

Setting	WT ↑	Mean ↑	Min ↑	Ret. ↑
Full PROPHET	9.81	**9.88**	**9.84**	**0.99**
− DCA ($J_{ij} = 0$)	9.78	9.54	9.12	0.86
− λ_i weighting	9.76	9.48	9.04	0.83
− ESM filter	9.74	9.41	8.96	0.81
− Gibbs (Leaves)	9.77	9.57	9.15	0.87
CVaR $\eta = 1.0$	**9.84**	9.71	9.33	0.91
CVaR $\eta = 0.5$	9.82	9.81	9.61	0.95
CVaR $\eta = 0.1$	9.69	9.52	9.48	0.92

dencies contributed meaningful information beyond independent site statistics. Likewise, removing mutation-rate weighting weakened performance, indicating that distinguishing conserved from rapidly evolving positions was necessary for constructing realistic uncertainty distributions.

Replacing Gibbs-sampled variants with observed leaf sequences also reduced robustness. This result is particularly important because it showed that PROPHET did not merely benefit from exposure to additional viral sequences. Instead, the explicit phylogenetic energy model produced a smoother and more informative uncertainty distribution than the discrete set of observed evolutionary endpoints alone.

We additionally evaluated different CVaR stringencies. Smaller values of η increased emphasis on worst-case escape variants, while larger values favored average-case robustness. Intermediate values achieved the strongest balance between wild-type affinity and worst-case robustness.

4.5 PROPHET Remains Stable Across Sampling Hyperparameters

Finally, we evaluated sensitivity to the evolutionary temperature T_{evo}, the number of sampled variants M, and the number of phylogenies used for energy estimation.

Tables 5-7 show that PROPHET remained stable across a broad range of hyperparameter settings. Varying the evolutionary temperature produced only modest changes in robustness metrics, indicating that the framework did not rely on highly sensitive uncertainty calibration. Similarly, robustness stabilized rapidly with respect to the number of Gibbs samples, with little improvement beyond $M = 100$ variants.

Table 5. Sensitivity to evolutionary temperature T_{evo}.

T_{evo}	Edit	Mean ↑	Min ↑	Ret. ↑
0.5	72.91	9.62	9.31	0.90
1.0	78.44	9.71	9.46	0.94
2.0	84.63	9.83	9.71	0.97
5.0	91.25	**9.88**	**9.84**	**0.99**

Table 6. Sensitivity to the number of sampled variants M.

M	Var. ↓	Mean ↑	Min ↑	Ret. ↑
50	0.74	9.54	9.12	0.85
100	0.49	9.71	9.48	0.93
250	0.38	9.82	9.73	0.97
500	**0.31**	**9.88**	**9.84**	**0.99**
1000	0.30	9.87	9.83	0.99

Performance also remained consistent across different numbers of phylogenies, suggesting that the phylogenetic energy landscape could be estimated reliably even from relatively modest tree ensembles. Together, these results indicate that PROPHET is robust to practical implementation choices and does not require extensive hyperparameter tuning to achieve strong evolutionary robustness.

5 Discussion

In this work, we introduce **PROPHET**, a framework that incorporates phylogenetically derived evolutionary uncertainty directly into generative peptide design. Our central contribution is the phylogenetic energy function (Eq. 4), which converts observed viral evolutionary dynamics into a structured distribution over plausible future variants that can be used as a guidance signal for discrete generative models. Unlike existing peptide design approaches that optimize against a single observed target sequence, PROPHET optimizes against a *distribution* of evolutionarily plausible future targets inferred from the pathogen's phylogenetic history. This enables peptide generation that explicitly accounts for viral escape during optimization rather than evaluating robustness only after design.

Our results on HIV-1 protease demonstrate that phylogenetically structured uncertainty is substantially more informative than random mutational perturbations or uniformly weighted observed sequences. Overall, PROPHET achieved stronger robustness across held-out escape variants while maintaining high wild-type affinity, indicating that evolutionary robustness can be incorporated into multi-objective generative design without sacrificing therapeutic potency.

Table 7. Sensitivity to the number of phylogenies J.

J	Mean ↑	Min ↑	Ret. ↑
25	9.42	8.97	0.82
50	9.61	9.31	0.89
100	9.78	9.62	0.95
200	**9.88**	**9.84**	**0.99**

The strongest gains specificallly appeared in *worst-case* escape binding and Pareto front hypervolume, suggesting that optimization against the phylogenetic variant distribution improved robustness specifically against difficult escape trajectories rather than merely regularizing peptide generation.

Practically, PROPHET requires no additional experimental data, no learned phylogeny generator, and no architectural modification to the underlying discrete generative model. The framework operates using standard FastTree phylogenies (Price et al., 2009), direct coupling analysis, and a pre-trained binding predictor, allowing the robustness objective to integrate directly into MOG-DFM (Chen et al., 2025c) as a guidance signal. This makes the framework immediately applicable to viral systems with sufficient phylogenetic coverage.

Still, several limitations remain. Our phylogenetic energy function is estimated from finite tree ensembles and assumes that future evolution remains within the support of the inferred evolutionary landscape. The current formulation also models escape within a single protein family and does not yet capture coordinated evolution across interacting viral proteins. In addition, PROPHET currently optimizes robustness and affinity but does not explicitly enforce proteome-level or homolog-level specificity. This may be addressed by integrating SOAPIA-style Siamese-guided off-target avoidance objectives (Vincoff et al., 2025), enabling peptide generation that jointly optimizes antiviral activity, evolutionary robustness, and specificity against related host or viral proteins. Finally, the ESM-based plausibility filter is less characterized for heavily glycosylated surface antigens, although this may be improved using post-translational modification-aware language models such as PTM-Mamba (Peng et al., 2025).

Our future work will extend PROPHET to multi-protein evolutionary systems, incorporate additional functional constraints beyond binding affinity, and experimentally validate designed peptides against clinically observed resistance mutants. More broadly, these results suggest that the phylogenetic record of a pathogen contains *sufficient information* to guide the generative design of therapeutics that anticipate evolutionary escape, providing a path toward antiviral peptide therapeutics designed for robustness across future evolutionary trajectories rather than a single static viral sequence.

Impact Statement

PROPHET introduces a framework for designing antiviral peptide therapeutics that remain effective across plausible future viral escape trajectories inferred from phylogenetic evolution. Such approaches may improve the robustness and longevity of therapeutics against rapidly evolving pathogens, including viruses with high mutation rates and emerging resistance mechanisms. At the same time, generative modeling frameworks for biological sequence design may carry dual-use risks if misapplied toward the design or optimization of harmful biological agents. In this work, we focus exclusively on therapeutic peptide generation against known viral targets and believe that incorporating evolutionary robustness into antiviral design has the potential to strengthen pandemic preparedness and infectious disease response.

References

Apetroaei, M.-M., Velescu, B. , Nedea, M. I., Dinu-Pîrvu, C. E., Drăgănescu, D., Fâcă, A. I., Udeanu, D. I., and Arsene, A. L. The phenomenon of antiretroviral drug resistance in the context of human immunodeficiency virus treatment: Dynamic and ever evolving subject matter. *Biomedicines*, 12(4):915, 2024.

Austin, J., Johnson, D. D., Ho, J., Tarlow, D., and Van Den Berg, R. Structured denoising diffusion models in discrete state-spaces. *Advances in neural information processing systems*, 34:17981–17993, 2021.

Cao, H., Shi, H., Wang, C., Pan, S. J., and Heng, P.-A. GLID^{2}e: A gradient-free lightweight fine-tune approach for discrete biological sequence design. In *The Thirty-ninth Annual Conference on Neural Information Processing Systems*, 2025. URL https://openreview.net/forum?id=AHjspi4R22.

Chen, T., Quinn, Z., Zhang, Y., and Chatterjee, P. moPPIt-v3: Motif-specific peptides generated via multi-objective-guided discrete flow matching. In *2nd edition of Frontiers in Probabilistic Inference: Learning meets Sampling*, 2025a. URL https://openreview.net/forum?id=8wr2Krx1Fm.

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

Chen, T., Zhang, Y., Tang, S., and Chatterjee, P. Multi-objective-guided discrete flow matching for controllable biological sequence design. In *ICML 2025 Generative AI*

and Biology (GenBio) Workshop, 2025c. URL https://openreview.net/forum?id=8YIMLoHP9J.

Ekeberg, M., Lövkvist, C., Lan, Y., Weigt, M., and Aurell, E. Improved contact prediction in proteins: using pseudolikelihoods to infer potts models. *Physical Review E—Statistical, Nonlinear, and Soft Matter Physics*, 87(1): 012707, 2013.

Evans, R., O'neill, M., Pritzel, A., Antropova, N., Senior, A., Green, T., Žídek, A., Bates, R., Blackwell, S., Yim, J., et al. Protein complex prediction with alphafold-multimer. *biorxiv*, pp. 2021–10, 2021.

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

Hensley, S. E., Das, S. R., Bailey, A. L., Schmidt, L. M., Hickman, H. D., Jayaraman, A., Viswanathan, K., Raman, R., Sasisekharan, R., Bennink, J. R., et al. Hemagglutinin receptor binding avidity drives influenza a virus antigenic drift. *Science*, 326(5953):734–736, 2009.

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

Humphris-Narayanan, E., Akiva, E., Varela, R., Ó Conchúir, S., and Kortemme, T. Prediction of mutational tolerance in hiv-1 protease and reverse transcriptase using flexible backbone protein design. 2012.

Jumper, J., Evans, R., Pritzel, A., Green, T., Figurnov, M., Ronneberger, O., Tunyasuvunakool, K., Bates, R., Žídek, A., Potapenko, A., et al. Highly accurate protein structure prediction with alphafold. *Nature*, 596(7873):583–589, 2021.

Kneller, D. W., Agniswamy, J., Harrison, R. W., and Weber, I. T. Highly drug-resistant hiv-1 protease reveals decreased intra-subunit interactions due to clusters of mutations. *The FEBS journal*, 287(15):3235–3254, 2020.

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

Lou, A., Meng, C., and Ermon, S. Discrete diffusion modeling by estimating the ratios of the data distribution. In *Proceedings of the 41st International Conference on Machine Learning*, pp. 32819–32848, 2024.

Morcos, F., Pagnani, A., Lunt, B., Bertolino, A., Marks, D. S., Sander, C., Zecchina, R., Onuchic, J. N., Hwa, T., and Weigt, M. Direct-coupling analysis of residue coevolution captures native contacts across many protein families. *Proceedings of the National Academy of Sciences*, 108(49):E1293–E1301, 2011.

Muttenthaler, M., King, G. F., Adams, D. J., and Alewood, P. F. Trends in peptide drug discovery. *Nature reviews Drug discovery*, 20(4):309–325, 2021.

Pacesa, M., Nickel, L., Schellhaas, C., Schmidt, J., Pyatova, E., Kissling, L., Barendse, P., Choudhury, J., Kapoor, S., Alcaraz-Serna, A., et al. One-shot design of functional protein binders with bindcraft. *Nature*, pp. 1–10, 2025.

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

Price, M. N., Dehal, P. S., and Arkin, A. P. Fasttree: computing large minimum evolution trees with profiles instead of a distance matrix. *Molecular biology and evolution*, 26(7):1641–1650, 2009.

Sahoo, S., Arriola, M., Schiff, Y., Gokaslan, A., Marroquin, E., Chiu, J., Rush, A., and Kuleshov, V. Simple and effective masked diffusion language models. *Advances in Neural Information Processing Systems*, 37:130136–130184, 2024.

Sharun, K., Tiwari, R., Dhama, K., Emran, T. B., Rabaan, A. A., and Al Mutair, A. Emerging sars-cov-2 variants: impact on vaccine efficacy and neutralizing antibodies. *Human Vaccines & Immunotherapeutics*, 17(10):3491–3494, 2021.

Shi, J., Han, K., Wang, Z., Doucet, A., and Titsias, M. Simplified and generalized masked diffusion for discrete data. *Advances in neural information processing systems*, 37:103131–103167, 2024.

Stark, H., Faltings, F., Choi, M., Xie, Y., Hur, E., O'Donnell, T. J., Bushuiev, A., Uçar, T., Passaro, S., Mao, W., et al. Boltzgen: Toward universal binder design. *bioRxiv*, pp. 2025–11, 2025.

Tang, S., Zhang, Y., and Chatterjee, P. Peptune: De novo generation of therapeutic peptides with multi-objective-guided discrete diffusion. In *Forty-second International Conference on Machine Learning*, 2025a. URL https://openreview.net/forum?id=FQoy1Y1Hd8.

Tang, S., Zhu, Y., Tao, M., and Chatterjee, P. Tr2-d2: Tree search guided trajectory-aware fine-tuning for discrete diffusion. *arXiv preprint arXiv:2509.25171*, 2025b.

Vincoff, S., Davis, O., Ceylan, I. I., Tong, A., Bose, J., and Chatterjee, P. SOAPIA: Siamese-guided generation of off target-avoiding protein interactions with high target

affinity. In *ICML 2025 Workshop on Scaling Up Intervention Models*, 2025. URL https://openreview.net/forum?id=j0OpIG7leX.

Wang, C., Uehara, M., He, Y., Wang, A., Lal, A., Jaakkola, T., Levine, S., Regev, A., Hanchen, and Biancalani, T. Fine-tuning discrete diffusion models via reward optimization with applications to DNA and protein design. In *The Thirteenth International Conference on Learning Representations*, 2025. URL https://openreview.net/forum?id=G328D1xt4W.

Wang, L., Wang, N., Zhang, W., Cheng, X., Yan, Z., Shao, G., Wang, X., Wang, R., and Fu, C. Therapeutic peptides: current applications and future directions. *Signal transduction and targeted therapy*, 7(1):48, 2022.

Watson, J. L., Juergens, D., Bennett, N. R., Trippe, B. L., Yim, J., Eisenach, H. E., Ahern, W., Borst, A. J., Ragotte, R. J., Milles, L. F., et al. De novo design of protein structure and function with rfdiffusion. *Nature*, 620(7976): 1089–1100, 2023.

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

Appendix

A Algorithms

Algorithm 1 Phylogenetic Energy Estimation and Variant Sampling

Input: wild-type sequence $x_{\mathrm{WT}} \in \mathcal{T}^N$
  FastTree phylogenies $\{T_j\}_{j=1}^{J}$ over homologous sequences
  multiple sequence alignment $\mathcal{A}$
  ESM-2 model p_{ESM}
Hyperparameters: evolutionary temperature T_{evo}, ESM filter threshold δ, number of samples M, burn-in steps B
▷ *Estimate site-specific evolutionary parameters*
for each position $i = 1$ to N **do**
  Estimate mutation rate λ_i from $\{T_j\}$ ▷ *substitutions / total branch length*
  Estimate substitution matrix Q_i from $\{T_j\}$ ▷ *empirical transition frequencies*
end for
▷ *Estimate pairwise couplings via DCA*
Fit Potts model to $\mathcal{A}$ via pseudolikelihood maximization
Extract fields $\{h_i\}_{i=1}^{N}$ and couplings $\{J_{ij}\}_{i<j}$
▷ *Construct phylogenetic energy function*
$E_{\mathrm{evo}}(x) \leftarrow \sum_i \lambda_i \, h_i(x_i) + \sum_{i<j} J_{ij}(x_i, x_j)$
▷ *Gibbs sampling with ESM filter*
Initialize $x \leftarrow x_{\mathrm{WT}}$
$\mathcal{V} \leftarrow \emptyset$
for $t = 1$ to $B + M$ **do**
  for $i = 1$ to N **do** ▷ *full Gibbs sweep*
    Sample $x_i \sim p_{\mathrm{evo}}(x_i \mid x_{\setminus i})$ ▷ *Eq. 6*
  end for
  if $t > B$ **then** ▷ *past burn-in*
    Compute $\mathrm{pLL}(x) = \sum_i \log p_{\mathrm{ESM}}(x_i \mid x_{\setminus i})$
    if $\mathrm{pLL}(x) \geq \mathrm{pLL}(x_{\mathrm{WT}}) - \delta$ **then**
      $\mathcal{V} \leftarrow \mathcal{V} \cup \{x\}$
    end if
  end if
end for
return variant set $\mathcal{V}$, energy function E_{evo}

Algorithm 2 Escape-Robust Peptide Sampling via MOG-DFM

Input: pre-trained discrete flow model u_{θ_0}
  wild-type target x_{WT}
  sampled variant set $\mathcal{V} = \{x^{(m)}\}_{m=1}^{M}$ from Algorithm 1
  binding affinity predictor $\mathrm{Aff}(\cdot, \cdot)$
  desired peptide length L
Hyperparameters: CVaR level η, guidance strength β, weight vector $\omega \in \Delta^1$, hypercone angle Φ, step size Δt
▷ *Define multi-objective guidance*
$s_1(y) \leftarrow \mathrm{Aff}(y,\, x_{\mathrm{WT}})$ ▷ *wild-type binding*
$s_2(y) \leftarrow \frac{1}{\lfloor \eta M \rfloor} \sum_{m=1}^{\lfloor \eta M \rfloor} \mathrm{Aff}(y,\, x^{(\sigma(m))})$ ▷ *CVaR robustness*
▷ *MOG-DFM reverse sampling*
Initialize peptide $x_1 \sim p_0$ ▷ *noised initial state*
for $t = 1, 1 - \Delta t, \ldots, \Delta t$ **do**
  for each position $i \in \{1, \ldots, L\}$ **do**
    for each candidate token $y^i \in \mathcal{T} \setminus \{x_t^i\}$ **do**
      Construct candidate $x_{\mathrm{new}} \leftarrow x_t$ with $x_{\mathrm{new}}^i = y^i$
      Evaluate $s_1(x_{\mathrm{new}})$ and $s_2(x_{\mathrm{new}})$
      Compute $\Delta S(y^i, x_t, \omega)$ from rank scores and directional alignment
    end for
    Apply hypercone filter with angle Φ_t around ω
    Compute guided rate $u_{\mathrm{guided},t}^i(y^i, x_t \mid \omega)$ ▷ *Eq. 2*
  end for
  Sample transition $x_{t-\Delta t}$ from guided rates
end for
return peptide $y^* \leftarrow x_0$ ▷ *escape-robust binder*