PEGASUS: QUERY-EFFICIENT MULTI-OBJECTIVE BIOMOLECULAR DESIGN WITH DISCRETE FLOW MAPS

Anonymous authors
Paper under double-blind review

ABSTRACT

Discrete flow maps compress generation into one or a few model evaluations, but multi-objective optimization can still require many expensive oracle evaluations because intermediate states provide limited information about reachable terminal quality. To address this, we introduce **PEGASUS** (**P**areto-**E**fficient **G**uidance for **A**ction **S**election **U**sing Future-Observability **S**cores), a query-efficient inference-time optimizer for pretrained discrete flow maps. PEGASUS introduces *future observability*, the conditional terminal objective potential of an intermediate state under the generator's continuation distribution. It estimates this potential through objective-independent finite-time Koopman transport and an online task readout, combining it with multiscale Hamming actions and Pareto Decision Observability (PDO) to focus exact evaluations on decision-relevant uncertainty. Under an exact decision-linear response model, the PDO acquisition bound depends on the intrinsic unresolved decision dimension and is independent of the candidate count. Empirically, PEGASUS achieves competitive or superior terminal Pareto quality with up to $99\times$ fewer unique oracle queries than state-of-the-art discrete multi-objective generative baselines. Overall, PEGASUS provides a powerful, sequence-based framework for query-efficient multi-property biomolecule design.

1 INTRODUCTION

Biomolecular design requires search over vast discrete sequence spaces for candidates that simultaneously satisfy competing functional and developability objectives. For example, therapeutic proteins and peptides require high affinity for target engagement, but also favorable stability, solubility, permeability, toxicity, and manufacturability (Hong et al., 2026; Chen et al., 2023). Evaluating candidate binding and developability can require repeated property predictions, structure modeling, and docking, while limited experimental throughput further constrains downstream validation (Ebrahimi & Samanta, 2023; Kuo et al., 2026). Newer generative models, such as discrete diffusion and discrete flows, enable direct generation in categorical sequence spaces and have supported multi-objective biological design (Tang et al., 2025a;b; Chen et al., 2025c;b;a; Vincoff et al., 2026). Recently, flow maps have also addressed complementary bottlenecks by learning finite-time transitions that compress long generative trajectories into one or a few model evaluations, with recent formulations extending this capability to categorical data (Boffi et al., 2024; Potaptchik et al., 2026; Roos et al., 2026). Overall, these advances make finite-time discrete generation attractive for biomolecular optimization, where sampling and evaluation costs *both* matter.

Finite-time generation removes intermediate states at which inference-time optimizers evaluate partial candidates, branch trajectories, and incorporate feedback. Methods such as PepTune and MOG-DFM repeatedly evaluate transitions or expand search trees, while DSearch estimates intermediate rewards through explicit lookahead (Tang et al., 2025a; Chen et al., 2025c; Li et al., 2025). Recovering feedback through repeated branching or rollouts can erode the computational savings of few-step generation. Continuous gradient-based flow-map guidance also does not transfer directly to categorical states, where first-order approximations can incur substantial error (Huang et al., 2026; Wan et al., 2026). **Few-step generation does not by itself establish few-query optimization.**

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

Closing this gap requires guidance that anticipates how intermediate choices affect final sequence quality while limiting the objective evaluations needed to make those choices. Current-state scores alone provide incomplete guidance because similarly scored states can lead to different terminal outcomes. We therefore introduce **future observability**, the terminal objective potential induced by the generator's continuation distribution from an online intermediate branch. Related methods obtain prospective information through reward-specific value estimation, explicit lookahead, or continuous reward propagation (Li et al., 2025; 2024; Huang et al., 2026). For multi-objective optimization under a limited oracle budget, we seek a finite-time representation that can support this prediction across objectives. Using these predictions efficiently also requires choosing how extensively to modify a sequence and which objective evaluations to acquire. Fine local edits support refinement, while coordinated changes can reach favorable regions beyond the local neighborhood. Additional queries are useful when their outcomes can change the selected Pareto action; uncertainty alone does not establish this need. Query-efficient discrete optimization therefore calls for a **coarse-to-fine process**: coarse search estimates reachable terminal potential and allocates continuations, while fine refinement combines multiscale actions with exact queries that can change the Pareto decision.

Here, we introduce **PEGASUS**, **P**areto-**E**fficient **G**uidance for **A**ction **S**election **U**sing Future-Observability **S**cores, for query-efficient multi-objective optimization of pretrained discrete flow maps. We use a learned objective-independent finite-time Koopman representation as a structured realization of future observability: it transports observables through the frozen generative dynamics (Otto & Rowley, 2021), while task-specific readouts fitted from the evaluated archive estimate terminal objective potential. PEGASUS uses these estimates to allocate continuation computation among stochastic branches, then applies multiscale sequence actions for fine refinement and **Pareto Decision Observability (PDO)** to allocate exact objective queries according to whether unresolved information can change the current decision. **Our key contributions are:**

- We introduce a **coarse search stage** that uses future observability, realized through objective-independent Koopman transport and task-specific readouts, to direct continuation computation toward promising terminal outcomes.
- We develop a **fine refinement stage** that couples adaptive multiscale actions with **Pareto Decision Observability (PDO)**, directing exact queries toward decision-relevant challengers. Under exact decision-linear responses, PDO's query bound depends on the intrinsic unresolved decision dimension rather than candidate count.
- PEGASUS achieves strong Pareto quality with far fewer oracle queries in six-objective peptide design across eight targets, and generalizes to three-objective DNA optimization.

2 Related Work

Multi-objective biological sequence design. PepTune uses tree-guided discrete diffusion (Tang et al., 2025a), while MOG-DFM and AReUReDi guide discrete flows (Chen et al., 2025c;b). MP2D and MosPro search over sequence edits (Kong et al., 2026; Luo et al., 2025); LaMBO applies Bayesian optimization (Stanton et al., 2022), LaMBO-2 introduces guided diffusion with edit constraints (Gruver et al., 2023), and TR2-D2 fine-tunes generators using optimized trajectories (Tang et al., 2025b). These approaches leave open how to jointly allocate sequence interventions and expensive objective evaluations using the terminal potential of intermediate generative states. PEGASUS addresses this problem under frozen discrete dynamics through prospective branch selection, multiscale actions, and selective oracle acquisition.

Generative Pareto optimization. ParetoFlow, Pareto-conditioned diffusion, and SPREAD use generative models for Pareto search or refinement (Yuan et al., 2025; Shrestha et al., 2026; Hotegni & Peitz, 2026), while Pareto-set learning amortizes preference-conditioned optimization (Lin et al., 2022). Their treatment of Pareto search leaves a distinct question of which intermediate discrete candidates warrant exact evaluation under a limited oracle budget. Our setting emphasizes query-efficient categorical search, where each exact evaluation returns the complete objective vector.

Finite-time generation and prospective guidance. Flow-map matching and discrete variants, including latent-kernel maps, enable few-step generation (Boffi et al., 2024; Potaptchik et al., 2026; Roos et al., 2026; Ahmed et al., 2026). DSearch, soft value-based decoding, reward-guided flow maps, and iterative biological refinement incorporate lookahead or reward information during generation (Li et al., 2025; 2024; Huang et al., 2026; Uehara et al., 2025; Wan et al., 2026). However,

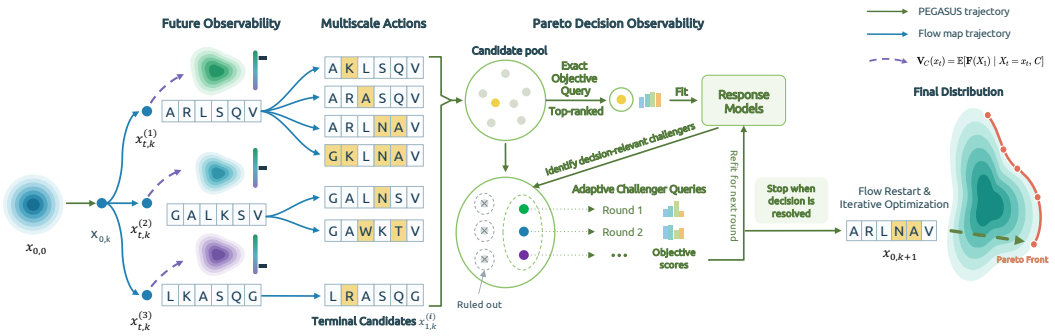


Figure 1: **Overview of PEGASUS.** Starting from $x_{0,0}$ sampled from the uniform categorical prior, each round k begins with **coarse search**: the frozen flow produces intermediate branches $x_{t,k}^{(i)}$, and **Future Observability** estimates their terminal objective potential to allocate continuations and verify promising terminal endpoints. **Fine refinement** then applies **Multiscale Actions** to generate terminal candidates at different edit scales. **Pareto Decision Observability (PDO)** exact-queries a top-ranked candidate, fits response models, and adaptively evaluates decision-relevant challengers while ruling out alternatives that cannot change the action decision. The best exactly evaluated improvement becomes $x_{0,k+1}$ and restarts the frozen flow. Repeating this process across rounds and preferences produces the final candidate library.

prospective guidance does not address how to compare multiscale actions and allocate exact evaluations under multi-objective uncertainty. PEGASUS instead estimates terminal objective potential across competing branches without retraining the base generator.

Koopman representations for prospective decisions. Koopman theory provides approximately linear evolution of learned observables (Otto & Rowley, 2021; Lusch et al., 2018) and has been applied to continuous generative sampling (Turan et al., 2025). The remaining challenge is to recover task-specific terminal objectives from objective-independent representations under stochastic discrete dynamics. PEGASUS combines finite-time Koopman transport with task-specific readouts to estimate reachable terminal quality.

Decision-aware acquisition and multiscale actions. Predictive entropy search targets Pareto-set uncertainty (Hernández-Lobato et al., 2016), multi-output linear bandits study Pareto identification (Kone et al., 2025), and variable-neighborhood search explores discrete edit scales (Duarte et al., 2015). These approaches do not address how multiscale actions should be coupled to decision-relevant exact evaluations. PEGASUS evaluates complete multiscale transformations and queries the exact oracle only when unresolved uncertainty could alter the selected action.

3 PEGASUS: Pareto-Efficient Guidance for Action Selection Using Future-Observability Scores

Background on multi-objective optimization, Latent-Kernel discrete flow maps (LKF), and Koopman operators is given in Appendix A. Let $\mathcal{X} = \mathcal{V}^L$ be the sequence space and $\mathbf{F}(x) \in \mathbb{R}^m$ the maximize-oriented objective vector. For preference $\mathbf{w}$, PEGASUS uses the augmented Tchebycheff utility $U_{\mathbf{w}}$, abbreviated U (Eq. 13), and one exact query evaluates the full vector $\mathbf{F}(x)$ for a previously unseen sequence. Figure 1 summarizes the two-stage workflow: coarse search uses future observability to allocate continuations and verify terminal candidates, followed by fine refinement with multiscale actions and PDO. Algorithm 1 in Appendix C gives the full procedure. We use t, s for flow time, k for optimization round, i for online branch/candidate, n for offline calibration or archive example, j for independent continuation, and ℓ for sequence position; thus online states are $x_{t,k}^{(i)}$, and sequences are written $x = (x[1], \ldots, x[L])$ with token $x[\ell]$.

3.1 Future Observability via Finite-Time Koopman Transport

For an online intermediate branch $x_{s,k}^{(i)}$ and a prescribed continuation C_i from time s to terminal time 1, define the **terminal objective potential**

$$\mathbf{V}_{C_i}(x_{s,k}^{(i)}) = \mathbb{E}\Big[\mathbf{F}(X_1) \mid X_s = x_{s,k}^{(i)}, C_i\Big]. \tag{1}$$

PEGASUS evaluates a branch through $U(\mathbf{V}_{C_i}(x_{s,k}^{(i)}))$ under the fixed preference, i.e., the utility of the conditional mean vector rather than the expected nonlinear utility. Directly estimating this potential for every branch would require repeated terminal rollouts and expensive objective evaluations. Instead, PEGASUS first estimates the expected terminal representation through objective-independent finite-time Koopman transport, then maps that representation to the current objectives using an online task-specific readout. Their composition provides a prospective estimate of terminal quality before committing exact oracle queries.

Objective-independent finite-time transport. To estimate how an intermediate state evolves without requiring downstream objective labels, we learn a Koopman observable ϕ_K using trajectories from the frozen LKF generator and calibrate finite-time operators $\{A_C\}$ (A.3). The defining **approximation of the expected terminal representation** is

$$\mathbb{E}[\phi_K(X_1, 1) \mid X_s = x_s, C] \approx A_C \phi_K(x_s, s). \tag{2}$$

Thus, the expected terminal representation can be estimated directly from an intermediate state without explicitly sampling its terminal continuations. The architecture and objective used to learn ϕ_K, including representation normalization and anti-collapse regularization, are given in Appendix B.3.

For each predefined continuation C, we calibrate its corresponding operator A_C from frozen LKF rollouts. Given offline intermediate states $x_s^{(n)}$ and J independent terminal continuations $x_{1,C}^{(n,1)}, \ldots, x_{1,C}^{(n,J)}$, let $\bar{\phi}_{1,C}^{(n)} = J^{-1} \sum_{j=1}^{J} \phi_K(x_{1,C}^{(n,j)}, 1)$. For exposition, the finite-time calibration is expressed in the effective observable coordinates as

$$A_C = \arg\min_A \sum_n \left\| \bar{\phi}_{1,C}^{(n)} - A\phi_K(x_s^{(n)}, s) \right\|_2^2 + \lambda_A \|A\|_F^2. \tag{3}$$

The implementation fits the operator in raw-anchor coordinates before applying the fixed information transform, as detailed in Appendix B.3. Here, ϕ_K provides objective-independent observable coordinates, while A_C transports those coordinates from the start of continuation C directly to their expected terminal value. Both trainings are objective-independent and frozen for PEGASUS.

Task-specific objective readout. Koopman transport estimates the expected terminal representation, but prospective branch selection requires objective values. PEGASUS learns an online readout from the exact evaluations collected during optimization. We index PEGASUS optimization rounds by k. Let $\mathcal{D}_k = \{(x_1^{(n)}, \mathbf{F}(x_1^{(n)}))\}_{n=1}^{N_k}$ denote the cumulative archive of all unique terminal sequences exact-queried up to round k, where $N_k = |\mathcal{D}_k|$. Using $\mathcal{D}_k$, we fit the current multi-output ridge readout

$$W_k = \arg\min_W \sum_{n=1}^{N_k} \left\| \mathbf{F}(x_1^{(n)}) - W^\top \phi_K(x_1^{(n)}, 1) \right\|_2^2 + \lambda_W \|W\|_F^2. \tag{4}$$

The discovery-set construction, feature standardization, and selection of λ_W are specified in Appendix B.5; measured readout and PDO committee fitting costs are reported in Table S28.

For any terminal sequence x_1, the readout predicts $\widehat{\mathbf{F}}_k(x_1) = W_k^\top \phi_K(x_1, 1)$. Because the readout is linear, it can also be applied to the expected terminal representation of an online branch estimated by Koopman transport (Eq. 2), giving the **estimated terminal objective potential**

$$\widehat{\mathbf{V}}_{k,C_i}(x_{s,k}^{(i)}) = W_k^\top A_{C_i} \phi_K(x_{s,k}^{(i)}, s). \tag{5}$$

Eq. 5 approximates the conditional terminal objective potential in Eq. 1 directly from an intermediate state, without generating terminal continuations solely to score that branch. The factorization separates reusable, objective-independent dynamics (ϕ_K and A_C) from the online task readout (W_k), so new exact evaluations update the objective estimates without retraining the transport.

Proposition 3.1 (Future-observability decision bound)**.** *Fix a finite set of compared branch–continuation pairs $\mathcal{B}_k = \{(x_{s,k}^{(i)}, C_i)\}_{i=1}^{b}$. Suppose, uniformly for every $(x_{s,k}, C) \in \mathcal{B}_k$, $\|\mathbb{E}[\phi_K(X_1, 1) \mid X_s = x_{s,k}, C] - A_C \phi_K(x_{s,k}, s)\|_2 \le \epsilon_K$ and $\mathbf{F}(x_1) = W_k^\top \phi_K(x_1, 1) + \mathbf{r}(x_1)$*

with $\|\mathbb{E}[\mathbf{r}(X_1) \mid X_s = x_{s,k}, C]\|_\infty \le \epsilon_R$. *Then*

$$\left\|\mathbf{V}_C(x_{s,k}) - \widehat{\mathbf{V}}_{k,C}(x_{s,k})\right\|_\infty \le \delta_{\mathcal{B}_k} := \|W_k^\top\|_{\infty\leftarrow 2}\epsilon_K + \epsilon_R. \tag{6}$$

For the augmented Tchebycheff utility, $|U(\mathbf{y}) - U(\mathbf{y}')| \le L_U\|\mathbf{y} - \mathbf{y}'\|_\infty$ *with* $L_U = \|\bar{\mathbf{w}}\|_\infty + \rho$. *If PEGASUS selects the branch maximizing* $U(\widehat{\mathbf{V}}_{k,C})$ *over* $\mathcal{B}_k$, *its true conditional-mean utility regret relative to the best pair in* $\mathcal{B}_k$ *is at most* $2L_U\delta_{\mathcal{B}_k}$.

The bound separates finite-time transport error from task-readout error and connects both directly to branch selection. The proof is given in Appendix D.3.

Coarse search: protected future-observability allocation. PEGASUS allocates continuation computation using estimated terminal potential. It samples N_{br} intermediate branches with the frozen LKF and scores branch i by $u_i = U(\widehat{\mathbf{V}}_{k,C_i}(x_{s,k}^{(i)}))$. Of the total continuation budget B, a protected fraction is allocated approximately uniformly across retained branches, subject to a minimum per-branch allocation. The remaining adaptive fraction follows $\pi_i \propto [\mathrm{rank}(u_i) + 1]^{-1}$. This preserves stochastic support and directs additional continuation computation toward branches with higher predicted terminal potential (Appendix C.2).

LKF then generates concrete terminal sequences, which PEGASUS ranks using $\widehat{\mathbf{F}}_k$ before fine-stage multiscale editing. The highest-ranked contender is queried exactly, and a small committee of task-readout refits assesses whether remaining candidates can outperform the best exactly evaluated candidate. Each query updates the cumulative archive and task readout. Another contender is queried only if a committee member predicts an improvement beyond the decision tolerance; the procedure stops otherwise or when the coarse query cap is reached. Only the best exactly evaluated improvement can replace the incumbent (Appendix C.3).

3.2 Adaptive Multiscale Actions

Future observability identifies promising stochastic futures during coarse search but does not determine the appropriate intervention scale. In fine refinement, PEGASUS therefore considers local substitutions and coordinated multiresidue edits within a single decision, enabling both precise refinement and larger moves across sequence space. For incumbent $x = (x[1], \ldots, x[L])$, let $T_x(a)$ denote the sequence obtained by action a. An action H_r changes exactly r distinct positions, i.e. with edit distance $d_H(T_x(a), x) = r$. For configured radii $\mathcal{R} = \{r_1, \ldots, r_S\}$, the action set is

$$\mathcal{A}_{\mathcal{R}}(x) = \{a_0\} \cup \bigcup_{r\in\mathcal{R}} \mathcal{A}_r^{\mathrm{feas}}(x), \tag{7}$$

where a_0 is the no-op and $\mathcal{A}_r^{\mathrm{feas}}(x)$ contains realizable exact-r edits.

H1 substitutions are exhaustively enumerated when tractable, while larger Hamming shells are sampled as fixed-size slates using distinct positions and legal replacement tokens. PEGASUS materializes complete sequences, removes duplicates, computes frozen features without objective queries, and retains a bounded model-ranked contender set per scale. All scales enter the same decision without assuming additive responses across substitutions (Appendix C.4).

3.3 Pareto Decision Observability

Fine-stage multiscale actions produce many candidate sequences, but exact evaluation is necessary only while unresolved objective information can change the action choice. For $a \in \mathcal{A}_{\mathcal{R}}(x)$, define its exact vector response as $\mathbf{Y}_x(a) = \mathbf{F}(T_x(a)) - \mathbf{F}(x)$.

Decision observability. Let $\mathfrak{P}$ denote a specified class of vector-response maps $\Psi : \mathcal{A}_{\mathcal{R}}(x) \to \mathbb{R}^m$. After q exact action queries, let $\mathcal{O}_q$ be the queried action set and define the version set

$$\mathcal{V}_q = \{\Psi \in \mathfrak{P} : \Psi(a) = \mathbf{Y}_x(a) \text{ for all } a \in \mathcal{O}_q\}. \tag{8}$$

Thus, $\mathcal{V}_q$ contains the response maps in $\mathfrak{P}$ that remain consistent with all exact responses observed for the current decision.

For $\Psi \in \mathcal{V}_q$, define the predicted utility gain $g_\Psi(a) = U(\mathbf{F}(x) + \Psi(a)) - U(\mathbf{F}(x))$ and its ϵ-good action set

$$E_\epsilon(\Psi) = \left\{ a \in \mathcal{A}_\mathcal{R}(x) : g_\Psi(a) \geq \max_{b \in \mathcal{A}_\mathcal{R}(x)} g_\Psi(b) - \epsilon \right\}. \tag{9}$$

Definition 3.2 (Pareto Decision Observability)**.** *The current action decision is ϵ-observable after q exact observations if*

$$\bigcap_{\Psi \in \mathcal{V}_q} E_\epsilon(\Psi) \neq \varnothing. \tag{10}$$

Any action in this intersection is ϵ-good under every response map that remains consistent with the observed exact responses.

PDO therefore asks whether the available exact information is sufficient to identify a robust action decision, without requiring every candidate objective vector to be resolved.

Intrinsic decision dimension. Each complete action a is represented by its frozen LKF feature displacement $z_x(a) = \phi_L(T_x(a)) - \phi_L(x)$. To characterize exact-query complexity, consider the decision-linear response model $\mathbf{Y}_x(a) = B^\star z_x(a)$, which applies to complete sequence transformations without assuming edit-wise additivity. Let $\mathcal{K}_0$ denote the subspace of displacements with known exact responses, and $\mathcal{S}_0$ the span of actions that are potentially ϵ-good under the initial response version set, including the no-op. The intrinsic unresolved decision dimension is

$$d_{\mathrm{PDO}} = \dim(\mathcal{K}_0 + \mathcal{S}_0) - \dim(\mathcal{K}_0). \tag{11}$$

Let Q_{acq} denote the number of exact oracle queries used to resolve previously unresolved displacement directions before the decision becomes ϵ-observable, and let Q_{fine} denote the total exact-query cost of the fine-decision stage, including final verification.

Theorem 3.3 (Intrinsic-dimension PDO query bound)**.** *Under realizable exact-linear responses, consistent version-set updates, and acquisition of an unresolved displacement direction before ϵ-observability,*

$$Q_{\mathrm{acq}} \leq d_{\mathrm{PDO}}. \tag{12}$$

Once the decision is ϵ-observable, every action in the common ϵ-good set has true regret at most ϵ. Including one final exact verification gives $Q_{\mathrm{fine}} \leq d_{\mathrm{PDO}} + 1$.

Thus, exact-query complexity depends on unresolved decision directions rather than directly on candidate count or the number of Hamming scales. Adding actions whose decision-relevant displacements lie in already resolved directions does not increase the bound. The formal assumptions, proof, and worst-case tightness construction are provided in Appendix D.4.

Practical empirical PDO. The exact-linear query bound motivates our acquisition strategy but does not directly guarantee the behavior of the empirical implementation. PEGASUS approximates the response version set using a small family of plausible objective-response maps constructed from frozen features and available exact evaluations (Appendix C.5). Each map predicts the objective changes of complete multiscale actions, allowing PEGASUS to assess whether the current action decision is sufficiently resolved.

To avoid premature termination from agreement among imperfect models, PEGASUS first exact-queries one protected exploit-first action from the full multiscale family. It then compares the action decisions across the empirical response maps and queries additional contenders only while unresolved responses can change the current decision. After each exact evaluation, the response models are updated and the decision is reassessed. Acquisition stops when no decision-relevant challenger remains or the fine query cap is reached (Appendix C.6). Only exactly evaluated actions that improve the incumbent utility beyond ϵ_{acc} are accepted. Algorithm 2 in Appendix C details the empirical PDO acquisition routine.

Flow restart. After coarse verification and PDO-controlled refinement, the accepted terminal candidate $x_{1,k}^{(i^\star)}$ becomes the initial state of the next optimization round, $x_{0,k+1} = x_{1,k}^{(i^\star)}$. PEGASUS

Table 1: **LKF achieves strong unconditional peptide generation quality with substantially lower sampling cost than existing discrete generative models.** Best values per length are bold.

Length	Method	External PLL ↑	Embedding MMD^2 ↓	Pairwise Hamming ↑	Native NFE ↓	Exact-length rate ↑	Effective NFE ↓
$L = 12$	LKF	**-2.963**	**0.00139**	0.9387	**8**	**100%**	**8**
	CFM	-2.998	0.00784	0.9385	16	3.02%	532
	PepDFM	-3.003	0.00219	0.9393	100	**100%**	100
	EvoDiff	-2.973	0.02718	0.9338	12	**100%**	12
	DPLM	-2.996	0.00722	**0.9420**	500	99.06%	982.9
$L = 40$	LKF	-2.842	0.02099	0.9210	**8**	**100%**	**8**
	CFM	-2.849	0.02568	0.9225	16	2.36%	680
	PepDFM	-2.860	0.02502	0.9233	100	**100%**	100
	EvoDiff	**-2.831**	**0.01415**	0.9219	40	**100%**	40
	DPLM	-2.918	0.01916	**0.9369**	500	98.98%	544.9

Table 2: **PEGASUS outperforms existing discrete MOO baselines in quality-efficiency across peptide lengths.** Best MOO values per length are bold. LKF Best-of-N is a non-MOO sampling control and its best runtime is underlined.

Length	Method	Terminal HV ↑	Mean utility ↑	Top-10 utility ↑	Runtime (min) ↓	Unique queries ↓
$L = 12$	PEGASUS	0.6867±0.0453	$\mathbf{-0.0449 \pm 0.0034}$	$\mathbf{-0.0373 \pm 0.0042}$	**7.95±0.28**	1303.8±67.6
	MOG-DFM	**0.6894±0.0561**	-0.0795 ± 0.0026	-0.0584 ± 0.0054	54.64±0.08	109949.4±1019.2
	AReUReDi	0.6416±0.0639	-0.0784 ± 0.0028	-0.0595 ± 0.0054	55.80±1.20	114842.0±1076.1
	LaMBO	0.4644±0.0683	-0.1196 ± 0.0029	-0.0752 ± 0.0055	271.8±56.4	**1278.8±1.1**
	PepTune+DPLM	0.2335±0.0216	-0.1459 ± 0.0021	-0.1108 ± 0.0044	23.65±1.43	1297.1±5.3
	LKF Best-of-N	0.3746±0.0207	-0.1316 ± 0.0014	-0.0889 ± 0.0038	<u>1.14±0.08</u>	1303.8±67.6
$L = 40$	PEGASUS	**0.5286±0.0528**	$\mathbf{-0.0631 \pm 0.0031}$	$\mathbf{-0.0524 \pm 0.0044}$	**15.08±0.55**	1568.2±118.7
	MOG-DFM	0.3629±0.0390	-0.1445 ± 0.0023	-0.0848 ± 0.0033	78.27±1.19	155016.7±820.3
	AReUReDi	0.3409±0.0402	-0.1485 ± 0.0031	-0.0862 ± 0.0035	78.94±0.51	161805.9±1157.0
	LaMBO	0.0479±0.0180	-0.1656 ± 0.0012	-0.1377 ± 0.0067	331.8±69.0	**1279.8±0.4**
	PepTune+DPLM	0.1150±0.0876	-0.1669 ± 0.0014	-0.1365 ± 0.0094	31.91±2.90	1585.5±49.5
	LKF Best-of-N	0.0418±0.0113	-0.1660 ± 0.0009	-0.1397 ± 0.0025	<u>1.39±0.06</u>	1568.2±118.7

then restarts the frozen LKF from this sequence, generating new stochastic branches for future-observability-guided exploration rather than continuing the previous flow trajectory. If neither stage yields an accepted improvement, the incumbent remains unchanged. This process repeats across independent lineages, with exact evaluations shared through the cumulative archive (Figure 1).

4 Results

We evaluated PEGASUS on $L = 12$ and $L = 40$ peptides under six predicted objectives (non-hemolysis, non-fouling, solubility, permeability, half-life, and target-conditioned affinity), and on 170-bp DNA under three condition-derived activity objectives. The peptide targets include three with known binders (3IDJ, 5AZ8, 7JVS) and five without (AMHR2, OX1R, DUSP12, EWS1::FLI1, MYC) (Chen et al., 2025b), testing established binding contexts and prospective discovery. Target selection, normalization, and evaluation protocols are detailed in Appendices E and F.

4.1 Peptide LKF Supports Accurate Generation and Koopman Transport

The frozen peptide LKF produced competitive unconditional generation quality with eight model evaluations (Table 1). On held-out branched trajectories, finite-time transport achieved cosine similarity of 0.989 from $s = 0.5$ to terminal time, compared with 0.827 after shuffling intermediate-to-terminal correspondences (Table S2). The calibration protocol and additional diagnostics are detailed in Appendix B.4.

4.2 PEGASUS Enables Query-Efficient Six-Objective Peptide Optimization

We compared PEGASUS with discrete generative optimizers, classical multi-objective search, LaMBO, and query-matched LKF Best-of-N across eight targets (Table 2; complete objective-level

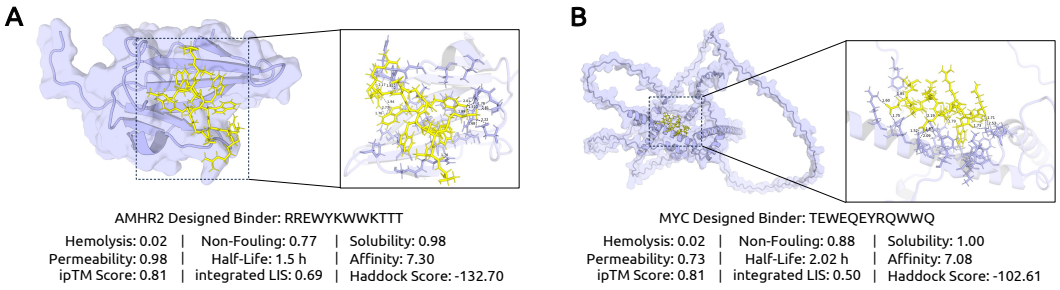


Figure 2: **HADDOCK-docked peptide-protein complexes generated for (A) structured protein AMHR2 and (B) intrinsically disordered protein MYC.** Proteins are shown in light blue and PEGASUS-designed peptides in yellow, with interface contacts enlarged in the insets.

results in Table S3). At $L = 12$, PEGASUS attained mean terminal-100 hypervolume 0.6867 versus 0.6894 for MOG-DFM, using $1{,}303.8$ versus $109{,}949.4$ unique full-vector queries. PEGASUS therefore matched the strongest Pareto-front quality with a small fraction of the oracle cost. At $L = 40$, PEGASUS attained 0.5286 versus 0.3629, using $1{,}568.2$ versus $155{,}016.7$ queries. PEGASUS also achieved the shortest runtime among the multi-objective methods at both lengths and the highest target-wise mean hypervolume on all eight targets at $L = 40$, demonstrating that its quality-efficiency advantage strengthened in the larger search space (Table S4). It surpassed LaMBO, PepTune+DPLM, and query-matched LKF Best-of-N, supporting gains beyond additional generator sampling. PEGASUS further achieved the highest mean and top-10 terminal utility at both lengths, confirming that its Pareto sets contained consistently high-quality leading candidates (Table S5).

The returned PEGASUS terminal libraries satisfy normalized pairwise Hamming separation $\delta = 0.20$ at both lengths, and substantial subsets remain separated at stricter thresholds (Table S8). Independent property predictors further corroborated favorable scores for selected PEGASUS designs, particularly for non-fouling and penetrance (Appendix G.1). Together, these analyses demonstrate that PEGASUS maintains sequence diversity while producing high-scoring candidates whose predicted properties are supported by orthogonal computational evaluations.

We assessed selected PEGASUS designs using five independent AlphaFold3 complex predictions and complementary HADDOCK docking. The predicted complexes achieved favorable ipTM and iLIS scores (Figs. 2 and S1), with representative AMHR2 and MYC designs attaining ipTM scores of 0.81 and iLIS scores of 0.69 and 0.50, respectively. Together with favorable docking interfaces and target contacts, these results provide orthogonal structural evidence supporting the potential of PEGASUS designs to engage both structured and intrinsically disordered proteins.

Several high-scoring $L = 12$ designs exhibited recurrent tryptophan-rich patterns. To investigate their contribution, we replaced all tryptophans in one selected binder per target with phenylalanine, tyrosine, or leucine, generating 24 counterfactual variants. All 24 variants exhibited lower six-objective utility under the benchmark oracle (Table S6). These consistent reductions across targets and substitution types support PEGASUS's ability to discover and exploit recurrent sequence patterns that contribute to favorable multi-objective performance.

4.3 Future Observability, Multiscale Actions, and PDO Drive PEGASUS Gains

Across eight targets and three seeds, removing future observability reduced mean terminal six-objective hypervolume at both lengths, with a larger effect at $L = 40$, supporting the contribution of prospective information to terminal Pareto quality (Table S9). Matched-observer controls and additional prediction analyses are reported in Appendix G.2. Compared with H1-only editing, the full multiscale ladder improved generated utility and accepted fine-stage gain at both lengths, with larger improvements at $L = 40$. H8 actions contributed a greater share of accepted gain and were selected more frequently for longer sequences (Table S10). Notably, adding H8 to the H1/H2/H4 ladder reduced realized exact-query expenditure at $L = 40$ while preserving generated utility (Table S16). Further analyses of action-scale contributions and query scaling are provided in Appendix G.3.

Table 3: **PEGASUS generalizes across different frozen unconditional flow generators.** Results use dynamics-matched finite-time observers, the same three objectives, and seed-matched terminal-100 libraries ($n = 3$ seeds). Native NFE excludes paid objective queries.

Frozen generator	Native full-path NFE	Unique objective queries	Clipped HV-AUC@2000↑	Final-library raw composite↑
LKF ($M = 1$)	16	$1{,}336 \pm 59$	0.96688 ± 0.0211	3.777 ± 0.133
CatFM	16	$1{,}627 \pm 145$	0.97095 ± 0.01178	2.442 ± 0.331
DirFM	99	$1{,}510 \pm 86$	0.96762 ± 0.01571	3.187 ± 0.030
FisherFM	100	$1{,}705 \pm 172$	0.96134 ± 0.02117	2.817 ± 0.039

Under matched fine-label allowances, PDO increased accepted fine-stage gain, while generated utility and terminal hypervolume showed positive point estimates whose target-aware confidence intervals included zero (Table S19). Without PDO's acquisition rule, unconstrained evaluation consumed substantially more queries and reached a different quality–cost operating point (Table S11). Frozen-state replay demonstrated sparse acquisition with low observed decision regret, while synthetic exact-linear experiments satisfied the intrinsic-rank query bound in all trials (Appendices G.4 and G.5). Additional online readout and exact-verification diagnostics are reported in Table S27.

4.4 PEGASUS Generalizes to Multi-Objective DNA Sequence Optimization

To evaluate generalization beyond peptides, we applied PEGASUS to 170-bp DNA design with three plantGREP-derived objectives: light-minus-dark response, warm-minus-cold response, and reduced maize-condition activity. We trained a DNA LKF on Arabidopsis open-chromatin sequences and used a separate frozen objective model for optimization. The implementation retains the same future-observability, multiscale-action, and PDO decision structure, with modality-specific generative dynamics and a calibrated observer (Appendix F).

The frozen DNA LKF achieved high clipped archive HV-AUC with the fewest unique objective queries and returned a 100-sequence library with high raw composite quality; its unconditional samples also achieved the lowest mean k-mer JS distance among the four generators (Tables 3, S32). To assess dependence on the underlying generator, we replaced LKF with CatFM, DirFM, and FisherFM, recalibrating the objective-independent finite-time observer for each frozen backbone. All four implementations achieved high clipped HV-AUC, with held-out transport diagnostics supporting the effectiveness of their dynamics-matched observers. These findings establish the cross-modality generalization of PEGASUS, with strong unconditional generation and effective finite-time future transport across distinct frozen DNA backbones.

The DNA ablations further support the contribution of multiscale actions: H1-only optimization attained lower clipped and upper-unclipped HV-AUC than the full configuration. Broadening the radius set further compounds this effect, with configurations with four or more scales outperforming the three-scale ladder on both clipped and upper-unclipped HV-AUC (Table S31). Together, these results demonstrate that the benefits of multiscale actions extend to longer DNA sequences, while PEGASUS preserves its core optimization mechanisms across distinct biological modalities and frozen generative dynamics.

5 Discussion

In this work, we introduce PEGASUS to combine future observability, adaptive multiscale actions, and decision-aware acquisition under frozen discrete generative dynamics. PEGASUS achieved strong Pareto quality with substantially fewer objective evaluations, with ablations supporting each mechanism and generalization beyond peptide modality. These findings support multi-objective search with shallow generators by separating generation depth, optimization horizon, and evaluation cost. To characterize query efficiency, we bounded PDO acquisition under exact decision-linear responses and observed efficient acquisition with the approximate model. These computational gains require experimental follow-up to establish binding and developability, despite additional assessment with independent predictors and structural analyses. Nonetheless, PEGASUS addresses a central bottleneck in therapeutic design across competing functional and developability objectives, enabling broader exploration of biological sequences under limited evaluation budgets.

AI USE STATEMENT

We acknowledge the use of AI to assist in polishing and editing parts of this manuscript. AI was used to refine phrasing, improve clarity, and ensure consistency of style across sections. All technical content, experiments, analyses, and conclusions were developed by the authors, with AI support limited to language refinement and editorial improvements.

REPRODUCIBILITY

We provide an anonymous code repository at `https://anonymous.4open.science/r/PEGASUS-5B68/`. The Appendix details the PEGASUS algorithms, training and calibration procedures, objective definitions, baseline configurations, query accounting, and evaluation protocols. Complete theoretical assumptions and proofs are also provided in the Appendix to support reproduction of the methodological and experimental results.

ETHICS

This study evaluates computational methods for peptide and DNA sequence optimization. All reported biological properties and structural interactions are computational predictions rather than experimentally established outcomes. Designed candidates require independent experimental assessment of binding, developability, and safety before any therapeutic application.

REFERENCES

Osama Abdin, Satra Nim, Han Wen, and Philip M Kim. Pepnn: a deep attention model for the identification of peptide binding sites. *Communications biology*, 5(1):503, 2022.

Josh Abramson, Jonas Adler, Jack Dunger, Richard Evans, Tim Green, Alexander Pritzel, Olaf Ronneberger, Lindsay Willmore, Andrew J Ballard, Joshua Bambrick, et al. Accurate structure prediction of biomolecular interactions with alphafold 3. *Nature*, 630(8016):493–500, 2024.

Mansoor Ahmed, Yue-Tsz Fan, Hemanth Venkateswara, and Murray Patterson. Latent-kernel discrete flow maps for few-step generation. *arXiv preprint arXiv:2607.27529*, 2026.

Sarah Alamdari, Nitya Thakkar, Rianne Van Den Berg, Neil Tenenholtz, Robert Strome, Alan M Moses, Alex X Lu, Nicolò Fusi, Ava P Amini, and Kevin K Yang. Protein generation with evolutionary diffusion: sequence is all you need. *BioRxiv*, pp. 2023–09, 2023.

Gonzalo Benegas, Sanjit Singh Batra, and Yun S. Song. DNA language models are powerful predictors of genome-wide variant effects. *Proceedings of the National Academy of Sciences*, 120(44): e2311219120, 2023. doi: 10.1073/pnas.2311219120.

Rayane Monique Bernardes-Loch, Gustavo de Oliveira Almeida, Igor Teixeira Brasiliano, Wagner Meira Jr, Douglas EV Pires, Maria Cristina Baracat-Pereira, and Sabrina de Azevedo Silveira. Perseucpp: a machine learning strategy to predict cell-penetrating peptides and their uptake efficiency. *Bioinformatics Advances*, 5(1):vbaf213, 2025.

Nicola Beume, Boris Naujoks, and Michael Emmerich. Sms-emoa: Multiobjective selection based on dominated hypervolume. *European journal of operational research*, 181(3):1653–1669, 2007.

Nicholas M Boffi, Michael S Albergo, and Eric Vanden-Eijnden. Flow map matching with stochastic interpolants: A mathematical framework for consistency models. *arXiv preprint arXiv:2406.07507*, 2024.

Anton Bushuiev, Roman Bushuiev, Petr Kouba, Anatolii Filkin, Marketa Gabrielova, Michal Gabriel, Jiri Sedlar, Tomas Pluskal, Jiri Damborsky, Stanislav Mazurenko, et al. Learning to design protein-protein interactions with enhanced generalization. In *International Conference on Learning Representations*, volume 2024, pp. 21010–21035, 2024.

Tianlai Chen, Lauren Hong, Vivian Yudistyra, Sophia Vincoff, and Pranam Chatterjee. Generative design of therapeutics that bind and modulate protein states. *Current Opinion in Biomedical Engineering*, 28:100496, 2023.

Tong Chen, Zachary Quinn, Yinuo Zhang, and Pranam Chatterjee. moppit-v3: Motif-specific peptides generated via multi-objective-guided discrete flow matching. In *NeurIPS 2025 AI for Science Workshop*, 2025a.

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

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

CA Coello Coello and Maximino Salazar Lechuga. Mopso: A proposal for multiple objective particle swarm optimization. In *Proceedings of the 2002 Congress on Evolutionary Computation. CEC'02 (Cat. No. 02TH8600)*, volume 2, pp. 1051–1056. IEEE, 2002.

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.

Abraham Duarte, Juan J Pantrigo, Eduardo G Pardo, and Nenad Mladenovic. Multi-objective variable neighborhood search: an application to combinatorial optimization problems. *Journal of Global Optimization*, 63(3):515–536, 2015.

Sasha B Ebrahimi and Devleena Samanta. Engineering protein-based therapeutics through structural and chemical design. *Nature communications*, 14(1):2411, 2023.

Zhun Fan. An evolutionary many-objective optimization algorithm using reference-point-based nondominated sorting approach, part i: solving problems with box constraints. *IEEE transactions on evolutionary computation*, 2014.

Veniamin Fishman, Yuri Kuratov, Aleksei Shmelev, Maxim Petrov, Dmitry Penzar, Denis Shepelin, Nikolay Chekanov, Olga Kardymon, and Mikhail Burtsev. GENA-LM: A family of open-source foundational dna language models for long sequences. *Nucleic Acids Research*, 53(2):gkae1310, 2025. doi: 10.1093/nar/gkae1310.

Nate Gruver, Samuel Stanton, Nathan Frey, Tim GJ Rudner, Isidro Hotzel, Julien Lafrance-Vanasse, Arvind Rajpal, Kyunghyun Cho, and Andrew G Wilson. Protein design with guided discrete diffusion. *Advances in neural information processing systems*, 36:12489–12517, 2023.

Chakradhar Guntuboina, Adrita Das, Parisa Mollaei, Seongwon Kim, and Amir Barati Farimani. Peptidebert: A language model based on transformers for peptide property prediction. *The Journal of Physical Chemistry Letters*, 14(46):10427–10434, 2023.

Daniel Hernández-Lobato, Jose Hernandez-Lobato, Amar Shah, and Ryan Adams. Predictive entropy search for multi-objective bayesian optimization. In *International conference on machine learning*, pp. 1492–1501. PMLR, 2016.

Lauren Hong, Sophia Vincoff, and Pranam Chatterjee. Ai-designed peptides as tools for biochemistry. *Biochemistry*, 65(8):1071–1083, 2026.

Rodrigo V Honorato, Mikael E Trellet, Brian Jiménez-García, Jörg J Schaarschmidt, Marco Giulini, Victor Reys, Panagiotis I Koukos, João PGLM Rodrigues, Ezgi Karaca, Gydo CP van Zundert, et al. The haddock2. 4 web server for integrative modeling of biomolecular complexes. *Nature protocols*, 19(11):3219–3241, 2024.

Sedjro Salomon Hotegni and Sebastian Peitz. Spread: Sampling-based pareto front refinement via efficient adaptive diffusion. In *International Conference on Learning Representations*, volume 2026, pp. 115284–115315, 2026.

Jerry Y Huang, Justin Lin, Sheel Shah, Kartik Nair, and Nicholas M Boffi. How to guide your flow: Few-step alignment via flow map reward guidance. *arXiv preprint arXiv:2604.27147*, 2026.

Himanshu Jain and Kalyanmoy Deb. An evolutionary many-objective optimization algorithm using reference-point based nondominated sorting approach, part ii: Handling constraints and extending to an adaptive approach. *IEEE Transactions on evolutionary computation*, 18(4):602–622, 2013.

Tobias Jores, Nicholas A. Mueth, Sayeh Gorjifard, Sebastian Triesch, Dominic Schirmer, Jackson Tonnies, Kerry L. Bubb, Josh T. Cuperus, Stanley Fields, and Christine Queitsch. Massively parallel characterization and deep learning of enhancers in plant genomes. *bioRxiv*, 2026. doi: 10.64898/2026.04.26.720828.

Ah-Ram Kim, Aram Comjean, Austin Veal, Jonathan Rodiger, Myeonghoon Han, Yanhui Hu, and Norbert Perrimon. Flypredictome: A structural atlas of predicted protein-protein interactions in drosophila. *bioRxiv*, 2026.

Cyrille Kone, Emilie Kaufmann, and Laura Richert. Bandit pareto set identification in a multi-output linear model. *arXiv preprint arXiv:2507.04255*, 2025.

Zitai Kong, Yifan Dong, Yixuan Wu, Zhaokang Liang, Jian Wu, and Hongxia Xu. Mp2d: Constrained monte carlo tree-guided diffusion for multi-objective protein sequence design. *arXiv preprint arXiv:2605.05829*, 2026.

Yun-Huai Kuo, Chloe N Brown, Ece Akin, Yunxuan Xie, and Peter M Tessier. Multi-objective antibody design and optimization using machine learning. *Nature Reviews Bioengineering*, pp. 1–15, 2026.

Xiner Li, Yulai Zhao, Chenyu Wang, Gabriele Scalia, Gokcen Eraslan, Surag Nair, Tommaso Biancalani, Shuiwang Ji, Aviv Regev, Sergey Levine, et al. Derivative-free guidance in continuous and discrete diffusion models with soft value-based decoding. *arXiv preprint arXiv:2408.08252*, 2024.

Xiner Li, Masatoshi Uehara, Xingyu Su, Gabriele Scalia, Tommaso Biancalani, Aviv Regev, Sergey Levine, and Shuiwang Ji. Dynamic search for inference-time alignment in diffusion models. *arXiv preprint arXiv:2503.02039*, 2025.

Xi Lin, Zhiyuan Yang, Xiaoyuan Zhang, and Qingfu Zhang. Pareto set learning for expensive multi-objective optimization. *Advances in neural information processing systems*, 35:19231–19247, 2022.

Jiaqi Luo, Kerr Ding, and Yunan Luo. Pareto-optimal sampling for multi-objective protein sequence design. *Iscience*, 28(3), 2025.

Bethany Lusch, J Nathan Kutz, and Steven L Brunton. Deep learning for universal linear embeddings of nonlinear dynamics. *Nature communications*, 9(1):4950, 2018.

Deepika Mathur, Sandeep Singh, Ayesha Mehta, Piyush Agrawal, and Gajendra PS Raghava. In silico approaches for predicting the half-life of natural and modified peptides in blood. *PloS one*, 13(6):e0196829, 2018.

Kaisa Miettinen et al. *Nonlinear multiobjective optimization*, volume 12. Kluwer Academic Publishers Boston, 1999.

Samuel E Otto and Clarence W Rowley. Koopman operators for estimation and control of dynamical systems. *Annual Review of Control, Robotics, and Autonomous Systems*, 4(1):59–87, 2021.

Peter Potaptchik, Jason Yim, Adhi Saravanan, Peter Holderrieth, Eric Vanden-Eijnden, and Michael S Albergo. Discrete flow maps. *arXiv preprint arXiv:2604.09784*, 2026.

Anand Singh Rathore, Nishant Kumar, Shubham Choudhury, Naman Kumar Mehta, and Gajendra PS Raghava. Prediction of hemolytic peptides and their hemolytic concentration. *Communications Biology*, 8(1):176, 2025.

Daan Roos, Oscar Davis, Floor Eijkelboom, Michael Bronstein, Max Welling, Ismail Ilkan Ceylan, Luca Ambrogioni, and Jan-Willem van de Meent. Categorical flow maps. *arXiv preprint arXiv:2602.12233*, 2026.

Jatan Shrestha, Santeri Heiskanen, Kari Hepola, Severi Rissanen, Pekka Jääskeläinen, and Joni Pajarinen. Pareto-conditioned diffusion models for offline multi-objective optimization. *arXiv preprint arXiv:2602.00737*, 2026.

Pietro Sormanni, Francesco A Aprile, and Michele Vendruscolo. The camsol method of rational design of protein mutants with enhanced solubility. *Journal of molecular biology*, 427(2):478–490, 2015.

Samuel Stanton, Wesley Maddox, Nate Gruver, Phillip Maffettone, Emily Delaney, Peyton Greenside, and Andrew Gordon Wilson. Accelerating bayesian optimization for biological sequence design with denoising autoencoders. In *International conference on machine learning*, pp. 20459–20478. PMLR, 2022.

Hannes Stärk, Bowen Jing, Chenyu Wang, Gabriele Corso, Bonnie Berger, Regina Barzilay, and Tommi S. Jaakkola. Dirichlet flow matching with applications to DNA sequence design. In *Proceedings of the 41st International Conference on Machine Learning*, volume 235 of *Proceedings of Machine Learning Research*, pp. 46495–46513. PMLR, 2024. URL `https://openreview.net/forum?id=syXFAVqx85`.

Sophia Tang, Yinuo Zhang, and Pranam Chatterjee. Peptune: De novo generation of therapeutic peptides with multi-objective-guided discrete diffusion. *ArXiv*, pp. arXiv–2412, 2025a.

Sophia Tang, Yuchen Zhu, Molei Tao, and Pranam Chatterjee. Tr2-d2: Tree search guided trajectory-aware fine-tuning for discrete diffusion. *ArXiv*, pp. arXiv–2509, 2025b.

Erkan Turan, Aristotelis Siozopoulos, and Maks Ovsjanikov. Unfolding generative flows with koopman operators: Fast and interpretable sampling. *arXiv e-prints*, pp. arXiv–2506, 2025.

Masatoshi Uehara, Xingyu Su, Yulai Zhao, Xiner Li, Aviv Regev, Shuiwang Ji, Sergey Levine, and Tommaso Biancalani. Reward-guided iterative refinement in diffusion models at test-time with applications to protein and dna design. *arXiv preprint arXiv:2502.14944*, 2025.

Sophia Vincoff, Howard Liu, Corinne M Linardic, and Pranam Chatterjee. Soapia: Specificity-guided generation of off-target-avoiding protein interactions with high target affinity. In *The 2026 Workshop on Generative and Agentic AI for Biology*, 2026.

Zhengyan Wan, Yidong Ouyang, Liyan Xie, Fang Fang, Hongyuan Zha, and Guang Cheng. Discrete guidance matching: Exact guidance for discrete flow matching. In *International Conference on Learning Representations*, volume 2026, pp. 148461–148493, 2026.

Xinyou Wang, Zaixiang Zheng, Fei Ye, Dongyu Xue, Shujian Huang, and Quanquan Gu. Diffusion language models are versatile protein learners. *arXiv preprint arXiv:2402.18567*, 2024.

Zhenji Wang, Minghao Liu, Fuming Lai, Qiangqiang Fu, Liang Xie, Yaping Fang, Qiangwei Zhou, and Guoliang Li. AraENCODE: A comprehensive epigenomic database of Arabidopsis thaliana. *Molecular Plant*, 16(7):1113–1116, 2023. doi: 10.1016/j.molp.2023.06.005.

Ye Yuan, Can Chen, Christopher Pal, and Xue Liu. Paretoflow: Guided flows in multi-objective optimization. In *International Conference on Learning Representations*, volume 2025, pp. 72594–72620, 2025.

Chengxin Zhang, Xi Zhang, Lydia Freddolino, and Yang Zhang. Biolip2: an updated structure database for biologically relevant ligand–protein interactions. *Nucleic acids research*, 52(D1): D404–D412, 2024.

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

Eckart Zitzler, Marco Laumanns, and Lothar Thiele. Spea2: Improving the strength pareto evolutionary algorithm. *TIK report*, 103, 2001.

A Background and Notation

A.1 Multi-Objective Discrete Sequence Optimization

Let $\mathcal{V}$ denote a finite vocabulary and $\mathcal{X} = \mathcal{V}^L$ the space of discrete sequences of length L. Each sequence $x = (x[1], \ldots, x[L]) \in \mathcal{X}$ has a maximize-oriented objective vector $\mathbf{F}(x) = (F_1(x), \ldots, F_m(x)) \in \mathbb{R}^m$. A sequence x Pareto-dominates x', written $x \succ x'$, if $F_o(x) \geq F_o(x')$ for every objective o and $F_o(x) > F_o(x')$ for at least one objective o. The Pareto set contains sequences that are not dominated by any other feasible sequence.

For preference-conditioned optimization, let $\mathbf{w} \in \mathbb{R}^m_{>0}$ and $\bar{\mathbf{w}} = \mathbf{w}/(\mathbf{1}^\top \mathbf{w})$. Given a reference vector $\mathbf{z}^\star$, PEGASUS uses the augmented Tchebycheff utility (Miettinen et al., 1999)

$$U_{\mathbf{w},\rho}(\mathbf{y}) = -\max_{o \in [m]} \bar{w}_o(z^\star_o - y_o) + \rho\, \bar{\mathbf{w}}^\top \mathbf{y}, \tag{13}$$

where $\rho \geq 0$ is the augmentation coefficient. We write U when $\mathbf{w}$ and ρ are fixed. One exact oracle query evaluates the complete vector $\mathbf{F}(x)$ for a previously unseen sequence; repeated occurrences reuse the cached vector and do not increase the unique-query count.

A.2 Finite-Time Discrete Flow Maps and LKF

A discrete generative process defines a stochastic trajectory $X_t \in \mathcal{X}$ for $t \in [0, 1]$. A finite-time flow map samples X_t from an intermediate state X_s through a learned transition $p_\theta(x_t \mid x_s, s, t)$ without explicitly simulating every intermediate time. We use $C = (\tau_0 = s < \tau_1 < \cdots < \tau_R = 1)$ to denote a prescribed continuation from time s to the terminal time. The terminal distribution under C is induced by the corresponding finite-time transitions. Our main peptide implementation uses a frozen Latent-Kernel discrete flow map (LKF) generator (Ahmed et al., 2026), and $\phi_L(x) \in \mathbb{R}^{d_L}$ denotes the frozen LKF feature used to describe explicit sequence transformations. All LKF parameters remain fixed during downstream PEGASUS optimization.

A.3 Koopman Operators for Finite-Time Dynamics

For an observable $g : \mathcal{X} \to \mathbb{R}^d$, the finite-time stochastic Koopman operator is $(\mathcal{K}_{s\to t} g)(x) = \mathbb{E}[g(X_t) \mid X_s = x]$ (Otto & Rowley, 2021). PEGASUS uses an objective-independent observable representation $\phi_K : \mathcal{X} \to \mathbb{R}^{d_K}$ and finite-time matrices A_C satisfying $\mathbb{E}[\phi_K(X_1) \mid X_s = x, C] \approx A_C \phi_K(x)$. The representations ϕ_L and ϕ_K serve different roles: ϕ_L parameterizes realized action displacements for fine-stage PDO, while ϕ_K supports finite-time prospective transport.

B Offline Model Learning and Task Initialization

B.1 Peptide Dataset and Preprocessing

The peptide generator uses the same corpus and split protocol as the wild-type peptide experiments in MOG-DFM and AReUReDi (Chen et al., 2025c;b). The corpus contains peptide sequences from PepNN (Abdin et al., 2022) and BioLiP2 (Zhang et al., 2024), together with PPIRef (Bushuiev et al., 2024) sequences of length 6 to 49 amino acids. The resulting dataset is divided into training, validation, and test partitions using an 80/10/10 split. PEGASUS operates on canonical amino-acid sequences; the $L = 12$ and $L = 40$ benchmark conditions use fixed-length generation from the same frozen peptide generator.

B.2 Peptide LKF Training

The final peptide backbone is a teacher-free Uniform-LKF trained from scratch on the peptide corpus. It does not initialize from PepDFM and does not use PepDFM-generated trajectories. For a corrupted intermediate sequence x_s and clean endpoint x_1, the LKF represents the conditional clean-sequence distribution as a mixture of M factorized categorical components,

$$p_\theta(x_1 \mid x_s, s) = \sum_{c=1}^{M} \pi_{\theta,c}(x_s, s) \prod_{\ell=1}^{L} p_{\theta,c,\ell}(x_1[\ell] \mid x_s, s), \tag{14}$$

and minimizes the negative log-likelihood of the clean endpoint under this mixture. Here M denotes the number of generator mixture components, not the number of coarse branches. The production peptide model uses $M = 8$ mixture components; the $M = 1$ model is retained as a mixture-component ablation reported in Appendix G.2. The frozen checkpoint is selected using validation performance before any downstream objective labels are introduced.

B.3 Objective-Independent Koopman Coordinates

PEGASUS constructs objective-independent Koopman coordinates from frozen LKF features and calibrates finite-time transport using generator trajectories. No peptide-property or target-affinity labels are used during this stage.

Frozen Gaussian information anchor. Let $h_L(x,t) \in \mathbb{R}^{d_h}$ denote the residue-pooled hidden feature from the frozen Uniform-LKF shared blocks at state x and time t, excluding CLS and EOS positions and including the configured time embedding when enabled. We define a fixed $d_K = 64$ dimensional information anchor

$$r(x,t) = P^\top h_L(x,t) \in \mathbb{R}^{d_K}, \qquad d_K = 64, \tag{15}$$

where $P \in \mathbb{R}^{d_h \times d_K}$ is sampled once from a Gaussian distribution with entries scaled by $1/\sqrt{d_K}$, using the fixed anchor seed 60042. The projection is stored as a non-trainable buffer and remains frozen throughout transport calibration and downstream optimization.

Information whitening. We calibrate a fixed information transform from an objective-independent bank of anchors obtained from frozen-LKF trajectories. Given calibration anchors $\{r_i\}_{i=1}^N$, let $\bar{r} = N^{-1}\sum_{i=1}^N r_i$ and compute the covariance C_r of the centered vectors $r_i - \bar{r}$. Eigen-decomposing $C_r = Q\,\mathrm{diag}(\lambda_j)Q^\top$, we retain eigenvectors whose eigenvalues exceed the relative eigenvalue floor and construct

$$T = \mathrm{diag}\Big((\lambda_j + \varepsilon_{\mathrm{wh}})^{-1/2}\Big)_{j \in \mathcal{I}_{\mathrm{keep}}} Q_{\mathrm{keep}}^\top, \tag{16}$$

where $\mathcal{I}_{\mathrm{keep}}$ indexes the retained modes, Q_{keep} contains their eigenvectors, and $\varepsilon_{\mathrm{wh}} > 0$ regularizes the inverse eigenvalue scaling. The resulting Koopman observable is

$$\phi_K(x,t) = Tr(x,t). \tag{17}$$

The implementation stores feature vectors in row form and applies the equivalent transformation $r(x,t)T^\top$. Centering is used to estimate the calibration covariance; runtime normalization applies the fixed transform T without subtracting the calibration mean or imposing per-sample unit norm. We write $\phi_K(x)$ when the associated time is clear from context.

Label-free finite-time transport. After fixing the anchor and information transform, we calibrate finite-time operators using trajectories from the frozen generator. Let $A_C^{(r)}$ denote the operator fitted in raw anchor coordinates for continuation C. The implemented terminal-observable prediction is

$$\widehat{\phi}_{K,C}(x_s,s) = TA_C^{(r)} r(x_s,s), \tag{18}$$

which approximates the conditional terminal observable mean without requiring downstream objective labels.

The main text uses A_C to denote effective transport in information-normalized coordinates, giving the compact approximation

$$\mathbb{E}[\phi_K(X_1,1) \mid X_s = x_s, C] \approx A_C \phi_K(x_s,s). \tag{19}$$

When T is invertible, the effective operator is $A_C = TA_C^{(r)}T^{-1}$, and the two representations are equivalent. If whitening removes eigenmodes, an effective operator can be defined as $A_C = TA_C^{(r)}T^\dagger$, where $T^\dagger$ is the Moore–Penrose pseudoinverse. In this case, the difference between the implemented prediction and the effective approximation is

$$TA_C^{(r)}\left(I - T^\dagger T\right) r(x_s,s). \tag{20}$$

This residual contributes to the transport approximation error in Proposition 3.1. The implementation evaluates Eq. 18 directly, retaining the raw-anchor information during transport.

The finite-time operators and subsequent controlled residual components are fitted using label-free generator trajectories. The frozen LKF, Gaussian projection, and calibrated information transform remain fixed during downstream PEGASUS optimization, while the task-specific objective readout is updated using the cumulative exact-evaluation archive.

B.4 Finite-Time Operator Calibration

After fixing ϕ_K, each operator A_C is calibrated from frozen-generator rollouts as described in Eq. 3. For each start state, multiple independent terminal continuations estimate the conditional terminal observable mean. The production calibration uses eight independent terminal continuations per start state and includes the $0.5 \to 1$ and $0.75 \to 1$ continuations used by the main coarse-stage optimizer. Operator fitting is objective-independent and uses ridge-regularized least squares. Held-out rollout states are reserved for finite-time transport diagnostics.

Held-out generation and transport diagnostics. Peptide LKF generated valid, novel peptides with competitive sequence and distributional quality using the fewest model evaluations compared with existing discrete generative baselines (Table 1). With LKF frozen, we learned objective-independent Koopman representations and finite-time operators. On held-out branched trajectories, predicted future features closely matched observed terminal features, achieving cosine similarity of 0.989 from $s = 0.5$ to terminal time. Shuffling intermediate-to-terminal correspondences reduced similarity to 0.827, supporting the presence of state-specific prospective information (Table S2). Additional transport controls, terminal-objective prediction, and branch ranking are reported in Appendix G.2.

B.5 Discovery Set and Online Objective Readout

Each downstream run begins with approximately 100 unique discovery sequences generated by the frozen physical generator and evaluated once by the complete objective oracle. These pairs initialize the cumulative exact archive $\mathcal{D}_0$. The task readout in Eq. 4 is a standardized multi-output ridge regression on ϕ_K. Feature centering and scaling are fitted from the discovery archive, the ridge coefficient is selected once by cross-validation over a fixed grid, and that coefficient is held fixed during later online refits. Every newly queried unique sequence is appended to the cumulative archive and remains available in all subsequent rounds. This update changes only the downstream readout; the LKF generator, ϕ_K, and all A_C remain frozen.

C Complete PEGASUS Algorithm and Implementation Details

C.1 PEGASUS Pseudocode and State Updates

Algorithm 1 gives the complete execution order for a fixed preference; Algorithm 2 expands its empirical fine-stage acquisition. Each lineage carries an exactly evaluated incumbent, while all lineages share a unique-sequence oracle cache $\mathcal{D}$ and its online task readout W. A call to $\mathrm{QueryOrCache}$ returns the cached complete objective vector when available; otherwise it evaluates $\mathbf{F}$ once, inserts the sequence into $\mathcal{D}$, and increments Q. The per-stage query caps count only fresh evaluations. The routines $\mathrm{NextCoarse}$, $\mathrm{DecisionRelevantChallengers}$, and $\mathrm{SelectProbe}$ implement the committee tests specified below.

Algorithm 1 PEGASUS: inference-time multi-objective sequence optimization

Require: Frozen generator p_θ, observables ϕ_K, ϕ_L, operators $\{A_C\}$, full-vector oracle $\mathbf{F}$, preference utility U, radii $\mathcal{R}$, protected share β, continuation budget B, coarse branch count N_{br}, fresh-query caps Q_c, Q_f, total budget $Q_{\max}$, lineage schedule

Ensure: Exactly evaluated, Hamming-separated terminal library $\mathcal{L}$

$\mathcal{D} \leftarrow \{(x^{(n)}, \mathbf{F}(x^{(n)}))\}_{n=1}^{N_0}$ from $\mathcal{X}_{\mathrm{disc}}$; $Q \leftarrow N_0$ ▷ Shared unique-query cache

Initialize exactly evaluated lineage incumbents from $\mathcal{D}$; $W \leftarrow \mathrm{Readout}(\mathcal{D})$

for scheduled lineage h and round k, while $Q < Q_{\max}$ **do**

  $x \leftarrow x_{0,k}$ for lineage h ▷ Exactly evaluated round incumbent

  Coarse stage: predict terminal potential, then verify realized endpoints.

  $\mathcal{B}_k \leftarrow \{(x_{s,k}^{(i)}, C_i)\}_{i=1}^{N_{\mathrm{br}}} \leftarrow \mathrm{Branches}(p_\theta, x)$

  $u_i \leftarrow U\Big(W^\top A_{C_i} \phi_K(x_{s,k}^{(i)})\Big), \quad \omega_i \leftarrow [1 + \mathrm{rank}(u_i)]^{-1}$

  $\mathbf{b} \leftarrow \mathrm{ProtectedAllocate}(B, \beta, \boldsymbol{\omega}, b_{\min})$ ▷ Uniform floor + rank-weighted remainder

  $\mathcal{T}_k \leftarrow \mathrm{Unique}\Big(\bigcup_i \mathrm{Continue}_{p_\theta}(x_{s,k}^{(i)}, C_i, b_i)\Big)$ ▷ $j = 1, \ldots, b_i$ independent continuations

  $\mathcal{T}_k \leftarrow \mathrm{Top}_{K_c}(\mathcal{T}_k;\ x' \mapsto U(W^\top \phi_K(x')))$ ▷ Reindex as terminal candidates $x_{1,k}^{(i)}$

  $\mathcal{E}_c \leftarrow \{x\}$; $Q_c^0 \leftarrow Q$

  while $Q - Q_c^0 < Q_c$ and $Q < Q_{\max}$ **do**

    $x' \leftarrow \mathrm{NextCoarse}(\mathcal{T}_k, \mathcal{E}_c, \mathcal{D}, W, \epsilon_c)$ ▷ Top first; then committee challengers

    if $x' = \varnothing$ **then**

      break

    end if

    $(\mathcal{D}, Q, \mathbf{F}_{\mathcal{D}}(x')) \leftarrow \mathrm{QueryOrCache}(x', \mathcal{D}, Q)$

    $\mathcal{E}_c \leftarrow \mathcal{E}_c \cup \{x'\}$ ▷ $Q - Q_c^0$ counts fresh coarse labels

    $W \leftarrow \mathrm{Readout}(\mathcal{D})$ ▷ Readout changes only on new labels

  end while

  $x^{\mathrm{c},\star} \leftarrow \arg\max_{x' \in \mathcal{E}_c} U(\mathbf{F}_{\mathcal{D}}(x'))$

  $x^{\mathrm{c}} \leftarrow x^{\mathrm{c},\star}$ if $U(\mathbf{F}_{\mathcal{D}}(x^{\mathrm{c},\star})) > U(\mathbf{F}_{\mathcal{D}}(x)) + \epsilon_{\mathrm{acc}}$; else $x^{\mathrm{c}} \leftarrow x$

  Fine stage: complete multiscale edits and decision-aware acquisition.

  $\mathcal{A} \leftarrow \{a_0\} \cup \bigcup_{r \in \mathcal{R}} \mathrm{Top}_{K_r}(\mathcal{A}_r^{\mathrm{feas}}(x^{\mathrm{c}}); U \circ \widehat{\mathbf{F}})$

  $z_{x^{\mathrm{c}}}(a) \leftarrow \phi_L(T_{x^{\mathrm{c}}}(a)) - \phi_L(x^{\mathrm{c}})$ for $a \in \mathcal{A}$

  $(\mathcal{E}_f, \mathcal{D}, Q, W) \leftarrow \mathrm{EmpiricalPDO}(x^{\mathrm{c}}, \mathcal{A}, \mathcal{D}, Q, W, Q_f, \epsilon_{\mathrm{dec}})$

  $x^{\mathrm{f}} \leftarrow \arg\max_{x' \in \mathcal{E}_f \cup \{x^{\mathrm{c}}\}} U(\mathbf{F}_{\mathcal{D}}(x'))$

  $x_{1,k}^{(i^\star)} \leftarrow x^{\mathrm{f}}$ if $U(\mathbf{F}_{\mathcal{D}}(x^{\mathrm{f}})) > U(\mathbf{F}_{\mathcal{D}}(x^{\mathrm{c}})) + \epsilon_{\mathrm{acc}}$; else $x_{1,k}^{(i^\star)} \leftarrow x^{\mathrm{c}}$

  $x_{0,k+1} \leftarrow x_{1,k}^{(i^\star)}$ for lineage h ▷ Restart a new flow trajectory next round

end for

$\mathcal{L} \leftarrow \mathrm{HammingPack}(\mathcal{D}, N, \delta)$ ▷ Final selection from paid archive

return $\mathcal{L}$

Algorithm 2 Empirical PDO: protected fine observation and selective exact acquisition

Require: Exactly evaluated x, action slate $\mathcal{A}$ (including no-op a_0) with displacements $z_x(a)$, shared $(\mathcal{D}, Q)$, readout W, caps $Q_f, Q_{\max}$, tolerance ϵ_{dec}
Ensure: Exactly evaluated fine candidates $\mathcal{E}_f$, updated $(\mathcal{D}, Q, W)$
if $Q \geq Q_{\max}$ or $\mathcal{A} = \{a_0\}$ **then**
  return $(\{x\}, \mathcal{D}, Q, W)$
end if
$\widehat{\Delta}_0(a) \leftarrow W^\top[\phi_K(T_x(a)) - \phi_K(x)]$ for $a \in \mathcal{A}$
$a^{\mathrm{anc}} \leftarrow \arg\max_{a \in \mathcal{A} \setminus \{a_0\}} U(\mathbf{F}_{\mathcal{D}}(x) + \widehat{\Delta}_0(a))$ ▷ Protected exploit-first anchor
$\mathcal{O} \leftarrow \varnothing$; $\mathcal{E}_f \leftarrow \{x\}$; $Q_f^0 \leftarrow Q$
$(\mathcal{D}, Q, \mathbf{F}_{\mathcal{D}}(T_x(a^{\mathrm{anc}}))) \leftarrow \mathrm{QueryOrCache}(T_x(a^{\mathrm{anc}}), \mathcal{D}, Q)$
$\mathcal{O} \leftarrow \mathcal{O} \cup \{a^{\mathrm{anc}}\}$; $\mathcal{E}_f \leftarrow \mathcal{E}_f \cup \{T_x(a^{\mathrm{anc}})\}$; $W \leftarrow \mathrm{Readout}(\mathcal{D})$
while $Q - Q_f^0 < Q_f$ and $Q < Q_{\max}$ **do**
  Refresh $\widehat{\Delta}_0(a)$ using the current W
  $\mathbf{r}_i \leftarrow \mathbf{F}_{\mathcal{D}}(T_x(a_i)) - \mathbf{F}_{\mathcal{D}}(x) - \widehat{\Delta}_0(a_i), \quad a_i \in \mathcal{O}$
  $\xi(a) \leftarrow \mathrm{ReducedCoordinates}(z_x(a))$ ▷ Normalized local SVD span
  $C_\lambda \leftarrow \arg\min_C \sum_{a_i \in \mathcal{O}} \|\mathbf{r}_i - C^\top \xi(a_i)\|_2^2 + \lambda \|C\|_F^2$
  $\mathfrak{P}_q \leftarrow \{\widehat{\Delta}_0,\ \widehat{\Psi}_\lambda(a) = \widehat{\Delta}_0(a) + C_\lambda^\top \xi(a)\}_{\lambda \in \Lambda}$
  Augment with leave-one-probe members; clamp observed actions to exact vectors
  $\mathcal{C} \leftarrow \mathrm{DecisionRelevantChallengers}(\mathfrak{P}_q, \mathcal{A}, \mathcal{O}, \epsilon_{\mathrm{dec}})$
  if $\mathcal{C} = \varnothing$ **then**
    break ▷ No unresolved decision-relevant action
  end if
  $a^\star \leftarrow \mathrm{SelectProbe}(\mathcal{C}, \mathfrak{P}_q)$ ▷ Common winner, no-op challenger, or information gain
  $(\mathcal{D}, Q, \mathbf{F}_{\mathcal{D}}(T_x(a^\star))) \leftarrow \mathrm{QueryOrCache}(T_x(a^\star), \mathcal{D}, Q)$
  $\mathcal{O} \leftarrow \mathcal{O} \cup \{a^\star\}$; $\mathcal{E}_f \leftarrow \mathcal{E}_f \cup \{T_x(a^\star)\}$; $W \leftarrow \mathrm{Readout}(\mathcal{D})$
end while
return $(\mathcal{E}_f, \mathcal{D}, Q, W)$

Both stages accept a move only after observing its exact objective vector. For the currently scheduled lineage, the selected terminal candidate $x_{1,k}^{(i^\star)}$ becomes the next-round incumbent, $x_{0,k+1} = x_{1,k}^{(i^\star)}$, and starts a new frozen-generator trajectory; it is not the next time step of the previous trajectory. Final library packing operates on the paid archive and does not change lineage-level decisions (Appendix C.7).

C.2 Protected Coarse Continuation Allocation

For N_{br} intermediate branches with prospective utilities $u_1, \ldots, u_{N_{\mathrm{br}}}$ and total continuation budget B, PEGASUS reserves a protected portion of the continuation budget for approximately uniform allocation across retained branches. The remaining adaptive portion is distributed using rank weights $\omega_i = [\mathrm{rank}(u_i) + 1]^{-1}$. The weights are normalized over retained branches, and integer continuation counts are produced by largest-remainder rounding while preserving the total budget and the configured minimum branch floor. The frozen peptide benchmark uses the AF$= 0.5$ configuration, corresponding to an equal protected and adaptive split in the final benchmark implementation. This allocation changes only model-side continuation computation and consumes no objective query.

C.3 Coarse Exploit-First Verification

After terminal endpoints are generated, PEGASUS retains a bounded top contender prefix under the current task readout and exact-queries the highest-ranked endpoint first. The coarse response committee contains the full-data readout and deterministic leave-one-block-out readouts; the production setting uses four held-out blocks, giving at most five coherent full-vector predictors. Queried candidates and the incumbent are clamped to their exact objective vectors in every committee member. An unqueried candidate remains eligible only if at least one member predicts utility above the best exact candidate by more than the coarse decision tolerance. Among eligible candidates, the next query is the strongest updated contender. Acquisition stops when no eligible challenger remains or the coarse fresh-query cap is reached. The production contender prefix is 32 and the maximum number of fresh coarse queries per slate is three.

C.4 Multiscale Candidate Construction

For incumbent x, an exact Hamming-r action chooses r distinct sequence positions and substitutes one legal token at each selected position. H1 is exhaustively enumerated when tractable. Larger shells are combinatorial, so PEGASUS samples a fixed-size slate of complete exact-r sequences, deduplicates identical targets, and evaluates their frozen features in batches before any exact objective call. The peptide benchmark uses $\mathcal{R} = \{1, 2, 4\}$ at $L = 12$ and $\mathcal{R} = \{1, 2, 4, 8\}$ at $L = 40$. The production setting samples 256 candidates for each non-H1 scale and retains a top contender prefix of 32 per scale for the short fine-stage verifier. These values control model-side opportunity exposure and do not impose one exact query per scale.

For each realized action a, PEGASUS computes the complete displacement $z_x(a) = \phi_L(T_x(a)) - \phi_L(x)$. A higher-order action enters PDO as one complete sequence transformation. The implementation does not sum single-edit objective responses to predict an H2, H4, or H8 response.

C.5 Empirical PDO Response Models

The global objective readout provides the parent-relative prior $\widehat{\Delta}_0(a) = \widehat{\mathbf{F}}_k(T_x(a)) - \widehat{\mathbf{F}}_k(x)$. For each queried action a_i, the exact local residual is $\mathbf{r}_i = \mathbf{Y}_x(a_i) - \widehat{\Delta}_0(a_i)$. PEGASUS normalizes the current displacement family using one global median nonzero norm, computes its empirical SVD span, and represents each action by a reduced coordinate $\xi(a) \in \mathbb{R}^{r_{\mathrm{loc}}}$. For ridge coefficient λ, the local correction solves

$$C_\lambda = \operatorname*{arg\,min}_{C \in \mathbb{R}^{r_{\mathrm{loc}} \times m}} \sum_{a_i \in \mathcal{O}_q} \|\mathbf{r}_i - C^\top \xi(a_i)\|_2^2 + \lambda \|C\|_F^2, \tag{21}$$

and the corresponding response map is $\widehat{\Psi}_\lambda(a) = \widehat{\Delta}_0(a) + C_\lambda^\top \xi(a)$. The production empirical family includes the uncorrected global prior, ridge corrections with $\lambda \in \{0.03, 0.1, 0.3, 1, 3\}$, and deterministic leave-one-probe sensitivity members when enough local observations are available. Each queried action is clamped to its exact response in every member. The family preserves coherent multi-objective vectors and does not construct Cartesian products of independent per-objective intervals.

C.6 PDO Acquisition and Stopping

Practical committee agreement is not assumed to provide formal coverage. Each active fine decision therefore receives one protected exploit-first exact observation from the full multiscale contender family before empirical PDO may terminate. After each exact response, PEGASUS recomputes the empirical ϵ-good sets. If all empirical response maps select the same unresolved non-noop winner, that action is queried directly. If no-op is the common decision, the strongest remaining non-noop challenger is tested. When members disagree, acquisition is restricted to their winners and ϵ-plausible alternatives. The implementation uses the ridge information matrix $G = \lambda I + \sum_{a \in \mathcal{O}_q} \xi(a)\xi(a)^\top$ and a rank-one variance-reduction score over disputed decision directions to select an informative probe. If the information score is numerically uninformative, the highest-value disputed action is queried. The process stops when the empirical decision is resolved or after eight fresh fine queries in the production peptide configuration.

The decision tolerance ϵ_{dec} controls empirical stopping, and the separate acceptance threshold ϵ_{acc} controls whether an exactly evaluated action replaces the incumbent. Their numerical values are experiment-specific configuration parameters.

C.7 Lineage Scheduling, Query Accounting, and Final Packing

Exact labels are shared across lineages through one unique-sequence cache, while each lineage forms its own coarse endpoint set, fine multiscale family, PDO decision, and acceptance update. Discovery queries, coarse verification queries, protected fine queries, and fine challenger queries all contribute to the same unique exact-query count. LKF sampling, Koopman propagation, feature computation, candidate enumeration, readout fitting, and response-model fitting are model-side computation and do not count as objective queries.

At termination, candidate sequences are ordered by the prescribed quality criterion and greedily admitted to the returned library only when their normalized Hamming distance from every already selected sequence meets the configured threshold. The main peptide benchmark returns 100 unique terminal sequences per target, seed, and length. The packing rule changes the reported library but does not retroactively alter optimization decisions.

C.8 Production Configuration and Computational Scaling

Table S1 summarizes the frozen peptide inference configuration used in the reported experiments.

Table S1: Production PEGASUS settings for the peptide benchmark.

Setting	Frozen peptide configuration
Discovery exact labels	approximately 100 unique sequences
Coarse continuations	$0.5 \to 1$ and $0.75 \to 1$
Intermediate starts per coarse chain	6
Protected/adaptive continuation split	$0.5/0.5$ (AF$= 0.5$)
Coarse contender prefix	32
Maximum fresh coarse queries per slate	3
L12 multiscale radii	$\{1, 2, 4\}$
L40 multiscale radii	$\{1, 2, 4, 8\}$
Candidates per sampled non-H1 scale	256
Fine contender prefix per scale	32
Maximum fresh fine queries per lineage/round	8
Coarse readout committee	full fit + 4 leave-block-out fits
Local PDO ridge grid	$\{0.03, 0.1, 0.3, 1, 3\}$
Objective preference in main benchmark	$\mathbf{w} = (1, 1, 1, 1, 1, 1)$
Tchebycheff augmentation	$\rho = 0.01$
Tchebycheff reference	$\mathbf{z}^\star = \mathbf{1}$
PDO decision tolerance	$\epsilon_{\mathrm{dec}} = 0.002$
Final library Hamming separation	$\delta = 0.20$

Let N_{br} be the number of realized coarse starts, B the continuation budget, K the number of active lineages, and C_r the model-side candidate count at radius r. Coarse physical generation scales with $N_{\mathrm{br}} + B$ finite-time transition samples, and Koopman look-ahead requires one finite-time observable propagation per start. Fine candidate construction scales with $\sum_{r \in \mathcal{R}} C_r$ model-side action and feature operations per lineage. Exact-query cost is governed by the protected anchors and unresolved challengers. Under the ideal PDO theorem, the acquisition component scales with d_{PDO} instead of the raw candidate count.

D Theoretical Analysis and Proofs

D.1 Pareto Soundness and Completeness of the Preference Scalarization

The following results concern the scalarization family, not the global convergence of PEGASUS. They establish that preference-conditioned augmented Tchebycheff optimization is compatible with Pareto-optimal trade-offs in a finite discrete design space.

Proposition D.1 (Pareto soundness)**.** *Let $\bar{\mathbf{w}} \in \mathbb{R}^m_{>0}$ and $\rho > 0$. Any global maximizer of $U_{\mathbf{w},\rho}$ over a finite feasible objective set is Pareto-optimal.*

Proof. Suppose $\mathbf{y}'$ Pareto-dominates $\mathbf{y}$. Then $z^\star_o - y'_o \leq z^\star_o - y_o$ for every objective o, so the weighted maximum gap in Eq. 13 cannot increase. Because all weights are positive and at least one coordinate improves strictly, $\bar{\mathbf{w}}^\top \mathbf{y}' > \bar{\mathbf{w}}^\top \mathbf{y}$. Hence $U_{\mathbf{w},\rho}(\mathbf{y}') > U_{\mathbf{w},\rho}(\mathbf{y})$, so a dominated point cannot be a global maximizer. □

Theorem D.2 (Pareto completeness of the augmented Tchebycheff family)**.** *Let $\mathcal{Y} \subset \mathbb{R}^m$ be a finite feasible objective set, and let $\mathbf{z}^\star$ strictly dominate every $\mathbf{y} \in \mathcal{Y}$ componentwise. For every Pareto-optimal objective vector $\mathbf{y}^\star \in \mathcal{Y}$, there exists a strictly positive preference vector $\mathbf{w}$ and a*

sufficiently small $\rho > 0$ such that $\mathbf{y}^\star$ is a global maximizer of $U_{\mathbf{w},\rho}$ over $\mathcal{Y}$. Any design sharing the same objective vector attains the same utility.

Proof. Choose unnormalized weights $w_o = (z_o^\star - y_o^\star)^{-1}$ and normalize them to obtain $\bar{\mathbf{w}}$. Then $\bar{w}_o(z_o^\star - y_o^\star) = c$ for the same constant $c > 0$ for every objective coordinate. For any distinct feasible objective vector $\mathbf{y} \neq \mathbf{y}^\star$, Pareto optimality of $\mathbf{y}^\star$ implies that $y_o < y_o^\star$ for at least one objective o. Hence $\max_o \bar{w}_o(z_o^\star - y_o) > c$, so $\mathbf{y}^\star$ uniquely minimizes the unaugmented weighted Tchebycheff gap among distinct objective vectors. Because $\mathcal{Y}$ is finite, the minimum positive gap margin γ between $\mathbf{y}^\star$ and all distinct competitors exists. Let $M = \max_{\mathbf{y}\in\mathcal{Y}} |\bar{\mathbf{w}}^\top(\mathbf{y} - \mathbf{y}^\star)|$. If $M = 0$, any $\rho > 0$ preserves the optimum. Otherwise choose $0 < \rho < \gamma/M$. The augmentation term then cannot reverse the strict unaugmented ordering, so $\mathbf{y}^\star$ maximizes $U_{\mathbf{w},\rho}$. □

The theorem establishes representability for a strictly dominating aspiration vector and sufficiently small augmentation. The implemented benchmark uses $\mathbf{z}^\star = \mathbf{1}$ with objective coordinates clipped at one, so that strict-domination premise does not hold for every feasible vector. We do not invoke the completeness theorem as a guarantee for this fixed implementation reference or its finite preference sweep. Pareto soundness applies to the implemented positive-weight, positive-augmentation utility.

D.2 Lipschitz Stability of the Augmented Tchebycheff Utility

Lemma D.3 (Tchebycheff utility is Lipschitz)**.** *For fixed* $\mathbf{w}$ *and* ρ, $|U(\mathbf{y}) - U(\mathbf{y}')| \leq L_U \|\mathbf{y} - \mathbf{y}'\|_\infty$ *with* $L_U = \|\bar{\mathbf{w}}\|_\infty + \rho$.

Proof. Let $M(\mathbf{y}) = \max_o \bar{w}_o(z_o^\star - y_o)$. Then $|M(\mathbf{y}) - M(\mathbf{y}')| \leq \|\bar{\mathbf{w}}\|_\infty \|\mathbf{y} - \mathbf{y}'\|_\infty$. In addition, $\rho|\bar{\mathbf{w}}^\top(\mathbf{y} - \mathbf{y}')| \leq \rho\|\bar{\mathbf{w}}\|_1 \|\mathbf{y} - \mathbf{y}'\|_\infty = \rho\|\mathbf{y} - \mathbf{y}'\|_\infty$. The triangle inequality gives the result. □

D.3 Proof of the Future-Observability Decision Bound

Proof of Proposition 3.1. Write $\mathbf{F}(X_1) = W_k^\top \phi_K(X_1) + \mathbf{r}(X_1)$. Conditioning on $X_s = x_{s,k}$ and continuation C gives

$$\mathbf{V}_C(x_{s,k}) - \widehat{\mathbf{V}}_{k,C}(x_{s,k}) = W_k^\top \left(\mathbb{E}[\phi_K(X_1) \mid X_s = x_{s,k}, C] - A_C \phi_K(x_{s,k})\right) + \mathbb{E}[\mathbf{r}(X_1) \mid X_s = x_{s,k}, C]. \tag{22}$$

The induced-norm inequality and the two assumptions in Proposition 3.1 yield $\|\mathbf{V}_C(x_{s,k}) - \widehat{\mathbf{V}}_{k,C}(x_{s,k})\|_\infty \leq \|W_k^\top\|_{\infty\leftarrow 2}\epsilon_K + \epsilon_R = \delta_{\mathcal{B}_k}$ uniformly over $\mathcal{B}_k$. By Lemma D.3, every compared pair has utility error at most $L_U\delta_{\mathcal{B}_k}$. Let $b^\star$ maximize the true conditional-mean utility over $\mathcal{B}_k$ and let $\hat{b}$ maximize its prediction. Writing $\mathbf{V}(b) = \mathbf{V}_C(x_{s,k})$ and $\widehat{\mathbf{V}}(b) = \widehat{\mathbf{V}}_{k,C}(x_{s,k})$ for $b = (x_{s,k}, C)$, we have

$$\begin{aligned} U(\mathbf{V}(b^\star)) - U(\mathbf{V}(\hat{b})) &\leq U(\mathbf{V}(b^\star)) - U(\widehat{\mathbf{V}}(b^\star)) && (23)\\ &\quad + U(\widehat{\mathbf{V}}(\hat{b})) - U(\mathbf{V}(\hat{b})) && (24)\\ &\leq 2L_U\delta_{\mathcal{B}_k}, && (25) \end{aligned}$$

where the predicted-maximizer property removes the intermediate difference $U(\widehat{\mathbf{V}}(b^\star)) - U(\widehat{\mathbf{V}}(\hat{b})) \leq 0$. The statement concerns the utility of the conditional-mean objective vector and does not bound the utility of a sampled terminal endpoint. □

D.4 Proof of the Intrinsic-Dimension PDO Bound

Proof of Theorem 3.3. After q exact action queries, exact consistency forces all response maps in the version set to agree on every displacement direction whose response lies in the known subspace $\mathcal{K}_q$. Suppose the decision is not ϵ-observable. If every initially decision-relevant displacement were already contained in $\mathcal{K}_q$, all remaining response maps would assign the same vector response to every action that can affect the decision and would therefore induce the same utility ordering on that set. A common ϵ-good action would exist, contradicting non-observability. Hence a non-observable decision admits an unresolved decision-relevant action with displacement outside $\mathcal{K}_q$. By

assumption, the acquisition rule queries an action that adds such a direction, increasing the resolved dimension within $\mathcal{K}_0 + \mathcal{S}_0$ by at least one. At most $d_{\mathrm{PDO}} = \dim(\mathcal{K}_0 + \mathcal{S}_0) - \dim(\mathcal{K}_0)$ strict increases are possible. Because the true response map remains in every exact version set, any action in the final common ϵ-good intersection has true regret at most ϵ. One optional final verification query gives $Q_{\mathrm{fine}} \le d_{\mathrm{PDO}} + 1$. □

D.5 Worst-Case Lower Bound

Theorem D.4 (Worst-case intrinsic-dimension lower bound)**.** *For every $d \ge 1$ and every $\epsilon < \Delta$, there exists an exact decision-linear instance with intrinsic decision dimension d for which any deterministic algorithm restricted to querying the finite candidate actions, and guaranteed to return an ϵ-optimal action, requires at least d linearly independent exact action queries in the worst case.*

Proof. Let the action displacements be $\{0, \pm e_1, \dots, \pm e_d\}$ and let the scalar decision gain be $g_\theta(a) = \theta^\top a$. After fewer than d linearly independent queries, some coordinate direction e_r remains outside the queried span. The two parameters $\theta^+ = \Delta e_r$ and $\theta^- = -\Delta e_r$ agree on every previously queried displacement, yet their unique optimal nonzero actions are e_r and $-e_r$, respectively. For $\epsilon < \Delta$, no single output is ϵ-optimal under both instances. Thus at least d independent queries are necessary in the worst case. □

D.6 Multiscale Expansion Corollary

Corollary D.5 (No mandatory query cost per added scale)**.** *Under the assumptions of Theorem 3.3, enlarging the finite action family with additional Hamming scales does not increase the PDO acquisition bound if the newly decision-relevant displacements lie in $\mathcal{K}_0 + \mathcal{S}_0$ without increasing d_{PDO}.*

Proof. Theorem 3.3 depends on the dimension increment d_{PDO} and contains no term for the number of actions or scale labels. If the enlarged family does not increase this dimension, the same upper bound applies. □

D.7 Approximate Decision-Linear Responses

Proposition D.6 (Misspecified response regret)**.** *Suppose $\mathbf{Y}_x(a) = B^\star z_x(a) + \mathbf{e}(a)$ with $\|\mathbf{e}(a)\|_\infty \le \eta$. Assume the queried displacements $z_x(q_1), \dots, z_x(q_r)$ are linearly independent and span the decision-relevant displacements, so that a linear map interpolating their exact responses exists. Write every relevant $z_x(a) = \sum_i \alpha_i(a) z_x(q_i)$ and set $\kappa_Q = \max_a \|\boldsymbol{\alpha}(a)\|_1$. Any such interpolating linear response map has vector error at most $\eta(1 + \kappa_Q)$ on the relevant action set. Its utility-maximizing action has regret at most $2L_U\eta(1 + \kappa_Q)$, plus any explicit ϵ stopping tolerance.*

Proof. For an interpolating map $\widehat{B}$, linearity gives $\widehat{B} z_x(a) = \sum_i \alpha_i(a) \mathbf{Y}_x(q_i) = B^\star z_x(a) + \sum_i \alpha_i(a) \mathbf{e}(q_i)$. Therefore $\|\mathbf{Y}_x(a) - \widehat{B} z_x(a)\|_\infty \le \eta + \eta \|\boldsymbol{\alpha}(a)\|_1 \le \eta(1 + \kappa_Q)$. Lemma D.3 converts this vector error to utility error, and comparison of the predicted and true maximizers contributes the factor of two. □

D.8 Exact Acceptance and No-Op Safety

Proposition D.7 (Exact monotone acceptance)**.** *Every accepted PEGASUS move $x \to y$ satisfies $U(\mathbf{F}(y)) > U(\mathbf{F}(x)) + \epsilon_{\mathrm{acc}}$.*

Proof. The implementation changes the incumbent only after obtaining the exact vector $\mathbf{F}(y)$ and applying the stated acceptance test. The result is independent of the future observer, task readout, and empirical PDO response family. □

Under the exact PDO assumptions, the no-op action has gain zero. If the optimal exact gain exceeds $\epsilon + \epsilon_{\mathrm{acc}}$, every ϵ-optimal action has gain above the acceptance threshold. If every action has nonpositive gain, no-op is optimal.

D.9 Protected Branch Support and Conditional Packing

Proposition D.8 (Protected branch support)**.** *Suppose a retained branch receives at least $b_{\min}$ independent terminal continuations and has probability $p_B > 0$ of reaching a sequence basin B. The probability that at least one protected continuation reaches B is $1 - (1 - p_B)^{b_{\min}}$.*

Proof. The probability that none of the $b_{\min}$ independent continuations reaches B is $(1 - p_B)^{b_{\min}}$. Taking the complement yields the claim. □

Proposition D.9 (Conditional Hamming packing)**.** *If the greedy final-packing procedure returns N sequences while admitting a candidate only when its normalized Hamming distance from every previously selected sequence is at least δ, then all returned pairs have normalized Hamming distance at least δ.*

Proof. The property holds after the first selected sequence. Each subsequent insertion is checked against every previously selected sequence, so induction preserves the pairwise threshold. The statement is conditional on the evaluated archive and the candidate ordering producing N admissible sequences. □

E Peptide Experimental Setup and Reproducibility

E.1 Targets, Sequence Lengths, Seeds, and Terminal Reporting

The primary terminal-100 benchmark comprises three structured targets with pre-existing peptide binders (3IDJ, 5AZ8, 7JVS), three structured targets without known binders (AMHR2, OX1R, DUSP12), and two intrinsically disordered targets without known binders (EWS1::FLI1, MYC) (Chen et al., 2025b). Including both binder-annotated and binder-unannotated targets probes optimization in established binding contexts and prospective discovery settings; the structured and disordered targets extend the test across protein classes. The separate controlled E4–E6 mechanism study uses eight targets without known binders: AMHR2, CLK1, DUSP12, EWS1::FLI1, MYC, PPP5, OX1R, and UCHL5. Its aggregate results are reported separately from the primary benchmark. Both studies include $L = 12$ and $L = 40$ and seeds 42, 73, and 101 unless a narrower experiment is specified. A primary E3 run returns 100 unique terminal peptides. Target, seed, and length define an independent benchmark run; the E4-A mechanism comparison contains 48 matched units and the E5 ladder experiment contains 144 target–seed–length–ladder cells.

E.2 Peptide Objective Oracle and Normalization

Every benchmark method uses the same frozen six-property oracle implementation. The objective order is Non-Hemolysis, Non-Fouling, Solubility, Permeability, Half-Life, and Affinity. Non-Hemolysis rewards lower predicted red-blood-cell lysis; Non-Fouling rewards resistance to nonspecific protein adsorption; Solubility rewards predicted aqueous solubility; and Permeability denotes predicted cell penetrance, not an experimentally measured permeability coefficient or a PAMPA/Caco-2 regression output. These four properties are classification scores. Half-Life denotes predicted persistence in human serum (raw output in hours), while Affinity denotes a target-conditioned, increasing binding-strength score on the predictor's negative-log-like scale. The affinity score is not a measured dissociation constant. These are computational predictions, not experimental measurements of the optimized peptides (Zhang et al., 2026). The first five objectives depend only on the peptide sequence; affinity also takes the target protein sequence. One counted exact query returns the complete six-dimensional vector for one unique peptide under the fixed target.

Let the raw benchmark outputs be y_{Hem}, y_{NF}, y_{Sol}, y_{Perm}, y_{HL}, and y_{Aff}. The optimization-facing scores are

$$\mathbf{F}(x) = \mathrm{clip}_{[0,1]}\left(1 - y_{\mathrm{Hem}},\ y_{\mathrm{NF}},\ y_{\mathrm{Sol}},\ y_{\mathrm{Perm}},\ \frac{\mathrm{clip}(y_{\mathrm{HL}}, 0, 2)}{2},\ \frac{y_{\mathrm{Aff}}}{10}\right). \tag{26}$$

Thus, a higher value is preferable for every coordinate: hemolysis is inverted, predicted half-life is capped at 2 hours and divided by 2, and the raw affinity score is divided by 10. The half-life cap is a benchmark normalization threshold, not a biological upper bound. The oracle wrapper applies

these transformations once; downstream methods do not renormalize scores per target or method. The main equal-preference benchmark uses $\mathbf{w} = (1, 1, 1, 1, 1, 1)$, $\rho = 0.01$, and $\mathbf{z}^\star = \mathbf{1}$.

The six predictors are frozen throughout optimization and across baselines.

E.3 Baseline Configuration and Query Fairness

MOG-DFM and AReUReDi use their published peptide optimization procedures (Chen et al., 2025c;b) with the same six-objective benchmark oracle. LKF Best-of-N draws from the same frozen physical generator as PEGASUS and is matched to PEGASUS in unique exact-query count. SMS-EMOA (Beume et al., 2007), NSGA-III (Jain & Deb, 2013; Fan, 2014), SPEA2 (Zitzler et al., 2001), and MOPSO (Coello & Lechuga, 2002) use 2,000 unique exact objective queries per run in the completed benchmark. The completed LaMBO baseline (Stanton et al., 2022) is a two-start portfolio for each target, seed, and sequence length. Its observed query counts include both neutral-generator warm starts and BO proposals and are not forced to match PEGASUS. The portfolio protocol and uncertainty estimates are reported in Appendix E.4.

All methods use the same target sequence, peptide length, frozen six-objective oracle, score orientation, and hypervolume definition. The hypervolume reference is the zero vector. The main E3 metric is computed on 100 returned sequences per run: PEGASUS returns a Hamming-separated library selected from its paid archive, MOG-DFM and AReUReDi use the first 100 unique completed terminals, and the classical methods select the final 100 from their 2,000-query archives. LKF Best-of-N is query-matched to PEGASUS. Distinct terminal-set construction rules and actual observed query costs are reported alongside the quality measurements. Query counts include one full-vector evaluation per newly queried sequence; repeated cache hits are excluded.

E.4 LaMBO Two-Start Portfolio

The completed LaMBO comparison covers eight peptide targets and seeds 42, 73, and 101 at both $L = 12$ and $L = 40$. For each target-seed pair, we form a portfolio from two independent 8-round LaMBO arms, each initialized with 512 neutral LKF warm-start sequences, to align the query budget with the other methods under comparison. All 1,024 warm-start evaluations count toward the paid unique-query budget, whereas only BO-proposed sequences are eligible for the reported terminal set of 100. The two arms have no sequence overlap in the supplied runs.

For computational feasibility, we used two independent 8-round arms rather than the upstream LaMBO default of one 64-round arm. Although both configurations provide comparable query budgets, a 64-round run was estimated to require at least 20.6 GPU-hours per target-seed pair, or approximately 493 GPU-hours across all 24 pairs. This made the full panel impractical and substantially more expensive than the other baseline runs, which generally required well under 10 GPU-hours per pair.

We report the hypervolume of the final 100 selected BO proposals, distinct from the full proposal archive hypervolume. The length-12 portfolio achieves final-100 HV 0.464 (hierarchical bootstrap 95% CI $[0.431, 0.505]$), full-proposal-archive HV 0.469 ($[0.435, 0.509]$), and $1{,}278.8$ unique queries ($[1{,}278.2, 1{,}279.3]$). The corresponding length-40 values are 0.048 ($[0.039, 0.058]$), 0.048 ($[0.040, 0.058]$), and $1{,}279.8$ ($[1{,}279.5, 1{,}280.0]$). The terminal-set pairwise and nearest-neighbor normalized Hamming distances are 0.912 and 0.333 at length 12, and 0.924 and 0.425 at length 40. Relative to the first 8-round arm, adding the independent second start increases final-100 HV by 0.050 ($[0.031, 0.083]$) at length 12 and 0.011 ($[0.007, 0.015]$) at length 40. Confidence intervals resample targets and then seeds within sampled targets; returned sequences within a run are not treated as independent replicates.

Tables S4 and S34 give the complete target-level portfolio results for 100 final top sequences. In particular, mean normalized non-fouling scores of the length-40 LaMBO terminal sets range from 0.054 to 0.078 across targets under the shared oracle.

E.5 Evaluation Metrics and Statistical Aggregation

The primary set-level quality metric is six-dimensional hypervolume of the 100 normalized maximize-oriented objective vectors using reference point $(0, 0, 0, 0, 0, 0)$. We also report the

mean of each objective, normalized pairwise Hamming diversity, target-wise hypervolume, and the number of unique exact queries. For final set $\{x^{(n)}\}_{n=1}^{N}$, pairwise Hamming diversity is $\binom{N}{2}^{-1} \sum_{1 \le n < n' \le N} d_H(x^{(n)}, x^{(n')})/L$. Sequence-space coverage additionally reports the largest retained subset satisfying prescribed minimum pairwise Hamming thresholds.

Run-level metrics are computed before aggregation. Main peptide summaries average over the eight targets and three seeds within each sequence length. Paired comparisons preserve target and seed pairing. Confidence intervals use the target-aware or hierarchical bootstrap specified by the corresponding experiment; the 100 terminal sequences within one run are not counted as independent experimental replicates.

E.6 Unconditional Generation Evaluation

The unconditional benchmark compares LKF with CFM, PepDFM, EvoDiff (Alamdari et al., 2023), and DPLM (Wang et al., 2024) at fixed peptide lengths $L = 12$ and $L = 40$. Reported metrics include external sequence pseudo-likelihood, embedding-distribution MMD, normalized pairwise Hamming diversity, native NFE, exact-length rate, and effective NFE required to obtain fixed-length samples.

Native NFE denotes the number of model forward evaluations in a method's native sampling procedure, whereas Effective NFE incorporates method-specific resampling costs. For CFM, samples that do not match the requested length are rejected; consequently, its Effective NFE is the Native NFE divided by the empirical exact-length acceptance rate. For DPLM, the Exact-length rate is instead the observed fraction of generated outputs matching the requested length and is not used as a rejection-acceptance probability. DPLM's Effective NFE is measured directly in batch-equivalent model forwards and includes the additional forwards incurred by its native anti-repeat resampling procedure. For methods that generate the requested length directly without an additional resampling adjustment, Effective NFE equals Native NFE. The benchmark therefore distinguishes native sampling depth from the effective computational cost of producing fixed-length samples.

E.7 Structural Evaluation Protocol

We evaluated selected PEGASUS peptide–target complexes with AlphaFold3 prediction (Abramson et al., 2024) and HADDOCK docking (Honorato et al., 2024). Each reported AlphaFold3 binder is evaluated across five independent runs. The main structural summary reports ipTM and integrated Local Interaction Score (iLIS) (Kim et al., 2026) for the configurations described in the Results. The $L = 12$ panel contains one PEGASUS binder per target. For $L = 40$, primary ipTM/iLIS reporting is restricted to targets with reproducible structural confidence in the completed analysis. Docking visualizations use the target protein in light blue and the peptide binder in yellow; tryptophan side chains are highlighted in light pink in the zoomed interface views. Structural predictions are computational corroboration of plausible target-contacting complexes and are not experimental binding measurements.

E.8 Orthogonal Property Evaluators

To test whether optimization-selected sequence features remain favorable under predictors absent from the PEGASUS loop, terminal peptides are rescored with HemoPI2 (Rathore et al., 2025) for hemolysis, PeptideBERT (Guntuboina et al., 2023) for non-fouling, PerseuCPP (Bernardes-Loch et al., 2025) for penetrance, CamSol (Sormanni et al., 2015) for intrinsic solubility, and PLifePred2 Model 1 (Mathur et al., 2018) for half-life. PeptideBERT evaluation at $L = 40$ is treated as exploratory because its training data are concentrated at shorter lengths. Binding affinity is excluded from this cross-model property panel because target engagement is analyzed separately with structural methods. The primary retention analysis is performed within target, seed, and length cells to avoid pooling different task contexts.

F DNA EXPERIMENTAL SETUP

F.1 DNA DATASET AND PREPROCESSING

The unconditional generators model fixed-length 170-bp DNA sequences over $\{\mathrm{A}, \mathrm{C}, \mathrm{G}, \mathrm{T}\}$. The source corpus comprises *Arabidopsis thaliana* open-chromatin regions collected from more than 250 ATAC-seq experiments in AraENCODE (Wang et al., 2023). Peaks were aligned to the TAIR10.1 reference genome, centered on their summits, and extracted as 170-bp windows; peaks within 25 bp of an annotated transcription start site were excluded to reduce promoter-proximal and genic enrichment. The merged (peak-level deduplicated) LKF input contains 189,676 sequences before the nucleotide filter. Removing seven sequences with non-ACGT bases gives chromosome-disjoint partitions of 117,574 sequences from chromosomes 1–3 for training, 29,330 from chromosome 4 for validation, and 42,765 from chromosome 5 for testing. The chromosome-level split is fixed before fitting to limit genomic-neighborhood leakage, consistent with evaluation practices in DNA sequence modeling (Benegas et al., 2023; Fishman et al., 2025); no exact duplicates were found across partitions. LKF uses reverse-complement augmentation on training sequences only after its first architecture-screening run. The DirFM and FisherFM experiments inherit the same source-corpus and chromosome-split specification from the accompanying dataset protocol, although their archived training FASTA files were unavailable for byte-level comparison with the LKF and CatFM inputs.

These sequences are unlabeled accessible-chromatin windows, not measured enhancer-activity examples. Enhancer activity is scored by a separate frozen plantGREP model trained on STARR-seq readouts (Jores et al., 2026); none of the plantGREP activity labels are used to train the unconditional generators.

F.2 DNA GENERATOR TRAINING AND FINITE-TIME CALIBRATION

The reported DNA PEGASUS production run uses a one-component (M = 1) DNA LKF rather than the $M = 8$ peptide setting. The available architecture-screening record includes a 256-wide, 12-layer, eight-head $M = 1$ LKF with four latent layers, trained with per-residue mixture negative log-likelihood under uniform categorical corruption. It reports an AdamW configuration with learning rate 3×10^{-4}, weight decay 0.05, batch size 128, and best-validation checkpoint selection; we report those architecture-screening settings separately from the production checkpoint used for all reported DNA results. The production LKF and three replacement generators are compared in the unconditional sequence-quality evaluation (Table S32); historical LKF architecture-screening measurements, resolved by k-mer length, are reported separately in Table S33.

CatFM (Roos et al., 2026) uses a two-time diffusion-transformer architecture with 21.6 million parameters, a hidden width of 512, a conditioning width of 128, six transformer blocks, eight attention heads, and dropout 0.1. The model learns to transport samples from an isotropic Gaussian prior to one-hot DNA sequences using a categorical vector-flow-matching cross-entropy objective. Lagrangian categorical self-distillation (CSD) loss is applied to 25% of the training batches. CatFM is trained with AdamW using a batch size of 128, weight decay of 10^{-5}, and seed 42. The initial training phase uses a learning rate of 10^{-4}. Training is then resumed from the best-validation checkpoint at step 22,032, to improve distribution matching, with the learning rate decaying from 5×10^{-5} to 10^{-5} by step 70,000 and from 10^{-5} to 2×10^{-6} by step 100,000. The PEGASUS experiments use the checkpoint saved at step 95,000, selected by mean k-mer Jensen–Shannon distance over $k = 1, \ldots, 7$ against the natural-sequence reference distribution. The final checkpoint is evaluated alongside the other selected backbones in Table S32.

DirFM (Stärk et al., 2024) and FisherFM (Davis et al., 2024) each use an approximately 13.3-million-parameter, 20-block one-dimensional convolutional network with 256 channels and a 256-dimensional Gaussian-Fourier time embedding. DirFM is trained with Adam using a learning rate of 5×10^{-4}, a batch size of 64, no dropout, and $\alpha_{\max} = 8$. Its selected trial-005 checkpoint was saved at epoch 154 after 284,890 optimizer steps and achieved a validation loss of 0.7928. FisherFM is trained for 800 epochs with AdamW using a learning rate of 10^{-3}, weight decay of 0.01, a batch size of 512, no dropout, gradient clipping at 1.0, and no learning-rate scheduler. Its selected checkpoint was saved at epoch 798 after 183,770 optimizer steps and attained the lowest validation loss recorded during that run. For DirFM and FisherFM, the inherited configuration was

selected from a ten-setting hyperparameter sweep using the held-out score labeled *test loss* in the original experimental report; within the selected configuration, checkpoint selection used validation loss. We retain this distinction rather than relabeling the sweep-selection metric as validation loss.

In the generator-replacement experiment, CatFM, DirFM, and FisherFM use their own continuous-state mechanics and their own objective-independent finite-time observer. Their full-path sampling costs in the downstream PEGASUS experiments are 16, 99, and 100 NFE, respectively; these should not be confused with the 16-step unconditional sampling evaluation in Table S32. The matched operator-calibration protocol draws 2,048 training parents, 512 validation parents, and 512 held-out parents from the plantGREP validation-sequence resource, with continuation chains $0.5 \to 1$ and $0.75 \to 1$. Each generator uses 64 anchor features; the ridge coefficient is selected independently for each chain from a fixed grid without plantGREP objective labels. This calibration split is distinct from the generator-training split.

F.3 DNA Objective Model and Normalization

One frozen plantGREP package, developed from STARR-seq enhancer-activity data rather than the AraENCODE generator-training corpus, predicts enhancer strength for a DNA sequence under five assay settings (Jores et al., 2026): light (L), dark (D), warm (W), and cold (C) conditions in tobacco, and a separate maize setting (M). These are five predictor outputs, not five optimization objectives. We derive three maximize-oriented objectives,

$$\mathbf{c}(x) = (L - D,\ W - C,\ -M). \tag{27}$$

The first objective favors greater predicted activity in light relative to dark; the second favors greater activity in warm relative to cold conditions; and the third penalizes predicted activity in maize. The contrast objectives are formed from the raw predictor outputs before normalization. The first two are *relative condition-response* objectives: a large difference does not necessarily imply high absolute activity under the favored condition. The anti-maize objective is $-M$, not a tobacco-minus-maize contrast, and by itself does not require high tobacco activity. All three are derived from predicted enhancer activity rather than measured activity of the generated sequences.

Each objective coordinate o uses fixed first- and ninety-ninth-percentile reference values estimated on a common 10,000-sequence reference set, followed by $z_o = \mathrm{clip}_{[0,1]}((c_o - q_{01,o})/(q_{99,o} - q_{01,o}))$. The reference values are $(-3.095779, 1.591785)$ for light-minus-dark, $(0.074640, 1.055210)$ for warm-minus-cold, and $(-2.304014, 2.489046)$ for anti-maize. The same frozen normalization and oracle package are used by PEGASUS and its component ablations and generator-replacement runs. One newly evaluated DNA sequence returns the complete three-vector and consumes one paid oracle query; cache hits consume none.

F.4 Optimization Protocol, Query Accounting, and Metrics

The frozen DNA comparison uses seeds 42, 73, and 101, a common 100-sequence exact discovery manifest for each seed, 12 PEGASUS lineages, an adaptive continuation fraction of 0.5, the two chains above, and the multiscale radii $\{1, 2, 4, 8\}$. The no-op remains available, and the protected final-library separation is normalized Hamming distance 0.20. The maximum paid-query budget is 2,000, but the scheduled full PEGASUS run ends at 1,331, 1,398, and 1,280 unique queries across the three seeds. Those early stops are completion of the fixed controller schedule, not a measured convergence certificate. The generator-replacement and component-ablation runs can have different actual stopping budgets. Each run returns 100 sequences from its own paid archive, and the archive, current population, nondominated subset, and returned library are treated as distinct evaluation objects.

The primary clipped three-objective hypervolume uses reference $(0, 0, 0)$ on the nondominated paid archive. HV-AUC@2000 integrates its value after each paid query, including the initial 100 discovery evaluations (consumed one at a time in manifest order), and carries the final value forward only after the true stopping point: $\mathrm{AUC}_{2000} = \{\sum_{q=1}^{Q} \mathrm{HV}_q + (2000 - Q)\mathrm{HV}_Q\}/2000$. No additional queries are imputed. Since $(1, 1, 1)$ alone attains clipped hypervolume one, we also analyze the same archived raw derived coordinates after removing only the upper normalization cap: $z_o^+ = \max\{0, (c_o - q_{01,o})/(q_{99,o} - q_{01,o})\}$. This upper-unclipped evaluation is post hoc and does not alter the clipped objectives on which any optimizer made decisions.

The returned library contains 100 sequences selected from each run's paid archive by the same clipped augmented Tchebycheff utility, subject to minimum pairwise normalized Hamming distance 0.20; ties are broken lexicographically. Its raw composite is the mean of $(L - D) + (W - C) - M$ over these selected sequences. We aggregate run-level metrics over the three seeds and distinguish the actual number of paid queries from the 2,000-query HV-AUC reporting horizon. The common discovery manifest contributes to every run's paid budget; selecting a returned sequence never incurs a new query.

G Additional Experiments and Diagnostics

G.1 Orthogonal Validation with Independent Peptide Property Predictors

Optimization against learned property models can produce sequences that score well under the guidance models without preserving the same ordering under alternative predictive formulations. We therefore rescore PEGASUS terminal designs using external peptide-property predictors that are never queried during optimization. Within each target, seed, and sequence-length cell, we identify the top quartile of PEGASUS designs according to each optimization objective and evaluate whether these designs remain highly ranked under the corresponding external predictor. We report median external percentile, the fraction retained above the external median, the fraction retained in the external top quartile, and enrichment over the empirical top-quartile background rate (Table S7).

The strongest cross-model retention occurs for non-fouling and penetrance, with additional directionally favorable agreement for non-hemolysis. Agreement is weaker for solubility and half-life. The solubility comparison changes from the learned PeptiVerse optimization model to the physicochemical CamSol score, while the half-life comparison uses PLifePred2, whose released output is not calibrated to the human-serum endpoint used by the optimization model. These results are computational cross-model corroboration; experimental measurements are required to establish the corresponding biological properties directly.

G.2 Alternative Realizations and Robustness of Future Observability

The held-out transport controls compare Koopman transport with identity-on-K, matched LKF-linear transport, and shuffled intermediate-to-terminal correspondences (Table S2). The online objective readout and decision-level branch-ranking diagnostics are reported in Tables S27 and S14, respectively. These diagnostics distinguish representation transport from objective prediction and decision ranking.

The main ablation supports prospective allocation but does not isolate the Koopman representation itself. We additionally compare the Koopman future observer with a dimension-matched generic finite-time linear observer trained from the same frozen LKF dynamics and objective-independent transition data, holding the downstream PEGASUS optimizer fixed. The generic observer supports strong downstream optimization and exceeds the Koopman realization on several pooled endpoint metrics, including $L = 40$ hypervolume, while exact-query cost is similar (Table S12). Thus, the measured closure of the learned Koopman observables and the downstream quality of a future observer are distinct properties. Koopman transport supplies a reusable, objective-independent finite-time factorization; neither this comparison nor the held-out aggregate transport errors establishes unique necessity, universal downstream superiority, or the uniform error assumptions of Proposition 3.1.

We also test sensitivity to the number of mixture components in the frozen LKF generator and its corresponding future-observability stack. Changing the production $M = 8$ mixture to $M = 1$ leaves downstream optimization competitive across both sequence lengths (Table S13). Here M is the generator mixture-component count, not the dimension of the 64-dimensional Koopman observable; this result supports robustness to backbone mixture complexity, not a claim that one Koopman coordinate suffices.

G.3 Detailed Multiscale Action Analysis

We compare three nested action ladders: $R1 = \{1\}$, $R124 = \{1, 2, 4\}$, and $R1248 = \{1, 2, 4, 8\}$. Expanding beyond H1 increases generated utility and accepted fine progress at both sequence lengths, with the largest Full-versus-H1 effect at $L = 40$ (Table S15). The intermediate $R124$ ladder captures a substantial part of the improvement, supporting the value of multiple intervention scales.

Adding H8 on top of H1/H2/H4 does not yield a statistically resolved change in final generated utility by itself (Table S16). Its contribution is clearer from action usage: within the full ladder, H8 supplies a larger fraction of accepted fine-stage gain and is selected more frequently at $L = 40$ (Table S10). At $L = 40$, H8 also reduces exact-query use relative to $R124$ in the completed comparison. Candidate exposure increases far more rapidly than exact fine-query demand as scales are added (Table S17), which is consistent with the PDO separation between model-side breadth and paid objective information.

G.4 PDO Acquisition Safeguards and Sparse Query Scaling

The production PDO procedure begins each fine decision with one protected exact observation. Removing this anchor reduces generated utility and hypervolume and decreases accepted fine progress (Table S18). The protected configuration spends more exact labels but fewer labels per accepted fine move, supporting the anchor as a practical safeguard against unsupported zero-query stopping.

Using 96 causally reconstructed historical optimization states, we also expand nested Hamming-1 contender sets and replay PDO without the production query cap. Median fresh-query demand remains one across all tested candidate-set sizes, while mean demand grows modestly (Table S20). Under the production cap of eight fresh observations, PDO avoids approximately 99.6 to 99.8% of the exhaustive Hamming-1 offline labels. Mean regret for the production prescreen is 0.0023 at $L = 12$ and 0.0019 at $L = 40$. Some positive opportunities are missed, and empirical committee resolution is not a certificate of true ϵ-optimality (Table S20).

G.5 PDO Decision Geometry and Exact-Linear Validation

Within-state fixed-effect estimates did not resolve an association of challenger queries with unresolved rank or candidate count. Pooled associations differed across these quantities, but describe variation among states and do not establish a within-state scaling law (Table S21).

The exact-linear synthetic experiment tests the theorem regime directly. Across 1,600 trials spanning candidate-set sizes from 64 to 4096 and response dimensions from 2 to 16, PDO identifies a zero-regret decision in every trial and satisfies the predicted rank bound in every trial (Table S21). Reconstruction error remains at floating-point precision, and query count tracks response dimension instead of candidate count.

G.6 Generator Substitution and Finite-Horizon Decision Geometry

We replace LKF with CFM and PepDFM while pairing each generator with a dynamics-matched future-observability representation. The objective oracle, preference, multiscale action set, PDO controller, and discovery states remain fixed. CFM gives similar utility to LKF at $L = 12$ and higher utility at $L = 40$, while PepDFM is weaker at both lengths in the completed comparison (Tables S22 and S23). These results show that the PEGASUS decision procedure can operate with different finite-time generators and that final optimization quality still depends on the physical dynamics supplied to the decision layer.

Both LKF- and PepDFM-matched observers predict conditional terminal means accurately (Table S24), yet their finite-horizon decision geometry differs (Table S25). Accurate conditional-mean prediction is therefore insufficient to guarantee identical branch-selection quality across generators.

G.7 Sequence-Space Coverage

The complete PEGASUS terminal libraries satisfy a normalized Hamming separation of $\delta = 0.20$ at both peptide lengths, and substantial subsets remain separated under stricter packing thresholds (Table S8). This remains true after restricting the archive to Pareto-nondominated terminal designs. The result supports broad sequence-space coverage of the returned libraries under the tested objectives; it does not imply that every objective landscape contains multiple distant high-quality basins.

G.8 Preference Steering and Pareto-Region Coverage

The equal-preference run is supplemented by six single-objective emphasis runs on representative $L = 12$ targets. Each emphasized preference assigns weight 0.5 to one objective and 0.1 to each of the other five, while all other PEGASUS settings remain fixed. Emphasizing non-fouling, solubility, permeability, half-life, or affinity increases the corresponding terminal objective mean in all nine target-seed cells, with the strongest shifts for permeability, non-fouling, and half-life (Table S26). Non-hemolysis changes are smaller and its confidence interval includes zero. Extreme-region enrichment is strongest for affinity, permeability, half-life, and non-fouling. These results show that the finite preference sweep can steer PEGASUS toward distinct Pareto-relevant regions. The scalarization completeness theorem in Appendix D concerns representability over all positive preferences and does not imply exhaustive coverage by these seven runs.

G.9 Structural and Tryptophan Counterfactual Analyses

AlphaFold3 and HADDOCK analyses provide out-of-loop structural context for selected PEGASUS designs (Abramson et al., 2024; Honorato et al., 2024). Figure 2 presents representative predicted complexes for AMHR2 and MYC, Figure S1 summarizes five-run ipTM and iLIS measurements, and Figure S2 presents additional predicted binding modes for the remaining six targets. For the recurrent tryptophan-rich $L = 12$ pattern, we replaced every tryptophan in one representative binder per target with phenylalanine, tyrosine, or leucine. All 24 counterfactual variants had lower six-objective utility than their corresponding original designs (Table S6). This supports the optimization relevance of the recurrent sequence pattern under the benchmark oracle but does not establish that individual tryptophan residues are physical interface hotspots.

G.10 DNA Supplementary Results

Unconditional sequence distribution and generator substitution. Table S32 compares the production LKF, CatFM, DirFM, and FisherFM generators using 40,000 generated 170-bp sequences per model under a 16-step unconditional sampling protocol. The arithmetic mean Jensen–Shannon distance over $k = 1, \ldots, 7$ is 0.017579 for LKF, 0.019795 for CatFM, 0.025746 for DirFM, and 0.025476 for FisherFM. LKF has the lowest observed mean k-mer distance in this single-run evaluation, whereas DirFM has the highest GPN/CNN mean $P(\mathrm{real})$; the metrics assess distinct aspects of sequence quality. The historical LKF architecture screen and its held-out train–test reference are reported separately in Table S33. The same PEGASUS decision procedure can operate with an LKF, CatFM, DirFM, or FisherFM backbone. Table 3 separates native generator NFE from paid oracle queries and reports clipped archive HV-AUC alongside mean raw composite in the returned 100-sequence library. The replacement generators have their own calibrated objective-independent finite-time operators. On held-out parents, the information-normalized transport errors improve over the no-transport identity control for both prescribed chains (Table S29). These are representation-transport results and must not be described as measured terminal-objective readout accuracy. Cross-generator upper-unclipped HV-AUC is unavailable because the original per-query raw archives for CatFM, DirFM, and FisherFM were not retained in the available analysis package; it is not reconstructed from final means.

Clipping sensitivity. Since the fixed objective normalization caps every coordinate at one, clipped hypervolume cannot distinguish raw improvements once an archive reaches the saturated vector $(1, 1, 1)$. We therefore retain an upper-unclipped post-hoc analysis of the saved PEGASUS and component-ablation archives in Table S30. These runs still optimized the original clipped objectives; the alternate metric is a reporting diagnostic, not an optimization rerun.

Table S2: Objective-independent finite-time transport for frozen peptide LKF (256 held-out roots, six continuations each). Cosine and relative ℓ_2 compare predicted with observed terminal features; shuffling averages three seeds.

Horizon	Transport	Cosine ↑	Relative ℓ_2 ↓
$0.5 \to 1$	Koopman	0.9886	0.1486
	Identity on ϕ_K	0.4634	0.9614
	Matched LKF-linear	0.9848	0.1718
	Shuffled Koopman	0.8270	0.5724
$0.75 \to 1$	Koopman	0.9940	0.1084
	Matched LKF-linear	0.9924	0.1213
	Shuffled Koopman	0.7341	0.7109

Table S3: Complete terminal-100 objective-level and diversity results for PEGASUS and all peptide baselines across eight targets and three seeds per length ($n = 24$). Entries are run-level mean ± SD. Higher objective scores, terminal hypervolume, and pairwise normalized Hamming distance are better; fewer unique queries are better. The best value within each length and metric is bold.

Length	Method	Non-hemolysis	Non-fouling	Solubility	Permeability	Half-life	Affinity	Terminal HV	Pairwise Hamming	Unique queries
$L = 12$	PEGASUS	**0.976±0.009**	0.818±0.014	0.961±0.012	0.797±0.032	**0.845±0.019**	**0.754±0.024**	0.6867±0.0453	0.6960±0.0209	1303.8±67.6
	MOG-DFM	0.968±0.002	**0.934±0.007**	**0.999±0.001**	0.923±0.006	0.603±0.015	0.615±0.027	**0.6894±0.0561**	0.8908±0.0058	109949.4±1019.2
	AReUReDi	0.956±0.002	0.926±0.006	0.998±0.001	**0.940±0.005**	0.611±0.018	0.610±0.027	0.6416±0.0639	0.8900±0.0056	114842.0±1076.1
	LKF Best-of-N	0.927±0.003	0.663±0.026	0.865±0.013	0.480±0.027	0.391±0.017	0.622±0.025	0.3746±0.0207	**0.9347±0.0027**	1303.8±67.6
	LaMBO	0.949±0.007	0.747±0.034	0.938±0.035	0.564±0.043	0.429±0.025	0.631±0.029	0.4644±0.0683	0.9119±0.0094	**1278.8±1.1**
	PepTune+DPLM	0.930±0.006	0.609±0.037	0.597±0.039	0.305±0.030	0.306±0.010	0.598±0.024	0.2335±0.0216	0.9133±0.0098	1297.1±5.3
	SMS-EMOA	0.918±0.018	0.610±0.056	0.811±0.087	0.578±0.068	0.505±0.039	0.703±0.034	0.5669±0.0662	0.8294±0.0212	2000.0±0.0
	NSGA-III	0.895±0.023	0.529±0.056	0.657±0.037	0.622±0.037	0.550±0.055	0.704±0.037	0.5396±0.0431	0.8406±0.0268	2000.0±0.0
	SPEA2	0.814±0.045	0.454±0.043	0.618±0.040	0.587±0.042	0.583±0.040	0.718±0.027	0.4976±0.0422	0.8091±0.0317	2000.0±0.0
	MOPSO	0.886±0.034	0.529±0.089	0.717±0.064	0.520±0.064	0.471±0.052	0.669±0.041	0.4393±0.0515	0.7573±0.0293	2000.0±0.0
$L = 40$	PEGASUS	0.882±0.028	**0.794±0.020**	**0.943±0.020**	**0.679±0.022**	**0.712±0.026**	0.661±0.022	**0.5286±0.0528**	0.8372±0.0143	1568.2±118.7
	MOG-DFM	**0.933±0.003**	0.265±0.023	0.696±0.061	0.432±0.016	0.610±0.016	0.699±0.046	0.3629±0.0390	**0.9474±0.0006**	155016.7±820.3
	AReURe Di	0.895±0.005	0.230±0.034	0.706±0.029	0.498±0.015	0.657±0.018	0.704±0.051	0.3409±0.0402	0.9326±0.0014	161805.9±1157.0
	LKF Best-of-N	0.753±0.011	0.079±0.008	0.381±0.025	0.380±0.020	0.402±0.016	0.698±0.035	0.0418±0.0113	0.9311±0.0013	1568.2±118.7
	LaMBO	0.797±0.018	0.067±0.012	0.384±0.042	0.367±0.031	0.419±0.016	0.708±0.044	0.0479±0.0180	0.9237±0.0038	**1279.8±0.4**
	PepTune+DPLM	0.792±0.014	0.064±0.015	0.272±0.020	0.358±0.026	0.343±0.017	0.709±0.051	0.1150±0.0876	0.9235±0.0067	1585.5±49.5
	SMS-EMOA	0.906±0.027	0.383±0.088	0.840±0.113	0.524±0.062	0.563±0.058	**0.755±0.044**	0.3602±0.0805	0.7212±0.0408	2000.0±0.0
	NSGA-III	0.865±0.035	0.399±0.069	0.706±0.067	0.458±0.051	0.470±0.032	0.697±0.038	0.2545±0.0510	0.8362±0.0200	2000.0±0.0
	SPEA2	0.766±0.029	0.278±0.041	0.559±0.041	0.483±0.036	0.518±0.020	0.727±0.029	0.2740±0.0439	0.8214±0.0163	2000.0±0.0
	MOPSO	0.746±0.030	0.180±0.045	0.515±0.052	0.383±0.041	0.410±0.039	0.697±0.034	0.1082±0.0350	0.8044±0.0206	2000.0±0.0

DNA component and radius ablations. The frozen full configuration uses $\{1, 2, 4, 8\}$; No-Future removes future-observability-guided allocation, No-PDO replaces decision-aware acquisition, and H1-only removes coordinated larger-radius actions. Table S30 reports both the historical clipped HV-AUC and the post-hoc upper-unclipped HV-AUC. No-PDO uses the entire 2,000-query allowance while Full uses approximately 1,336, so its higher upper-unclipped AUC is not a matched-budget estimate of the effect of PDO. The three-seed comparisons measured lower H1-only performance under both evaluation normalizations. We report No-Future and No-PDO comparisons descriptively because their actual paid-query counts differ. Additional radius sets through H16 and H32 are given in Table S31; because every added scale also changes candidate allocation under the common controller, these are whole-policy comparisons, not isolated estimates of an H16 or H32 mutation's contribution. The final-library Hamming floor acts between returned sequences, not between a proposal and its parent, and therefore cannot be used to rule out H32 edits.

Table S4: Target-wise terminal-100 hypervolume (mean across three seeds). Best value per target and length is bold; Mean averages eight targets.

Length	Method	3IDJ	5AZ8	7JVS	AMHR2	OX1R	DUSP12	EWS::FLI1	MYC	Mean
$L = 12$	PEGASUS	0.6585	**0.6619**	**0.6887**	0.7131	0.7779	0.6251	**0.6893**	**0.6789**	0.6867
	MOG-DFM	**0.6933**	0.6489	0.6812	**0.7300**	**0.8081**	**0.6273**	0.6602	0.6663	**0.6894**
	AReUReDi	0.6399	0.6007	0.6302	0.6716	0.7653	0.5869	0.6148	0.6235	0.6416
	LKF Best-of-N	0.3635	0.3672	0.3817	0.3934	0.4037	0.3519	0.3634	0.3724	0.3746
	LaMBO	0.494	0.423	0.444	0.497	0.560	0.407	0.460	0.428	0.464
	PepTune+DPLM	0.2079	0.2251	0.2275	0.2500	0.2326	0.2409	0.2524	0.2314	0.2335
	SMS-EMOA	0.5028	0.5131	0.5775	0.6381	0.6127	0.5168	0.5727	0.6013	0.5669
	NSGA-III	0.4984	0.5141	0.5220	0.5893	0.5856	0.5127	0.5495	0.5450	0.5396
	SPEA2	0.4754	0.4965	0.4799	0.5122	0.5441	0.4953	0.4797	0.4976	0.4976
	MOPSO	0.4399	0.4233	0.4054	0.4516	0.5103	0.4070	0.4563	0.4205	0.4393
$L = 40$	PEGASUS	**0.5250**	**0.5530**	**0.5155**	**0.5205**	**0.5707**	**0.4837**	**0.4970**	**0.5635**	**0.5286**
	MOG-DFM	0.3420	0.3356	0.3600	0.3486	0.4240	0.3202	0.3818	0.3913	0.3629
	AReUReDi	0.3342	0.3136	0.3262	0.3526	0.4077	0.3136	0.3365	0.3427	0.3409
	LKF Best-of-N	0.0431	0.0414	0.0433	0.0358	0.0509	0.0311	0.0442	0.0446	0.0418
	LaMBO	0.040	0.037	0.035	0.066	0.049	0.045	0.062	0.049	0.048
	PepTune+DPLM	0.1019	0.1939	0.0852	0.0804	0.1730	0.0945	0.1303	0.0605	0.1150
	SMS-EMOA	0.2870	0.3511	0.3791	0.2919	0.4930	0.3315	0.3185	0.4292	0.3602
	NSGA-III	0.2327	0.2338	0.2419	0.2803	0.3316	0.2068	0.2565	0.2526	0.2545
	SPEA2	0.2295	0.2606	0.2896	0.3188	0.3050	0.2573	0.2659	0.2651	0.2740
	MOPSO	0.0927	0.0980	0.1045	0.1433	0.1247	0.0847	0.0893	0.1285	0.1082

Table S5: Terminal-100 peptide utility under equal-weight augmented Tchebycheff ($\rho = 0.01$). Mean $\pm$ SD across 24 target–seed runs per length; top-10 averages the ten highest-utility sequences in each returned library. Best values are bold.

Method	L12 mean utility ↑	L12 top-10 utility ↑	L40 mean utility ↑	L40 top-10 utility ↑
PEGASUS	$\mathbf{-0.0449 \pm 0.0034}$	$\mathbf{-0.0373 \pm 0.0042}$	$\mathbf{-0.0631 \pm 0.0031}$	$\mathbf{-0.0524 \pm 0.0044}$
MOG-DFM	-0.0795 ± 0.0026	-0.0584 ± 0.0054	-0.1445 ± 0.0023	-0.0848 ± 0.0033
AReUReDi	-0.0784 ± 0.0028	-0.0595 ± 0.0054	-0.1485 ± 0.0031	-0.0862 ± 0.0035
LKF Best-of-N	-0.1316 ± 0.0014	-0.0889 ± 0.0038	-0.1660 ± 0.0009	-0.1397 ± 0.0025
LaMBO	-0.1196 ± 0.0029	-0.0752 ± 0.0055	-0.1656 ± 0.0012	-0.1377 ± 0.0067
PepTune+DPLM	-0.1459 ± 0.0021	-0.1108 ± 0.0044	-0.1669 ± 0.0014	-0.1365 ± 0.0094
SMS-EMOA	-0.1192 ± 0.0053	-0.0674 ± 0.0075	-0.1309 ± 0.0092	-0.0854 ± 0.0100
NSGA-III	-0.1271 ± 0.0037	-0.0752 ± 0.0062	-0.1382 ± 0.0059	-0.0999 ± 0.0068
SPEA2	-0.1331 ± 0.0041	-0.0793 ± 0.0073	-0.1486 ± 0.0039	-0.0990 ± 0.0075
MOPSO	-0.1351 ± 0.0052	-0.0866 ± 0.0075	-0.1605 ± 0.0034	-0.1238 ± 0.0068

Table S6: Tryptophan substitution ablation on one $L = 12$ PEGASUS binder per target. W→F/Y/L variants are rescored by the same six-objective utility (higher is better).

Target	PEGASUS binder	WT utility ↑	W→F utility ↑	W→Y utility ↑	W→L utility ↑
7JVS	LQYQQKWWWYTE	-0.050710	-0.120203	-0.095909	-0.087803
3IDJ	RTERQQWWRWQL	-0.040456	-0.122577	-0.104165	-0.095874
MYC	TEWEQEYRQWWQ	-0.049859	-0.152875	-0.092729	-0.158613
EWS::FLI1	TPARQRRERWWW	-0.045686	-0.109868	-0.103116	-0.096355
DUSP12	EYTYQEYWEWYW	-0.051878	-0.110330	-0.093343	-0.097183
OX1R	TPGWWYPYWTKK	-0.046626	-0.100745	-0.112417	-0.087134
5AZ8	TEAQWQYYWWRQ	-0.047842	-0.086154	-0.067591	-0.130776
AMHR2	RREWYKWWKTTT	-0.046341	-0.086405	-0.101757	-0.078303
Mean	n/a	-0.047425	-0.111145	-0.096378	-0.104005

All 24 variants have lower utility than their original binder. The summary row has no pooled sequence (n/a).

Table S7: Independent-predictor retention for PEGASUS-high terminal peptides (optimization top quartile). Entries are medians across 24 target–seed cells per length with hierarchical-bootstrap 95% CIs; $\mathrm{HSR}_{25\to50}$ and $\mathrm{HSR}_{25\to25}$ denote retention above the external median and top quartile. The upper block shows percentile and median retention; the lower block shows top-quartile retention and enrichment.

Length	Property	Evaluator	External percentile [CI]	$\mathrm{HSR}_{25\to50}$ [CI]
$L = 12$	Non-fouling	PeptideBERT	0.7677 [0.7172,0.8081]	0.7600 [0.7200,0.8400]
	Penetrance	PERSEUcpp	0.6616 [0.6162,0.7273]	0.6600 [0.6200,0.7200]
	Hemolysis	HemoPI2	0.6414 [0.5404,0.7273]	0.6600 [0.6000,0.7308]
	Solubility	CamSol	0.5303 [0.5051,0.6162]	0.5400 [0.5200,0.6400]
	Half-life	PLifePred2	0.4343 [0.3838,0.4697]	0.4400 [0.3600,0.4800]
$L = 40$	Non-fouling†	PeptideBERT	0.6061 [0.5707,0.6364]	0.6000 [0.5600,0.6800]
	Penetrance	PERSEUcpp	0.7323 [0.7071,0.7677]	0.7600 [0.7200,0.8000]
	Hemolysis	HemoPI2	0.5657 [0.4848,0.6364]	0.5600 [0.4800,0.6400]
	Solubility	CamSol	0.4949 [0.3838,0.5556]	0.5000 [0.4000,0.5800]
	Half-life	PLifePred2	0.5354 [0.4747,0.5657]	0.5200 [0.4800,0.5600]

Length	Property	$\mathrm{HSR}_{25\to25}$ [CI]	Enrichment [CI]
$L = 12$	Non-fouling	0.5200 [0.4800,0.5600]	2.080 [1.920,2.240]
	Penetrance	0.4000 [0.3600,0.4600]	1.600 [1.440,1.840]
	Hemolysis	0.3600 [0.2400,0.4400]	1.440 [0.960,1.760]
	Solubility	0.3200 [0.2400,0.3600]	1.280 [0.960,1.440]
	Half-life	0.1600 [0.1200,0.2400]	0.640 [0.480,0.960]
$L = 40$	Non-fouling†	0.3200 [0.2600,0.3600]	1.280 [1.040,1.440]
	Penetrance	0.4800 [0.4400,0.5200]	1.920 [1.760,2.080]
	Hemolysis	0.3000 [0.2200,0.3600]	1.200 [0.880,1.440]
	Solubility	0.2200 [0.1600,0.2800]	0.880 [0.640,1.120]
	Half-life	0.2800 [0.2200,0.3200]	1.120 [0.880,1.280]

†PeptideBERT was trained on substantially shorter peptides; its L40 non-fouling evaluation is therefore treated as exploratory rather than primary evidence. Affinity is excluded because binding is evaluated separately using structural analyses.

Table S8: Hamming-separated coverage of PEGASUS terminal libraries (eight targets, three seeds; mean ± SD). P_{all} and P_{ND} are maximal retained counts among all and Pareto-nondominated designs, respectively.

	All terminal designs			Pareto-nondominated designs	
Length	$\mathbf{P}_{\mathrm{all}}(\mathbf{0.20})$	$\mathbf{P}_{\mathrm{all}}(\mathbf{0.30})$	$\mathbf{P}_{\mathrm{all}}(\mathbf{0.40})$	$\mathbf{P}_{\mathrm{ND}}(\mathbf{0.30})$	$\mathbf{P}_{\mathrm{ND}}(\mathbf{0.40})$
$L = 12$	100.0 ± 0.0	67.5 ± 4.2	49.2 ± 5.3	40.2 ± 4.8	31.0 ± 3.7
$L = 40$	100.0 ± 0.0	85.9 ± 2.9	72.9 ± 2.9	52.0 ± 8.2	44.0 ± 7.0

For sequence length L, the separation constraint is implemented using integer Hamming distance $k = \lceil \delta L \rceil$.

Table S9: Future-observability ablation across eight mechanism-study targets and three seeds per length. ΔHV is Full minus No-Future with target-cluster 95% CI; W/T/L counts paired wins/ties/losses.

Length	Full HV	No-Future HV	ΔHV [95% CI]	Queries (Full / No-Future)	W/T/L
$L = 12$	0.6675	0.6630	$+0.0044$ $[-0.0028, +0.0117]$	$\approx 1{,}325/1{,}324$	15/0/9
$L = 40$	0.5351	0.5011	$+0.0340$ **[**$+0.0169, +0.0503$**]**	$\approx 1{,}511/1{,}521$	18/0/6
Pooled	0.6013	0.5820	$+0.0192$ **[**$+0.0099, +0.0287$**]**	$\approx 1{,}418/1{,}423$	33/0/15

Table S10: Multiscale versus H1-only actions across eight mechanism-study targets. Paired differences are $R1248 - R1$ with target-aware 95% CIs; H8 statistics refer to the full ladder.

Metric	$L = 12$	$L = 40$
Generated mean utility, $R1248 - R1 \uparrow$	$+0.0022$ $[0.0013, 0.0031]$	$+0.0057$ $[0.0038, 0.0077]$
Accepted fine gain, $R1248 - R1 \uparrow$	$+0.0704$ $[0.0526, 0.0888]$	$+0.1805$ $[0.1501, 0.2154]$
H8 accepted-gain share ↑	0.193 ± 0.008	0.326 ± 0.013
H8 accepted moves / run	5.17	18.25

Table S11: PDO ablation across eight mechanism-study targets, three seeds, and two lengths ($n = 48$ pairs; AF$= 0.5$). Δ is Full minus No-PDO; matched rows share fresh fine-label allowances, whereas natural-budget rows have unequal realized query counts.

Panel	Metric	Full PEGASUS	No-PDO	Δ (Full−No-PDO)	Paired 95% CI	Target-aware 95% CI
Matched	Generated mean utility	-0.054696	-0.055475	+0.000779	[+0.000158, +0.001380]	[-0.000163, +0.001591]
	Top-10 utility	-0.045892	-0.046447	+0.000555	[-0.000394, +0.001458]	[-0.000883, +0.001824]
	Best utility	-0.041843	-0.042068	+0.000225	[-0.000909, +0.001303]	[-0.001308, +0.001640]
	Hypervolume	0.597522	0.582645	+0.014877	[+0.002120, +0.027997]	[-0.003761, +0.034411]
	Pareto count	59.875	58.333	+1.542	[-1.376, +4.542]	[-3.459, +5.833]
	Accepted fine gain	0.630926	0.595905	+0.035020	[+0.021233, +0.049772]	[+0.016790, +0.053182]
	Pairwise Hamming	0.761748	0.767050	-0.005302	[-0.010023, -0.000532]	[-0.012172, +0.001315]
	Total exact queries	1446.27	1460.13	-13.85	[-17.83, -10.08]	[-19.17, -8.85]
Natural	Generated mean utility	-0.0547	-0.0501	-0.0046	[-0.0053, -0.0039]	[-0.005547, -0.003550]
	Top-10 utility	-0.0459	-0.0403	-0.0056	[-0.0071, -0.0043]	[-0.007287, -0.004122]
	Hypervolume	0.5975	0.6095	-0.0120	[-0.0283, +0.0038]	[-0.030792, +0.005268]
	Accepted fine gain	0.6309	0.7114	-0.0805	[-0.1002, -0.0621]	[-0.104255, -0.058617]
	Total exact queries	1446.3	2377.9	-931.6	[-983.9, -879.0]	[-1016.629, -844.083]
	Fine queries	536.6	1461.4	-924.8	[-977.6, -871.9]	[-1010.856, -837.413]
	Queries / accepted fine move	16.62	19.00	-2.37	[-3.06, -1.71]	[-3.173, -1.604]

Table S12: Koopman versus matched generic linear future observer under the same frozen LKF and PEGASUS protocol. Values are paired Koopman-minus-generic differences with 95% CIs.

Metric	All runs: Δ [95% CI]	$L = 12$: Δ [95% CI]	$L = 40$: Δ [95% CI]
Generated mean utility ↑	$-0.0027\,[-0.0038, -0.0015]$	$-0.0006\,[-0.0016, +0.0003]$	$-0.0047\,[-0.0065, -0.0027]$
Top-10 utility ↑	$-0.0026\,[-0.0041, -0.0011]$	$+0.0002\,[-0.0008, +0.0010]$	$-0.0054\,[-0.0078, -0.0030]$
Hypervolume ↑	$-0.0326\,[-0.0508, -0.0152]$	$+0.0072\,[-0.0006, +0.0144]$	$-0.0724\,[-0.0990, -0.0451]$
Accepted fine gain ↑	$-0.0123\,[-0.0397, +0.0129]$	$+0.0470\,[+0.0297, +0.0635]$	$-0.0716\,[-0.1093, -0.0363]$
Unique exact queries ↓	$+4.1\,[-26.1, +32.6]$	$+34.5\,[-0.7, +69.0]$	$-26.4\,[-71.9, +16.4]$

Table S13: LKF mixture-component ablation ($M = 1$ versus $M = 8$) under a common PEGASUS configuration. Mean ± SD over 24 runs per length; M is not the Koopman feature dimension.

	$L = 12$		$L = 40$	
Metric	$M = 1$	$M = 8$	$M = 1$	$M = 8$
Unique exact queries	1322 ± 75	1307 ± 75	1554 ± 144	1586 ± 134
Accepted moves	51.2 ± 7.6	47.9 ± 7.4	28.7 ± 3.4	30.3 ± 6.3
Strong hits	7.17 ± 2.30	3.25 ± 1.82	10.4 ± 2.2	8.54 ± 2.32
Verification positive rate	0.0645 ± 0.0098	0.0603 ± 0.0094	0.0350 ± 0.0040	0.0367 ± 0.0076
Pairwise Hamming	0.699 ± 0.028	0.686 ± 0.024	0.851 ± 0.012	0.837 ± 0.015
Affinity mean	0.744 ± 0.035	0.748 ± 0.032	0.674 ± 0.039	0.658 ± 0.026
Affinity best	0.812 ± 0.040	0.812 ± 0.040	0.787 ± 0.062	0.779 ± 0.060

Table S14: Branch-ranking diagnostics for the E2 feature representations (higher is better). These static decision scores are distinct from finite-time transport accuracy.

Representation / control	Spearman ↑	Top-1 agreement ↑
Identity-style control	0.265	0.316
Matched LKF-linear	0.254	0.288
Koopman representation	0.207	0.261

Recorded utility differences were approximately $+3 \times 10^{-4}$ for Koopman transport relative to uniform continuation and -8×10^{-5} relative to the matched LKF-linear control.

Table S15: Multiscale action-ladder outcomes across eight mechanism-study targets and three seeds ($n = 24$ per length and ladder; AF$= 0.5$). Utilities and accepted gain are higher-is-better; query count is lower-is-better.

	$L = 12$			$L = 40$		
Ladder	**Generated utility ↑**	**Exact queries ↓**	**Accepted gain ↑**	**Generated utility ↑**	**Exact queries ↓**	**Accepted gain ↑**
$R1$	-0.0481 ± 0.0009	**1206.8**	0.3724	-0.0692 ± 0.0012	**1552.0**	0.6385
$R124$	-0.0457 ± 0.0011	1302.5	0.4272	-0.0638 ± 0.0008	1749.9	**0.8221**
$R1248$	-0.0459 ± 0.0009	1306.8	**0.4428**	-0.0635 ± 0.0006	1585.7	0.8190

Table S16: Marginal effect of H8 relative to H1/H2/H4: paired $R1248 - R124$ differences with 95% CIs across target–seed runs (AF$= 0.5$).

Length	**Metric**	Δ	**95% CI**
$L = 12$	Generated mean utility	-0.0001	[-0.0009, 0.0005]
$L = 40$	Generated mean utility	+0.0002	[-0.0009, 0.0014]
$L = 12$	Exact queries	+4.4	[-20.5, 30.2]
$L = 40$	Exact queries	-164.2	[-281.1, -64.0]

Table S17: Candidate exposure versus exact fine-query cost for the multiscale ladders. Exposure counts model-side evaluations, not oracle calls.

	$L = 12$		$L = 40$	
Ladder	**Candidate exposures**	**Fine queries**	**Candidate exposures**	**Fine queries**
$R1$	175,101	297.5	583,680	593.7
$R124$	568,311.5	407.9	976,896	817.2
$R1248$	764,919	412.4	1,173,504	660.8

Table S18: Protected exact-observation ablation ($n = 48$ paired target–seed–length runs). Δ is Full minus No-Protected, with 95% CIs.

Metric	**Full PEGASUS**	**No Protected**	Δ	**95% CI**
Generated mean utility ↑	-0.0547	-0.0555	+0.0008	[0.0002, 0.0015]
Top-10 utility ↑	-0.0459	-0.0463	+0.0004	[-0.0004, 0.0013]
Hypervolume ↑	0.5975	0.5847	+0.0128	[0.0018, 0.0250]
Accepted fine gain ↑	0.6309	0.5964	+0.0345	[0.0200, 0.0495]
Unique exact queries ↓	1446.3	1379.7	+66.6	[28.0, 106.6]
Fine queries ↓	536.6	469.9	+66.7	[29.2, 105.5]
Queries / accepted fine move ↓	16.62	18.24	-1.62	[-2.20, -1.04]

Table S19: PDO versus fixed acquisition under matched fresh fine-label allowances ($n = 48$ pairs; AF$= 0.5$). Differences are Full minus No-PDO with target-aware 95% CIs; non-fine queries may differ.

Metric	Difference	Target-aware 95% CI
Accepted fine gain ↑	$+0.035020$	$[+0.016790, +0.053182]$
Generated mean utility ↑	$+0.000779$	$[-0.000163, +0.001591]$
Terminal hypervolume ↑	$+0.014877$	$[-0.003761, +0.034411]$
Total unique queries ↓	-13.854	$[-19.167, -8.854]$

Table S20: PDO acquisition on 96 frozen decision states. Uncapped rows vary contender size; capped rows permit eight fresh exact queries. Offline exhaustive labels are excluded from online query counts.

Panel	Setting	Mean Q	Median / max Q	Mean regret	Resolved	ϵ-optimal rate / missed of present	Labels saved
Uncapped	L12, $N = 64$	1.00	1 / 11	0.00246	1.00	0.56	99.6%
	L12, $N = 128$	1.00	1 / 11	0.00264	1.00	0.52	99.6%
	L40, $N = 64$	2.52	1 / 29	0.00267	1.00	0.50	99.7%
	L40, $N = 128$	2.60	1 / 36	0.00292	1.00	0.44	99.7%
	L40, $N = 256$	2.69	1 / 60	0.00335	1.00	0.40	99.6%
	L40, $N = 512$	3.38	1 / 83	0.00386	1.00	0.35	99.6%
Cap $= 8$	Full H1, L12	0.94	1 / 8	0.0027	0.979	36 / 44	99.6%
	Full H1, L40	1.46	1 / 8	0.0040	0.958	27 / 48	99.8%
	Production prescreen, L12	0.94	1 / 8	0.0023	0.979	33 / 41	99.6%
	Production prescreen, L40	1.50	1 / 8	0.0019	0.938	26 / 46	99.8%

In uncapped rows, the seventh column is the true ϵ-optimal fraction; in capped rows it counts missed improving opportunities among those present offline. Labels saved are $1 - Q_{\text{total,fresh}}/Q_{\text{exhaustive,fresh}}$.

Table S21: Decision-geometry diagnostics: 96 frozen states and 1,600 synthetic exact-linear trials (16 settings $\times$ 100). Fixed-effects estimates are the primary real-state analysis. n/a denotes a quantity not applicable to that diagnostic; it is not missing numerical data.

Panel	Analysis / quantity	Estimate	R^2 / rate	95% CI / result
Real	FE: $Q_{\text{challenger}} \sim d_{\text{unresolved}}$	-0.370	0.008	[-0.676, +0.183]
	FE: $Q_{\text{challenger}} \sim \log N$	+0.354	0.007	[-0.497, +1.664]
	FE: joint model	rank -0.469; log$N$ $+0.467$	0.019	both CIs include zero
	Pooled Spearman: Q vs unresolved rank	0.977	n/a	[0.947, 0.993]
	Pooled Spearman: Q vs $N/\log N$	0.040	n/a	[-0.041, 0.118]
Synthetic	Trials	1600	n/a	16 grid cells $\times$ 100
	Exact-linear correctness	n/a	1600/1600	100%
	Rank-bound pass	n/a	1600/1600	100%
	Mean max reconstruction error	$\sim 10^{-16}$	n/a	numerical precision
	Mean regret / query scaling	0	n/a	$Q \approx d + 1$

Table S22: Generator substitution with dynamics-matched future observers and fixed PEGASUS decisions. Best value per length and metric is bold. Mean over eight targets and three seeds ($n = 24$ per length); queries are unique exact objective-vector requests.

Length	Generator	Best utility ↑	Strong hits ↑	Unique queries ↓
$L = 12$	LKF	-0.03569	3.25	1307
	CFM	**-0.03513**	4.96	**1298**
	PepDFM	-0.03679	**5.00**	1357
$L = 40$	LKF	-0.04800	**8.54**	**1586**
	CFM	**-0.03772**	6.88	1760
	PepDFM	-0.05306	4.04	1715

Table S23: Paired generator-minus-LKF utility differences (95% bootstrap CIs). Query-matched values use the largest exact-query budget shared within each paired run ($n = 24$ per length).

Group	Generator vs. LKF	Final ΔU [95% CI]	Query-matched ΔU [95% CI]
$L = 12$	CFM	$+0.00055\,[-0.00035,\ +0.00154]$	$+0.00046\,[-0.00062,\ +0.00155]$
	PepDFM	$-0.00111\,[-0.00208,\ -0.00018]$	$-0.00128\,[-0.00227,\ -0.00032]$
$L = 40$	CFM	$\mathbf{+0.01028}\,[\mathbf{+0.00850},\ \mathbf{+0.01214}]$	$\mathbf{+0.00863}\,[\mathbf{+0.00673},\ \mathbf{+0.01064}]$
	PepDFM	$-0.00506\,[-0.00678,\ -0.00317]$	$-0.00548\,[-0.00735,\ -0.00341]$
All	CFM	$\mathbf{+0.00542}\,[\mathbf{+0.00372},\ \mathbf{+0.00715}]$	$\mathbf{+0.00455}\,[\mathbf{+0.00297},\ \mathbf{+0.00619}]$
	PepDFM	$-0.00308\,[-0.00424,\ -0.00192]$	$-0.00338\,[-0.00462,\ -0.00213]$

Table S24: Finite-time terminal-feature prediction for dynamics-matched LKF and PepDFM observers. Δ is PepDFM minus LKF with target-cluster 95% CIs (5,000 bootstrap replicates).

Length	Progress	Metric	LKF	PepDFM	Δ	95% CI for Δ
12	0.50	Cosine ↑	0.998578	0.998390	-0.000189	$[-0.000255, -0.000132]$
12	0.75	Cosine ↑	0.999278	0.998277	-0.001001	$[-0.001058, -0.000944]$
40	0.50	Cosine ↑	0.999355	0.999498	+0.000144	$[+0.000111, +0.000179]$
40	0.75	Cosine ↑	0.999635	0.999498	-0.000137	$[-0.000155, -0.000117]$
12	0.50	Relative ℓ_2 ↓	0.055122	0.058091	+0.002969	$[+0.002094, +0.004071]$
12	0.75	Relative ℓ_2 ↓	0.038673	0.059989	+0.021317	$[+0.020109, +0.022466]$
40	0.50	Relative ℓ_2 ↓	0.037726	0.032438	-0.005288	$[-0.006561, -0.004073]$
40	0.75	Relative ℓ_2 ↓	0.027201	0.033005	+0.005804	$[+0.005208, +0.006349]$

Table S25: Finite-horizon branch-decision geometry under dynamics-matched LKF and PepDFM observers. Δ is PepDFM minus LKF with target-cluster 95% CIs (5,000 replicates).

Length	Progress	Metric	LKF	PepDFM	Δ	95% CI for Δ
12	0.50	DecisionSNR ↑	0.2586	0.1988	-0.0598	$[-0.1067, -0.0193]$
12	0.50	Rank Spearman ↑	0.4412	0.3298	-0.1114	$[-0.1996, -0.0288]$
12	0.50	Selection regret ↓	0.00785	0.00887	+0.00103	$[+0.00021, +0.00201]$
12	0.75	DecisionSNR ↑	0.3990	0.3501	-0.0489	$[-0.1220, +0.0300]$
12	0.75	Rank Spearman ↑	0.5779	0.4828	-0.0950	$[-0.1829, -0.0010]$
12	0.75	Selection regret ↓	0.00632	0.00771	+0.00139	$[+0.00011, +0.00267]$
40	0.50	DecisionSNR ↑	0.1690	0.1335	-0.0355	$[-0.0752, +0.0046]$
40	0.50	Rank Spearman ↑	0.3038	0.2483	-0.0555	$[-0.1846, +0.0920]$
40	0.50	Selection regret ↓	0.00255	0.00222	-0.00034	$[-0.00091, +0.00031]$
40	0.75	DecisionSNR ↑	0.2635	0.2507	-0.0128	$[-0.1025, +0.0627]$
40	0.75	Rank Spearman ↑	0.4429	0.4548	+0.0120	$[-0.0376, +0.0606]$
40	0.75	Selection regret ↓	0.00243	0.00180	-0.00062	$[-0.00098, -0.00027]$

Table S26: Objective-emphasized versus equal-preference peptide optimization ($L = 12$; three targets, three seeds). Δ is the paired score change (95% CI); enrichment is relative to uniform preference share $1/7$.

Objective	Equal	Emphasized	Δ (95% CI)	Positive cells	Top-10% share	Top-10% enrichment	Top-20% share
Non-Hemolysis	0.971	0.977	$+0.0058$ (-0.0011, $+0.0131$)	7/9	0.12	$0.84\times$	0.16
Non-Fouling	0.821	0.965	$+0.1447$ ($+0.1361$, $+0.1548$)	9/9	0.57	$4.0\times$	0.62
Solubility	0.967	0.996	$+0.0292$ ($+0.0228$, $+0.0361$)	9/9	0.18	$1.3\times$	0.22
Permeability	0.803	0.955	$+0.1516$ ($+0.1372$, $+0.1644$)	9/9	0.85	$6.0\times$	0.69
Half-Life	0.845	0.978	$+0.1338$ ($+0.1226$, $+0.1464$)	9/9	0.75	$5.2\times$	0.66
Affinity	0.759	0.854	$+0.0946$ ($+0.0850$, $+0.1039$)	9/9	1.00	$7.0\times$	0.71

Table S27: Online objective-readout and PDO-response diagnostics across three $L = 12$ targets and three seeds ($n = 9$). Results are mean $\pm$ SD unless noted.

Diagnostic	Result
Global objective readout	
Mean objective Spearman correlation	0.588 ± 0.016
Utility pairwise ranking accuracy	0.664 ± 0.021
PDO local response models and exact verification	
False-positive proposals blocked	$\mathbf{357/357}$ $(\mathbf{100}\%)$
False-positive proposals accepted	$\mathbf{0}$
Median observed decision regret	0.004 ± 0.006

Table S28: Median global-readout and PDO-committee fitting times at median/maximum labeled-set sizes. Cumulative fitting time is reconstructed from size-binned timings, not end-to-end runtime.

	Global readout		PDO committee		Oracle	Estimated fitting
Length	$n_{\text{med}}/n_{\max}$	Time (ms)†	$n_{\text{med}}/n_{\max}$	Time (ms)†	ms/vector	s/run
$L = 12$	876 / 1,449	0.37 / 0.59	774 / 1,450	1.51 / 2.58	46.54	$\sim$17.4
$L = 40$	1,086 / 1,878	0.73 / 1.21	935 / 1,879	3.01 / 5.78	27.80	$\sim$30.1

†Median latency at the median-size / maximum-size training input; n counts labeled examples. Each PDO timing includes the complete committee fit. The estimated cumulative times include reconstructed joint-view rebuilds and depend on the measured hardware and workload. Feature extraction, generation, and candidate scoring are excluded.

Table S29: Held-out objective-independent transport on three replacement DNA generators (512 parents per continuation). Relative ℓ_2 uses within-generator information-normalized coordinates; identity leaves intermediate features unchanged.

Generator	Continuation	Koopman cosine$\uparrow$	Identity cosine$\uparrow$	Koopman rel. $\ell_2 \downarrow$	Identity rel. $\ell_2 \downarrow$
CatFM	$0.5 \to 1$	0.99615	0.95673	0.08432	0.29149
CatFM	$0.75 \to 1$	0.99759	0.92725	0.06510	0.39078
DirFM	$0.5 \to 1$	0.99926	0.99804	0.03784	0.06151
DirFM	$0.75 \to 1$	0.99981	0.99912	0.01937	0.04137
FisherFM	$0.5 \to 1$	0.99819	0.98508	0.05987	0.17081
FisherFM	$0.75 \to 1$	0.99878	0.99010	0.05065	0.13932

Table S30: DNA ablations under the frozen LKF controller (three seeds). Upper-unclipped HV-AUC is a post-hoc archive evaluation; unequal native-stop query counts preclude matched-budget interpretation.

Variant	Queries	Clipped HV-AUC↑	Upper-unclipped HV-AUC↑
Full PEGASUS	1,336	0.96688 ± 0.0211	1.34385 ± 0.1457
No Future	1,419	0.96365 ± 0.0282	1.43087 ± 0.2037
No PDO	2,000	0.96487 ± 0.0233	1.52735 ± 0.2537
H1 only	992	0.95295 ± 0.0251	1.22499 ± 0.1068

Table S31: DNA Hamming-radius sweep (24 paired replicates per radius set: three discovery manifests $\times$ eight controller seeds). Upper-unclipped HV-AUC is evaluated post hoc; adding radii changes proposal allocation, not only individual move availability.

Radius set	Scales	Clipped HV-AUC↑	Upper-unclipped HV-AUC↑
$\{1\}$	1	0.94704 ± 0.0313	1.22340 ± 0.1382
$\{1, 2, 4\}$	3	0.95814 ± 0.0256	1.29405 ± 0.1732
$\{1, 2, 4, 8\}$ (frozen)	4	0.96286 ± 0.0213	1.37180 ± 0.1926
$\{1, 2, 4, 16\}$	4	0.96037 ± 0.0237	1.31370 ± 0.1654
$\{1, 2, 4, 8, 16\}$	5	0.96208 ± 0.0252	1.34100 ± 0.1572
$\{1, 2, 4, 8, 32\}$	5	0.96315 ± 0.0223	1.35911 ± 0.1738
$\{1, 2, 4, 8, 16, 32\}$	6	0.96450 ± 0.0175	1.37211 ± 0.1291

Table S32: Unconditional quality of frozen DNA generators (40,000 length-170 samples each, 16 sampling steps). Mean k-mer JS distance averages $k = 1, \ldots, 7$ against a common reference; these are single-run results, not downstream PEGASUS NFE. Bold indicates the best observed metric; GC has no monotonic direction.

Generator	Mean GC	Diversity	Novelty	Mean k-mer JS ↓	Mean $P(\text{real})$ ↑	Median $P(\text{real})$ ↑
LKF (production, $M = 1$)	0.365206	0.546246	0.435076	**0.017579**	0.463911	0.495486
CatFM (95k)	0.365389	**0.547557**	0.433729	0.019795	0.380758	0.356644
DirFM	0.353954	0.542133	0.432341	0.025746	**0.503738**	**0.554763**
FisherFM	0.367097	0.544862	**0.436624**	0.025476	0.467018	0.501596

Table S33: Historical LKF architecture screen: generated-versus-training k-mer JS distance ($k = 1, \ldots, 7$; 16 NFE; lower is better). The train–test row is a held-out reference; screening models are not the frozen downstream generators.

LKF screening configuration	$k=1$	$k=2$	$k=3$	$k=4$	$k=5$	$k=6$	$k=7$	Mean ↓
Width 512, $M = 8$	0.017206	0.027129	0.034828	0.043105	0.052802	0.063701	0.076329	0.045014
Width 256, $M = 8$	0.012968	0.021651	0.028792	0.037566	0.047770	0.059227	0.072474	0.040064
Width 128, $M = 8$	0.016720	0.024951	0.032062	0.040284	0.050316	0.061501	0.074488	0.042903
Width 256, $M = 1$	0.005516	0.014587	0.021547	0.029690	0.040080	0.051977	0.065889	0.032755
Train–test reference	0.001492	0.002619	0.003756	0.005130	0.008178	0.014257	0.026569	0.008857

Table S34: LaMBO two-start portfolio: normalized objective mean ± sample SD across three seed-level means of 100 returned proposals per target. Target-wise hypervolume appears in Table S4.

Length	Target	Non-hemolysis	Non-fouling	Solubility	Permeability	Half-life	Affinity
12	3IDJ	0.950±0.009	0.761±0.034	0.943±0.027	0.591±0.042	0.430±0.026	0.607±0.015
12	5AZ8	0.947±0.008	0.748±0.038	0.965±0.023	0.590±0.060	0.438±0.035	0.619±0.006
12	7JVS	0.944±0.013	0.772±0.037	0.946±0.032	0.603±0.015	0.430±0.004	0.644±0.008
12	AMHR2	0.954±0.007	0.748±0.041	0.927±0.047	0.544±0.070	0.412±0.024	0.680±0.007
12	OX1R	0.952±0.004	0.728±0.017	0.912±0.054	0.563±0.020	0.429±0.047	0.661±0.017
12	DUSP12	0.948±0.004	0.725±0.024	0.924±0.033	0.548±0.025	0.441±0.008	0.592±0.006
12	EWS::FLI1	0.950±0.003	0.767±0.049	0.965±0.018	0.548±0.022	0.424±0.003	0.621±0.002
12	MYC	0.945±0.010	0.725±0.029	0.927±0.037	0.528±0.036	0.427±0.042	0.625±0.006
40	3IDJ	0.801±0.019	0.062±0.014	0.386±0.029	0.389±0.058	0.419±0.006	0.698±0.004
40	5AZ8	0.794±0.004	0.069±0.015	0.388±0.051	0.361±0.024	0.419±0.019	0.682±0.006
40	7JVS	0.787±0.004	0.054±0.009	0.378±0.036	0.377±0.026	0.428±0.015	0.707±0.002
40	AMHR2	0.815±0.020	0.071±0.010	0.389±0.024	0.354±0.026	0.414±0.023	0.753±0.008
40	OX1R	0.783±0.021	0.065±0.016	0.385±0.047	0.375±0.020	0.403±0.016	0.794±0.008
40	DUSP12	0.797±0.022	0.064±0.003	0.396±0.078	0.372±0.045	0.415±0.007	0.643±0.005
40	EWS::FLI1	0.798±0.025	0.078±0.010	0.363±0.042	0.357±0.021	0.421±0.023	0.688±0.004
40	MYC	0.801±0.023	0.075±0.005	0.389±0.059	0.349±0.032	0.430±0.002	0.699±0.010

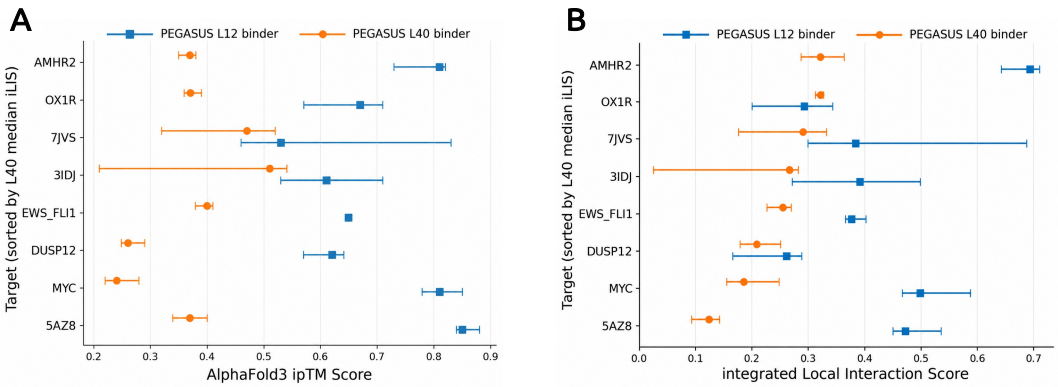


Figure S1: **(A)** AlphaFold3 predicted interface TM-score (ipTM) and **(B)** integrated Local Interaction Score (iLIS) for PEGASUS peptide–target complexes across the eight primary benchmark targets. Blue squares correspond to L12 designs and orange circles to L40 designs. Each point is the median across five independent AlphaFold3 runs, with horizontal error bars indicating their standard deviation across five seeds. All eight L40 series are plotted for context; primary L40 interpretation is restricted to 3IDJ, 7JVS, AMHR2, and OX1R. These are predicted structural-confidence metrics, not measured binding affinities.

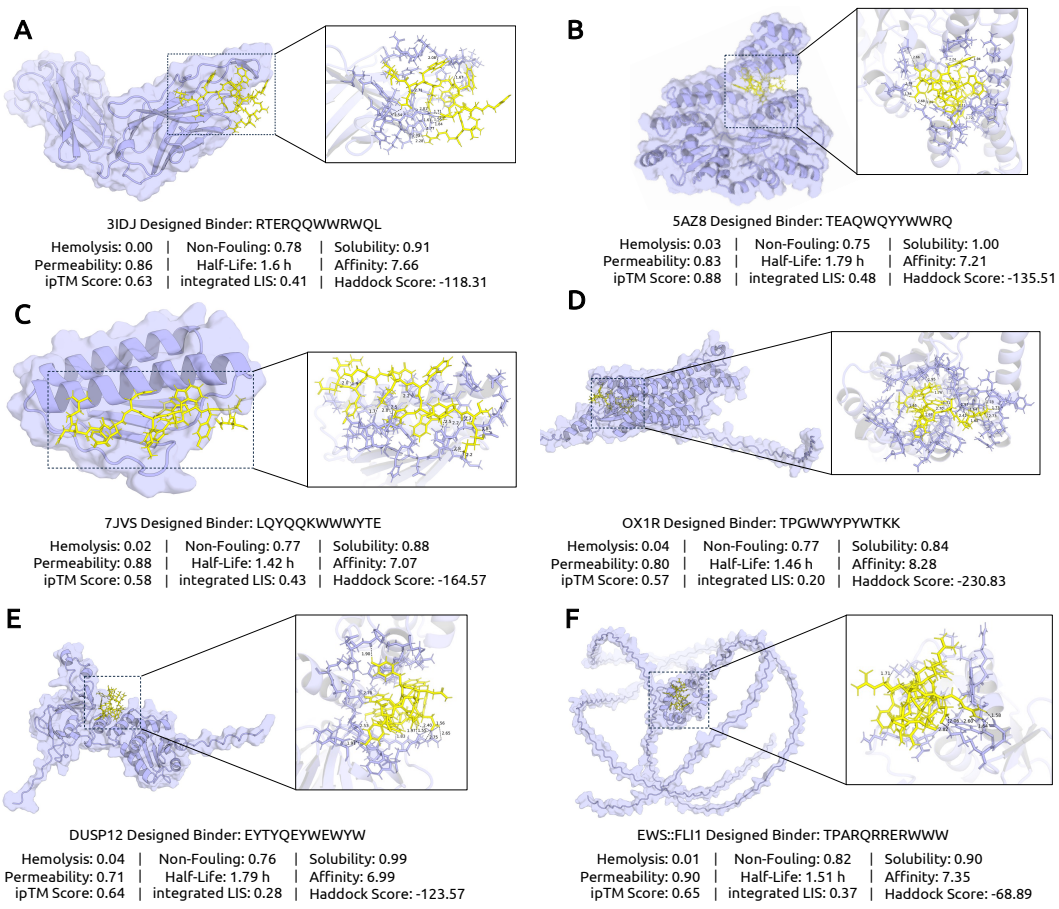


Figure S2: **Predicted binding modes of PEGASUS-designed peptide binders across the remaining six benchmark targets.** Target proteins are shown in light blue and designed peptides in yellow, with enlarged views of the predicted binding interfaces. Predicted peptide properties, structural confidence metrics, and docking scores are reported in each panel.