# SHAPE: LEARNING TO SEARCH WITH AN EVOLVING HARNESS FOR TARGET-BINDING ANTIMICROBIAL PEPTIDE DESIGN

**Anonymous authors**
Paper under double-blind review

## ABSTRACT

Designing deployable antimicrobial peptides (AMPs) is inherently multi-objective, requiring strong antimicrobial activity while balancing safety, AMP likelihood, and sequence diversity. Requiring effective engagement with biologically relevant targets, including drug-resistance and tissue-repair targets, adds another coupled objective and makes the search problem substantially harder. Recent computational methods provide strong peptide design priors, while agentic methods enable iterative search through LLM reasoning and tool feedback. Yet generative methods remain bounded by learned priors, and agentic search typically uses a fixed harness, so accumulated experience improves individual trajectories without improving the search system itself. This raises a broader challenge: how can a peptide discovery system improve its own search process as peptide search proceeds?

We address this challenge with SHAPE, a multi-agent system coupling target-aware peptide search with harness evolution. To evaluate this setting, we introduce a target-aware AMP design benchmark that jointly measures predicted potency, safety, AMP likelihood, sequence diversity, and target binding. SHAPE distills search trajectories into executable condition–action operators, validates them by matched equal-budget comparisons against the current harness, and reuses them when applicable under a shared scorer budget. Across the benchmark, SHAPE improves the composite score by **23.99%** over the strongest baseline under the same budget. Operator-guided edits achieve Pareto improvements in **42%** of evaluated cases, versus **27%** for LLM-reflection-driven edits. Evolved operators generalize across eleven generators: when frozen and reused for three rounds, they achieve a mean cumulative gain of **12.03%** under an external composite score, with all sources finishing above their original means. Together, these results show that SHAPE improves both the designed peptides and the search harness itself by turning search experience into reusable search capabilities.

To enable full reproducibility, we publicly release the SHAPE framework at `https://anonymous.4open.science/r/SHAPE-AMP-Design-D44D`.

## 1 INTRODUCTION

Antimicrobial peptides (AMPs) hold considerable potential for infection control and the development of new antimicrobial agents (Zasloff, 2002; Hancock & Sahl, 2006; Mookherjee et al., 2020). AI-based peptide design offers a powerful way to navigate the vast sequence space, accelerate candidate discovery, and generate experimentally active AMPs (Das et al., 2021; Szymczak et al., 2023; Wang et al., 2025a;b).

Designing strong candidates, however, requires coordinating multiple peptide-level and target-interaction objectives under limited evaluation budgets (Liu et al., 2025; Torres et al., 2026). In wound healing, peptide therapeutics may need to support both infection control and tissue repair (Mookherjee et al., 2020; Thapa et al., 2020). The VEGF-A-binding peptide PR1P provides a concrete precedent, enhancing angiogenesis and accelerating wound closure in mice (Adini et al., 2017). Motivated by this dual role, we study *target-aware AMP design* across three bacterial targets (PBP2a,



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.

GyrA, AceA) and three repair-associated targets (VEGFA, FGF2, ANGPT1), extending AMP design from peptide-level properties to biologically relevant target interactions.

Recent advances span generative and property-guided models, foundation-model approaches, and agentic iterative search. Established methods such as HydrAMP, PepGLAD, and PepZOO provide sequence-, structure-, or multi-objective search priors (Szymczak et al., 2023; Kong et al., 2024; Liu et al., 2025). Foundation-model approaches, represented by AMP-Designer Full, use an LLM-based model for peptide design (Wang et al., 2025b). Agentic workflows add iterative tool use and feedback: PepCraft + AMP-GAN v3 combines conditional generation with agent coordination (Jung et al., 2026), while Pepti-Agent coordinates multiple peptide-design tools (Chen et al., 2026), PDAgent provides protein-mutation strategies that we further adapt to AMP design (Ouyang et al., 2026), and MAC-AMP updates reward functions and generators through multi-agent feedback (Zhou et al., 2026). We additionally evaluate LLM Single-stage (GPT-6), which directly prompts an LLM to generate candidate sequences, and an AutoResearch workflow adapted to peptide search and design (Karpathy, 2026). Generative methods provide increasingly strong peptide design priors, while agentic methods further enable iterative search through LLM reasoning, tools, and feedback. Yet the search harness itself typically remains fixed, leaving accumulated search experience within individual trajectories rather than using it to improve subsequent search. This motivates a broader question: *how can a peptide discovery system improve its own search process as peptide search proceeds?* We decompose this question into two interconnected challenges.

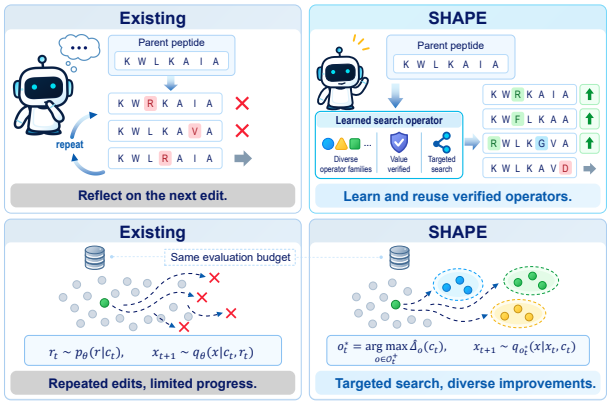



Figure 1: **Learning reusable search operators.** Program evolution turns AMP trajectories into conditional operators. Matched evaluation supports their inheritance and context-dependent use through harness assembly.

**First, jointly optimizing AMP properties and target binding.** There are trade-offs between improving peptide properties and strengthening target binding. Search must therefore coordinate these objectives while maintaining diversity across candidates.

**Second, learning and validating context-appropriate search operators.** In repeated edit–score cycles, sequence edits determine which part of the search space to explore, shaping search efficiency and candidate diversity. Effective editing strategies depend on the current parent, target, and limiting properties. The challenge is to distill reflection-derived experience into reusable procedures with explicit applicability conditions, validate their added value over the current search system, and incorporate them into subsequent search.

We introduce **SHAPE**, a multi-agent system coupling target-aware peptide search with *harness evolution* (Figure 1). Its peptide-search path generates, edits, and evaluates candidates, while its operator-evolution path distills trajectories into executable programs that specify when and how to edit. Validated operators are incorporated into the search harness and reused when their conditions match the current context. The two paths share trajectory evidence and a single scorer budget, including operator validation.

Our contributions are threefold:

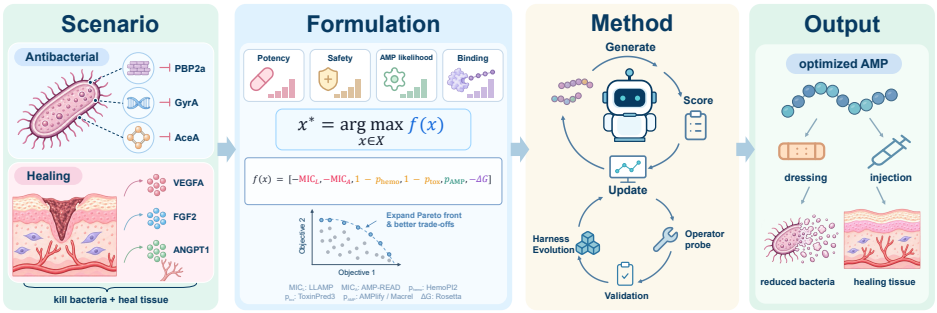



Figure 2: **Target-aware multi-objective AMP design.** Bacterial and repair-associated targets supply molecular contexts for potency, safety, AMP likelihood, and predicted target interactions. LLAMP (Bae et al., 2025), AMP-READ (Wang et al., 2025a), Macrel (Santos-Júnior et al., 2020), ToxinPred3 (Rathore et al., 2024), AMPlify (Li et al., 2022), and Rosetta (RosettaCommons, 2016) provide the depicted scoring components. Peptide search and operator evaluation share the discovery budget; candidate-set diversity also enters the reported composite score. The output panel presents computationally prioritized candidates alongside prospective applications and functional evaluation.

1. **Target-aware multi-objective AMP design.** We formulate a peptide-design benchmark that jointly optimizes traditional AMP properties and predicted target binding while maintaining sequence diversity across six bacterial and repair-associated targets under a common evaluation protocol.
2. **Harness evolution through validated conditional operators.** SHAPE learns executable operators from search experience and validates each by comparing the search system with and without it from the same starting state under equal scoring budgets. Operator-guided edits achieve Pareto improvements over their parents with an empirical probability of **42%**, compared with **27%** for reflection-driven edits.
3. **Strong benchmark performance and transferable search capabilities.** SHAPE achieves the best performance among the evaluated methods, improving composite score by **23.99%** over the strongest baseline (Table 1). When the evolved operators are frozen and applied to peptides from all ten baseline generators over three rounds, the external composite score increases by **12.03%** on average, with every source exceeding its original mean by round 3 (Table 3). This demonstrates that harness evolution produces reusable and transferable editing capabilities.

We publicly release the SHAPE framework, agent prompts, and scorer protocol at `https://anonymous.4open.science/r/SHAPE-AMP-Design-D44D`, which is fully reproducible.

## 2 Target-Aware AMP Design

The six searches cover bacterial AceA, GyrA, and PBP2a and repair-associated VEGFA, FGF2, and ANGPT1 (Figure 2, Appendix C). The benchmark evaluates four candidate-level dimensions: potency, safety, AMP likelihood, target binding and set-level sequence diversity. **potency** uses LLAMP and AMP-READ predicted MIC ($\mu$M) (Bae et al., 2025; Wang et al., 2025a); **safety** uses ToxinPred3 (Rathore et al., 2024) and HemoPI2 (Rathore et al., 2025); **AMP likelihood** uses AMPlify (Li et al., 2022) and Macrel AMP probability (Santos-Júnior et al., 2020); **target binding** uses Rosetta $\Delta G$ (REU) (RosettaCommons, 2016); and **diversity** is mean pairwise Levenshtein edit distance. Target-interaction scores serve as molecular design objectives; functional endpoints are detailed in Appendix C.

The task combines peptide-level potency, safety, AMP likelihood, and target binding with candidate-set diversity. Benchmark composite score equally averages the eight direction-aligned scores as $C = \mathrm{mean}(L, R, A, T, H_{\mathrm{hemo}}, M, E, D)$; normalization and reporting cohorts are specified in Appendix C.2. Section 3 formulates how SHAPE learns and deploys search operators to pursue these objectives under a shared scorer budget.

# 3 SHAPE

SHAPE couples target-aware peptide search with harness evolution. The search harness is the reusable system that guides candidate generation and editing. Its peptide-search path uses reflection and active operators to explore peptide candidates, while its operator-evolution path converts trajectory evidence into executable editing procedures. Each proposed operator is validated against the current harness before being incorporated into subsequent search. Both paths share evidence and a scorer-call budget.

## 3.1 Candidates and Search Objective

Let $\mathcal{Z}$ denote the set of protein targets. For a target $z \in \mathcal{Z}$ and a peptide amino-acid sequence $x$, the scalar utility $u_z(x)$ evaluates the candidate for search, with larger values preferred. The peptide-property and target-binding evaluations described in Section 2 provide the search feedback.

An archive $P$ stores evaluated peptide sequences together with their scores. For a nonempty archive, $A_K(P)$ denotes the subset retained by the archive-selection rule, containing between one and $K$ peptides, where $K$ is a positive integer capacity. The number of retained peptides is $|A_K(P)|$, and $\mathrm{Div}(A_K(P))$ measures their sequence diversity. With diversity weight $\eta$, the archive objective $J_z(P)$ combines mean candidate utility and diversity, assigning zero to the empty archive $\varnothing$:

$$J_z(P) = \frac{1}{|A_K(P)|} \sum_{x \in A_K(P)} u_z(x) + \eta \, \mathrm{Div}(A_K(P)), \qquad J_z(\varnothing) = 0. \tag{1}$$

Search uses $J_z$ to evaluate candidate sets; benchmark reporting uses the Composite Score defined in Appendix C.2.

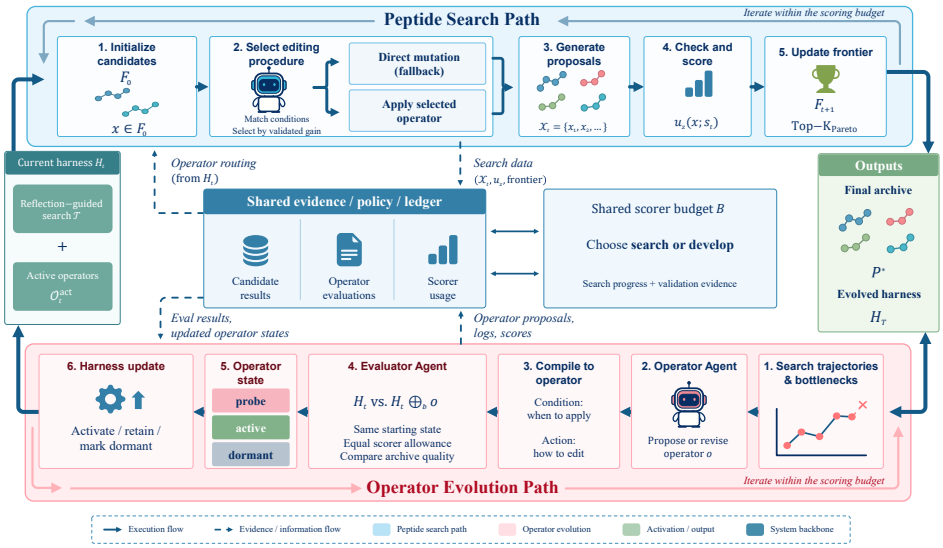



Figure 3: **Harness evolution couples peptide search with operator learning.** The Search Agent routes candidate generation through applicable validated operators or reflection-guided mutation, while scored candidates update the archive and Pareto frontier. The Operator Agent distills trajectory evidence into executable condition–action operators, which are validated by matched evaluation and inherited by the evolving harness. Both paths share evidence and a unified scorer budget; ESM2 (Lin et al., 2023) and GFlowNet (Bengio et al., 2021) are shown as example operator components.

## 3.2 Peptide Search

At step $t$, the search state $s_t$ contains the target, the evaluated archive $P_t$, the working frontier, and the trajectory context. The evaluated archive $P_t$ is the set of peptides that have already been evaluated

and retained. The working frontier supplies parent peptides for further edits. The trajectory context records recent proposals, scores, and edit histories used for reflection. The current harness $H_t$ combines a reflection-guided search routine with active operators. An operator $o$ is a function that maps a search state $s$ and a scorer-call allowance $b$ to a peptide proposal set, i.e., $o(s, b)$ returns the candidate peptides generated by $o$ under $s$ and $b$. It is implemented as an executable edit program with an applicability condition and an editing action. Active operators are those currently available in $H_t$.

**Operator selection.** The Search Agent first identifies the active operators whose conditions match $s_t$, forming the applicable set $\mathcal{C}_t(s_t)$. For each operator $o$, $\widehat{\Delta}_t(o)$ denotes its stored estimate of incremental gain from matched validation, i.e., $\widehat{\Delta}_t(o) = \widehat{\Delta}_o(s_t; H_t)$. The stored value is updated only when matched validation is performed. When the applicable set is nonempty, the agent selects an operator $o_t$ with the highest stored gain, $o_t \in \arg\max_{o \in \mathcal{C}_t(s_t)} \widehat{\Delta}_t(o)$, breaking ties in a fixed order. When none applies, it uses reflection-guided direct mutation.

Under the allotted scorer-call budget $b_t$, the selected operator's executable action $a_{o_t}$ generates the peptide proposal set $U_t$ as $U_t = a_{o_t}(s_t, b_t)$. Both operator execution and direct mutation feed the same checking, scoring, and archive-update workflow. The working frontier is refreshed for subsequent edits, and candidate results and trajectories are recorded as evidence for further search and operator development.

### 3.3 Operator Evolution

**Learning conditions and editing actions.** The Operator Agent synthesizes and revises executable operators using search trajectories, candidate scores, and observed bottlenecks. Feedback shapes both when an operator applies and how it constructs candidate sequences. For example, feedback on fragment recombination produced a successor that selects activity-limited parents, draws donors with complementary property profiles, and constrains charge and hydrophobic runs (Appendix E.2). The resulting program turns these editing decisions into a conditional procedure available for validation and reuse.

**Matched validation.** The Evaluator Agent tests a proposed operator $o$ through paired search blocks. Each block comprises candidate generation and evaluation under a fixed scorer-call allowance. Both branches begin from the same search state $s$, containing the starting archive $P$, and receive the same allowance $b$. The control uses the current harness $H_t$. The trial uses $H_t \oplus_b o$, where $\oplus_b$ integrates $o$ into the harness under that allowance.

For a block returning the evaluated peptide set $\mathcal{X}$, its outcome $Y_s(\mathcal{X}) = J_z(P \cup \mathcal{X})$ is the archive quality after adding those candidates to the starting archive. Let $n_o$ be the positive number of matched probe pairs. For pair $i = 1, \ldots, n_o$, the sets $\mathcal{X}_{o,i}^{\mathrm{ctrl}}$ and $\mathcal{X}_{o,i}^{\mathrm{trt}}$ contain the evaluated candidates returned by the control and trial branches, respectively. The estimated incremental gain of $o$, denoted by $\widehat{\Delta}_o(s; H_t)$, is the mean paired difference:

$$\widehat{\Delta}_o(s; H_t) = \frac{1}{n_o} \sum_{i=1}^{n_o} \left[ Y_s\left(\mathcal{X}_{o,i}^{\mathrm{trt}}\right) - Y_s\left(\mathcal{X}_{o,i}^{\mathrm{ctrl}}\right) \right]. \tag{2}$$

Positive values indicate improved archive quality under the operator-enabled harness. Validation informs whether an operator is activated, retained, or marked dormant. Active operators become available to subsequent search, and their stored validation gains guide routing when their conditions match the current state. These stored validation results are used by operator selection in Section 3.2.

**Budget and alternation.** Peptide search and operator development share the same scorer-call budget $B$, including calls used for matched validation. Based on the accumulated search evidence and remaining budget, SHAPE alternates between candidate search and operator development. It finally returns the peptide archive $P^*$ and evolved harness $H_T$.

# 4 EVALUATION

We evaluate SHAPE and eleven baselines on 300 peptides per method, with 50 peptides for each of three bacterial and three repair-associated targets. Target-agnostic AMP generators and target-conditioned binding designers contribute complementary capabilities under a common evaluation protocol. Agent methods receive 1,200 scorer calls per run, while generators follow their standard workflows. Appendix C.5 details the method-specific protocols and scorer-call accounting.

## 4.1 CANDIDATE QUALITY ACROSS TARGETS

Table 1: Overall means: 300 peptides per method, final Top 50 per target. Best/second-best: **bold**/underlined. Scoring and cohorts: Appendix C.2.

| Method | LLAMP MIC ($\mu$M) ↓ | AMP-READ MIC ($\mu$M) ↓ | AMPlify ↑ | ToxinPred3 ↓ | HemoPI2 ↓ | Macrel AMP ↑ | Rosetta $\Delta G$ (REU) ↓ | diversity (norm.) ↑ | composite score ↑ |
|---|---|---|---|---|---|---|---|---|---|
| Pepti-Agent | 63.60 | 766.17 | 0.494 | 0.149 | 0.199 | 0.236 | -37.17 | 0.757 | 0.425 |
| AMP-Designer Full | 22.66 | 47.35 | 0.905 | 0.206 | 0.272 | 0.663 | -43.73 | 0.550 | 0.632 |
| LLM Single-stage (GPT-6) | 62.53 | 459.08 | 0.376 | 0.141 | 0.195 | 0.237 | -43.35 | 0.769 | 0.461 |
| PepCraft + AMP-GAN v3 | 32.98 | 18.04 | 0.958 | 0.115 | 0.230 | **0.720** | -44.31 | 0.534 | 0.687 |
| PDAgent | 30.34 | 80.74 | 0.734 | 0.117 | **0.181** | 0.451 | -42.45 | 0.633 | 0.634 |
| MAC-AMP | 40.59 | 570.84 | 0.697 | 0.117 | 0.260 | 0.524 | -41.79 | 0.647 | 0.541 |
| AutoResearch | 24.21 | 41.11 | 0.674 | 0.127 | 0.293 | 0.395 | -34.93 | 0.842 | 0.553 |
| PepZOO | 73.81 | 294.37 | 0.851 | 0.188 | 0.426 | 0.529 | 2.07 | 0.773 | 0.470 |
| HydrAMP | 47.30 | 29.41 | 0.953 | 0.856 | 0.248 | 0.548 | 21.32 | 0.010 | 0.450 |
| RFdiffusion | 105.07 | 821.19 | 0.287 | 0.156 | 0.294 | 0.264 | 101.87 | 0.132 | 0.206 |
| PepGLAD | 144.09 | 1637.78 | 0.206 | 0.192 | 0.265 | 0.173 | 1077.00 | 0.252 | 0.177 |
| SHAPE | **4.94** | **10.59** | **0.976** | **0.000** | 0.201 | 0.554 | **-64.65** | **0.850** | **0.852** |

SHAPE improves the overall composite score by **23.99%** over PepCraft + AMP-GAN v3, the strongest baseline (Table 1). It leads six of eight components, combining strong predicted potency and target interactions with broad sequence diversity. Across the six targets, its composite-score advantage over the strongest baseline ranges from **21.43% to 26.55%**. These complementary gains support target-aware peptide design across both antibacterial and repair-associated objectives. SHAPE also achieves **42%** improvement success over 778 proposals, compared with **27%** over AutoResearch's 1,200, linking learned editing procedures to more frequent successful proposals (Appendix C.3).

We further analyze the top 15 peptides per target (90 total) using trajectory-averaged interface geometry, interaction occupancies, and endpoint MM/GBSA energies (Massova & Kollman, 2000). AceA shows the lowest mean iRMSD, highest native-contact retention, and negative MM/GBSA, while GyrA is also negative; VEGFA and FGF2 retain native-contact fractions of 0.69 and 0.67. Hydrogen-bond occupancy reaches 100% across targets, while salt bridges remain target specific, together indicating preserved interfaces with target-dependent interaction patterns.

## 4.2 GENERALIZATION OF EVOLVED OPERATORS

Generalization of evolved operators directly tests whether harness evolution has distilled search experience into reusable capabilities beyond SHAPE's own trajectories. The complete operator library is frozen and applied for three rounds to 100 parents from each of ten generator and agent sources. All parents match at least one operator condition, and the 100 outputs from each round proceeds to the next without selection.

To evaluate generalization independently of SHAPE's own scorer, we use an external six-component panel: PepGLAD's Rosetta score and BLOSUM62-clustering diversity (Kong et al., 2024), AMP Scanner v2 AMP probability (Veltri et al., 2018), Macrel hemolysis probability (Santos-Júnior et al., 2020), PALM aggregation score (Eschbach et al., 2026), and Peptide Dashboard sol-rnn solubility (Ansari & White, 2023). Component directions are aligned before aggregation (Appendix F).

The frozen harness achieves a mean cumulative gain of **12.03%** by R3 (Table 3), with all eleven source means finishing above their originals and improving again in both R2 and R3. The consistency across heterogeneous sources shows that the learned operators capture reusable search structure rather than source-specific corrections, while the continued gains across rounds indicate that

Table 2: **MD interface and MM/GBSA metrics.** Means ± SD across the top 15 peptides per target (90 total). iRMSD and native contacts are trajectory means; H-bond and salt-bridge values are aggregate occupancies. Arrows indicate preferred direction.

| Target | iRMSD ↓ (nm) | Native contacts ↑ (fraction) | H-bond ↑ (%) | Salt bridge ↑ (%) | MM/GBSA ↓ (kcal/mol) |
|---|---|---|---|---|---|
| AceA | $0.26 \pm 0.02$ | $0.70 \pm 0.08$ | $100.00 \pm 0.00$ | $100.00 \pm 0.00$ | $-137.0 \pm 37.4$ |
| GyrA | $0.32 \pm 0.09$ | $0.57 \pm 0.06$ | $100.00 \pm 0.00$ | $97.33 \pm 4.62$ | $-85.8 \pm 24.1$ |
| PBP2a | $0.38 \pm 0.24$ | $0.61 \pm 0.05$ | $100.00 \pm 0.00$ | $99.66 \pm 0.59$ | $-100.3 \pm 14.9$ |
| VEGFA | $0.31 \pm 0.11$ | $0.69 \pm 0.11$ | $100.00 \pm 0.00$ | $97.19 \pm 2.58$ | $-115.9 \pm 33.8$ |
| FGF2 | $0.39 \pm 0.11$ | $0.67 \pm 0.12$ | $100.00 \pm 0.00$ | $87.95 \pm 19.11$ | $-95.5 \pm 15.7$ |
| ANGPT1 | $0.57 \pm 0.36$ | $0.50 \pm 0.08$ | $100.00 \pm 0.00$ | $91.83 \pm 14.16$ | $-81.0 \pm 27.5$ |

Table 3: **Generalization of evolved operators.** Each source starts with 100 parents; all satisfy an operator condition, and all 100 outputs from each round proceed to the next without selection. gain is relative to the original mean, and improved denotes the percentage above the preceding round's mean. Best and second-best results are **bold** and underlined.

| | original | R1 | | | R2 | | | R3 | | |
|---|---|---|---|---|---|---|---|---|---|---|
| Method | external score | external score | improved (%) | gain (%) | external score | improved (%) | gain (%) | external score | improved (%) | gain (%) |
| AMP-Designer Full | 0.4394 | 0.4492 | 48 | +2.2 | 0.4924 | 71 | +12.1 | 0.5257 | 62 | **+19.6** |
| HydrAMP | 0.5504 | **0.6051** | **94** | **+9.9** | **0.6371** | **81** | **+15.8** | **0.6504** | 59 | +18.2 |
| LLM Single-stage (GPT-6) | 0.5091 | 0.5038 | 44 | -1.0 | 0.5219 | 61 | +2.5 | 0.5317 | 51 | +4.4 |
| MAC-AMP | 0.4832 | 0.5214 | 59 | +7.9 | 0.5290 | 49 | +9.5 | 0.5433 | 54 | +12.4 |
| PDAgent | 0.4757 | 0.5029 | 63 | +5.7 | 0.5299 | 60 | +11.4 | 0.5545 | 58 | +16.5 |
| PepCraft + AMP-GAN v3 | 0.4505 | 0.4453 | 33 | -1.2 | 0.4793 | 67 | +6.4 | 0.5132 | **68** | +13.9 |
| PepGLAD | **0.6244** | 0.5834 | 28 | -6.6 | 0.6147 | 69 | -1.5 | 0.6398 | 62 | +2.5 |
| Pepti-Agent | 0.5470 | 0.5910 | 72 | +8.0 | 0.6167 | 64 | +12.7 | 0.6441 | 64 | +17.7 |
| PepZOO | 0.5611 | 0.5647 | 50 | +0.6 | 0.5818 | 64 | +3.7 | 0.5975 | 57 | +6.5 |
| AutoResearch | 0.4321 | 0.4159 | 28 | -3.7 | 0.4575 | 65 | +5.9 | 0.4602 | 59 | +6.5 |
| RFdiffusion | 0.4956 | 0.5265 | 69 | +6.2 | 0.5453 | 62 | +10.0 | 0.5645 | 66 | +13.9 |
| **avg baseline** | 0.5062 | 0.5190 | 53.5 | +2.56 | 0.5460 | 64.8 | +8.03 | 0.5659 | 60.0 | +12.03 |

they remain useful as candidate states change. This is especially clear for initially resistant sources: PepGLAD moves from $-6.6\%$ in R1 to $+2.5\%$ in R3, and PepCraft + AMP-GAN v3 from $-1.2\%$ to $+13.9\%$, showing that immediate response can underestimate the cumulative value of repeated conditional reuse. Higher predicted solubility and lower aggregation across all sources further show that these gains are accompanied by improved developability-related properties. Together, the results demonstrate successful harness evolution: trajectory experience is distilled into validated operators that transfer across sources and continue improving already edited candidates without further operator learning.

### 4.3 Checkpoint comparison and mutation effects

This matched checkpoint experiment isolates the capability learned through harness evolution: from identical parents and search states, SHAPE's evolved operators produce stronger one-step edits than AutoResearch reasoning/reflection. Figure 4A–C shows that SHAPE moves farther toward the reference basin, improves 2.87 scorers per edit versus 1.30, and raises eligibility from 31% to 58%. Higher family coverage and sequence diversity indicate that this advantage reflects broader usable editing behavior rather than optimization of a narrow mutation pattern. Panel D further shows structured, position- and property-dependent mutation preferences.

Panel E provides a PBP2a molecular-dynamics case study of the resulting structural behavior. The SHAPE edit retains the parental R10–D93 salt contact, whereas the AutoResearch edit loses the corresponding interaction, while also maintaining more interface contacts and lower late peptide C$\alpha$ RMSD. Together, these results demonstrate that harness evolution improves the search system itself: from identical starting states, SHAPE produces broader and more eligible edits.

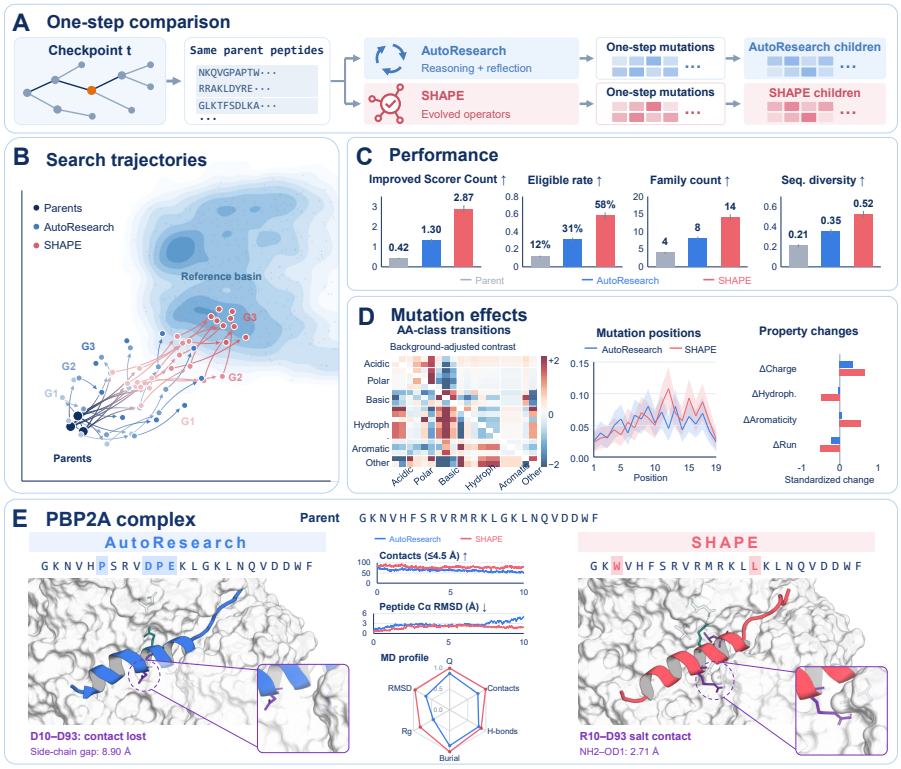



Figure 4: **Matched checkpoint comparison and mutation effects.** **(A)** One-step editing from identical checkpoints, parent peptides, and search states: AutoResearch uses reasoning and reflection, whereas SHAPE applies evolved operators; all children are evaluated with the same scoring pipeline. **(B)** Parent-to-child trajectories relative to the reference basin. **(C)** Improved-scorer count, eligibility rate, peptide-family count, and sequence diversity. **(D)** Background-adjusted amino-acid-class transitions, positional mutation frequencies, and sequence-property changes. **(E)** Representative PBP2a case showing sequence edits, the R10–D93 interaction, interface contacts, peptide C$\alpha$ RMSD, and MD profiles. Together, the panels compare the one-step editing behavior of evolved operators and reflection-guided search from matched starting states.

### 4.4 Search-path analysis

Checkpoint performance alone does not reveal whether the search system itself improves over time. We therefore trace a full search on a common scorer-budget axis, jointly recording best-so-far candidate quality, operator lineage, and operator usage to test whether harness evolution converts trajectory evidence into validated capabilities and increasingly reuses them.

Figure 5A shows the resulting change in search efficiency. SHAPE reaches the target score of 0.65 at 850 scorer calls, 200 calls before PaperResearch, and continues to 0.86 by the end of the budget, compared with 0.69 and 0.59 for PaperResearch and AutoResearch. Panels B and C explain this trajectory: operators formed early are selectively promoted, reused, merged, or retired, while validated families such as sequence-aware edits and successful-edit reuse become recurrent components of later search rather than remaining trajectory-specific discoveries.

The system-level shift is visualized in Panel D. From the early to late windows, new-family discovery falls from 3.00 to 0.75 while positive-effect gains per 100 scorer calls rise from 1.00 to 3.75.

Importantly, this reduced discovery coincides with continued improvement in Panel A, indicating a transition from acquiring new procedures to extracting more value from validated ones rather than search stagnation. Panel E further shows that gain is concentrated in repeatedly deployed capabilities, particularly sequence-aware edits and successful-edit reuse, while several other operator families also contribute. Together, the path analysis shows harness evolution changing *how* the budget is spent: early trajectory evidence becomes a reusable operator library, and later search increasingly exploits the capabilities that have demonstrated value.

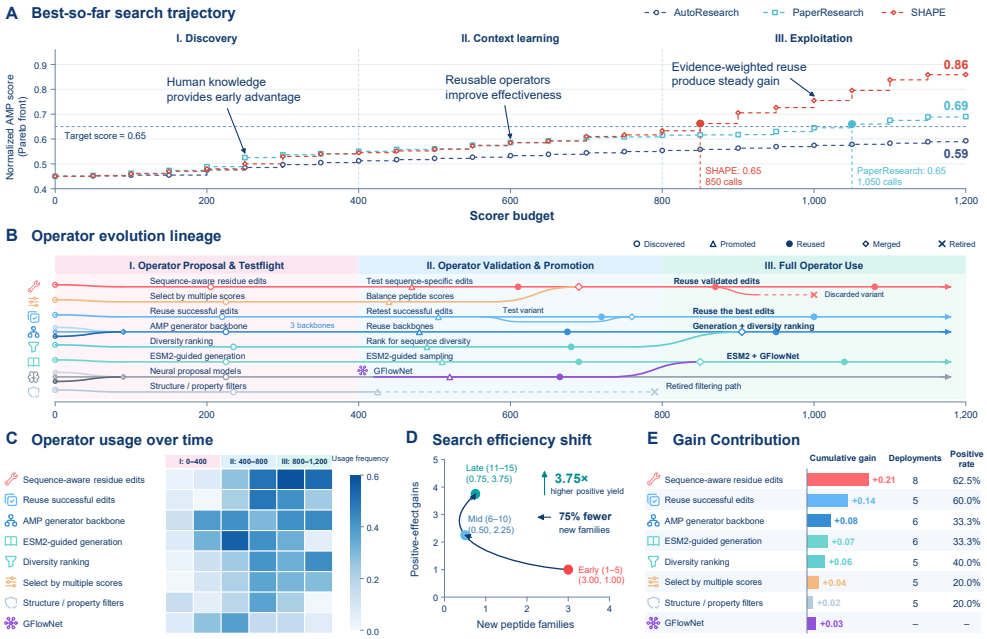



Figure 5: **Search trajectories and operator evolution.** **(A)** Best-so-far normalized AMP scores on the Pareto front across discovery, context learning, and exploitation. **(B)** Operator lineages on the same scorer-budget axis, showing discovery, promotion, reuse, merging, and retirement. **(C)** Usage frequencies for eight operator families across six temporal bins, two per search phase. **(D)** Early-, mid-, and late-window family discovery and positive-effect gains per 100 scorer calls. **(E)** Single-case composite gains by operator family, deployment counts, and positive rates. Together, these views connect search progress with operator refinement and increasingly productive reuse.

## 5 Conclusion

Designing deployable antimicrobial peptides is inherently multi-objective, requiring antimicrobial potency to be balanced with safety, AMP likelihood, and sequence diversity. Adding engagement with biologically relevant targets, including bacterial drug-resistance and tissue-repair-associated targets, further couples these objectives and makes the search substantially harder. We introduced SHAPE to address this setting by combining *target-aware peptide optimization* with *harness evolution*. Rather than keeping the search procedure fixed, SHAPE turns search experience into validated, reusable condition-action operators and incorporates them back into subsequent search under the same scorer budget. Across six targets, SHAPE consistently improves candidate quality over the strongest evaluated baseline, while the evolved operators produce more effective edits and remain useful when frozen and transferred to peptides from heterogeneous sources.

More broadly, SHAPE suggests a view of molecular design in which the optimizer itself is part of what is learned. As objectives become more heterogeneous and evaluations more costly, a fixed search strategy is unlikely to remain effective throughout optimization. Harness evolution provides a mechanism for adapting how search is performed, including which procedures to apply, in which contexts, and when to reuse or revise them. Extending this idea to richer structural signals, higher-

fidelity evaluation, and experimental feedback could move molecular design toward systems that improve both the molecules they discover and the search processes used to discover them.

## AI USE STATEMENT

AI systems were used in both the research workflow and manuscript preparation. Codex served as the agent backbone for SHAPE, and GPT-6 was included as a baseline for evaluating direct peptide generation. ChatGPT was used to assist with language polishing and presentation of the manuscript. All experimental design, analyses, results, and scientific conclusions were determined and verified by the authors.

## REFERENCES

Avner Adini, Irit Adini, Zai-Long Chi, Ratmir Derda, Amy E. Birsner, Benjamin D. Matthews, and Robert J. D'Amato. A novel strategy to enhance angiogenesis in vivo using the small VEGF-binding peptide PR1P. *Angiogenesis*, 20(3):399–408, 2017. doi: 10.1007/s10456-017-9556-7. URL https://doi.org/10.1007/s10456-017-9556-7.

Mehrad Ansari and Andrew D. White. Serverless prediction of peptide properties with recurrent neural networks. *Journal of Chemical Information and Modeling*, 63(8):2546–2553, 2023. doi: 10.1021/acs.jcim.2c01317.

Daehun Bae, Minsang Kim, Jiwon Seo, and Hojung Nam. AI-guided discovery and optimization of antimicrobial peptides through species-aware language model. *Briefings in Bioinformatics*, 26 (4):bbaf343, 2025. doi: 10.1093/bib/bbaf343. URL https://doi.org/10.1093/bib/bbaf343.

Emmanuel Bengio, Moksh Jain, Maksym Korablyov, Doina Precup, and Yoshua Bengio. Flow network based generative models for non-iterative diverse candidate generation. *Advances in Neural Information Processing Systems*, 34:27381–27394, 2021. URL https://papers.nips.cc/paper/2021/hash/e614f646836aaed9f89ce58e837e2310-Abstract.html.

Boltz developers. Boltz: Official repository for the boltz biomolecular interaction models. https://github.com/jwohlwend/boltz, 2026. Official software repository.

Houxu Chen, Achuth Chandrasekhar, and Amir Barati Farimani. Pepti-Agent: An AI agent for peptide design and optimization, 2026. URL https://arxiv.org/abs/2606.15422.

Chung-Hyun Cho et al. COMP-angiopoietin-1 promotes wound healing through enhanced angiogenesis, lymphangiogenesis, and blood flow in a diabetic mouse model. *Proceedings of the National Academy of Sciences*, 103(13):4946–4951, 2006. doi: 10.1073/pnas.0506352103. URL https://pubmed.ncbi.nlm.nih.gov/16543381/.

Payel Das, Tom Sercu, Kahini Wadhawan, Inkit Padhi, Sebastian Gehrmann, Flaviu Cipcigan, Vijil Chenthamarakshan, Hendrik Strobelt, Cicero Dos Santos, Pin-Yu Chen, Yi Yan Yang, Jeremy P. K. Tan, James Hedrick, Jason Crain, and Aleksandra Mojsilovic. Accelerated antimicrobial discovery via deep generative models and molecular dynamics simulations. *Nature Biomedical Engineering*, 5(6):613–623, 2021. doi: 10.1038/s41551-021-00689-x. URL https://www.nature.com/articles/s41551-021-00689-x.

Ethan Eschbach, Kristine Deibler, Deepa Korani, and Sebastian Swanson. Predicting peptide aggregation with protein language model embeddings. *Journal of Chemical Information and Modeling*, 66(9):5206–5216, 2026. doi: 10.1021/acs.jcim.6c00577.

X. Fu, Z. Shen, Y. Chen, J. Xie, Z. Guo, M. Zhang, and Z. Sheng. Randomised placebo-controlled trial of use of topical recombinant bovine basic fibroblast growth factor for second-degree burns. *The Lancet*, 352(9141):1661–1664, 1998. doi: 10.1016/S0140-6736(98)01260-4. URL https://pubmed.ncbi.nlm.nih.gov/9853438/.

R. D. Galiano, O. M. Tepper, C. R. Pelo, K. A. Bhatt, M. Callaghan, N. Bastidas, S. Bunting, H. G. Steinmetz, and G. C. Gurtner. Topical vascular endothelial growth factor accelerates diabetic wound healing through increased angiogenesis and by mobilizing and recruiting bone marrow-derived cells. *The American Journal of Pathology*, 164(6):1935–1947, 2004. doi: 10.1016/S0002-9440(10)63754-6. URL https://pubmed.ncbi.nlm.nih.gov/15161630/.

Robert E. W. Hancock and Hans-Georg Sahl. Antimicrobial and host-defense peptides as new anti-infective therapeutic strategies. *Nature Biotechnology*, 24(12):1551–1557, 2006. doi: 10.1038/nbt1267. URL https://doi.org/10.1038/nbt1267.

Steven Henikoff and Jorja G. Henikoff. Amino acid substitution matrices from protein blocks. *Proceedings of the National Academy of Sciences*, 89(22):10915–10919, 1992. doi: 10.1073/pnas.89.22.10915.

Jay Jung, Xiaohan Zhang, Shenghan Song, Mahmoud Sayedahmed, Chijian Xiang, Yunong Xu, Ahmed AbdelKhalek, Severin T. Schneebeli, Matthew J. Wargo, Jianing Li, and Safwan Wshah. Agentic discovery of non-canonical antimicrobial peptides with AMPGAN v3. In *GenBio Workshop at the International Conference on Machine Learning*, 2026. URL https://arxiv.org/abs/2606.17127.

Andrej Karpathy. autoresearch: AI agents running research on single-GPU nanochat training automatically. GitHub repository, 2026. URL https://github.com/karpathy/autoresearch. Accessed: 2026-09-21.

Xiangzhe Kong, Yinjun Jia, Wenbing Huang, and Yang Liu. Full-atom peptide design with geometric latent diffusion. In *Advances in Neural Information Processing Systems*, volume 37, 2024. URL https://proceedings.neurips.cc/paper_files/paper/2024/hash/88ad9774ffcb7a272457e9396f793a07-Abstract-Conference.html.

Chenkai Li, Darcy Sutherland, S. Austin Hammond, Chen Yang, Figali Taho, Lauren Bergman, Simon Houston, René L. Warren, Titus Wong, Linda M. N. Hoang, Caroline E. Cameron, Caren C. Helbing, and Inanc Birol. Amplify: attentive deep learning model for discovery of novel antimicrobial peptides effective against who priority pathogens. *BMC Genomics*, 23:77, 2022. doi: 10.1186/s12864-022-08310-4.

Zeming Lin, Halil Akin, Roshan Rao, Brian Hie, Zhongkai Zhu, Wenting Lu, Nikita Smetanin, Robert Verkuil, Ori Kabeli, Yaniv Shmueli, Allan dos Santos Costa, Maryam Fazel-Zarandi, Tom Sercu, Salvatore Candido, and Alexander Rives. Evolutionary-scale prediction of atomic-level protein structure with a language model. *Science*, 379(6637):1123–1130, 2023. doi: 10.1126/science.ade2574.

Fei Liu, Tong Xialiang, Mingxuan Yuan, Xi Lin, Fu Luo, Zhenkun Wang, Zhichao Lu, and Qingfu Zhang. Evolution of heuristics: Towards efficient automatic algorithm design using large language model. In *Proceedings of the 41st International Conference on Machine Learning*, volume 235 of *Proceedings of Machine Learning Research*, pp. 32201–32223. PMLR, 2024. URL https://proceedings.mlr.press/v235/liu24bs.html.

Xianliang Liu, Jiawei Luo, Xinyan Wang, Yang Zhang, and Junjie Chen. Directed evolution of antimicrobial peptides using multi-objective zeroth-order optimization. *Briefings in Bioinformatics*, 26(1):bbae715, 2025. doi: 10.1093/bib/bbae715. URL https://academic.oup.com/bib/article/26/1/bbae715/7952006.

Yuchen Ma, Yue Huang, Han Bao, Haomin Zhuang, Swadheen Shukla, Michel Galley, Xiangliang Zhang, and Stefan Feuerriegel. SkillGen: Verified inference-time agent skill synthesis, 2026. URL https://arxiv.org/abs/2605.10999.

Irina Massova and Peter A. Kollman. Combined molecular mechanical and continuum solvent approach (MM-PBSA/GBSA) to predict ligand binding. *Perspectives in Drug Discovery and Design*, 18(1):113–135, 2000. doi: 10.1023/A:1008763014207. URL https://doi.org/10.1023/A:1008763014207.

Neeloffer Mookherjee, Marilyn A. Anderson, Henk P. Haagsman, and Donald J. Davidson. Antimicrobial host defence peptides: functions and clinical potential. *Nature Reviews Drug Discovery*, 19(5):311–332, 2020. doi: 10.1038/s41573-019-0058-8. URL https://doi.org/10.1038/s41573-019-0058-8.

Song Ouyang, Zhijie Dong, Yong Luo, Kehua Su, Huangxuan Zhao, Miaojing Shi, and Bo Du. PDAgent: An LLM-driven autonomous agent framework towards in silico protein design via

directed mutation. International Conference on Machine Learning (ICML), 2026. URL https://icml.cc/virtual/2026/poster/63768. OpenReview record: https://openreview.net/forum?id=TwTvDVj6cH.

Anand Singh Rathore, Shubham Choudhury, Akanksha Arora, Purva Tijare, and Gajendra P. S. Raghava. Toxinpred 3.0: An improved method for predicting the toxicity of peptides. *Computers in Biology and Medicine*, 179:108926, 2024. doi: 10.1016/j.compbiomed.2024.108926.

Anand Singh Rathore, Nishant Kumar, Shubham Choudhury, Naman Kumar Mehta, and Gajendra P. S. Raghava. Prediction of hemolytic peptides and their hemolytic concentration. *Communications Biology*, 8:176, 2025. doi: 10.1038/s42003-025-07615-w. URL https://www.nature.com/articles/s42003-025-07615-w.

Barak Raveh, Nir London, and Ora Schueler-Furman. Sub-angstrom modeling of complexes between flexible peptides and globular proteins. *Proteins: Structure, Function, and Bioinformatics*, 78(9):2029–2040, 2010. doi: 10.1002/prot.22716.

RosettaCommons. InterfaceAnalyzer documentation. Rosetta software documentation, 2016. URL https://docs.rosettacommons.org/docs/latest/application_documentation/analysis/interface-analyzer. Accessed: 2026-09-21.

Célio Dias Santos-Júnior, Shaojun Pan, Xing-Ming Zhao, and Luis Pedro Coelho. Macrel: antimicrobial peptide screening in genomes and metagenomes. *PeerJ*, 8:e10555, 2020. doi: 10.7717/peerj.10555.

Noah Shinn, Federico Cassano, Ashwin Gopinath, Karthik Narasimhan, and Shunyu Yao. Reflexion: Language agents with verbal reinforcement learning. In *Advances in Neural Information Processing Systems*, volume 36, 2023. URL https://proceedings.neurips.cc/paper_files/paper/2023/hash/1b44b878bb782e6954cd888628510e90-Abstract-Conference.html.

Niranjan Srinivas, Andreas Krause, Sham M. Kakade, and Matthias Seeger. Gaussian process optimization in the bandit setting: No regret and experimental design. In *Proceedings of the 27th International Conference on Machine Learning*, pp. 1015–1022, 2010. URL https://icml.cc/Conferences/2010/abstracts.html.

Paulina Szymczak, Marcin Możejko, Tomasz Grzegorzek, Radosław Jurczak, Marta Bauer, Damian Neubauer, Karol Sikora, Michał Michalski, Jacek Sroka, Piotr Setny, Wojciech Kamysz, and Ewa Szczurek. Discovering highly potent antimicrobial peptides with deep generative model HydrAMP. *Nature Communications*, 14:1453, 2023. doi: 10.1038/s41467-023-36994-z. URL https://www.nature.com/articles/s41467-023-36994-z.

Raj Kumar Thapa, Dzung B. Diep, and Hanne Hjorth Tønnesen. Topical antimicrobial peptide formulations for wound healing: current developments and future prospects. *Acta Biomaterialia*, 103:52–67, 2020. doi: 10.1016/j.actbio.2019.12.025. URL https://pubmed.ncbi.nlm.nih.gov/31874224/.

Marcelo D. T. Torres, Yimeng Zeng, Fangping Wan, Natalie Maus, Jacob Gardner, and Cesar de la Fuente-Nunez. A generative artificial intelligence approach for peptide antibiotic optimization. *Nature Machine Intelligence*, 8:841–856, 2026. doi: 10.1038/s42256-026-01237-5. URL https://www.nature.com/articles/s42256-026-01237-5.

Daniel Veltri, Uday Kamath, and Amarda Shehu. Deep learning improves antimicrobial peptide recognition. *Bioinformatics*, 34(16):2740–2747, 2018. doi: 10.1093/bioinformatics/bty179.

B. Wang, P. Lin, Y. Zhong, X. Tan, Y. Shen, Y. Huang, et al. Explainable deep learning and virtual evolution identifies antimicrobial peptides with activity against multidrug-resistant human pathogens. *Nature Microbiology*, 10(2):332–347, 2025a. doi: 10.1038/s41564-024-01907-3. URL https://www.nature.com/articles/s41564-024-01907-3.

Jike Wang, Jianwen Feng, Yu Kang, Peichen Pan, Jingxuan Ge, Yan Wang, Mingyang Wang, Zhenxing Wu, Xingcai Zhang, Jiameng Yu, Xujun Zhang, Tianyue Wang, Lirong Wen, Guangning Yan, Yafeng Deng, Hui Shi, Chang-Yu Hsieh, Zhihui Jiang, and Tingjun Hou. Discovery of antimicrobial peptides with notable antibacterial potency by an LLM-based foundation model. *Science Advances*, 11(10):eads8932, 2025b. doi: 10.1126/sciadv.ads8932. URL https://www.science.org/doi/10.1126/sciadv.ads8932.

Joseph L. Watson, David Juergens, Nathaniel R. Bennett, Brian L. Trippe, Jason Yim, Helen E. Eisenach, Woody Ahern, Andrew J. Borst, Robert J. Ragotte, Lukas F. Milles, et al. De novo design of protein structure and function with RFdiffusion. *Nature*, 620:1089–1100, 2023. doi: 10.1038/s41586-023-06415-8. URL https://www.nature.com/articles/s41586-023-06415-8.

Tianxin Wei, Zhan Shi, Minhua Lin, Bing He, Zewen Liu, Yisi Sang, Yuanchen Bei, Xuying Ning, Jiaru Zou, Ting-Wei Li, Xiao Lin, Yanjun Zhao, Chi Wang, Benoit Dumoulin, Dakuo Wang, Jingrui He, and Hanqing Lu. Evo-Harness: Context-to-harness skill compilation for self-evolving agents, 2026. URL https://arxiv.org/abs/2608.15071.

Michael Zasloff. Antimicrobial peptides of multicellular organisms. *Nature*, 415(6870):389–395, 2002. doi: 10.1038/415389a. URL https://pubmed.ncbi.nlm.nih.gov/11807545/.

Baoyun Zhao, He Wang, and Liang Zeng. G-LNS: Generative large neighborhood search for LLM-based automatic heuristic design, 2026. URL https://arxiv.org/abs/2602.08253.

Gen Zhou, Sugitha Janarthanan, Lianghong Chen, and Pingzhao Hu. MAC-AMP: A closed-loop multi-agent collaboration system for multi-objective antimicrobial peptide design. In *International Conference on Learning Representations*, 2026. URL https://proceedings.iclr.cc/paper_files/paper/2026/hash/0d230e3015b806e7c9cc9afd5ac2c4a9-Abstract-Conference.html.

# A Related Work

**AMP design and target-aware applications.** AMPs provide a basis for infection control and host-directed peptide therapeutics (Zasloff, 2002; Hancock & Sahl, 2006; Mookherjee et al., 2020). Generative design and predictor-guided optimization have yielded active peptides (Das et al., 2021; Szymczak et al., 2023; Wang et al., 2025a;b), while multi-objective methods coordinate competing design criteria (Liu et al., 2025; Torres et al., 2026). Wound-healing applications connect infection control with tissue repair (Mookherjee et al., 2020; Thapa et al., 2020); the VEGF-A-binding peptide PR1P enhances angiogenesis and accelerates wound closure in mice (Adini et al., 2017). This motivates coordinating AMP properties with bacterial and repair-associated target interactions.

**Traditional and property-guided peptide generation.** Deep generative models have enabled controlled peptide sampling and computational screening (Das et al., 2021). HydrAMP performs conditional latent generation for new peptides and analogues (Szymczak et al., 2023), while PepZOO combines learned representations with predictor-guided multi-objective zeroth-order optimization (Liu et al., 2025). Structure-conditioned design extends these approaches to three-dimensional contexts: RFdiffusion generates protein backbones (Watson et al., 2023), and PepGLAD performs full-atom peptide generation conditioned on target binding sites (Kong et al., 2024). Related peptide-optimization systems also combine generative priors with property or interaction feedback (Torres et al., 2026). AMPGAN v3 provides a multi-objective conditional generator that is used within the PepCraft agent workflow (Jung et al., 2026). Together, these methods provide sequence-, property-, and structure-conditioned proposal mechanisms represented in our benchmark.

**LLM-based peptide design.** AMP-Designer Full uses an LLM-based foundation model with prompt tuning and property-guided optimization (Wang et al., 2025b). Our LLM Single-stage (GPT-6) baseline tests direct prompted peptide generation. More generally, language models can serve both as generators and as controllers that reason over candidate histories, model outputs, and external tools.

**Agent-based design.** Agentic peptide-design systems introduce iterative decisions around generation and evaluation. Pepti-Agent coordinates peptide generation, prediction, and mutation tools (Chen et al., 2026); PDAgent performs agent-directed protein mutation and is adapted here to AMP search (Ouyang et al., 2026); and MAC-AMP couples multi-agent feedback with reward and generator updates (Zhou et al., 2026). PepCraft coordinates generation, filtering, and verification around AMPGAN v3 (Jung et al., 2026). AutoResearch provides a general autonomous research workflow that we adapt to iterative peptide search (Karpathy, 2026). These systems motivate treating the search procedure itself as an object that can accumulate useful experience.

**Operator learning and harness evolution.** Automatic algorithm-design methods provide related mechanisms for improving search procedures. EoH evolves heuristic descriptions and executable code (Liu et al., 2024), while G-LNS evolves operators for large-neighborhood search (Zhao et al., 2026). Reflexion retains linguistic feedback across attempts (Shinn et al., 2023); SkillGen synthesizes and evaluates reusable agent skills (Ma et al., 2026); and Evo-Harness compiles experience into reusable harness components (Wei et al., 2026). SHAPE brings this idea to target-aware peptide design by distilling trajectory evidence into conditional editing operators, evaluating their incremental value against the current harness, and incorporating validated operators into later search under the same scorer budget.

# B SHAPE Implementation Details

## B.1 Agent Control and an Instantiated Search Operator

Algorithm 1 summarizes SHAPE's agent control flow and expands one operator invocation. The left panel coordinates peptide search, operator development, and evaluation. The right panel instantiates an ESM2-based lower-confidence-bound proposal operator.

**Algorithm 1** SHAPE agent actions and an example search operator

**Control flow.** S/O/E denote the Search, Operator, and Evaluator Agents. $s$ contains the target, scored archive $P$, frontier, and search context; $\mathcal{E}$ is shared evidence and $H$ the harness initialized by reflection-guided routine $\mathcal{T}$. $B$, $b_{\text{rem}}$, and $b$ are total, remaining, and assigned scorer allowances, with $b$ per probe branch. For proposed operator $o$, $\widehat{\Delta}_o$ is its mean paired archive-quality gain. $\tau$ is the highest-gain applicable operator's action, with fixed tie-breaking and direct mutation as fallback. $U$, $\mathcal{X}$, and $\mathcal{X}_p$ are proposed, scored, and pooled probe peptide sets; $k$ counts complete-block calls, including nested queries. $P^*$ and $H_T$ are the final archive and harness at step $T$.

**Example operator.** $\phi$ produces ESM2-8M embeddings; GP is a Gaussian process fitted to AMP-READ observations $\mathcal{D}_{\text{train}}$ with log-MIC targets $y_{\text{logMIC}}$ (Lin et al., 2023; Srinivas et al., 2010). For parent peptide $p$, $I_p$ marks protected positions, $\text{Mut}_1$ generates single-residue edits, and $\mathcal{N}_p$ is the filtered neighborhood. $\mathcal{V}$ and $Q$ exclude history/training sequences and earlier-branch selections; `history` and seq supply recorded and training sequences. For candidate $x$, $d_x$ is its Euclidean embedding distance from $p$, $m_x$ and $\sigma_x$ its predictive mean and standard deviation, and $\ell_x$ its lower-confidence-bound (LCB) score. $r_p$ is the parent-specific radius and $R$ stores retained scores paired with sequences. Sorting uses score then sequence; Unique returns up to four distinct, unexcluded proposals.

**SHAPE control flow**

Initialize $s, \mathcal{E}$; $H \leftarrow \mathcal{T}, b_{\text{rem}} \leftarrow B$
**while** budget covers the next block:
  **if** develop$(s, \mathcal{E}, b_{\text{rem}})$:
    O: $o \leftarrow \text{LearnOperator}(s, \mathcal{E})$
    E: $(\widehat{\Delta}_o, \mathcal{X}_p, k) \leftarrow \text{Probe}_b(H, o; s)$
    $b_{\text{rem}} \leftarrow b_{\text{rem}} - k$; update$(s, \mathcal{E}; \mathcal{X}_p)$
    $H \leftarrow \text{Assemble}(H, o, \widehat{\Delta}_o)$; **end if**
  **if** no search block fits $b_{\text{rem}}$: **break**
  S: $\tau \leftarrow \text{Route}(H, s, \mathcal{E})$
  **S:** $U \leftarrow \tau(s, b)$
  E: $(\mathcal{X}, k) \leftarrow \text{CheckScore}(U)$
  $b_{\text{rem}} \leftarrow b_{\text{rem}} - k$; S: update$(s, \mathcal{E}; \mathcal{X})$
**end while**
**return** $(P, H)$ as $(P^*, H_T)$

**Example of a Proposed Operator**

$\phi \leftarrow$ ESM2-8M; $R \leftarrow \varnothing$
$\text{GP} \leftarrow \text{Fit}(\phi(\mathcal{D}_{\text{train}}), y_{\text{logMIC}})$
$\mathcal{V} \leftarrow \text{history} \cup \text{seq}(\mathcal{D}_{\text{train}})$
$Q \leftarrow$ earlier-branch selections
**for** parent $p$ with fixed positions $I_p$:
  $\mathcal{N}_p \leftarrow \text{Mut}_1(p; I_p) \setminus \mathcal{V}$; skip if empty
  $d_x \leftarrow \|\phi(x) - \phi(p)\|_2, \quad x \in \mathcal{N}_p$
  $r_p \leftarrow \max\left(\text{median}_{x \in \mathcal{N}_p} d_x, 10^{-9}\right)$
  $(m_x, \sigma_x) \leftarrow \text{GP}(\phi(x))$
  $\ell_x \leftarrow m_x - 1.5\sigma_x$
  $R \leftarrow R \cup \{(\ell_x, x) : x \in \mathcal{N}_p,\ d_x \le r_p\}$
**end for**
$U \leftarrow \text{Unique}\left(\text{Sort}_{\uparrow(\ell,x)} R; \mathcal{V} \cup Q\right)$
**return** $U$

### B.2 Execution Interfaces

ROUTE identifies active operators whose conditions match the current search state and selects the operator with the highest stored validation gain, using a fixed tie order. Reflection-guided mutation supplies the fallback when no active operator applies.

PROBE compares the current harness with an operator-enabled version from the same starting state under equal scorer allowances. It returns the mean paired gain, pooled evaluated candidates, and consumed scorer calls. ASSEMBLE uses the validation result to activate, retain, or mark the operator dormant, updating the procedures available to subsequent search. After each search or probe block, UPDATE incorporates scored candidates into the archive, working frontier, and shared evidence ledger.

## C Task Semantics and Reporting Cohorts

**Targets and endpoints.** AceA, GyrA, and PBP2a provide bacterial interaction targets, while VEGFA, FGF2, and ANGPT1 provide repair-associated interaction targets. The repair setting is motivated by studies of growth-factor-mediated wound healing (Fu et al., 1998; Galiano et al., 2004; Cho et al., 2006). Candidate evaluation combines peptide potency, safety, AMP likelihood, target interaction, and sequence diversity.

**Scorer interpretation.** LLAMP and AMP-READ report predicted MIC in $\mu$M. ToxinPred3 and HemoPI2 provide toxicity and hemolysis outputs (Rathore et al., 2024; 2025); AMPlify and Macrel provide AMP probabilities (Li et al., 2022; Santos-Júnior et al., 2020); and Rosetta reports target-

specific interaction energy in REU. Mean pairwise Levenshtein distance supplies the sequence-diversity statistic.

### C.1 Archive Utility and Composite Score

Matched operator validation evaluates both branches from the same starting archive using the archive objective in Section 3.1. For a search block returning scored candidates $\mathcal{X}$, its archive outcome is the quality after adding $\mathcal{X}$ to that starting archive. Paired differences are averaged to obtain the stored operator gain used by subsequent routing. Benchmark Composite Scores follow Appendix C.2.

### C.2 Benchmark Normalization and Aggregation

Let $L, R, A, T, H_{\text{hemo}}, M, E, D$ denote the direction-aligned benchmark components for LLAMP, AMP-READ, AMPlify, ToxinPred3, HemoPI2, Macrel, Rosetta binding, and sequence diversity. The benchmark composite score gives all eight components equal weight:

$$C = \text{mean}\left(L, R, A, T, H_{\text{hemo}}, M, E, D\right). \tag{3}$$

For each reporting panel $z$, the eleven baselines form the calibration set $\mathcal{B}$. For each non-binding metric $k$, let $q_{mzk}$ be the reported value of method $m$, with baseline mean $\mu_{zk}$ and sample standard deviation $s_{zk}$. Its normalized direction-aligned score is

$$S_{mzk} = \Phi\left(d_k \frac{q_{mzk} - \mu_{zk}}{s_{zk}}\right), \tag{4}$$

where $\Phi$ is the standard normal CDF. The direction is $d_k = +1$ for larger-is-better quantities and $d_k = -1$ for lower-is-better quantities. A zero baseline standard deviation maps to $0.5$. The same panel-specific calibration is applied to all methods. Sequence diversity uses mean pairwise Levenshtein distance before this normalization.

**Reported cohorts.** All eleven baselines and SHAPE contribute 300 peptides per method, with the final Top 50 for each of six targets, yielding $12 \times 6 \times 50 = 3{,}600$ evaluation results. Overall results pool the six target cohorts for each method; target-specific rows summarize 50 peptides. Table 1 uses the pooled cohorts.

**Relative gain.** The Overall improvement over the strongest baseline is

$$\text{gain} = 100\left(\frac{C_{\text{SHAPE}}}{\max_{b \in \mathcal{B}} C_b} - 1\right)\%. \tag{5}$$

This gives **23.99%** Overall. The corresponding target-specific gains are $26.55\%$ (AceA), $21.81\%$ (GyrA), $22.28\%$ (PBP2a), $22.59\%$ (VEGFA), $21.43\%$ (FGF2), and $23.58\%$ (ANGPT1).

### C.3 Improvement Success Rate

Each trial is a logged proposal event with a child peptide and its parent. A proposal is successful when the child is at least as good as its parent on every direction-aligned candidate objective and strictly improves at least one. Let $I_{m,i} = 1$ for a successful proposal and 0 otherwise. For the $N_m$ proposals generated by method $m$,

$$\text{ISR}_m = \frac{\sum_{i=1}^{N_m} I_{m,i}}{N_m} \times 100\%. \tag{6}$$

SHAPE records $N_{\text{SHAPE}} = 778$ proposal events and an ISR of $42\%$; AutoResearch records $N_{\text{AutoResearch}} = 1{,}200$ proposal events and an ISR of $27\%$.

### C.4 Evaluation Protocol

**Evaluators and score directions.** AMP-likelihood evaluation uses AMPlify and AMP Scanner v2 where specified (Li et al., 2022; Veltri et al., 2018). Safety and developability evaluation uses ToxinPred3, Macrel, PALM, and the Peptide Dashboard sol-rnn predictor in the experiments where they

are reported (Rathore et al., 2024; Santos-Júnior et al., 2020; Eschbach et al., 2026; Ansari & White, 2023). Structural analyses use the pinned Boltz source, Rosetta FlexPepDock, and MM/GBSA where indicated (Boltz developers, 2026; Raveh et al., 2010; Massova & Kollman, 2000). ESM-2 supplies sequence embeddings for search components that use protein-language-model representations (Lin et al., 2023).

**Matched evaluation accounting.** Treatment and control start from the same search state and receive equal scorer-call allowances. Each matched block contributes the difference between treatment and control archive quality, and these differences are averaged across the probe pairs. Scored candidate calls determine budget consumption.

**Molecular-dynamics validation.** We analyze the top 15 peptides per target, giving 90 complexes across six targets. Table 2 reports means and standard deviations for interface RMSD, native-contact retention, hydrogen-bond and salt-bridge occupancies, and endpoint MM/GBSA energies.

### C.5 Method Budgets and Scorer-Call Accounting

Each agent-based method receives $1{,}200$ scorer calls per run. One scorer call is one candidate–target evaluation request to the common scoring workflow; a batch of $q$ candidates uses $q$ calls. SHAPE's peptide-search and operator-evolution paths share the same $1{,}200$-call budget, including matched probes and nested scorer queries. The five Rosetta decoys are internal replicates of one candidate's structural endpoint.

| Method | Scorer-budget regime | Final outputs |
|---|---|---|
| Pepti-Agent | 1,200/run | 50/target |
| AMP-Designer Full | Standard generation | 50/target |
| LLM Single-stage (GPT-6) | Standard generation | 50/target |
| PepCraft + AMP-GAN v3 | 1,200/run | 50/target |
| PDAgent | 1,200/run | 50/target |
| MAC-AMP | 1,200/run | 50/target |
| AutoResearch | 1,200/run | 50/target |
| PepZOO | Standard generation | 50/target |
| HydrAMP | Standard generation | 50/target |
| RFdiffusion | Standard generation | 50/target |
| PepGLAD | Standard generation | 50/target |
| SHAPE | 1,200/run, shared paths | 50/target |

Every method contributes its final Top 50 outputs for each of six targets, or 300 peptides per method. LLM and machine-learning generators follow their standard generation procedures and then enter the common evaluation workflow. PaperResearch, used in Figure 5, also receives $1{,}200$ scorer calls.

**Target-agnostic methods.** For methods without a protein-target input, peptide-level properties are shared across target assignments and target binding is evaluated against the assigned protein. This produces the repeated non-binding summaries for Pepti-Agent, AMP-Designer Full, PepCraft, PDAgent, and MAC-AMP in the six-target tables.

**PaperResearch.** PaperResearch mines peptide- and AMP-related literature and uses retrieved design knowledge in candidate proposal and reflection. It is compared with AutoResearch and SHAPE under the same scorer-call allowance in Figure 5.

### C.6 Target-Specific Rosetta Evaluation

Each candidate keeps its peptide sequence, target identity, and fixed base seed throughout structural evaluation. Boltz predicts a receptor–peptide complex with receptor chain A and peptide chain B using one sample, three recycling steps, 200 sampling steps, `use_msa_server=false`, `use_potentials=true`, and `no_kernels=true`.

The complex undergoes one FlexPepDock prepack (`-flexpep_prepack`). Five refinement decoys are then generated from the same prepacked structure using base-seed off-

sets 1 through 5, `-pep_refine`, `-ex1`, `-ex2aro`, `-use_input_sc`, and `ref2015`. Rosetta InterfaceAnalyzer evaluates interface `A_B` with `ref2015`, `pack_input=false`, and `pack_separated=false`. The candidate-level binding score is the median `dG_separated` across the five decoys in REU. The same workflow is used for benchmark candidates and frozen-transfer parents and children.

# D Baseline Sources

The eleven evaluated baselines span peptide generators, structure-conditioned design, direct LLM generation, and agent-based search. Pepti-Agent (Chen et al., 2026), AMP-Designer (Wang et al., 2025b), PepCraft + AMP-GAN v3 (Jung et al., 2026), PDAgent (Ouyang et al., 2026), MAC-AMP (Zhou et al., 2026), and the AutoResearch adaptation (Karpathy, 2026) cover tool-using and iterative workflows; LLM Single-stage (GPT-6) supplies direct prompted proposals. PepZOO (Liu et al., 2025), HydrAMP (Szymczak et al., 2023), RFdiffusion (Watson et al., 2023), and PepGLAD (Kong et al., 2024) provide optimization, sequence-generation, and structure-conditioned baselines. Each method contributes 300 final peptides, with Top 50 outputs for each of six targets. Budget accounting is specified in Appendix C.5, and cohort aggregation and normalization follow Appendix C.2.

# E Checkpoint Analysis and Structural Case

Figure 4 compares AutoResearch and SHAPE from identical parent peptides and search states, applying one editing step and the same scoring pipeline. Mutation profiles connect changes in sequence properties to candidate quality and family coverage. For the PBP2a case, parent and child complexes are compared through residue contacts and molecular-dynamics trajectories, linking the sequence edits to interface retention and peptide stability.

## E.1 Concrete Search-Operator Implementation

The example operator in Algorithm 1 enumerates substitutions at the first six editable positions, using up to eight deterministic cyclic amino-acid substitutions per position while preserving protected residues. After removing history and training sequences, ESM2-8M embeddings define a parent-specific neighborhood using the median candidate distance as its radius.

A Gaussian process fitted to AMP-READ log-MIC observations provides predictions and uncertainty estimates. Candidates within the embedding radius are pooled across parents and ordered by ascending LCB score, with sequence-based tie-breaking. After excluding earlier-branch selections and duplicates, the operator returns up to four proposals for evaluation.

## E.2 From Feedback to Conditional Fragment-Recombination Operators

Cross-front fragment recombination illustrates how search feedback changes an operator's applicability and editing action. The initial program transferred fixed middle or quartile segments between peptides drawn from complementary property fronts.

Multi-property feedback produced a conditional successor. Parents were pooled from the Macrel, LLAMP, and AMP-READ activity fronts, with editing triggered below their upper-median Macrel AMP probability, 0.505 in the recorded execution. Donors came from AMP-activity, low-toxicity, and stability fronts. The action transferred contiguous segments of three to seven residues while preserving the parent sequence outside the window. Candidate generation capped hydrophobic runs at two and restricted the charge proxy to 2 through 10. The charge proxy counts lysine and arginine residues minus aspartate and glutamate residues.

Each invocation performs one editing generation and returns selected candidates for scoring. The validated program was compiled into the search harness and invoked again in the next round, carrying the learned conditions and editing procedure into subsequent search.

# F GENERALIZATION OF EVOLVED OPERATORS: EXPERIMENTAL SETUP

**Complete-cohort transfer.** Generalization of evolved operators uses 100 original peptides from each of eleven generator and agent sources. After SHAPE forms and validates its operator library, that library is frozen and applied for three rounds without further operator learning. Each parent is routed to an operator whose condition matches its current state. Condition coverage is 100% for every source cohort, and every evolved output becomes a parent in the next round, giving 100% round-to-round carry-forward.

**External evaluation panel.** The external panel uses four released predictors and two components from the PepGLAD evaluation pipeline. All six components are direction-aligned so that larger values are preferred.

- **AMP Scanner v2:** AMP probability from the released original-paper network and preprocessing, using `OriginalPaper_081917_FULL_MODEL.h5` and pinned source `933052e` (Veltri et al., 2018). Source: AMP Scanner v2.
- **Macrel:** hemolysis probability from version 1.6.1 and the released `Hemo.onnx.gz` model, with lower hemolysis preferred (Santos-Júnior et al., 2020). Source: Macrel.
- **PALM:** sequence-level amyloidogenic aggregation score from the mean of the five released folds in v1.1.0, source `f907ec3`, using ESM2 embeddings; lower aggregation is preferred (Eschbach et al., 2026). Source: PALM.
- **Peptide Dashboard sol-rnn:** released soluble-class score for the E. coli expression endpoint, using the published `sol-rnn` weights and residue-frequency encoding, source `b02d011`; higher solubility is preferred (Ansari & White, 2023). Source: Peptide Dashboard.
- **PepGLAD diversity and Rosetta score:** BLOSUM62-based sequence-clustering diversity, computed as the number of clusters divided by cohort size, and the Rosetta-based score from the released PepGLAD evaluation pipeline, source `bad015c` (Henikoff & Henikoff, 1992; Kong et al., 2024). Higher diversity is preferred. The Rosetta endpoint is converted to the normalized, direction-aligned PepGLAD Rosetta score $E_{\mathrm{PG}}$, with larger values preferred, using the same calibration rule across all sources and rounds. Source: PepGLAD.

**External composite score.** For peptide $x$, let $A(x)$ be AMP Scanner probability, $H(x)$ Macrel hemolysis probability, $P(x)$ PALM aggregation score, and $Q(x)$ Dashboard sol-rnn score. Let $D_{m,r}$ be PepGLAD clustering diversity for source $m$ in round $r$, and define $E_{\mathrm{PG}}(x) = \mathrm{clip}(-\Delta G_{\mathrm{PG}}(x)/100, 0, 1)$ from the PepGLAD Rosetta interface energy in REU, with the same scaling across sources and rounds. The six equally weighted components define

$$C_{m,r}(x) = \mathrm{mean}(A(x),\, 1 - H(x),\, 1 - P(x),\, Q(x),\, D_{m,r},\, E_{\mathrm{PG}}(x))\,. \tag{7}$$

**Reported statistics and Avg baseline.** For each source and round, the external score is the cohort mean of $C_{m,r}(x)$. The cumulative gain compares the round mean $S_{m,r}$ with the same source's original mean $S_{m,0}$:

$$G_{m,r} = 100\left(\frac{S_{m,r}}{S_{m,0}} - 1\right). \tag{8}$$

The reported *improved* percentage is the fraction of evolved peptides whose external composite exceeds the preceding round's corresponding parents. The Avg baseline row averages the source-level results with equal weight.