

000 001 002 003 004 005 006 007 008 009 010 011 012 013 014 015 016 017 018 019 020 021 022 023 024 025 026 027 028 029 030 031 032 033 034 035 036 037 038 039 040 041 042 043 044 045 046 047 048 049 050 051 052 053 PEPCOMPASS: NAVIGATING PEPTIDE EMBEDDING SPACES USING RIEMANNIAN GEOMETRY

Anonymous authors

Paper under double-blind review

ABSTRACT

Antimicrobial peptide discovery is challenged by the astronomical size of peptide space and the relative scarcity of active peptides. Generative models provide continuous latent “maps” of peptide space, but conventionally ignore decoder-induced geometry and rely on flat Euclidean metrics, rendering exploration and optimization distorted and inefficient. Prior manifold-based remedies assume fixed intrinsic dimensionality, which critically fails in practice for peptide data. Here, we introduce **PepCompass**, a geometry-aware framework for peptide exploration and optimization. At its core, we define a **Union of κ -Stable Riemannian Manifolds** \mathbb{M}^κ , a family of decoder-induced manifolds that captures local geometry while ensuring computational stability. We propose two local exploration methods: **Second-Order Riemannian Brownian Efficient Sampling**, which provides a convergent second-order approximation to Riemannian Brownian motion, and **Mutation Enumeration in Tangent Space**, which reinterprets tangent directions as discrete amino-acid substitutions. Combining these yields Local Enumeration Bayesian Optimization (**LE-BO**), an efficient algorithm for local activity optimization. Finally, we introduce Potential-minimizing Geodesic Search (**PoGS**), which interpolates between prototype embeddings along property-enriched geodesics, biasing discovery toward seeds, i.e. peptides with favorable activity. *In-vitro* validation confirms the effectiveness of PepCompass: PoGS yields four novel seeds, and subsequent optimization with LE-BO discovers 25 highly active peptides with broad-spectrum activity, including against resistant bacterial strains. These results demonstrate that geometry-informed exploration provides a powerful new paradigm for antimicrobial peptide design.

1 INTRODUCTION

Efficient exploration of peptide space is notoriously difficult. At the global level, there are more than 3.3×10^{32} combinatorially possible amino acid sequences of length at most 25. At the local level, each peptide of length 25 has nearly 1000 neighbors within an edit radius of one. Moreover, only a small fraction of amino acid sequences correspond to antimicrobial peptides (AMPs), which have high activity against bacteria Szymczak & Szczerk (2023); Szymczak et al. (2025). This extreme combinatorial complexity renders AMP discovery by brute-force exploration intractable. To address this challenge, we turn to one of the most prolific inventions of humankind: maps.

Since the dawn of civilization, maps have provided a structured way to support both local and global navigation, driving scientific discovery. In modern machine learning, latent-space generative models such as VAEs, GANs, WAEs, and normalizing flows (Bond-Taylor et al., 2021) enable building continuous latent representations—maps—of peptides. Such maps have already facilitated the discovery of promising new AMPs (Szymczak et al., 2023; Oort et al., 2021; Wang et al., 2022; Das et al., 2020). The standard workflow assumes that once the model is trained, its latent space together with the decoder properly models a set of valid, synthetizable peptides. The latent space is typically chosen to be \mathbb{R}^d with a flat Euclidean metric, enabling direct application of existing exploration and optimization algorithms. However, such flat representations suffer from a significant flaw: they ignore the differential geometry induced by the decoder, leading to distortions in distances.

Typical approaches attempting to circumvent this problem assume the *manifold hypothesis* (Bengio et al., 2012) and use the pullback metric (Arvanitidis et al., 2018). However, recent work (Loaiza-

Ganem et al., 2024; Brown et al., 2022; Wang & Wang, 2024) has shown that the manifold hypothesis does not withstand empirical scrutiny for image data, where sets of images are better modeled as unions or CW-complexes of manifolds with varying low dimensionality (we refer to Appendix A for Related Work). We show that peptide spaces suffer from a similar issue and introduce decoder-derived **Union of κ -Stable Riemannian Manifolds** \mathbb{M}^κ , which captures both the complex structure of peptide space and its local geometry, with computational stability controlled by a parameter κ . Intuitively, we cut the globally distorted map into a set of charts that enable efficient and distortion-free exploration and optimization.

Building upon the union-of-manifolds structure, we introduce **PepCompass**, a geometry-informed framework for peptide exploration and optimization at both global and local levels. At the *global level*, we propose Potential-minimizing Geodesic Search (**PoGS**), which models geodesic curves between known prototype peptides to identify promising seeds for further optimization (Figure 1A). We represent geodesics as energy-minimizing curves in peptide space and augment them with a potential function encoding antimicrobial activity. This biases exploration toward regions not only similar to the starting prototypes but also exhibiting higher activity. In doing so, our method extends standard local analogue search around a single prototype into a bi-prototype, controllable regime.

For *local* search on a single manifold from the family \mathbb{M}^κ , we designed two geometry-informed approaches: Second-Order Riemannian Brownian Efficient Sampling (**SORBES**) and Mutation Enumeration in Tangent Space (**MUTANG**). SORBES is a provably convergent, second-order approximation of the Riemannian Brownian motion (Schwarz et al., 2022; Herrmann et al., 2023), serving as a Riemannian analogue of local Gaussian search. MUTANG addresses the discrete nature of peptide space by reinterpreting the local tangent space not as continuous vectors but as discrete mutations, directly corresponding to amino-acid substitutions. This reinterpretation provides both interpretability and efficiency, enabling enumeration of a given peptide’s neighbours. We further combine SORBES and MUTANG into an iterative **Local Enumeration** procedure that densely populates the neighbourhood of a given peptide with valid, diverse neighbours. Finally, integrating this enumeration with a Bayesian optimization scheme yields an efficient Local Enumeration Bayesian Optimization procedure (**LE-BO**; see Figure 1B,C).

In vitro microbiological assays demonstrated unprecedented, 100% success rate of PepCompass in AMP optimization. Using PoGS we derived four peptide seeds, all of which showed significant antimicrobial activity. Further optimization of these seeds with LE-BO yielded 25/25 highly active peptides with broad-spectrum activity, including activity against multi-resistant bacterial strains. Code is available at <https://anonymous.4open.science/r/pep-compass-2ABF>.

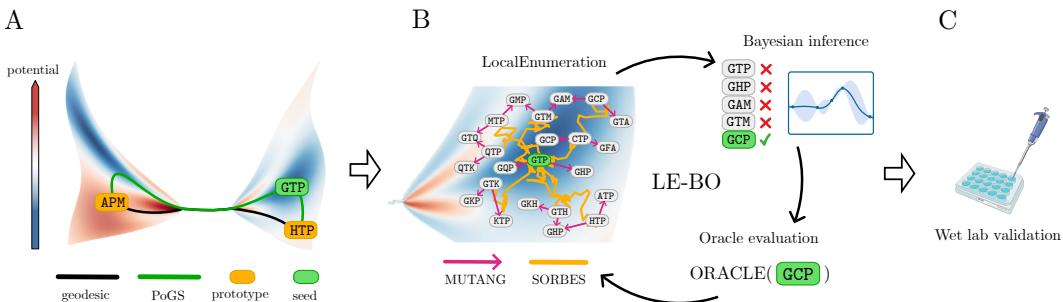


Figure 1: **PepCompass** overview.

2 METHODS

2.1 BACKGROUND

Let $\text{Dec}_\theta \in C^\infty : \mathcal{Z} \rightarrow \mathcal{X}$ be the deterministic decoder mapping latent vectors $z \in \mathbb{R}^d$ to position-factorized peptide probabilities $\text{Dec}_\theta(z) \in \mathbb{R}^{L \times A}$ (with L the maximum peptide length and A the size of the amino acid alphabet \mathcal{A} extended with a padding token pad , and $\mathbb{R}^{L \times A}$ - set of matrices of shape (L, A)). Define

$$X = \text{Dec}_\theta(\mathcal{Z}) \subset \mathbb{R}^{L \times A}, \quad p(z) = \text{argmax}(\text{Dec}_\theta(z), \text{dim} = 1) \in \mathcal{A}^L,$$

108 where $p(z)$ denotes the decoded peptide sequence. Intuitively, X is a continuous probabilistic
 109 approximation of the peptide space (with \mathcal{Z} as its map), from which the concrete peptides are decoded
 110 back using the p operator. For clarity, we drop explicit parameter dependence and simply write Dec
 111 instead of Dec_θ . We additionally introduce $\hat{\text{Dec}} : \mathcal{Z} \rightarrow \mathbb{R}^{LA}$ and $\hat{X} = \hat{\text{Dec}}(\mathcal{Z})$, i.e. the flattened
 112 versions of Dec and X .

113 A common approach is to equip \mathcal{Z} with the standard Euclidean inner product $\langle \cdot, \cdot \rangle_{\mathbb{R}^d}$, and use the
 114 associated Euclidean distance as a base for exploration. However, this ignores the geometry induced
 115 by the decoder, which, under the manifold hypothesis, can be accounted for using the *pullback*
 116 *metric* (Bengio et al., 2012; Arvanitidis et al., 2018).
 117

118 **Pullback metric** Under the manifold hypothesis, $\hat{\text{Dec}}$ is full rank (i.e. $\text{rank } J_{\hat{\text{Dec}}}(z) = d$ for all
 119 $z \in \mathcal{Z}$, where $J_{\hat{\text{Dec}}}(z) \in \mathbb{R}^{(LA) \times d}$ is the decoder Jacobian), and $(\mathcal{Z}, G_{\hat{\text{Dec}}})$ is a d -dimensional
 120 Riemannian manifold (do Carmo, 1992) where $G_{\hat{\text{Dec}}}$ is a natural, decoder-induced pullback metric
 121 on \mathcal{Z} given by:

$$122 \quad G_{\hat{\text{Dec}}}(z) = J_{\hat{\text{Dec}}}(z)^\top J_{\hat{\text{Dec}}}(z) \in \mathbb{R}^{d \times d}, \quad (1)$$

124 for $z \in \mathcal{Z}$. To simplify notation, let us set $G_{\text{Dec}} = G_{\hat{\text{Dec}}}$. Now, for a tangent space $T_z \mathcal{Z}$ at a point z
 125 and tangent vectors $u, v \in T_z \mathcal{Z}$,

$$126 \quad \langle u, v \rangle_z^{\text{Dec}} = u^\top G_{\text{Dec}}(z) v. \quad (2)$$

127 Note that $\hat{\text{Dec}} : (\mathcal{Z}, G_{\text{Dec}}) \rightarrow (\hat{X}, \langle \cdot, \cdot \rangle_{\mathbb{R}^{LA}})$ is an isometric diffeomorphism.
 128

129 **When the manifold hypothesis fails: union of manifolds** Previous work demonstrates that the
 130 manifold hypothesis often fails for complex data such as images (Loaiza-Ganem et al., 2024; Brown
 131 et al., 2022; Wang & Wang, 2024), and the data is better represented as unions of local manifolds
 132 of varying dimension, typically lower than that of the latent space. However, the previous methods
 133 defined the submanifolds based on pre-specified datasets and could not generalize to new data.
 134

135 2.2 UNION OF κ -STABLE RIEMANNIAN MANIFOLDS

137 Assuming that the manifold hypothesis is indeed violated and that a given generative model has
 138 learned to faithfully capture the lower-dimensional structure in the data, it should be reflected in the
 139 decoder having rank strictly smaller than the latent dimensionality. We verified this phenomenon for
 140 antimicrobial peptides data in two state-of-the-art latent generative models (Das et al., 2018; Szym-
 141 czak et al., 2023) (see Appendix B). This implies that the G_{Dec} is not of full rank, and consequently
 142 the pair $(\mathcal{Z}, G_{\text{Dec}})$ does not constitute a Riemannian manifold.
 143

144 To address this, we equip each point $z \in \mathcal{Z}$ with a local, potentially lower-dimensional manifold,
 145 which we further enrich with a Riemannian structure from the pullback metric. In contrast to pre-
 146 vious methods, we adapt a decoder-dependent approach in the submanifold definition, enabling
 147 generalization to any point encoded in the latent space. Namely, we decompose \mathcal{Z} as a union of
 148 locally κ -stable Riemannian submanifolds ($\kappa \geq 0$)

$$148 \quad \mathbb{M}^\kappa = \{M_z^\kappa : z \in \mathcal{Z}\}, \quad M_z^\kappa = (W_z^\kappa, G_{\text{Dec}}),$$

149 where each $W_z^\kappa \ni z$ is an open affine submanifold of \mathcal{Z} of *maximal dimension* (denoted as k_z^κ and
 150 referred to as κ -stable dimension), such that the pullback metric G_{Dec} restricted to W_z^κ , denoted
 151 $G_{\text{Dec}}|_{W_z^\kappa}$, has full rank and satisfies the κ -stability condition

$$152 \quad \inf_{\substack{v \in T_z W_z^\kappa \\ \langle v, v \rangle_{\mathbb{R}^d} = 1}} \langle v, v \rangle_z^{\text{Dec}} > \kappa^2.$$

155 Intuitively, W_z^κ removes degenerated (non-active) directions of the decoder, ensuring that all eigen-
 156 values of $G_{\text{Dec}}|_{W_z^\kappa}$ are bounded below by κ . This guarantees numerical stability for geometric
 157 computations requiring inversion of G_{Dec} . Note, that similarly to the full-rank case, $\hat{\text{Dec}}|_{W_z^\kappa} : (W_z^\kappa, G_{\text{Dec}}) \rightarrow (\hat{\text{Dec}}(W_z^\kappa), \langle \cdot, \cdot \rangle_{\mathbb{R}^{LA}})$ is an isometric diffeomorphism.
 158

159 The explicit SVD-based construction of W_z^κ is deferred to Appendix C. For stability near bound-
 160 aries, we also use contracted domains $\hat{W}_z^\kappa(\alpha) = \{z + \alpha(v - z) : v \in W_z^\kappa\}$ and $M_z^\kappa(\alpha) =$
 161 $(W_z^\kappa(\alpha), G_{\text{Dec}})$, with $\alpha \in (0, 1)$.
 162

162 2.3 SORBES - SECOND-ORDER RIEMANNIAN BROWNIAN EFFICIENT SAMPLING
163

164 Having a stable Riemannian approximation M_z^κ of X in the vicinity of a point $z \in \mathcal{Z}$, we now
165 describe how to explore it efficiently within the local neighbourhood of z . Our goal is to simulate
166 a Riemannian Brownian motion for a time T starting from z that is a Riemannian equivalent of a
167 local Gaussian perturbation $z + \epsilon$, $\epsilon \sim \mathcal{N}(0, T)$. To this end, we introduce the SORBES (Second-
168 Order Riemannian Brownian efficient Sampling) procedure, described in Algorithm 1. SORBES
169 improves the flat Gaussian noise by exploring only active local subspace W_z^κ , it is isotropic w.r.t. to
170 the decoder geometry, and respects the local curvature of M_z^κ .

171 **Algorithm 1** SORBES
172

173 **Require:** $z \in \mathcal{Z}$, $\kappa \geq 0$, step size ϵ , diffusion time T , $\alpha = 0.99$
174 1: Initialize $z_0^\epsilon \leftarrow z$, stopped \leftarrow False, $\sigma \leftarrow 0$ (σ tracks diffusion time)
175 2: $W_z^\kappa, G_{\text{Dec}} \leftarrow M_z^\kappa$
176 3: **for** $i = 1$ to $\lfloor \frac{T}{\epsilon^2} \rfloor$ **do**
177 4: Sample a unit tangent direction $\bar{v} \in S_z^\kappa = \{u \in T_z M_z^\kappa : \langle u, u \rangle_z^{\text{Dec}} = 1\}$
178 5: Set $v \leftarrow \sqrt{k_z^\kappa} \bar{v}$
179 6: **if** not stopped **then**
180 7: Update (Γ denotes the local Christoffel symbol for M_z^κ)
181
$$z_i^\epsilon = z_{i-1}^\epsilon + \underbrace{\epsilon v}_{\substack{\text{first-order} \\ \text{geodesic approximation}}} - \underbrace{\epsilon^2 \Gamma(z_{i-1}^\epsilon)[v, v]}_{\substack{\text{second-order} \\ \text{geodesic approximation}}},$$

182
183
184 8: $\sigma \leftarrow \sigma + \epsilon^2$ (diffusion time update)
185 9: **if** $z_i^\epsilon \notin W_z^\kappa(\alpha)$ **then**
186 10: stopped \leftarrow True
187 11: **end if**
188 12: **else**
189 13: $z_i^\epsilon \leftarrow z_{i-1}^\epsilon$ (absorbed state)
190 14: **end if**
191 15: **end for**
192 16: **return** $(z_i^\epsilon)_{0 \leq i \leq \lfloor \frac{T}{\epsilon^2} \rfloor}, \sigma$
193

194 Before introducing the key theoretical property of this algorithm, let's recall the crucial notation.
195 For $A \subset M_z^\kappa$, A^c is the complement of A in M_z^κ . Let $d_{M_z^\kappa}$ be the geodesic distance on M_z^κ w.r.t. to
196 the pullback metric, and $d_{M_z^\kappa}(x, A) = \inf_{y \in A} d_{M_z^\kappa}(x, y)$ for $x \in M_z^\kappa$ and $A \subset M_z^\kappa$. Let Ric be the
197 Ricci curvature. Then the key theoretical property of the Algorithm 1 is summarized by:

198 **Theorem 1.** *Let $(Z_i^\epsilon)_{i \geq 0}$ be the sequence produced by Algorithm 1, for $M_z^\kappa(\alpha)$ with $\alpha \in (0, 1)$ and
199 diffusion horizon $T > 0$, and define its continuous-time interpolation*

$$Z^\epsilon(t) := Z_{\lfloor \epsilon^{-2} t \rfloor}^\epsilon, \quad t \geq 0.$$

200 Let $R_z^\kappa = d_{M_z^\kappa}(z, (W_z^\kappa)^c)$, and suppose $L \geq 1$ satisfies

$$\sup_{x \in M_z^\kappa(\alpha)} \text{Ric}_{M_z^\kappa}(x) \geq -L^2.$$

201 Then for $T < \frac{(R_z^\kappa)^2}{4k_z^\kappa L}$, as $\epsilon \rightarrow 0$, the process Z^ϵ converges in distribution to Riemannian Brownian
202 motion stopped at the boundary of $M_z^\kappa(\alpha)$, with respect to the Skorokhod topology, on a set $C_{\kappa, z}^T \subset$
203 Ω such that

$$\mathbb{P}(C_{\kappa, z}^T) \geq 1 - \exp\left(-\frac{(R_z^\kappa)^2}{32T}\right).$$

204 For the proof, see Appendix D. Intuitively, in the small-step limit, our algorithm converges to Rie-
205 manian Brownian motion on $M_z^\kappa(\alpha)$, stopped at the boundary, with the deviation probability de-
206 caying exponentially in the inverse time horizon. Theorem 1 extends the main convergence result
207 of Schwarz et al. (2022) to possibly non-compact manifolds. Importantly, Schwarz et al. (2022)

216 showed that achieving this convergence requires a *second-order correction* term (capturing the ef-
 217 fect of Christoffel symbols), rather than the commonly used naive first-order update. This motivates
 218 our use of the second-order scheme in Algorithm 1. In Appendix E (Algorithm 4), we describe
 219 SORBES-SE (*Stable/Efficient*), an implementation of SORBES that approximates the second-
 220 order correction using finite differences. It employs an adaptive step size ϵ that adjusts to the local
 221 curvature of the space, while still preserving the convergence guarantees.

223 2.4 MUTANG - MUTATION ENUMERATION IN TANGENT SPACE

225 To further exploit the manifold structure of M_z^κ , modelling the neighborhood of a point $z \in \mathcal{Z}$, let us
 226 observe that the ambient tangent space $T_{\text{Dec}(z)} \hat{\text{Dec}}(W_z^\kappa)$ identifies directions in peptide space along
 227 which the decoder output is the most sensitive. We interpret these directions as defining a *mutation*
 228 *space* for the decoded peptide $p(z)$, providing candidate amino-acid substitutions.

229 Formally, let $U^\kappa(z)$ (see Equation 8) denote an orthonormal basis of $T_{\text{Dec}(z)} \hat{\text{Dec}}(W_z^\kappa)$ in the am-
 230 bient space (Figure 2A–B), and let $u_j \in \mathbb{R}^{LA}$ be the j -th basis vector. We reshape u_j into matrix
 231 form

$$232 \Delta \text{Dec}^{(j)}(z) = \text{reshape}(u_j, (L, A)) \in \mathbb{R}^{L \times A}. \quad (3)$$

234 Intuitively, each entry $\Delta \text{Dec}^{(j)}(z)_{\ell, a}$ measures the first-order sensitivity of the probability assigned
 235 to amino acid \mathcal{A}_a at position ℓ , thereby suggesting a possible substitution. To extract candidate
 236 mutations, we introduce a sensitivity threshold $\theta_{\text{mut}} > 0$ and declare that

$$238 |\Delta \text{Dec}^{(j)}(z)_{\ell, a}| \geq \theta_{\text{mut}} \Rightarrow \text{add mutation } p(z)_\ell \rightarrow a, \quad (4)$$

240 where $p(z)_\ell$ is the current residue at position ℓ . Applying this rule across all $j = 1, \dots, k_z^\kappa$ yields a
 241 *mutation pool*

$$242 \mathcal{P} \subseteq \{1, \dots, L\} \times \mathcal{A}.$$

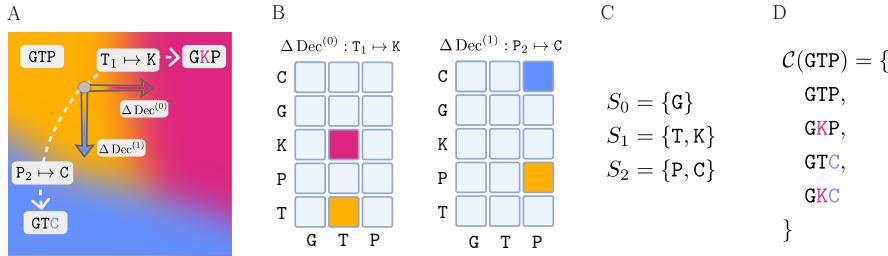
243 To enumerate candidate peptides, for each sequence position ℓ we define the set of admissible
 244 residues (Figure 2C) as

$$245 S_\ell = \{ \mathcal{A}_a \mid (\ell, a) \in \mathcal{P} \} \cup \{p(z)_\ell\},$$

246 i.e., all suggested mutations together with the identity residue. The complete candidate set is then
 247 obtained as the Cartesian product (Figure 2D):

$$249 \mathcal{C}(p(z)) = \prod_{\ell=1}^L S_\ell = \{ y \in \Sigma^L : y_\ell \in S_\ell \ \forall \ell \}. \quad (5)$$

253 The details of MUTANG are provided in Appendix F (Algorithm 5).



264 Figure 2: **Tangent space as mutation space and local enumeration.** **A)** Around an example
 265 peptide GTP we consider two orthogonal peptide-space tangent directions $\Delta \text{Dec}^{(j)}$ obtained from
 266 the SVD of the decoder Jacobian at the peptide code. Each direction suggests a specific substitution:
 267 $T_1 \rightarrow K$ and $P_2 \rightarrow C$. **B)** Each $\Delta \text{Dec}^{(j)}$ is reshaped into an $L \times A$ map (rows: amino acids; columns:
 268 positions). **C)** Thresholded entries define per-position sets of admissible residues (identity always
 269 included). **D)** The candidate set is the Cartesian product $\mathcal{C}(GTP) = \prod_\ell S_\ell$.

270 2.5 LOCAL ENUMERATION
271

272 We aim to densely populate the neighbourhood of a single prototype peptide with valid, diverse
273 candidates. Starting from a seed peptide p with latent code z_0 , we launch multiple Riemannian ran-
274 dom walk trajectories using the SORBES algorithm (Sec. 2.3). At the start and after each step, the
275 current latent state is decoded into a peptide and augmented with additional variants generated by
276 MUTANG (Sec. 2.4). Additionally, to account for local dimension variability, we re-estimate the
277 κ -stable submanifold M_z^κ at each step of the random walk. The union of all decoded walk steps and
278 tangent-space mutations yields a compact, high-quality *local candidate set* around p . LOCALENU-
279 MERATION is presented in Algorithm 2.
279

280 **Algorithm 2** LOCALENUMERATION
281

282 **Require:** seed peptide p , $\kappa_{\text{SORBES}}, \kappa_{\text{MUTANG}} \geq 0$, number of trajectories M , walk time budget
283 T_{walk} , nominal step size ϵ , mutation threshold θ_{mut}

284 1: Encode p to latent z_0 ; $\mathcal{C} \leftarrow \{p\}$
285 2: **for** $m = 1$ to M **do**
286 3: $z \leftarrow z_0$; $t \leftarrow 0$
287 4: $\mathcal{C} \leftarrow \mathcal{C} \cup \text{MUTANG}(z, \kappa_{\text{MUTANG}}, \theta_{\text{mut}})$
288 5: **while** $t < T_{\text{walk}}$ **do**
289 6: $((\cdot, z), \sigma) \leftarrow \text{SORBES-SE}(z, \kappa_{\text{SORBES}}, \epsilon, \text{STEP}_{\text{max}}=1)$ (single step of SORBES-SE)
290 7: $t \leftarrow t + \sigma^2$; $\mathcal{C} \leftarrow \mathcal{C} \cup \{p(z)\}$
291 8: $\mathcal{C} \leftarrow \mathcal{C} \cup \text{MUTANG}(z, \kappa_{\text{MUTANG}}, \theta_{\text{mut}})$
292 9: **end while**
293 10: **end for**
11: **return** \mathcal{C} (local candidate set)

294 2.6 LE-BO - LOCAL ENUMERATION BAYESIAN OPTIMIZATION
295

296 Finally, we integrate our LOCALENUMERATION algorithm into a Bayesian optimization (Garnett,
297 2023) framework for peptide design, which we term Local Enumeration Bayesian Optimization
298 (LE-BO). Instead of performing costly optimization of the acquisition function in the latent space,
299 which typically relies on Euclidean-distance kernels and ignores both the latent geometry and the
300 discrete nature of peptides, we use surrogate Gaussian process models (Seeger, 2004), defined di-
301 rectly in the peptide space. The acquisition function is optimized by locally enumerating peptides
302 in the vicinity of the most promising candidates and then selecting the peptide that maximizes the
303 acquisition value. To further encourage exploration and increase the diversity of discovered can-
304 didates, we employ the ROBOT scheme (Maus et al., 2023), which promotes searching across a
305 broader set of promising regions. Details of LE-BO are presented in Algorithm 3.
305

306 2.7 POGS - POTENTIAL-MINIMIZING GEODESIC SEARCH
307

308 Given two prototype peptides with latent vectors z_a and z_b , we aim to generate *seeds*, i.e., *analogues*
309 *that are jointly similar to both vectors and have high predicted activity*. To this end, we construct
310 a discrete geodesic-like curve connecting z_a and z_b , interpreted physically as a system with *kinetic*
311 *energy* (geometric term) and an added *potential energy* (property term). This provides a natural
312 tradeoff between similarity to both seeds and the desired molecular property.
313

314 Because the decoder Jacobian may have varying rank, we avoid intrinsic pullback computations
315 and work in the ambient peptide-probability space \mathbb{R}^{LA} . For a sequence of latent waypoints $Z =$
316 $\{z_0 = z_a, z_1, \dots, z_N = z_b\}$, we define their decoded logits $X_k = \log(\hat{\text{Dec}}(z_k))$, and approximate
317 curve length using *chord distances* $\|X_{k+1} - X_k\|_2$. This extrinsic metric serves as a first-order
318 surrogate for geodesic energy, bypassing costly Christoffel evaluations and remaining stable under
319 rank variability.
319

320 We define the total energy of a discrete path Z as

$$\mathcal{E}_{\lambda, \mu}(Z) = \underbrace{\sum_{k=0}^{N-1} \|X_{k+1} - X_k\|_2^2}_{\text{kinetic term: geometric similarity}} + \lambda \underbrace{\sum_{k=0}^N \Phi(X_k)}_{\text{potential term: property bias}} + \mu \underbrace{\sum_{k=0}^{N-1} \|z_{k+1} - z_k\|_2^2}_{\text{latent regularizer}}, \quad (6)$$

324

Algorithm 3 LE-BO

325

Require: ORACLE function to be optimized; seed p_{seed} ; maximum budget B_{max} ; trust region distance d_{trust} ; number of ROBOT evaluations per iteration k_{ROBOT} ; diversity threshold d_{ROBOT} ; a surrogate Gaussian Process GP model with GP.acquisition function.

326

```

1:  $p_{\text{current}} := p_{\text{best}} := p_{\text{seed}}$ ,  $\mathcal{D} := \{p_{\text{seed}}\}$ ,  $\mathcal{E} = \{(p_{\text{seed}}, \text{ORACLE}(p_{\text{seed}}))\}$ 
2: for  $\text{iter} := 1$  to  $\lfloor B_{\text{max}}/k_{\text{ROBOT}} \rfloor$  do
3:    $\mathcal{D} := \mathcal{D} \cup \text{LOCALENUMERATION}(p_{\text{current}})$  Explore the neighborhood of  $p_{\text{current}}$ 
4:    $\text{GP} := \text{GP}. \text{fit}(\mathcal{E})$  Fit the surrogate GP model
5:    $\mathcal{D}_{\text{trust}} := \{p \in \mathcal{D} \mid \text{Levenshtein}(p, p_{\text{best}}) \leq d_{\text{trust}}\}$  Define the trust region
6:   for  $i := 1$  to  $k_{\text{ROBOT}}$  do
7:      $p_{\text{ROBOT}}^i := \arg \max_{p \in \mathcal{D}_{\text{trust}}} \text{GP}. \text{acquisition}(p)$ 
8:      $\mathcal{E} := \mathcal{E} \cup \{(p_{\text{ROBOT}}^i, \text{ORACLE}(p_{\text{ROBOT}}^i))\}$ 
9:      $\mathcal{D}_{\text{trust}} := \mathcal{D}_{\text{trust}} \setminus \{p \in \mathcal{D} \mid \text{Levenshtein}(p, p_{\text{ROBOT}}^i) \leq d_{\text{ROBOT}}\}$  ROBOT diversity filtering
10:  end for
11:   $p_{\text{current}} := \arg \min_{1 \leq i \leq k_{\text{ROBOT}}} \text{ORACLE}(p_{\text{ROBOT}}^i)$ 
12:  if  $\text{ORACLE}(p_{\text{current}}) \geq \text{ORACLE}(p_{\text{best}})$  then
13:     $p_{\text{best}} := p_{\text{current}}$ 
14:  end if
15: end for
16: return  $p_{\text{best}}$ 

```

344

345

where Φ is the property prediction (e.g. negative log MIC), $\lambda \geq 0$ balances geometry vs. property, and $\mu \geq 0$ regularizes latent jumps to discourage large chords in \mathcal{Z} . The first term corresponds to kinetic energy (favoring smooth, short ambient curves), the second to a potential energy that biases toward low Φ , and the third acts as a stabilizer ensuring robustness of the chord points in the latent space.

350

351

To perform optimization and search, we initialize Z by straight-line interpolation in latent space, and later optimize $\mathcal{E}_{\lambda, \mu}(Z)$ w.r.t. Z using ADAM solver. During optimization, only the interior points z_1, \dots, z_{N-1} are updated. Given an optimized path Z , we decode each z_i to a peptide p_i and remove consecutive duplicates, obtaining a *peptide path* $(p'_0, \dots, p'_{N'})$ of length N' . We call a peptide p'_k a *seed* if its potential satisfies $\Phi(p'_k) \leq \theta_{\text{pot}}$ for a threshold θ_{pot} . A peptide p'_k is called a *well* if it is a seed and also a local minimum of the potential Φ along the peptide path. The detailed algorithm is presented in Appendix G (Algorithm 6).

357

358

3 RESULTS

359

3.1 POGS EVALUATION

360

361

To evaluate the PoGS procedure, we applied it for the HydrAMP model (Szymczak et al., 2023), using average standardized MIC predictions against 3 *Escherichia coli* strains (*E. coli* ATCC11775, AIG221 and AIG222) of a APEX-derived transformer prediction (Wan et al., 2024) (Appendix H) as the potential function. 300 prototype pairs (z_a, z_b) were drawn from the Veltri dataset (Veltri et al., 2018), restricted to active peptides (average MIC $\leq 32 \mu\text{g}/\text{ml}$ against 3 *E. coli* strains) with edit distance between prototypes ≥ 10 .

368

369

For each pair z_a and z_b , we compared three paths between z_a and z_b : straight Euclidean interpolation, PoGS without potential, and full PoGS ; Sec. 2.7), measuring chord latent and ambient lengths, decoded peptide path lengths, and property-based counts of seeds and wells. PoGS hyperparameters and metrics are described in Appendix G. PoGS achieved shorter ambient paths, and substantially more seeds and wells (Table 1), showing that property-aware potentials enrich trajectories for active candidates.

374

375

3.2 LE-BO EVALUATION

376

377

We next evaluated our LE-BO algorithm on a black-box peptide optimization task with a budget of 1400 evaluations. Optimization was initialized from four seed peptides derived using PoGS

378
 379
 380
 381
 382 **Table 1: PoGS results.** Comparison of straight interpolation, geodesic (no potential), and property-
 383 extended geodesic. Reported are latent length, ambient length, peptide path per length, and counts
 384 of seeds and wells.

Method	Latent length	Ambient length	Peptide path length	Potential	Seeds	Wells
Straight interpolation	6.43 \pm 0.13	1601.56 \pm 57.04	71.25 \pm 6.12	-501.00 \pm 23.92	11.70 \pm 5.54	2.50 \pm 1.90
PoGS w.o. potential ($\lambda = 0$)	7.98 \pm 0.18	1240.03 \pm 40.72	66.68 \pm 4.23	-603.34 \pm 15.12	18.90 \pm 8.40	6.40 \pm 3.20
PoGS ($\lambda = 0.01$)	9.02 \pm 0.09	1432.55 \pm 19.98	77.12 \pm 6.12	-785.45 \pm 34.12	22.70 \pm 10.49	12.60 \pm 4.45

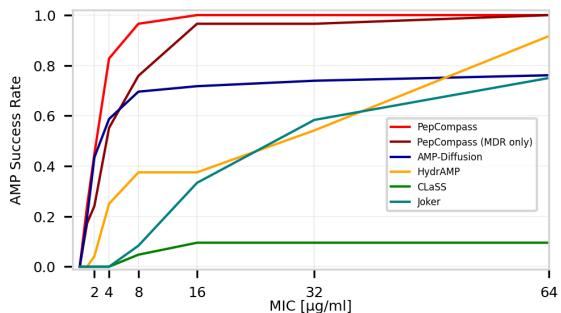
385
 386
 387
 388
 389
 390
 391 for $\lambda = 0.01$ and $\mu = 0.1$ (Sec. 2.7): *KY14*, *KF16*, *KK16*, *FL14*, as well as two previously
 392 described AMPs (*mammuthusin-3*, *hydrodamin-2*) (Wan et al., 2024). The optimization procedure
 393 was repeated 10 times, and for each run the best value achieved across all optimization steps was
 394 recorded and reported in Table 3.3. LE-BO hyperparameters are described in Appendix I.

395 We compared LE-BO against a diverse set of state-of-the-art methods, including generative models:
 396 HydrAMP (Szymczak et al., 2023) with different creativity hyperparameter τ and PepCVAE (Das
 397 et al., 2018); the diffusion-based model LaMBO-2 (Gruver et al., 2023); evolutionary strategies: La-
 398 tent CMA-ES and Relaxed CMA-ES (Hansen, 2016), where the former operates in the HydrAMP
 399 latent space and the latter optimizes directly on a continuous relaxation of one-hot sequence en-
 400 codings, AdaLead (Sinai et al., 2020), PEX (Ren et al., 2022); GFlowNet-based approaches: GFN-
 401 AL (Jain et al., 2022) and GFN-AL- δ CS (Kim et al., 2025); the reinforcement learning method
 402 DyNA-PPO (Angermueller et al., 2020); probabilistic methods: CbAS (Brookes et al., 2019) and
 403 Evolutionary BO (Sinai et al., 2020); the insertion-based Joker method (Porto et al., 2018); and
 404 greedy Random Mutation (González-Duque et al., 2024). Comparison to competitor latent BO
 405 methods SAASBO (Eriksson & Jankowiak, 2021), Hvarfner’s Vanilla BO (Hvarfner et al., 2024),
 406 and LineBO (Kirschner et al., 2019) was not feasible due to their relative computational inefficiency:
 407 while LE-BO finalized within 2 hours, for all those methods, a single run of the same number of
 408 1400 steps within HydrAMP latent space did not finish within a week.

409 Additionally, we conducted an ablation
 410 study on LE-BO variants to isolate the
 411 effects of random walks and mutation
 412 enumeration. These variants modify the
 413 ENUMERATELOCAL sub-procedure in Al-
 414 gorithm 2. The *Euclidean walk* vari-
 415 ant replaces SORBES with a naive Eu-
 416 clidean random walk in the latent space.
 417 The *mutation-disabled* variant omits MU-
 418 TANG. Finally, the *walk-disabled* variant
 419 uses only a single MUTANG without ran-
 420 dom walks.

421 As the optimized black-box function, we
 422 used the APEX MIC regressor (Wan et al.,
 423 2024), with the objective of minimization.
 424 We minimize the average \log_2 MIC across
 425 three *E. coli* strains, reporting the mean
 426 over the best results from 10 repeated op-
 427 timization runs (Table 3.3).

428 LE-BO achieved the best performance on
 429 five out of six prototypes, outperforming all baselines except for *FL14*, where it ranked second.
 430 The Ablated LE-BO variant with Euclidean walk and enabled mutations achieved the second-best
 431 performance on four out of six prototypes, underscoring the importance of mutation enumeration
 432 in the optimization process. These ablations further confirmed that both Riemannian random walks
 433 and mutations are essential for consistently achieving low MIC values.



432 **Figure 3: Antimicrobial peptide success rates across**
 433 **MIC thresholds.** Success rate is defined as the fraction
 434 of generated peptides with MIC below the specified
 435 threshold against at least one tested strain. Results
 436 are based on experimental validation against 19 bacte-
 437 rial strains, including 8 MDR isolates.

432 Table 2: Minimal \log_2 MIC values achieved by each optimization method. Reported values are the
 433 mean and standard deviation over 10 runs. The top row shows the predicted \log_2 MIC of seeds
 434 before optimization. The first block reports baseline results, the second block shows ablations, and
 435 the last row presents our method. **Bold**: best overall value for a prototype. Underline: second-best.
 436

Method	KY14	KF16	KK16	FL14	mammuthusin-3	hydrodamin-2
log ₂ MIC value (seed)	2.88	3.39	2.71	4.00	4.30	6.96
GFN-AL	2.73 \pm 0.27	3.20 \pm 0.19	2.63 \pm 0.17	3.73 \pm 0.47	4.30 \pm 0.00	3.85 \pm 0.76
GFN-AL- δ CS	1.91 \pm 0.19	1.74 \pm 0.20	1.79 \pm 0.08	2.06 \pm 0.42	3.02 \pm 0.19	1.77 \pm 0.41
PEX	1.39 \pm 0.07	1.48 \pm 0.09	1.54 \pm 0.14	1.27 \pm 0.27	2.65 \pm 0.09	1.14 \pm 0.11
Joker	4.66 \pm 2.04	3.98 \pm 1.52	3.79 \pm 1.61	4.22 \pm 1.08	6.49 \pm 1.75	6.28 \pm 0.33
Random Mutation	1.23 \pm 0.26	1.17 \pm 0.13	0.97 \pm 0.40	1.02 \pm 0.59	2.12 \pm 0.80	0.69 \pm 0.19
LaMBO-2	1.88 \pm 0.25	2.11 \pm 0.20	1.72 \pm 0.18	1.76 \pm 0.32	2.75 \pm 0.15	2.17 \pm 0.35
Relaxed CMA-ES	1.92 \pm 0.13	1.83 \pm 0.28	1.87 \pm 0.30	1.93 \pm 0.40	2.86 \pm 0.29	1.66 \pm 0.34
Latent CMA-ES	1.81 \pm 0.33	2.07 \pm 0.56	1.72 \pm 0.29	1.72 \pm 0.20	2.17 \pm 0.60	1.87 \pm 0.44
CbAS	2.88 \pm 0.00	3.39 \pm 0.00	2.71 \pm 0.00	4.00 \pm 0.00	4.30 \pm 0.00	5.50 \pm 0.93
DyNAPPO	1.45 \pm 0.36	1.31 \pm 0.22	1.34 \pm 0.20	0.73 \pm 0.53	2.42 \pm 0.54	0.82 \pm 0.39
Evolutionary BO	1.65 \pm 0.23	1.45 \pm 0.42	1.50 \pm 0.12	1.70 \pm 0.19	2.68 \pm 0.12	1.28 \pm 0.33
AdaLead	0.87 \pm 0.49	1.01 \pm 0.28	0.93 \pm 0.24	0.51 \pm 0.38	2.30 \pm 0.28	0.66 \pm 0.21
PepVAE	3.66 \pm 0.00	2.87 \pm 0.01	3.14 \pm 0.11	2.12 \pm 0.00	4.30 \pm 0.00	6.74 \pm 0.06
HydrAMP $\tau = 5.0$	2.35 \pm 0.08	2.03 \pm 0.12	1.88 \pm 0.10	2.19 \pm 0.12	3.19 \pm 0.27	2.39 \pm 0.38
HydrAMP $\tau = 2.0$	2.60 \pm 0.03	2.27 \pm 0.02	2.11 \pm 0.11	2.81 \pm 0.30	3.99 \pm 0.01	5.02 \pm 0.28
HydrAMP $\tau = 1.0$	2.86 \pm 0.01	2.27 \pm 0.00	2.35 \pm 0.00	3.72 \pm 0.35	4.27 \pm 0.10	6.09 \pm 0.01
Walk		Mutation				
Ablated LE-BO	Euclidean	–	1.37 \pm 0.27	1.33 \pm 0.23	1.24 \pm 0.18	1.29 \pm 0.17
	SORBES-SE	–	1.37 \pm 0.22	1.42 \pm 0.22	1.12 \pm 0.35	1.18 \pm 0.20
	–	✓	1.71 \pm 0.20	1.46 \pm 0.40	1.82 \pm 0.13	1.24 \pm 0.13
Euclidean		✓	<u>0.65 \pm 0.18</u>	<u>0.71 \pm 0.18</u>	<u>0.83 \pm 0.32</u>	0.87 \pm 0.22
LE-BO	SORBES-SE	✓	0.50 \pm 0.24	0.60 \pm 0.29	0.50 \pm 0.14	0.60 \pm 0.22
						0.50 \pm 0.38
						0.58 \pm 0.34

455 3.3 WET-LAB VALIDATION

456 To validate the computational predictions from PoGS and LE-BO, we conducted comprehensive
 457 *in vitro* antimicrobial testing of the generated peptides. A total of 29 novel peptides were exper-
 458 imentally evaluated: 4 seed peptides discovered through PoGS bi-prototype geodesics with property-
 459 aware potentials, and 25 analogs derived from these 4 seeds through LE-BO optimization for *E. coli*
 460 activity. These peptides were tested against a panel of 19 bacterial strains, including 8 multidrug-
 461 resistant (MDR) isolates (Appendix J), to assess both broad-spectrum activity and efficacy against
 462 clinically relevant resistant pathogens. We compared the success rate of PepCompass to previ-
 463 ous methods that were also validated experimentally (HydrAMP (Szymczak et al., 2023), AMP-
 464 Diffusion (Torres et al., 2025), CLaSS (Das et al., 2020), Joker (Porto et al., 2018)). To this end, for
 465 each activity threshold, we computed the fraction of tested peptides that were active against at least
 466 one bacterial strain with this activity threshold.

467 As demonstrated in Figure 3, with unprecedented 100% success rate for the standard activity thresh-
 468 old of $32\mu\text{g}/\text{ml}$ and 82% rate at much more demanding threshold of $4\mu\text{g}/\text{ml}$, PepCompass achieved
 469 superior performance across various MIC thresholds compared to previous methods, maintaining
 470 high success rates even when evaluated specifically against MDR strains. The experimental results
 471 confirmed the expected activity increase from prototypes to seeds to analogs for Gram-negative bac-
 472 teria, directly validating our optimization strategy that targeted *E. coli* activity. Indeed, while the
 473 generated peptides showed some activity against Gram-positive bacteria (Figure 8), the clear en-
 474 hancement from optimization was primarily observed against Gram-negative pathogens (Figure 9).
 475

477 CONCLUSIONS

478 By leveraging Riemannian latent geometry, interpretable tangent-space mutations, and potential-
 479 augmented geodesics, PepCompass enables efficient navigation of and optimization within peptide
 480 space across global and local scales. One of possible limitation of our approach is the reliance on the
 481 Euclidean-distance based metric on decoder outputs in the ambient space. However, any other met-
 482 rics that would better capture the ambient manifold could easily be incorporated. Already now, our
 483 computational experiments demonstrate superior performance over state-of-the-art baselines, and
 484 wet-lab validation confirms unprecedented success rates, with all tested peptides showing activity
 485 *in vitro*, including activity against multidrug-resistant pathogens. These results establish geometry-

486 aware exploration as a powerful new paradigm for controlled generative design in vast biological
 487 spaces.
 488

489 **REPRODUCIBILITY STATEMENT**
 490

491 All proofs, together with their explanations and underlying assumptions, are provided in Appen-
 492 dices B, C and D. All implementation details and hyperparameters used in the experiments are listed
 493 in Appendices E, F, G, H and I. Additionally, we release the source code and experimental con-
 494 figurations necessary to reproduce the key results. Full details of wet-lab validation procedure are
 495 described in Appendix J.

496 **ETHICS STATEMENT**
 497

498 This work includes methods generally applicable to peptide design. An example of potential mali-
 499 cious use of PepCompass would include optimization of peptide toxicity. However, the intention of
 500 this paper is to instead provide tools facilitating the design of therapeutic peptides.
 501

502 **LLM USAGE**
 503

504 Large Language Models (LLMs) were used in this work to improve the clarity and structure of the
 505 text. Their use was limited to rephrasing and stylistic refinement. In addition, LLMs were employed
 506 to support the search of related work, helping to verify the accuracy of claims about prior research.
 507

508 **REFERENCES**
 509

- 510 Alexander Alemi et al. Diffeomorphic explanations with normalizing flows. *arXiv preprint*, 2023.
 511 URL <https://arxiv.org/abs/2301.12844>.
- 512 Christof Angermueller, David Dohan, David Belanger, Ramya Deshpande, Kevin Murphy, and Lucy
 513 Colwell. Model-based reinforcement learning for biological sequence design. In *International
 514 Conference on Learning Representations (ICLR)*, 2020. URL <https://openreview.net/pdf?id=Hk1xbgBKvr>.
 515
- 516 Georgios Arvanitidis, Lars Kai Hansen, and Søren Hauberg. Latent space oddity: On the curvature
 517 of deep generative models. In *International Conference on Learning Representations (ICLR)*,
 518 2018. URL <https://arxiv.org/abs/1710.11379>.
- 519 Georgios Arvanitidis, Søren Hauberg, et al. Geometrically enriched latent spaces. *arXiv preprint
 520 arXiv:2008.00565*, 2020.
- 521 Yoshua Bengio, Aaron Courville, and Pascal Vincent. Representation learning: A review and new
 522 perspectives. *arXiv preprint arXiv:1206.5538*, 2012. doi: 10.48550/arXiv.1206.5538. URL
 523 <https://arxiv.org/abs/1206.5538>.
- 524 Mathieu Blondel, Kimon Fountoulakis, et al. (deep) generative geodesics. *arXiv preprint
 525 arXiv:2407.11244*, 2024.
- 526 Sam Bond-Taylor, Adam Leach, Yang Long, and Chris G. Willcocks. Deep generative modelling:
 527 A comparative review of vaes, gans, normalizing flows, energy-based and autoregressive models.
 528 *arXiv preprint arXiv:2103.04922*, 2021.
- 529 David Brookes, Hahnbeom Park, and Jennifer Listgarten. Conditioning by adaptive sampling for
 530 robust design. In *Proceedings of the 36th International Conference on Machine Learning (ICML)*,
 531 volume 97 of *Proceedings of Machine Learning Research*, pp. 773–782, 2019. URL <https://proceedings.mlr.press/v97/brookes19a.html>.
- 532 Bradley C. A. Brown, Anthony L. Caterini, Brendan L. Ross, Jesse C. Cresswell, and Gabriel
 533 Loaiza-Ganem. Verifying the union of manifolds hypothesis for image data. *arXiv preprint
 534 arXiv:2207.02862*, 2022.

- 540 T. Cao et al. Ancestral protein sequence reconstruction with a tree-structured ornstein-uhlenbeck
 541 vae. In *ICLR (OpenReview)*, 2022. URL <https://openreview.net/forum?id=FZoZ7a31GCW>.
- 543
- 544 Alice Capecchi, Daniel Probst, and Jean-Louis Reymond. One molecular fingerprint to rule them
 545 all: drugs, biomolecules, and the metabolome. *Journal of Cheminformatics*, 12(1):43, 2020. doi:
 546 10.1186/s13321-020-00445-4.
- 547
- 548 Payel Das, Kahini Wadhawan, Oscar Chang, Tom Sercu, Cicero dos Santos, Matthew Riemer, Vi-
 549 jil Chenthamarakshan, Inkit Padhi, and Aleksandra Mojsilovic. Pepcvae: Semi-supervised tar-
 550 getted design of antimicrobial peptide sequences. *arXiv preprint arXiv:1810.07743*, 2018. URL
 551 <https://arxiv.org/abs/1810.07743>. Revised version v3, Nov. 2018.
- 552
- 553 Payel Das, Tom Sercu, Kahini Wadhawan, Inkit Padhi, Sebastian Gehrmann, Flaviu Cipcigan,
 554 Vijil Chenthamarakshan, Hendrik Strobelt, Cicero dos Santos, Pin-Yu Chen, Yi Yan Yang,
 555 Jeremy Tan, James Hedrick, Jason Crain, and Aleksandra Mojsilovic. Accelerating antimicro-
 556 bial discovery with controllable deep generative models and molecular dynamics. *arXiv preprint*
 557 *arXiv:2005.11248*, 2020. URL <https://arxiv.org/abs/2005.11248>. CLaSS (Con-
 558 trolled Latent Attribute Space Sampling).
- 559
- 560 Valentin De Bortoli, James Thornton, et al. Riemannian score-based generative modeling. In
 561 *NeurIPS*, 2022.
- 562
- 563 Nicki S. Detlefsen, Søren Hauberg, and Wouter Boomsma. Learning meaningful represen-
 564 tations of protein sequences. *Nature Communications*, 13(1):1914, 2022. doi: 10.1038/
 565 s41467-022-29443-w.
- 566
- 567 Manfredo P. do Carmo. *Riemannian Geometry*. Birkhäuser, 1992. Translated by Francis Flaherty,
 568 reprint edition.
- 569
- 570 Weitao Du and Elton P. Hsu. Reflecting brownian motion and the gauss-bonnet-chern theorem,
 571 2021. URL <https://arxiv.org/abs/2106.10825>.
- 572
- 573 Jerome Eberhardt, Aidan Lees, Markus Lill, and Torsten Schwede. Combining bayesian optimiza-
 574 tion with sequence-or structure-based strategies for optimization of protein-peptide binding, 2024.
 575 URL <https://doi.org/10.26434/chemrxiv-2023-b7181-v2>.
- 576
- 577 David Eisenberg, Robert M. Weiss, Thomas C. Terwilliger, and William Wilcox. Hydrophobic
 578 moments and protein structure. *Faraday Symposia of the Chemical Society*, 17:109–120, 1982.
- 579
- 580 David Eriksson and Martin Jankowiak. High-dimensional bayesian optimization with sparse axis-
 581 aligned subspaces. In *Uncertainty in Artificial Intelligence*, pp. 493–503. PMLR, 2021.
- 582
- 583 Zhou Feng et al. Mario plays on a manifold: Generating functional content in latent space. In *ICML*,
 584 2021. URL <http://proceedings.mlr.press/v139/feng21g/feng21g.pdf>.
- 585
- 586 Roman Garnett. *Bayesian Optimization*. Cambridge University Press, 2023. ISBN 9781108348973.
- 587
- 588 Gary W. Gibbons. The jacobi metric for timelike geodesics in static spacetimes. *Classical and*
 589 *Quantum Gravity*, 33, 2015. URL <https://arxiv.org/abs/1508.06755>.
- 590
- 591 Miguel González-Duque, Richard Michael, Simon Bartels, Yevgen Zainchkovskyy, Søren Hauberg,
 592 and Wouter Boomsma. A survey and benchmark of high-dimensional bayesian optimization of
 593 discrete sequences. *NeurIPS 2024 (Datasets & Benchmarks Track)*, 2024. doi: 10.48550/arXiv.
 594 2406.04739. URL <https://arxiv.org/abs/2406.04739>.
- 595
- 596 Nate Gruver, Samuel Stanton, Nathan Frey, Tim G. J. Rudner, Julien Lafrance-Vanassee, Isidro
 597 Hotzel, Arvind Rajpal, Kyunghyun Cho, and Andrew Gordon Wilson. Protein design with
 598 guided discrete diffusion. In *Advances in Neural Information Processing Systems (NeurIPS)*,
 599 2023. URL https://proceedings.neurips.cc/paper_files/paper/2023/file/29591f355702c3f4436991335784b503-Paper-Conference.pdf.

- 594 Jiahui Guan, Peilin Xie, Dian Meng, Lantian Yao, Dan Yu, Ying Chih Chiang, Tzong Yi Lee, and
 595 Junwen Wang. Toxipep: Peptide toxicity prediction via fusion of context-aware representation
 596 and atomic-level graph. *Computational and Structural Biotechnology Journal*, 27:2347–2358,
 597 2025. doi: 10.1016/j.csbj.2025.05.039.
- 598 Nikolaus Hansen. The cma evolution strategy: A tutorial. *arXiv*, 2016. doi: 10.48550/arXiv.1604.
 599 00772. URL <https://arxiv.org/abs/1604.00772>.
- 600 Michael Herrmann, Pit Neumann, Simon Schwarz, Anja Sturm, and Max Wardetzky. Sub-
 601 riemannian random walks: From connections to retractions. *arXiv preprint arXiv:2311.17289*,
 602 2023.
- 603 Elton P. Hsu. *Stochastic Analysis on Manifolds*, volume 38 of *Graduate Studies in Mathematics*.
 604 American Mathematical Society, Providence, R.I.; Great Britain, 2002a. ISBN 0821808028.
- 605 Elton P Hsu. *Stochastic analysis on manifolds*. Number 38. American Mathematical Soc., 2002b.
- 606 Carl Hvarfner, Erik Orm Hellsten, and Luigi Nardi. Vanilla bayesian optimization performs great in
 607 high dimensions. *arXiv preprint arXiv:2402.02229*, 2024.
- 608 Moksh Jain, Emmanuel Bengio, Alex Hernandez-Garcia, Jarrid Rector-Brooks, Bonaventure F. P.
 609 Dossou, Chanakya Ekbot, Jie Fu, Tianyu Zhang, Michael Kilgour, Dinghuai Zhang, Lena
 610 Simine, Payel Das, and Yoshua Bengio. Biological sequence design with gflownets. In *Proceed-
 611 ings of the 39th International Conference on Machine Learning (ICML)*, volume 162 of *Proceed-
 612 ings of Machine Learning Research*, pp. 9845–9895, 2022. URL <https://proceedings.mlr.press/v162/jain22a.html>.
- 613 Dimitris Kalatzis, David Eklund, Georgios Arvanitidis, and Søren Hauberg. Variational autoen-
 614 coders with riemannian brownian motion priors. In *ICML*, 2020. arXiv:2002.05227.
- 615 Hyeonah Kim, Minsu Kim, Taeyoung Yun, Sanghyeok Choi, Emmanuel Bengio, Alex Hernández-
 616 García, and Jinkyoo Park. Improved off-policy reinforcement learning in biological sequence
 617 design. In *Proceedings of the 42nd International Conference on Machine Learning (ICML) —
 618 Poster*, 2025. URL <https://openreview.net/forum?id=0TY51hhdZm>.
- 619 Johannes Kirschner, Mojmir Mutny, Nicole Hiller, Rasmus Ischebeck, and Andreas Krause. Adap-
 620 tive and safe bayesian optimization in high dimensions via one-dimensional subspaces. In *Inter-
 621 national Conference on Machine Learning*, pp. 3429–3438. PMLR, 2019.
- 622 Gabriel Loaiza-Ganem, Brendan Leigh Ross, Rasa Hosseinzadeh, Anthony L. Caterini, and Jesse C.
 623 Cresswell. Deep generative models through the lens of the manifold hypothesis: A survey and
 624 new connections. *arXiv preprint arXiv:2404.02954*, 2024. doi: 10.48550/arXiv.2404.02954. URL
 625 <https://arxiv.org/abs/2404.02954>. Also certified at TMLR (2024).
- 626 Shuwen Lou. Discrete approximation to brownian motion with varying dimension in bounded do-
 627 mains. *Kyoto Journal of Mathematics*, 63(3):641–667, 2023. doi: 10.1215/21562261-10607383.
- 628 Natalie Maus, Kaiwen Wu, David Eriksson, and Jacob Gardner. Discovering many diverse sol-
 629 lutions with bayesian optimization. In *Proceedings of the 26th International Conference on
 630 Artificial Intelligence and Statistics (AISTATS)*, volume 206 of *Proceedings of Machine Learn-
 631 ing Research*, pp. 1779–1798. PMLR, April 2023. doi: 10.48550/arXiv.2210.10953. URL
 632 <https://proceedings.mlr.press/v206/maus23a/maus23a.pdf>.
- 633 Colin M. Van Oort, Jonathon B. Ferrell, Jacob M. Remington, Safwan Wshah, Jianing Li, et al.
 634 Ampgan v2: Machine learning-guided design of antimicrobial peptides. *Journal of Chemical
 635 Information and Modeling*, 61(5):2198–2207, May 2021. doi: 10.1021/acs.jcim.0c01441.
- 636 Amir Pandi, David Adam, Amir Zare, Van Tuan Trinh, Stefan L. Schaefer, Marie Burt, Björn
 637 Klabunde, Elizaveta Bobkova, Manish Kushwaha, Yeganeh Foroughijabbari, Peter Braun,
 638 Christoph Spahn, Christian Preußen, Elke Pogge von Strandmann, Helge B. Bode, Heiner von
 639 Buttlar, Wilhelm Bertrams, Anna Lena Jung, Frank Abendroth, Bernd Schmeck, Gerhard Hum-
 640 mer, and Tobias J. ... Erb. Cell-free biosynthesis combined with deep learning accelerates
 641 de novo-development of antimicrobial peptides. *Nature Communications*, 14, 2023. doi:
 642 10.1038/s41467-023-42434-9.

- 648 Seonghyeon Park et al. Low-rank subspaces in gans. In *NeurIPS*, 2021.
649
- 650 Paraskevas Pegios, Aasa Feragen, Andreas A. Hansen, and Georgios Arvanitidis. Counterfactual
651 explanations via riemannian latent space traversal. *arXiv preprint arXiv:2411.02259*, 2024.
- 652 William F. Porto, Isabel C. M. Fensterseifer, Suzana M. Ribeiro, and Octavio L. Franco. Joker: An
653 algorithm to insert patterns into sequences for designing antimicrobial peptides. *Biochimica et
654 Biophysica Acta (BBA) - General Subjects*, 1862(9):2043–2052, 2018. doi: 10.1016/j.bbagen.
655 2018.06.011. URL <https://pubmed.ncbi.nlm.nih.gov/29928920/>.
- 656
- 657 Zhizhou Ren, Jiahua Li, Fan Ding, Yuan Zhou, Jianzhu Ma, and Jian Peng. Proximal exploration
658 for model-guided protein sequence design. In *Proceedings of the 39th International Conference
659 on Machine Learning (ICML)*, volume 162 of *Proceedings of Machine Learning Research*, pp.
660 18520–18536, 2022. URL <https://proceedings.mlr.press/v162/ren22a.html>.
- 661
- 662 Steffen Schneider et al. Internal representations of vision models through the lens of frames on data
663 manifolds. *arXiv preprint*, 2022. URL <https://arxiv.org/abs/2211.10558>.
- 664
- 665 Simon Schwarz, Michael Herrmann, Anja Sturm, and Max Wardetzky. Efficient random walks on
666 riemannian manifolds. *arXiv preprint arXiv:2202.00959*, 2022. URL <https://arxiv.org/abs/2202.00959>.
- 667
- 668 Matthias Seeger. Gaussian processes for machine learning. *International journal of neural systems*,
669 14(02):69–106, 2004.
- 670
- 671 Hao Shao, Nicha C. Dvornek, et al. The riemannian geometry of deep generative models. In *CVPR
672 Workshops*, 2018.
- 673
- 674 Yujun Shen et al. Interpreting the latent space of gans for semantic face editing. In *CVPR*, 2020.
- 675
- 676 Sam Sinai, Richard Wang, Alexander Whatley, Stewart Slocum, Elina Locane, and Eric D. Kelsic.
677 Adalead: A simple and robust adaptive greedy search algorithm for sequence design. *arXiv*, 2020.
678 doi: 10.48550/arXiv.2010.02141. URL <https://arxiv.org/abs/2010.02141>.
- 679
- 680 Sándor Szedmak and Eric Bach. On the generalization of tanimoto-type kernels to real valued
681 functions. *arXiv preprint arXiv:2007.05943*, 2025. URL <https://arxiv.org/abs/2007.05943>. v3 (originally posted 2020).
- 682
- 683 P. Szymczak and E. Szczerk. Artificial intelligence-driven antimicrobial peptide discovery. *Current
684 Opinion in Structural Biology*, 83:102733, December 2023. ISSN 0959-440X. doi: 10.1016/j.
685 sbi.2023.102733. Epub 2023 Nov 21.
- 686
- 687 P. Szymczak, W. Zarzecki, J. Wang, Y. Duan, J. Wang, L. P. Coelho, C. de la Fuente-Nunez, and
688 E. Szczerk. Ai-driven antimicrobial peptide discovery: Mining and generation. *Accounts of
689 Chemical Research*, 58(12):1831–1846, June 2025. ISSN 0001-4842. doi: 10.1021/acs.accounts.
690 0c00594. Epub 2025 Jun 3.
- 691
- 692 Paulina Szymczak, Marcin Możejko, Tomasz Grzegorzek, Radosław Jurczak, Marta Bauer, Damian
693 Neubauer, Karol Sikora, Michał Michalski, Jacek Sroka, Piotr Setny, Wojciech Kamysz, and
694 Ewa Szczerk. Discovering highly potent antimicrobial peptides with deep generative model
695 hydram. *Nature Communications*, 14(1):5129, 2023. doi: 10.1038/s41467-023-36994-z. URL
696 <https://doi.org/10.1038/s41467-023-36994-z>.
- 697
- 698 Marcelo DT Torres, Leo Tianlai Chen, Fangping Wan, Pranam Chatterjee, and Cesar de la Fuente-
699 Nunez. Generative latent diffusion language modeling yields anti-infective synthetic peptides.
700 *Cell Biomaterials*, 2025.
- 701
- 702 Daniel Veltri, Uday Kamath, and Amarda Shehu. Deep learning improves antimicrobial pep-
703 tide recognition. *Bioinformatics*, 34(16):2740–2747, 2018. doi: 10.1093/bioinformatics/
704 bty179. URL <https://academic.oup.com/bioinformatics/article/34/16/2740/4953367>.

- 702 Fangping Wan, Marcelo D. T. Torres, Jacqueline Peng, Cesar de la Fuente-Nunez, et al. Deep-
 703 learning-enabled antibiotic discovery through molecular de-extinction. *Nature Biomedical Engineering*, 8(7):854–871, 2024. doi: 10.1038/s41551-024-01201-x. URL <https://doi.org/10.1038/s41551-024-01201-x>.
- 704
- 705
- 706 Danqing Wang, Zeyu Wen, Fei Ye, Lei Li, and Hao Zhou. Accelerating antimicrobial peptide
 707 discovery with latent structure model (Issamp). *arXiv preprint arXiv:2212.09450*, 2022. URL
 708 <https://arxiv.org/abs/2212.09450>.
- 709
- 710 Xiaodong Wang. Compactifications of complete riemannian manifolds and their applications. 12
 711 2010.
- 712
- 713 Yi Wang and Zhiren Wang. Cw complex hypothesis for image data. In Ruslan Salakhutdinov, Zico
 714 Kolter, Katherine Heller, Adrian Weller, Nuria Oliver, Jonathan Scarlett, and Felix Berkenkamp
 715 (eds.), *Proceedings of the 41st International Conference on Machine Learning*, volume 235 of
 716 *Proceedings of Machine Learning Research*, pp. 51596–51613. PMLR, Jul 2024. URL <https://proceedings.mlr.press/v235/wang24bs.html>.
- 717
- 718 Yixin Wang et al. Understanding the latent space of diffusion models through the lens of riemannian
 719 geometry. *arXiv preprint*, 2024. URL <https://arxiv.org/abs/2307.12868>.
- 720
- 721 Robin Winter et al. Navigating chemical space with latent flows. *arXiv preprint arXiv:2206.06072*,
 722 2022.
- 723
- 724 Siyu Zhong et al. Chemspace: Interpretable and interactive chemical space exploration. *ChemRxiv*,
 725 2022. URL <https://chemrxiv.org/engage/chemrxiv/article-details/633ba619ea6a22b0ba07e841>.
- 726
- 727
- 728
- 729 **APPENDIX**
- 730
- 731 **A RELATED WORK**
- 732
- 733 **Riemannian latent geometry.** Deep generative models can be endowed with Riemannian struc-
 734 ture (do Carmo, 1992) by pulling back the ambient metric through the decoder Jacobian (Arvanitidis
 735 et al., 2018; 2020; Shao et al., 2018). This allows distances and geodesics to reflect data geometry
 736 rather than Euclidean latent coordinates. However, most approaches assume a single smooth man-
 737 ifold of fixed dimension. Evidence from both theory and experiments shows that data often lie on
 738 unions of manifolds with varying intrinsic dimension or CW-complex structures (Lou, 2023; Brown
 739 et al., 2022; Wang & Wang, 2024). In practice, existing methods either restrict the latent to very
 740 low dimensions (≤ 8) or inflate the metric with variance terms to ensure full rank (Arvanitidis et al.,
 741 2020; Detlefsen et al., 2022), but these ignore extrinsic geometry. Our model instead *decomposes*
 742 *the latent into Riemannian submanifolds of varying dimensions*, enabling principled geometry across
 743 heterogeneous regions.
- 744
- 745 **Brownian motion and random walks.** Latent Brownian motion has been used as a prior for VAEs
 746 (Kalatzis et al., 2020) and for Riemannian score-based modeling (De Bortoli et al., 2022). But these
 747 approaches rely on first-order updates. Convergence results for geodesic random walks show that
 748 correct Riemannian and sub-Riemannian Brownian motion requires second-order approximations
 749 (Schwarz et al., 2022; Herrmann et al., 2023). Our method explicitly incorporates this requirement,
 750 yielding diffusion-consistent walks where previous methods diverge.
- 751
- 752 **Tangent spaces and interpretability.** Tangent-space analysis has mostly been applied in vision,
 753 where interpretable latent directions are discovered in GANs or diffusion models via Jacobian or
 754 eigen decompositions (Shen et al., 2020; Park et al., 2021; Wang et al., 2024; Alemi et al., 2023).
 755 Frames induced by augmentations provide another lens on local tangent geometry (Schneider et al.,
 2022). We are the first to provide an *interpretable tangent space in peptide sequence models*, where
 tangent vectors correspond directly to biologically meaningful mutations.

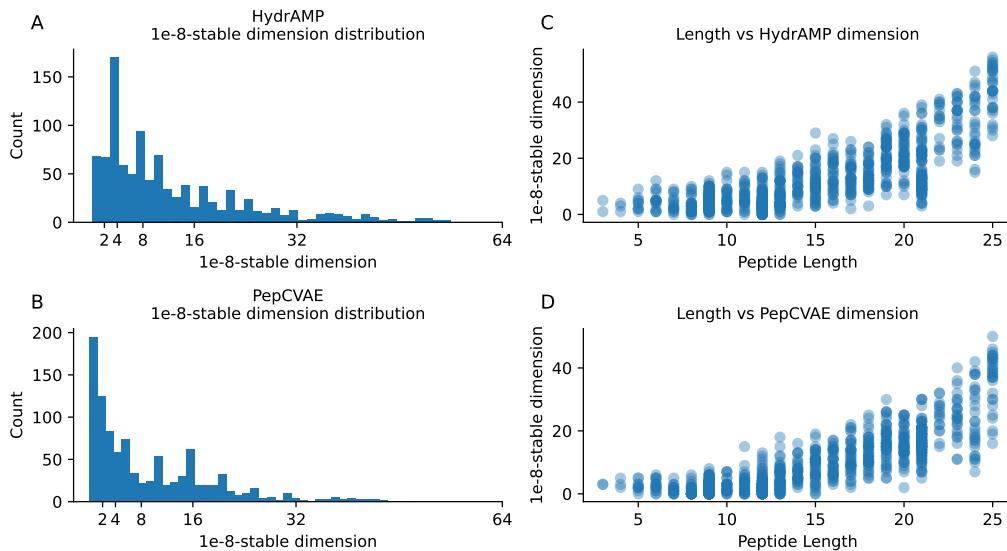
756 **Geodesics and potentials.** Geodesics are widely used for interpolation and counterfactual reasoning in latent space (Pegios et al., 2024; Blondel et al., 2024). Yet these are typically free geodesics. 757 We extend the concept with *potentials*, leveraging the Jacobi metric (Gibbons, 2015) so that peptide 758 traversals account for both geometry and biochemical preferences. 759

760 **Applications in molecules and proteins.** Geometry-aware latent models have been used for 761 chemical-space exploration (Zhong et al., 2022; Winter et al., 2022), molecule optimization (Feng 762 et al., 2021), and protein sequence modeling (Cao et al., 2022; Detlefsen et al., 2022). Our approach 763 complements these by combining: (i) varying-dimension latent decomposition, (ii) second- 764 order consistent Brownian walks, (iii) interpretable tangent spaces via mutations, and (iv) potential- 765 augmented geodesics tailored to peptide design. 766

767

768 B κ -STABLE DIMENSION OF THE PEPCVAE (DAS ET AL., 2018) AND 769 HYDRAMP SZYMCZAK ET AL. (2023) MODELS

770



791 **Figure 4: Stable rank and peptide statistics.** **A–B)** Distributions of the κ -stable dimension ($\kappa = 792 10^{-8}$) for HydrAMP (A) and PepCVAE (B) across sampled peptides. **C–D)** Scatter plots of peptide 793 length versus κ -stable dimension for HydrAMP (C) and PepCVAE (D), revealing a clear positive 794 correlation: longer peptides tend to yield higher stable dimensions. 795

796 To quantify how the κ -stable dimension varies across the latent space, we sampled 1,000 points 797 from the HydrAMP (Szymczak et al., 2023) training set and computed their κ -stable dimensions 798 under both the PepCVAE (Das et al., 2018) and HydrAMP (Szymczak et al., 2023) models. We 799 set $\kappa = 10^{-8}$, corresponding to the precision of the `float32` format commonly used in neural 800 network computations. This choice ensures that no eigenvalue of the inverse metric tensor exceeds 801 10^8 , thereby avoiding numerical instabilities. 802

803 As shown in Figure 4A, the κ -stable dimension was always strictly below the latent dimensionality 804 (64) of both models. This indicates that the effective local dimensionalities of the peptide spaces 805 are substantially smaller than the nominal latent dimension. Furthermore, when comparing peptide 806 length (Figure 4B) to κ -stable dimension (Figure 4C), we observe a clear positive correlation: longer 807 peptides systematically yield higher κ -stable dimensions. Intuitively, this suggests that longer 808 sequences admit more locally meaningful perturbations, which naturally translate into a richer set of 809 candidate substitutions. This observation further justifies our MUTANG strategy (§2.4), as it 810 allocates a larger and more diverse mutation pool precisely where biological sequence length provides 811 greater combinatorial flexibility. 812

810 **C CONSTRUCTION OF κ -STABLE MANIFOLDS**
 811

812 In this section we will introduce a construction of κ -stable Riemannian submanifolds, namely
 813

814
$$\mathbb{M}^\kappa = \{M_z^\kappa : z \in \mathcal{Z}\}, \quad M_z^\kappa = (W_z^\kappa, G_{\text{Dec}}),$$

815 where each $W_z^\kappa \ni z$ is an open affine submanifold of *maximal dimension* (denoted k_z^κ) through z ,
 816 such that the pullback metric G_{Dec} restricted to W_z^κ , denoted $G_{\text{Dec}}|_{W_z^\kappa}$, has full rank and satisfies
 817 the κ -stability condition

818
$$\inf_{\substack{v \in T_z W_z^\kappa \\ \langle v, v \rangle_{\mathbb{R}^d} = 1}} \langle v, v \rangle_z^{\text{Dec}} > \kappa^2.$$

 819
 820

821 For this, we will use the truncated-SVD of a flattened decoder Jacobian $J_{\hat{\text{Dec}}}$. Let
 822

823
$$J_{\hat{\text{Dec}}}(z) = U(z) \Sigma(z) V(z)^\top$$

824 be the thin SVD of the decoder Jacobian, with singular values $\sigma_0(z) \geq \sigma_1(z) \geq \dots \geq 0$, and
 825 $\Sigma(z) = \text{diag}((\sigma_0, \sigma_1, \dots, \sigma_{d-1}))$. Now let us note that:
 826

827
$$G_{\text{Dec}}(z) = J_{\hat{\text{Dec}}}(z)^\top J_{\hat{\text{Dec}}}(z) = V(z) \Sigma^2(z) V(z)^\top, \quad (7)$$

828 and define the κ -stable dimension:
 829

830
$$k_z^\kappa = \#\{i : \sigma_i(z)^2 > \kappa\}.$$

832 Truncating SVD-decomposition to first k_z^κ singular values gives
 833

834
$$U^\kappa(z) \in \mathbb{R}^{(LA) \times k_z^\kappa}, \quad \Sigma_z^\kappa \in \mathbb{R}^{k_z^\kappa \times k_z^\kappa}, \quad V^\kappa(z) \in \mathbb{R}^{d \times k_z^\kappa}, \quad (8)$$

835 with truncated Jacobian:
 836

837
$$J_{\hat{\text{Dec}}}^\kappa = U(z)^\kappa \Sigma(z)^\kappa V(z)^\kappa \in \mathbb{R}^{LA \times d}.$$

838 We then define the affine subspace (together with its parametrization)

839
$$\mathcal{V}_z^\kappa = \{\phi_z^\kappa(x) : x \in \mathbb{R}^{k_z^\kappa}\}, \quad \phi_z^\kappa(x) = z + V^\kappa(z)x.$$

 840

841 Restricting the decoder to \mathcal{V}_z^κ gives $\text{Dec}_z^\kappa = \hat{\text{Dec}} \circ \phi_z^\kappa$, with Jacobian
 842

843
$$J_{\text{Dec}_z^\kappa}(0) = U^\kappa(z) \Sigma^\kappa(z), \quad (9)$$

844 which has full column rank k_z^κ . By the inverse function theorem, there exists a neighborhood $\hat{W}_z^\kappa =$
 845 $B(0, r_z^\kappa)$ such that $\text{Dec}_z^\kappa(\hat{W}_z^\kappa)$ is a smooth k_z^κ -dimensional manifold. Its pullback metric is
 846

847
$$G_{\text{Dec}_z^\kappa}(0) = (\Sigma^\kappa(z))^2, \quad (10)$$

 848

849 with eigenvalues $\{\sigma_i(z)^2 : \sigma_i(z)^2 > \kappa\}$, all $\geq \kappa$.
 850

851 Finally, set $W_z^\kappa = \phi_z^\kappa(\hat{W}_z^\kappa)$. Then
 852

853
$$M_z^\kappa = (W_z^\kappa, G_{\text{Dec}}),$$

 854

855 and $(\phi_z^\kappa)^{-1}$ is a diffeomorphic isometry between $(W_z^\kappa, G_{\text{Dec}})$ and $(\hat{W}_z^\kappa, G_{\text{Dec}_z^\kappa})$. From this and
 856 Eq. 10 it follows that M_z^κ satisfies the κ -stability condition.
 857

858 The maximality of k_z^κ follows from the fact that if there existed an affine subspace \bar{W}_z^κ with
 859 $\dim(\bar{W}_z^\kappa) > k_z^\kappa$ such that $G_{\text{Dec}}|_{\bar{W}_z^\kappa}$ has full rank and satisfies the κ -stability condition, then we
 860 could define
 861

862
$$\bar{W}_z^\kappa \cap (W_z^\kappa)^\perp = \{w \in \bar{W}_z^\kappa : \forall v \in W_z^\kappa, \langle v, w \rangle^{\text{Euc}} = 0\} \subset \bar{W}_z^\kappa,$$

 863

864 as the subspace of \bar{W}_z^κ orthogonal to W_z^κ . Since $\dim(\bar{W}_z^\kappa) > \dim(W_z^\kappa)$, it follows that
 865 $\dim(\bar{W}_z^\kappa \cap (W_z^\kappa)^\perp) > 0$. By construction
 866

867
$$\bar{W}_z^\kappa \cap (W_z^\kappa)^\perp \subset \text{span}\{V(z)_{:, k_z^\kappa}, \dots, V(z)_{:, d}\}.$$

 868

869 what implies that for all $v \in \bar{W}_z^\kappa \cap (W_z^\kappa)^\perp$ it holds that
 870

864

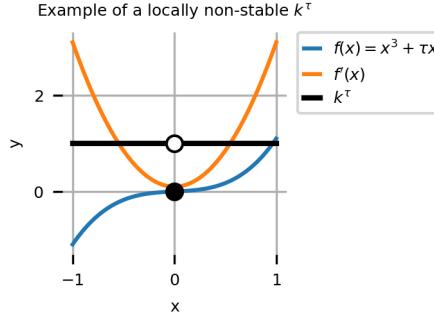
865
866
$$v = a_{k_z^\kappa} V(z)_{:,k_z^\kappa} + \cdots + a_{d-1} V(z)_{:,d-1},$$

867
868
869 for some $a_{k_z^\kappa}, \dots, a_{d-1} \in \mathbb{R}$. Now take $v \in \bar{W}_z^\kappa \cap (W_z^\kappa)^\perp$ such that $\langle v, v \rangle^{\text{Euc}} = 1$. Equation 7 then
implies

870
871
872
$$\langle v, v \rangle_z^{\text{Dec}} = a_{k_z^\kappa}^2 \sigma_{k_z^\kappa}^2 \|V(z)_{:,k_z^\kappa}\|^2 + \cdots + a_{d-1}^2 \sigma_{d-1}^2 \|V(z)_{:,d-1}\|^2 \quad (11)$$

873
874
$$\leq \kappa \left(a_{k_z^\kappa}^2 \|V(z)_{:,k_z^\kappa}\|^2 + \cdots + a_{d-1}^2 \|V(z)_{:,d-1}\|^2 \right) \quad (12)$$

875
876
$$= \kappa, \quad (13)$$

877
878 which contradicts the assumption that \bar{V}_z^κ satisfies the κ -stability condition.879
880 **Note.** This construction cannot be replaced by a direct Frobenius theorem argument, since $k_z^\kappa = \ell$
at a point does not imply that k^κ is constant in a neighborhood (see Appendix C.1).881
882
883
884
885 C.1 ON THE LOCAL INSTABILITY OF κ -STABLE DIMENSION886
887
888 Recall that $k_z^\kappa = \#\{i : \sigma_i(z) > \sqrt{\kappa}\}$ counts the number of singular values of $J_{\text{Dec}}(z)$ exceeding
889 the threshold $\sqrt{\kappa}$. While k_z^κ is well defined at every point z , it need not be locally constant. In
890 particular, singular values of $J_{\text{Dec}}(z)$ depend continuously on z , but they can cross the threshold $\sqrt{\kappa}$
891 arbitrarily close to a given point. Hence, even if $k_z^\kappa = \ell$ at some z , there may exist nearby points z'
892 with $k_{z'}^\kappa > \ell$ (see Figure 5).

- 918 1. $B_0 = x \in M'$ almost surely, for some fixed starting point $x \in M$.
 919
 920 2. The sample paths $t \mapsto B_t$ are almost surely continuous and adapted to the filtration $\{\mathcal{F}_t\}$.
 921
 922 3. For every smooth function $f \in C^\infty(M)$, the process

923
$$f(B_t) - f(B_0) - \frac{1}{2} \int_0^t (\Delta_g f)(B_s) ds$$

 924
 925

926 is a real-valued local martingale, where Δ_g denotes the Laplace–Beltrami operator associated
 927 with g .
 928

- 929 4. The generator of B_t is $\frac{1}{2}\Delta_g$, i.e.
 930

931
$$\lim_{t \rightarrow 0} \frac{\mathbb{E}[f(B_t)] - f(x)}{t} = \frac{1}{2}(\Delta_g f)(x), \quad \forall f \in C^\infty(M).$$

 932
 933

934 **Extension to the non-compact case: smooth spherical-cap compactification.** Our manifolds
 935 of interest, $M_z^\kappa(\alpha)$, $\alpha \in (0, 1)$, $z \in \mathcal{Z}$, are open subsets of M_z^κ and therefore non-compact.
 936 To define Brownian motion in this setting, one needs to control the behaviour of paths near the
 937 boundary. Classical approaches include: (i) compactification (Wang, 2010), (ii) stopping the process
 938 at the boundary (Hsu, 2002b), or (iii) reflecting it (Du & Hsu, 2021). In our work we adopt a
 939 compactification strategy via a smooth spherical cap (see Figure 6), followed by stopping on the
 940 boundary of a natural embedding of $M_z^\kappa(\alpha)$.
 941

942 Concretely, let k_z^κ be the κ -stable dimension and consider the unit sphere

943
$$S_1^{k_z^\kappa} \subset \mathbb{R}^{k_z^\kappa+1}.$$

 944
 945

946 Take an atlas of this sphere consisting of two charts (U_1, ψ_1) , (U_2, ψ_2) such that $U_1, U_2 \subset \mathbb{R}^{k_z^\kappa}$
 947 $\psi_1(U_1) \cup \psi_2(U_2) = S_1^{k_z^\kappa}$, with $\alpha \bar{W}_z^\kappa \subset U_1$ ($\alpha \bar{W}_z^\kappa = \{\alpha v : v \in \bar{W}_z^\kappa\}$) and
 948

949
$$\psi_1 \circ \psi_2^{-1}(U_2) \cap \bar{W}_z^\kappa = \emptyset.$$

 950

951 Let G_{ψ_i} be the pullback metric on U_i induced by ψ_i .
 952

953 Choose a smooth bump function $b \in C^\infty$ such that

954
$$b \equiv 1 \quad \text{on } (\frac{2}{3}\alpha + \frac{1}{3})\bar{W}_z^\kappa, \quad b \equiv 0 \quad \text{on } (\frac{1}{3}\alpha + \frac{2}{3})\bar{W}_z^\kappa.$$

 955

956 We then define the compactified manifold $\overline{M_z^\kappa}(\alpha)$ with charts $\{(U_1, \psi_1), (U_2, \psi_2)\}$ and Riemannian
 957 metric on U_1 given by
 958

959
$$G = b G_{\text{Dec}_z^\kappa} + (1 - b) G_{\psi_1},$$

 960

961 where

962
$$\text{Dec}_z^\kappa(x) = \begin{cases} \text{Dec}_z^\kappa(x), & x \in \alpha \bar{W}_z^\kappa, \\ 0, & \text{otherwise.} \end{cases}$$

 963
 964

965 By construction, there is a natural isometric embedding
 966

967
$$M_z^\kappa(\alpha) \hookrightarrow \overline{M_z^\kappa}(\alpha).$$

 968

969 This compactification allows us to invoke convergence results for Brownian motion on compact
 970 manifolds, while ensuring that in the region of interest the geometry coincides with that of the
 971 original κ -stable manifold.

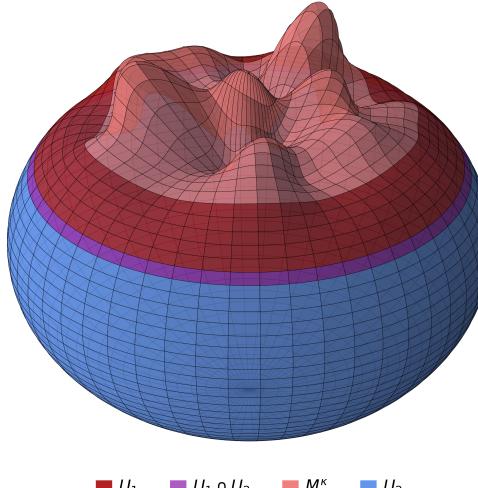


Figure 6: **Sphere-cap manifold compactification** $\overline{M_z^\kappa}$. The figure illustrates how the M_z^κ manifold is compactified by capping off the spherical domain with the appropriate boundary conditions.

Theorem (Main Theorem). *Let $(Z_i^\epsilon)_{i \geq 0}$ be the sequence produced by Algorithm 1, for $M_z^\kappa(\alpha)$ with $\alpha \in (0, 1)$ and diffusion horizon $T > 0$, and define its continuous-time interpolation*

$$Z^\epsilon(t) := Z_{\lfloor \epsilon^{-2}t \rfloor}^\epsilon, \quad t \geq 0.$$

Let $R_z^\kappa = d_{M_z^\kappa}(z, \alpha W_z^\kappa)^c$, and suppose $L \geq 1$ satisfies

$$\sup_{x \in M_z^\kappa(\alpha)} \text{Ric}_{M_z^\kappa}(x) \geq -L^2.$$

Then for $T < \frac{(R_z^\kappa)^2}{4k_z^\kappa L}$, as $\epsilon \rightarrow 0$, the process X^ϵ converges in distribution to Riemannian Brownian motion stopped at the boundary of $M_z^\kappa(\alpha)$, with respect to the Skorokhod topology, on a set $C_{\kappa,z}^T \subset \Omega$ such that

$$\mathbb{P}(C_{\kappa,z}^T) \geq 1 - \exp\left(-\frac{(R_z^\kappa)^2}{32T}\right).$$

Proof. Let $\overline{M_z^\kappa(\alpha)}$ be the spherical-cap compactification of $M_z^\kappa(\alpha)$. Schwarz et al. (2022) showed that the non-stopped version of Algorithm 1, extended to $\overline{M_z^\kappa(\alpha)}$, produces a process \overline{Z}^ϵ that converges to B in the Skorokhod topology on $\overline{M_z^\kappa(\alpha)}$.

For a closed set A define the exit time

$$T_A^X = \inf\{t \in [0, T] : X(t) \notin A\} \wedge T$$

Both $T_{\alpha W_z^\kappa}^{\overline{X}^\epsilon}$ and $T_{\alpha W_z^\kappa}^B$ are valid stopping times (by right-continuity and because $(W_z^\kappa)^c$ is closed), and (from the construction of the Algorithm 1):

$$Z^\epsilon = \overline{Z}^{\overline{X}^\epsilon}_{\cdot \wedge \alpha W_z^\kappa}$$

Convergence in Skorokhod topology does not automatically imply convergence of stopping times (see Appendix D.1 for a counterexample). However, for the high-probability event

$$C_{\kappa,z}^T = \{\omega \in \Omega : \forall t \in [0, T], B_t(\omega) \in \alpha W_z^\kappa\},$$

1026 we have that $\exists \epsilon_0$ such that $\forall \epsilon > \epsilon_0$,

$$1027 \quad T_{\alpha W_z^\kappa}^{\overline{X^\kappa}} \equiv T, \quad T_{\alpha W_z^\kappa}^B \equiv T,$$

1029 and thus convergence holds on $C_{\kappa, z}^T$.

1030 It remains to lower-bound $\mathbb{P}(C_{\kappa, z}^T)$. Observe that

$$1032 \quad \Omega \setminus C_{\kappa, z}^T \subset D_{\kappa, z}^T = \left\{ \omega \in \Omega : \exists t \in [0, T], d_{M_z^\kappa}(B_t(\omega), z) \geq R_z^\kappa \right\}.$$

1033 Hence

$$1034 \quad \mathbb{P}(C_{\kappa, z}^T) \geq 1 - \mathbb{P}(D_{\kappa, z}^T).$$

1035 **Lemma 1** (Exit-time bound). *Suppose $L \geq 1$ satisfies*

$$1037 \quad \sup_{x \in M_z^\kappa(\alpha)} \text{Ric}_{M_z^\kappa}(x) \geq -L^2.$$

1038 Let $T_{R_z^\kappa}$ be the first exit time of Riemannian Brownian motion from

$$1039 \quad B_{M_z^\kappa}(B_0, R_z^\kappa) = \{x \in M_z^\kappa(\alpha) : d_{M_z^\kappa}(B_0, x) < R_z^\kappa\}.$$

1040 Then

$$1041 \quad \mathbb{P}(T_{R_z^\kappa} \leq T) \leq \exp\left(-\frac{(R_z^\kappa)^2}{8T} \left(1 - \frac{2Tk_z^\kappa L}{(R_z^\kappa)^2}\right)^2\right).$$

1042 *Proof.* Let $r(x) = d_{M_z^\kappa}(B_0, x)$ and write $r_t := r(B_t)$. On M_z^κ the function r is smooth. The
1043 semimartingale decomposition of r_t (see, e.g., (Hsu, 2002a, Eq. (3.6.1))) gives

$$1044 \quad r_t^2 \leq 2 \int_0^t r_s d\beta_s + \int_0^t r_s \Delta r(B_s) ds + t,$$

1045 where β is a real Brownian motion adapted to B , and we have used that the local time term is
1046 nonnegative and can be dropped to obtain an inequality.

1047 By the Laplacian comparison theorem ($\text{Ricci}(\cdot) \geq -(k_z^\kappa - 1)L^2$), for $r > 0$,

$$1048 \quad \Delta r \leq (k_z^\kappa - 1)L \coth(Lr) \leq (k_z^\kappa - 1)\left(L + \frac{1}{r}\right),$$

1049 hence

$$1050 \quad r \Delta r \leq (k_z^\kappa - 1)(Lr + 1).$$

1051 Up to the first exit time $T_{R_z^\kappa} := \inf\{t \geq 0 : r_t \geq R_z^\kappa\}$ we have $r_s \leq R_z^\kappa$, so

$$1052 \quad r_s \Delta r(B_s) \leq (k_z^\kappa - 1)(LR_z^\kappa + 1) \leq k_z^\kappa L + k_z^\kappa \leq 2k_z^\kappa L,$$

1053 using $L \geq 1$. Therefore, for $t = T_{R_z^\kappa} \wedge T$,

$$1054 \quad r_t^2 \leq 2 \int_0^t r_s d\beta_s + 2k_z^\kappa L t + t \leq 2 \int_0^t r_s d\beta_s + 2k_z^\kappa L t + t. \quad (14)$$

1055 On the event $\{T_{R_z^\kappa} \leq T\}$ we have $t = T_{R_z^\kappa}$ and $r_t \geq R_z^\kappa$, hence from equation 14

$$1056 \quad (R_z^\kappa)^2 \leq 2 \int_0^{T_{R_z^\kappa}} r_s d\beta_s + 2k_z^\kappa L T.$$

1057 Rearranging,

$$1058 \quad \int_0^{T_{R_z^\kappa}} r_s d\beta_s \geq \frac{(R_z^\kappa)^2 - 2k_z^\kappa L T}{2}.$$

1059 Set $M_t := \int_0^t r_s d\beta_s$, a continuous martingale with quadratic variation $\langle M \rangle_t = \int_0^t r_s^2 ds \leq (R_z^\kappa)^2 t$
1060 up to time $T_{R_z^\kappa}$. By the Dambis–Dubins–Schwarz theorem there exists a standard Brownian motion
1061 W such that $M_{T_{R_z^\kappa}} = W_\eta$ with $\eta = \langle M \rangle_{T_{R_z^\kappa}} \leq (R_z^\kappa)^2 T_{R_z^\kappa} \leq (R_z^\kappa)^2 T$ on $\{T_{R_z^\kappa} \leq T\}$. Thus,

$$1062 \quad \{T_{R_z^\kappa} \leq T\} \subset \left\{ W_\eta \geq \frac{(R_z^\kappa)^2 - 2k_z^\kappa L T}{2} \right\}.$$

1063 Using the Gaussian tail bound together with $\eta \leq (R_z^\kappa)^2 T$ (and the reflection principle),

$$1064 \quad \mathbb{P}(T_{R_z^\kappa} \leq T) \leq \exp\left(-\frac{((R_z^\kappa)^2 - 2k_z^\kappa L T)^2}{8(R_z^\kappa)^2 T}\right) = \exp\left(-\frac{(R_z^\kappa)^2}{8T} \left(1 - \frac{2k_z^\kappa L T}{(R_z^\kappa)^2}\right)^2\right),$$

1065 which is the claimed bound. \square

1080 Applying the lemma, if $T < \frac{(R_z^\kappa)^2}{4k_z^\kappa L}$, then
 1081

$$1082 \mathbb{P}(D_{\kappa,z}^T) = \mathbb{P}(T_{R_z^\kappa} \leq T) \leq \exp\left(-\frac{(R_z^\kappa)^2}{32T}\right),$$

1083 which yields
 1084

$$1085 \mathbb{P}(C_{\kappa,z}^T) \geq 1 - \exp\left(-\frac{(R_z^\kappa)^2}{32T}\right).$$

□

1088 **Remark.** The lemma shows that the probability of exiting the ball of radius R_z^κ before time T
 1089 decays exponentially in $\frac{(R_z^\kappa)^2}{T}$, up to curvature- and rank-dependent constants. Intuitively, this means
 1090 that with overwhelming probability the Riemannian Brownian motion (and hence our random walk
 1091 in the $\epsilon \rightarrow 0$ limit) remains confined inside $B_{M_z^\kappa}(z, R_z^\kappa)$ for all $t \leq T$. This high-probability
 1092 control is what allows us to restrict attention to the event $C_{\kappa,z}^T$ in the proof of Theorem 1.
 1093

1094 D.1 ON STOPPING TIMES AND SKOROHOD CONVERGENCE

1096 An important subtlety in the proof of Theorem 1 is that convergence of processes in the Skorohod
 1097 topology does not, in general, imply convergence of associated stopping times. Figure 7 illustrates
 1098 this phenomenon with a simple deterministic example.

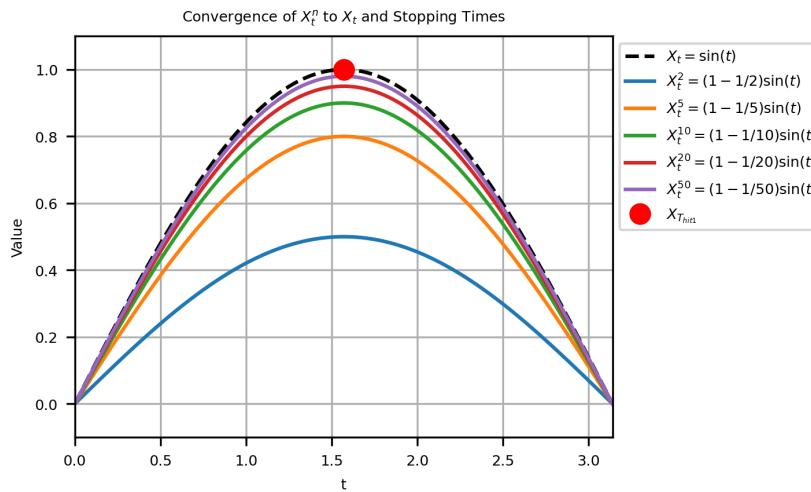
1099 Let $X_t = \sin(t)$, and consider the approximating sequence of processes
 1100

$$1101 X_t^n = \left(1 - \frac{1}{n}\right) \sin(t).$$

1102 Clearly $X^n \rightarrow X$ uniformly on compact time intervals, hence also in the Skorohod topology. How-
 1103 ever, the stopping time defined as

$$1104 T = \inf\{t \geq 0 : X_t = 1\}$$

1105 does not converge along this sequence. Indeed, $T = \pi/2$ for X , but for every finite n , the process
 1106 X^n never reaches 1 and therefore $T^n = \infty$. Thus, despite $X^n \rightarrow X$ in Skorohod topology, we have
 1107 $T^n \not\rightarrow T$.



1125 **Figure 7: Skorohod convergence does not imply convergence of stopping times.** The black
 1126 dashed curve shows $X_t = \sin(t)$, which reaches 1 at $t = \pi/2$ (red dot). The colored curves show
 1127 approximations $X_t^n = (1 - \frac{1}{n}) \sin(t)$, converging uniformly to X_t . However, none of the X_t^n
 1128 ever reach 1, so their stopping times for hitting level 1 are infinite. This illustrates that Skorohod
 1129 convergence of processes does not guarantee convergence of stopping times defined by hitting closed
 1130 sets.

1131 In the proof of Theorem 1 we avoid this issue by restricting to the high-probability set $C_{\kappa,z}^T$ where
 1132 the Brownian path remains in the interior of the ball. On this event the stopping times agree with T ,
 1133 ensuring consistency with the limiting process.

1134 E SORBES IMPLEMENTATION DETAILS
11351136 We now provide practical details for the implementation of the SORBES algorithm.
11371138 **Approximating Jacobian.** The decoder Jacobian $J_{\hat{\text{Dec}}}$ can be obtained exactly using one decoder
1139 forward pass and LA backward passes, where LA is a dimensionality of ambient space. To reduce
1140 computational cost, we approximate it via finite differences. Specifically, the i -th column of the
1141 Jacobian is approximated as
1142

1143
$$\frac{\partial \hat{\text{Dec}}}{\partial z_i}(z) \approx \frac{\hat{\text{Dec}}(z + \varepsilon e_i) - \hat{\text{Dec}}(z)}{\varepsilon}, \quad (15)$$

1144

1145 where e_i denotes the i -th standard basis vector in the latent space \mathcal{Z} and $\varepsilon > 0$ is a small pertur-
1146 bation parameter. This approximation requires only $d + 1$ decoder forward passes, where d is a
1147 dimensionality of latent space, providing a substantial reduction in computational overhead.
11481149 **Sampling a unit tangent direction.** We adopt the efficient implementation of Schwarz et al.
1150 (2022), which exploits a thin SVD of the decoder Jacobian to orthogonalize tangent directions.
11511152 **Approximating $\Gamma(z)[v, v]$.** Computing Christoffel symbols directly requires evaluating first
1153 derivatives of the metric, which is computationally expensive and numerically unstable. Instead,
1154 we use an extrinsic approach: the covariant derivative of a curve can be obtained from its Euclidean
1155 acceleration in the ambient space, projected back onto the tangent space (do Carmo, 1992). Con-
1156 cretely, for a point $z \in \mathcal{Z}$, $z' \in W_z^\kappa$ and a probe radius $\rho > 0$, we approximate the extrinsic
1157 acceleration of a decoded curve along $v \in T_z M_z^\kappa$ by a second-order central difference:
1158

1159
$$a_{\text{ex}}^{z, \kappa}(v; z', \rho) \approx \frac{\hat{\text{Dec}}(z' + \rho v) - 2 \hat{\text{Dec}}(z') + \hat{\text{Dec}}(z' - \rho v)}{\rho^2}. \quad (16)$$

1160

1161 Projecting back to the latent tangent space using the Moore–Penrose pseudoinverse of the truncated
1162 Jacobian yields an efficient approximation of the Christoffel correction:
1163

1164
$$\Gamma(z')[v, v] \approx c(v; z', \rho) = J_{\hat{\text{Dec}}}^\kappa(z')^+ a_{\text{ex}}(v; z', \rho), \quad c(v; z', \rho) \in \mathbb{R}^{k_z^\kappa}. \quad (17)$$

1165

1166 **Adaptive step size ϵ .** Since the computation of Christoffel symbols involves the (pseudo)inverse
1167 of the Jacobian, small values of κ may amplify numerical noise. In this case, the update
1168

1169
$$z_i^\epsilon = z_{i-1}^\epsilon + \epsilon v - \epsilon^2 \Gamma[v, v] = z_{i-1}^\epsilon + \Delta(\epsilon),$$

1170

1171 may become unreasonably large in the ambient Euclidean metric on U_z^κ , making the algorithm
1172 unstable.1173 To control this, we adapt the step size ϵ so that the update norm never exceeds a predefined threshold
1174 Δ_{max} :

1175
$$\epsilon_{\Delta_{\text{max}}} = \min \left\{ \epsilon' > 0 : \|\Delta(\epsilon')\| \geq \Delta_{\text{max}} \right\} \vee \epsilon,$$

1176

1177 where $a \vee b = \max\{a, b\}$. This guarantees that the step size is never greater than the nominal ϵ , but
1178 shrinks adaptively whenever the second-order correction is large.
11791180 **SORBES (Stable/Efficient).** We refer to the resulting algorithm with adaptive step size control
1181 and extrinsic approximation of Christoffel symbols (Sec. 2.3) as SORBES-SE. This variant is nu-
1182 mERICALLY stable in ill-conditioned regions of the decoder geometry while preserving the efficiency
1183 of the original scheme.
11841185 F MUTANG - MUTATION ENUMERATION IN TANGENT SPACE
1186

1187 The detailed description of MUTANG algorithm is presented in Algorithm 5.

1188

Algorithm 4 SORBES-SE

1189 **Require:** $z \in \mathcal{Z}$, $\kappa \geq 0$, step size ϵ , diffusion time T , maximum number of steps STEP_{\max} ,
 1190 $\alpha = 0.99$, stability threshold $\Delta_{\max} = 0.5$

1191 1: $W_z^\kappa, G_{\text{Dec}} \leftarrow M_z^\kappa$

1192 2: $z_0^\epsilon \leftarrow z$,

1193 3: $\text{stopped} \leftarrow \text{False}$,

1194 4: $\sigma \leftarrow 0$, (σ tracks diffusion time)

1195 5: $\text{step} = 0$,

1196 6: **while** $\sigma < T$ and $\text{step} < \text{STEP}_{\max}$ **do**

1197 7: Sample a unit tangent direction $\bar{v} \in S_z^\kappa = \{u \in T_z M_z^\kappa : \langle u, u \rangle_z^{\text{Dec}} = 1\}$

1198 8: Set $v \leftarrow \sqrt{k_z^\kappa} \bar{v}$

1199 9: **if** not stopped **then**

1200 10: Compute trial update $\Delta(\epsilon) \leftarrow \epsilon v - \epsilon^2 \Gamma(z)[v, v]$

1201 11: **Adaptive adjustment:** If $\|\Delta(\epsilon)\| > \Delta_{\max}$, shrink step size:
 $\epsilon \leftarrow \min\{\epsilon' > 0 : \|\Delta(\epsilon')\| \leq \Delta_{\max}\}$

1202 12: and recompute $\Delta(\epsilon)$.

1203 13: Update latent coordinate: $z_i^\epsilon \leftarrow z_{i-1}^\epsilon + \Delta(\epsilon)$ (update of diffusion time)

1204 14: **if** $z_i^\epsilon \notin W_z^\kappa(\alpha)$ **then**

1205 15: $\text{stopped} \leftarrow \text{True}$

1206 16: **end if**

1207 17: **else**

1208 18: $z_i^\epsilon \leftarrow z_{i-1}^\epsilon$ (absorbing state)

1209 19: **end if**

1210 20: $\text{step} \leftarrow \text{step} + 1$ (step update)

1211 21: **end while**

1212 22: **return** $(z_i^\epsilon)_{0 \leq i \leq \text{step}}, \sigma$

1217

Algorithm 5 MUTANG

1218 **Require:** latent $z \in \mathcal{Z}$; $\kappa \geq 0$; token threshold θ_{tok} .

1219 1: Set U_z^κ, k_z^κ as in Equation 8

1220 2: $\mathbf{p} \leftarrow \mathbf{p}(z)$

1221 3: $\mathcal{P} \leftarrow \emptyset$

1222 4: **for** $j = 1$ to k_z^κ **do**

1223 5: $\Delta_{\text{Dec}}^{(j)} \leftarrow \text{reshape}(U^\kappa(z)_{:,j}, (L, A))$

1224 6: **for** $\ell = 1$ to L **do**

1225 7: **for each** $a \in \mathcal{A}$ **do**

1226 8: **if** $|\Delta_{\text{Dec}}_{\ell,a}^{(j)}| \geq \theta_{\text{tok}}$ **then**

1227 9: $\mathcal{P} \leftarrow \mathcal{P} \cup \{(\ell, a)\}$

1228 10: **end if**

1229 11: **end for**

1230 12: **end for**

1231 13: **end for**

1232 14: **for** $\ell = 1$ to L **do**

1233 15: $S_\ell \leftarrow \{a : (\ell, a) \in \mathcal{P}\} \cup \{\mathbf{p}_\ell\}$ (identity included)

1234 16: **end for**

1235 17: $\mathcal{C}(\mathbf{p}(z)) \leftarrow \prod_{\ell=1}^L S_\ell$

1236 18: **return** $\mathcal{C}(\mathbf{p}(z))$

1238

1239

G POGS

1240

1241 Below we present the details of PoGS training and evaluation metrics:

- 1242 PoGS hyperparameters:
 1243
- 1244 • PoGS without potential: $\lambda = 0$ and $\mu = 0.1$,
 - 1245 • Full PoGS: $\lambda = 0.01$ and $\mu = 0.1$.
 - 1246 • All: $\theta_{\text{pot}} = 5$.

1248 PoGS metrics:

- 1250
- 1251 • chord ambient length:
- $$1252 \quad \sum_{k=0}^{N-1} \|X_{k+1} - X_k\|_2$$
- 1253
- 1254 • chord latent length:
- $$1255 \quad \sum_{k=0}^{N-1} \|z_{k+1} - z_k\|_2$$
- 1256

1260 For computation of seeds and wells, we excluded first and last 20% of a peptide path were excluded
 1261 to avoid trivial rediscovery.

1262 For each pair, the chord length N was determined dynamically as

$$1264 \quad N = \lfloor \rho \cdot \|z_a - z_b\|_2 \rfloor,$$

1265 where ρ is the point density hyperparameter (set to $\rho = 90$ in our experiments). This construction
 1266 guarantees that longer trajectories in the latent space are sampled more densely than shorter ones,
 1267 preserving a uniform resolution across geodesics of varying length. The geodesic points $\{z_i\}_{i=1}^n$
 1268 were optimized using the Adam optimizer with learning rate $\eta = 10^{-3}$ and weight decay 10^{-5} . We
 1269 applied a ReduceLROnPlateau scheduler, which decreased the learning rate by a factor of 0.8
 1270 whenever no improvement in the loss was observed for a number of iterations equal to the patience
 1271 hyperparameter. The endpoints z_a and z_b were kept fixed throughout the optimization by zeroing
 1272 their gradients at every step.

1274 Algorithm 6 PoGS

1275 **Require:** seeds z_a, z_b , potential function Φ , nb of segments N , weights λ, μ , steps T ,

- 1276 1: Initialize $z_0 \leftarrow z_a, z_N \leftarrow z_b, z_{1:N-1}$ by linear interpolation in latent space
 - 1277 2: **for** $t = 1$ to T **do**
 - 1278 3: $X_k \leftarrow \log(\text{Dec})(z_k)$ for $k = 0..N$
 - 1279 4: Compute energy $\mathcal{E}_\lambda(Z)$ as in equation 6
 - 1280 5: Take a gradient step on $z_{1:N-1}$ to minimize $\mathcal{E}_\lambda(Z)$
 - 1281 6: **end for**
 - 1282 7: **return** $\{\text{p}(z_k)\}_{k=0}^N$
-

1285 H APEX-POTENTIAL FOR POGS

1287 The APEX predictor (Wan et al., 2024) estimates minimum inhibitory concentration (MIC) values
 1288 against 11 bacterial strains, but it operates on concrete peptide *sequences*. In Potential-Minimizing
 1289 Geodesic Search (PoGS), optimization proceeds over *latent-space chords*, i.e., intermediate points
 1290 z that decode to *position-factorized distributions* over peptides rather than single sequences:

$$1292 \quad \text{Dec}(z) \in \mathbb{R}^{L \times A},$$

1293 where L is the maximum peptide length and $A = 21$ is the amino-acid alphabet augmented with
 1294 padding. To enable PoGS, we first *distill* the sequence-level APEX potential into a surrogate that
 1295 accepts peptide *distributions*.

1296 **Dataset construction.** Peptides from the HydrAMP training set (Szymczak et al., 2023) were
 1297 encoded into latent codes z . In order to obtain multiple distributions from a single peptide, we then
 1298 created four clones z' of latent codes z . We applied a 2×2 perturbation scheme: two clones were
 1299 injected with Gaussian noise $N(0, 0.05)$, and two were left unchanged. Finally, these four latent
 1300 codes were decoded to $\text{Dec}(z')$, using a softmax scaling with temperature of 1.0 for one pair (noisy
 1301 and non-noisy) and a temperature of 1.5 to the other pair, resulting in four distributions per peptide.
 1302 This yielded 1,060,000 peptide distributions in total. For each $\text{Dec}(z')$, we enumerated the $N = 20$
 1303 most-probable sequences

$$(P_0(z'), \dots, P_{N-1}(z')) \quad \text{with probabilities} \quad (p_0(z'), \dots, p_{N-1}(z')),$$

1304 1305 applied APEX to each $P_i(z')$ to obtain MIC *vectors* $\text{MIC}_{P_i(z')} \in \mathbb{R}^{11}$, and defined the distribution's
 1306 1307 *expected* MIC vector via the probability-weighted average

$$\Phi_{\text{MIC}}^{\text{true}}(\text{Dec}(z')) = \sum_{i=0}^{N-1} \text{MIC}_{P_i(z')} \cdot \frac{p_i(z')}{\sum_{j=0}^{N-1} p_j(z')} \in \mathbb{R}^{11}.$$

1311 **Training protocol and standardization.** We split the dataset into 80% train, 10% validation, and
 1312 10% test in such a way that no two sets contain distributions originating from the same peptide. Let
 1313 $\mu, \sigma \in \mathbb{R}^{11}$ be the per-strain mean and standard deviation computed *on the training set*. Targets
 1314 were z-scored componentwise:

$$y^z = \frac{y - \mu}{\sigma}.$$

1316 1317 We trained an encoder-only transformer that *operates on distributions* $\text{Dec}(z) \in \mathbb{R}^{L \times A}$ and predicts
 1318 1319 z-scored MIC vectors in \mathbb{R}^{11} :

$$\Phi_{\text{MIC}}^{\text{model}} : \text{Dec}(z) \mapsto \mathbb{R}^{11}.$$

1320 1321 Architecture: three transformer encoder layers (four heads), embedding dimension 128, feed-
 1322 1323 forward dimension 256, dropout 0.05. Optimization used Adam (learning rate 10^{-4}) for 15 epochs
 1324 with mean-squared error (MSE) loss on z-scored targets:

$$\mathcal{L}_{\text{MSE}} = \frac{1}{11} \|\Phi_{\text{MIC}}^{\text{model}}(\text{Dec}(z)) - y^z\|_2^2.$$

1326 **Final potential used by PoGS.** PoGS operates on flattened *log-probabilities*. Let $X \in \mathbb{R}^{L \cdot A}$ be
 1327 1328 the flattened log-probability vector. We reconstruct a valid distribution using PyTorch-style opera-
 1329 1330 tions:

$$P(X) = \text{softmax}(\text{X. reshape}(L, A), \text{dim} = 1) \in [0, 1]^{L \times A},$$

1331 where `dim=1` is the amino-acid dimension. The surrogate outputs a z-scored MIC vector

$$\hat{m}^z(X) = \Phi_{\text{MIC}}^{\text{model}}(P(X)) \in \mathbb{R}^{11}.$$

1333 1334 Restricting to the three target *E. coli* strains (index set $\mathcal{I}_{E. \text{coli}}$), the scalar property potential used by
 1335 1336 PoGS is

$$\Phi(X) = \mathbf{1}^\top [\hat{m}^z(X)]_{\mathcal{I}_{E. \text{coli}}}.$$

I LE-BO HYPERPARAMETERS

1339 1340 We enumerate all hyperparameter values of our optimization algorithm LE-BO and all its sub-
 1341 1342 algorithms.

- 1342 • Algorithm 3 LE-BO - Local Enumeration Bayesian Optimization
 - 1343 – Trust region distance $d_{\text{trust}} = 2$.
 - 1344 – Number of ROBOT evaluations per iteration $k_{\text{ROBOT}} = 3$.
 - 1345 – Diversity threshold $d_{\text{ROBOT}} = 2$.
 - 1346 – Following the approach of Eberhardt et al. (2024), as a surrogate model, we use
 1347 a Gaussian Process GP with the Tanimoto similarity kernel Szedmak & Bach (2025),
 1348 applied to the MAP4 fingerprints of peptides Capecchi et al. (2020).
 - 1349 – Aquisition function GP. acquistion was chosen to be Log Expected Improvement.

- 1350 • Algorithm 2 LOCALENUMERATION
 1351 – $\kappa_{SORBES} = 0.01$.
 1352 – $\kappa_{MUTANG} = 10^{-6}$.
 1353 – Number of trajectories $M = 10$.
 1354 – Walk time budget $T_{\text{walk}} = 0.1$.
 1355 – Nominal step size $\epsilon = 0.1$.
 1356 – Mutation threshold θ_{mut}
 1357
 1358 • A probe radius $\rho > 0$ in the second-order central difference approximation of the extrinsic
 1359 acceleration (Equation 16) $\rho = 0.05$.
 1360
 1361 • Step of the finite-difference approximation of the decoder Jacobian (Equation 15) $\varepsilon = 0.05$.
 1362

1363 J WET-LAB VALIDATION
 1364

1365 J.1 PEPTIDE SYNTHESIS AND CHARACTERIZATION

1366 Peptides were synthesized on an automated peptide synthesizer (Symphony X, Gyros Protein Technologies) by standard Fmoc-based solid-phase peptide synthesis (SPPS) on Fmoc-protected amino acid-Wang resins (100–200 mesh). The following preloaded resins were employed with their respective loading capacities (100 μmol scale): Fmoc-Asn(Trt)-Wang Resin (0.510 mmol g^{-1}), Fmoc-His(Trt)-Wang Resin (0.480 mmol g^{-1}), Fmoc-Leu-Wang Resin (0.538 mmol g^{-1}), Fmoc-Lys(Boc)-Wang Resin (0.564 mmol g^{-1}), Fmoc-Phe-Wang Resin (0.643 mmol g^{-1}), Fmoc-Thr(tBu)-Wang Resin (0.697 mmol g^{-1}), Fmoc-Trp(Boc)-Wang Resin (0.460 mmol g^{-1}), Fmoc-Tyr(tBu)-Wang Resin (0.520 mmol g^{-1}). In addition to preloaded resins, standard Fmoc-protected amino acids were employed for chain elongation, including: Fmoc-Ala-OH, Fmoc-Cys(Trt)-OH, Fmoc-Glu(OtBu)-OH, Fmoc-Phe-OH, Fmoc-Gly-OH, Fmoc-His(Trt)-OH, Fmoc-Ile-OH, Fmoc-Lys(Boc)-OH, Fmoc-Leu-OH, Fmoc-Met-OH, Fmoc-Asn(Trt)-OH, Fmoc-Arg(Pbf)-OH, Fmoc-Ser(tBu)-OH, Fmoc-Thr(tBu)-OH, Fmoc-Val-OH, Fmoc-Trp(Boc)-OH, and Fmoc-Tyr(tBu)-OH. N,N-Dimethylformamide (DMF) was used as the primary solvent throughout synthesis. Stock solutions included: 500 mmol L^{-1} Fmoc-protected amino acids in DMF, a coupling mixture of HBTU (450 mmol L^{-1}) and N-methylmorpholine (NMM, 900 mmol L^{-1}) in DMF, and 20% (v/v) piperidine in DMF for Fmoc deprotection. After synthesis, peptides were deprotected and cleaved from the resin using a cleavage cocktail of trifluoroacetic acid (TFA)/triisopropylsilane (TIS)/dithiothreitol (DTT)/water (92.8% v/v, 1.1% v/v, 0.9% w/v, 4.8% w/w) for 2.5 hours with stirring at room temperature. The resin was removed by vacuum filtration, and the peptide-containing solution was collected. Crude peptides were precipitated with cold diethyl ether and incubated for 20 min at -20°C , pelleted by centrifugation, and washed once more with cold diethyl ether. The resulting pellets were dissolved in 0.1% (v/v) aqueous formic acid and incubated overnight at -20°C , followed by lyophilization to obtain dried peptides. For characterization, peptides were dried, reconstituted in 0.1% formic acid, and quantified spectrophotometrically. Peptide separations were performed on a Waters XBridge C₁₈ column (4.6 \times 50 mm, 3.5 μm , 120 \AA) at room temperature using a conventional high-performance liquid chromatography (HPLC) system. Mobile phases were water with 0.1% formic acid (solvent A) and acetonitrile with 0.1% formic acid (solvent B). A linear gradient of 1–95% B over 7 min was applied at 1.5 mL min⁻¹. UV detection was monitored at 220 nm. Eluates were analyzed on Waters SQ Detector 2 with electrospray ionization in positive mode. Full scan spectra were collected over m/z 100–2,000. Selected Ion Recording (SIR) was used for targeted peptides. Source conditions were capillary voltage 3.0 kV, cone voltage 25–40 V, source temperature 120 $^\circ\text{C}$, and desolvation temperature 350 $^\circ\text{C}$. Mass spectra were processed with MassLynx software. Observed peptide masses were compared with theoretical values, and quantitative analysis was based on integrated SIR peak areas.

1399 J.2 BACTERIAL STRAINS AND GROWTH CONDITIONS
 1400

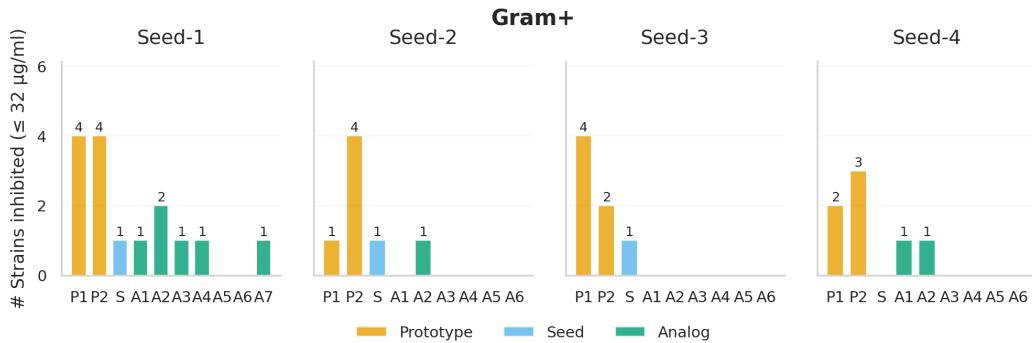
1401 The bacterial panel utilized in this study consisted of the following pathogenic strains: *Acinetobacter baumannii* ATCC 19606; *A. baumannii* ATCC BAA-1605 (resistant to ceftazidime, gentamicin, ticarcillin, piperacillin, aztreonam, cefepime, ciprofloxacin, imipenem, and meropenem);

1404 *Escherichia coli* ATCC 11775; *E. coli* AIC221 [MG1655 phnE_2::FRT, polymyxin-sensitive control];
 1405 *E. coli* AIC222 [MG1655 pmrA53 phnE_2::FRT, polymyxin-resistant]; *E. coli* ATCC BAA-
 1406 3170 (resistant to colistin and polymyxin B); *Enterobacter cloacae* ATCC 13047; *Klebsiella pneumo-*
 1407 *niae* ATCC 13883; *K. pneumoniae* ATCC BAA-2342 (resistant to ertapenem and imipenem);
 1408 *Pseudomonas aeruginosa* PAO1; *P. aeruginosa* PA14; *P. aeruginosa* ATCC BAA-3197 (resistant to
 1409 fluoroquinolones, β -lactams, and carbapenems); *Salmonella enterica* ATCC 9150; *S. enterica*
 1410 subsp. *enterica* Typhimurium ATCC 700720; *Bacillus subtilis* ATCC 23857; *Staphylococcus au-*
 1411 *reus* ATCC 12600; *S. aureus* ATCC BAA-1556 (methicillin-resistant); *Enterococcus faecalis* ATCC
 1412 700802 (vancomycin-resistant); and *Enterococcus faecium* ATCC 700221 (vancomycin-resistant).
 1413 *P. aeruginosa* strains were propagated on *Pseudomonas* Isolation Agar, whereas all other species
 1414 were maintained on Luria-Bertani (LB) agar and broth. For each assay, cultures were initiated from
 1415 single colonies, incubated overnight at 37 °C, and subsequently diluted 1:100 into fresh medium to
 1416 obtain cells in mid-logarithmic phase.

J.3 MINIMAL INHIBITORY CONCENTRATION (MIC) DETERMINATION

1419 MIC values were established using the standard broth microdilution method in untreated 96-well
 1420 plates. Test peptides were dissolved in sterile water and prepared as twofold serial dilutions ranging
 1421 from 1 to 64 $\mu\text{mol L}^{-1}$. Each dilution was combined at a 1:1 ratio with LB broth containing
 1422 4×10^6 CFU mL $^{-1}$ of the target bacterial strain. Plates were incubated at 37 °C for 24 h, and
 1423 the MIC was defined as the lowest peptide concentration that completely inhibited visible bacterial
 1424 growth. All experiments were conducted independently in triplicate.

J.4 DETAILED WET-LAB VALIDATION RESULTS



1440 **Figure 8: Antimicrobial activity against Gram-positive bacterial strains by seed family.** Bar
 1441 chart shows the number of Gram-positive strains (out of 5 total) against which each peptide achieved
 1442 MIC $\leq 32 \mu\text{g/ml}$, organized by seed family (Seed-1 through Seed-4). Within each family, results
 1443 are shown for prototypes (P1, P2; orange), seeds (S; blue), and analogs (A1-A7; green). Numbers
 1444 above bars indicate the count of active strains for each peptide.

K REBUTTAL-RELATED ANALYSES

K.1 SENSITIVITY ANALYSIS OF PoGS

K.1.1 SENSITIVITY TO THE CHOICE OF METRIC

1452 To assess sensitivity of PoGS to the choice of the metric as well as to the choice of prototypes,
 1453 we compared our implementation with two alternatives: using amino-acid probabilities instead of
 1454 logits, and straight (Euclidean) interpolation, on a new set of 60 prototypes (Table K.1.1). Logits
 1455 yielded better performance in terms of potential, counts of seeds and wells, supporting our choice.
 1456 Importantly, both metrics outperformed straight interpolation, and PoGS achieved even better results
 1457 on these new prototypes than those reported in Table 1. These results show that PoGS improves
 1458 results regardless of the metric used or the prototype set.

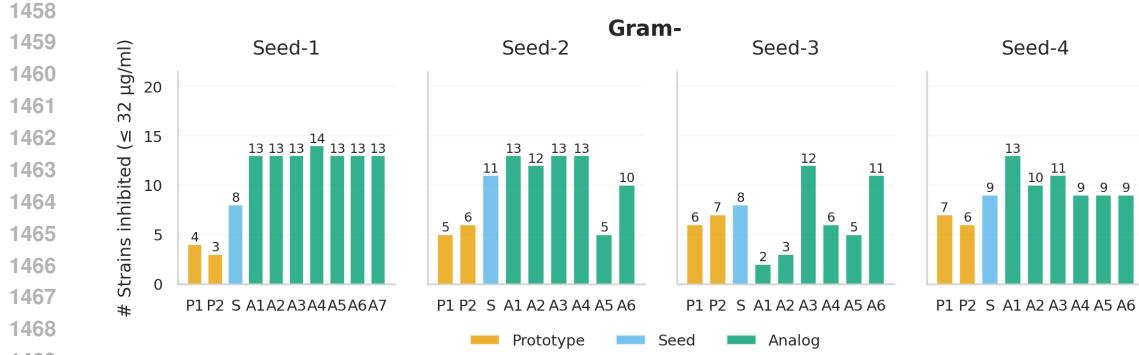


Figure 9: **Antimicrobial activity against Gram-negative bacterial strains by seed family.** Bar chart shows the number of Gram-negative strains (out of 14 total) against which each peptide achieved $\text{MIC} \leq 32 \mu\text{g/ml}$, organized by seed family (Seed-1 through Seed-4). Within each family, results are shown for prototypes (P1, P2; orange), seeds (S; blue), and analogs (A1-A7; green). Numbers above bars indicate the count of active strains for each peptide.

Method	Latent Length	Ambient Length (orig.)	Ambient Length (probs)	Peptide Path Length	Potential	Seeds	Wells
Straight interpolation	6.50 ± 1.34	3448.7 ± 403.8	4.76 ± 4.94	47.0 ± 20.9	-1732.8 ± 272.7	117 ± 42	41 ± 14
PoGS (with probs metric)	11.48 ± 3.36	17227.93 ± 147.8	1.67 ± 1.16	42.7 ± 18.1	-1729.9 ± 271.4	176 ± 43	67 ± 14
PoGS (original)	14.12 ± 3.07	5001.4 ± 236.7	9.97 ± 4.96	63.3 ± 29.3	-2329.0 ± 420.0	188 ± 29	80 ± 20

Table 3: Table PoGS with the original metric compared to PoGS with a metric on decoded probabilities, as well as to straight interpolation, for a different set of prototypes than in Table 1.

K.1.2 SENSITIVITY TO THE CHOICE OF POTENTIAL

To evaluate PoGS’s sensitivity to the choice of potential and its applicability to multi-objective settings, we applied it to jointly maximize two physicochemical properties by assigning weight α to the hydrophobicity and $1 - \alpha$ to charge (Table K.1.2). Across all values of α , PoGS outperformed straight-line (Euclidean) interpolation. For this experiment, we selected a new set of 60 prototypes from the GRAMPA and DRAMP datasets, demonstrating that PoGS achieves superior performance regardless of both the potential used and the prototype set.

Method	Latent Length	Ambient Length	Peptide Path Length	Max Hydrophobicity	Max Charge	Maximized Multiobjective Potential
PoGS w. $\alpha = 0.9$	12.07 ± 2.57	4573.57 ± 19	66.86 ± 24.14	8.40 ± 0.72	15.81 ± 2.31	24.21 ± 2.42
PoGS w. $\alpha = 0.5$	10.04 ± 2.05	3624.69 ± 185.09	57.66 ± 27.08	7.64 ± 0.63	15.75 ± 2.15	23.39 ± 2.24
PoGS w. $\alpha = 0.1$	10.61 ± 2.13	3651.19 ± 190.39	55.93 ± 25.80	7.91 ± 0.64	14.76 ± 2.26	22.67 ± 2.35
Straight line (Euclidean)	6.51 ± 1.34	3448.70 ± 403.80	47.00 ± 20.90	7.90 ± 0.64	13.76 ± 2.22	21.66 ± 2.35

Table 4: Comparison of PoGS multiobjective optimization across different α values and Euclidean straight-line baselines.

K.2 LE-BO TIME AND MEMORY PROFILING

To quantify the computational gains of our method, we measured the average runtime of LE-BO over 10 iterations and 6 seeds, and compared it to SAASBO (Eriksson & Jankowiak (2021)) under the same conditions (Table K.2). LE-BO required 8 \times less time per iteration and used 1.5 \times less memory. Moreover, Local Enumeration accounted for only 35% of the total runtime, underscoring its efficiency. Average execution time of a single optimization run of LE-Bo with 1400 iterations was 1h20m (± 30 m). All computations were performed for the HydRAMP model with dimension 64 and measured on a Mac Mini M4Pro machine with 24GB of RAM.

Method	Local Enumeration Time / Iteration (s)	Total Iteration Time (s)	Memory (MB)
LE-BO	1.03 ± 0.31	2.89 ± 2.21	1490 ± 210
SAASBO (Eriksson & Jankowiak (2021))	N/A	23.31 ± 9.94	2248 ± 12

Table 5: Runtime and memory comparison between LE-BO and SAASBO (Eriksson & Jankowiak (2021)).

1512 **K.3 LE-BO ABLATION STUDY**
 1513

1514 **K.3.1 ANALYSIS OF RESULTS FOR DIFFERENT κ_{SORBES} AND κ_{MUTANG}**
 1515

1516 Additional ablation study show that alternative κ values can even outperform those originally se-
 1517 lected, demonstrating that LE-BO’s performance can be further enhanced through targeted hyper-
 1518 parameter tuning *K.3.1*. An ablation with respect to α is unnecessary because, with only a single
 1519 SORBES-SE step, α has no effect on the search process.

κ_{mutang}	κ_{sorbes}	FL14	KY14	KF16	KK16	mammuthusin-3	hydrodamin-2	Avg. Diff. from LE-BO
10^{-8}	10^{-3}	0.74 ± 0.28	0.83 ± 0.44	0.56 ± 0.29	0.58 ± 0.30	0.44 ± 0.29	0.38 ± 0.13	0.04 ± 0.18
10^{-4}	10^{-3}	0.53 ± 0.26	0.65 ± 0.39	0.66 ± 0.11	0.36 ± 0.09	0.60 ± 0.69	0.45 ± 0.14	-0.01 ± 0.12
10^{-8}	10^{-1}	0.41 ± 0.11	0.65 ± 0.44	0.38 ± 0.14	0.73 ± 0.34	0.58 ± 0.21	0.45 ± 0.11	-0.01 ± 0.19
10^{-4}	10^{-1}	0.83 ± 0.20	0.54 ± 0.31	0.72 ± 0.29	0.66 ± 0.31	0.60 ± 0.45	0.63 ± 0.40	0.12 ± 0.07
10^{-6} (orig.)	10^{-2} (orig.)	0.604 ± 0.22	0.502 ± 0.24	0.600 ± 0.29	0.498 ± 0.14	0.498 ± 0.38	0.581 ± 0.34	0.0 ± 0.0

1524 **Table 6: Table LE-BO performance for different hyperparameter values. Best values of minimized**
 1525 **MIC as predicted by APEX are bolded. The last column reports the average difference between the**
 1526 **alternative hyperparameter setting and the original used in the manuscript.**

1528 **Analysis of results of LE-BO without ROBOT (Maus et al. (2023))**
 1529

1530 Ablation analysis of the model without ROBOT Maus et al. (2023) confirms that ROBOT improves
 1531 LE-BO for 4 out of 6 seeds (with average difference in log MIC of -0.03).

Method	FL14	KY14	KF16	KK16	mammuthusin-3	hydrodamin-2	Avg. diff. from LE-BO
LE-BO w/o ROBOT	0.75 ± 0.41	0.52 ± 0.17	0.56 ± 0.11	0.63 ± 0.28	0.57 ± 0.61	0.42 ± 0.20	0.03 ± 0.12
LE-BO	0.604 ± 0.22	0.502 ± 0.24	0.600 ± 0.29	0.498 ± 0.14	0.498 ± 0.38	0.581 ± 0.34	0.000 ± 0.000

1535 **Table 7: Comparison of LE-BO with and without ROBOT across six peptides and mean deviation**
 1536 **from LE-BO baseline.**

1538 **K.3.2 ANALYSIS OF RESULTS FOR OTHER SEEDS**
 1539

1540 To show consistent gains compared to ablations across different seeds, we additionally performed
 1541 an LE-BO ablation study on three further seeds (LL13, RC16, and KI21) from the PoGS optimiza-
 1542 tion (Table *K.3.2*). The results confirm that clear advantage of enabling mutations and of using
 1543 SORBES-SE as the random walk method can be observed regardless of the chosen seeds. Taken
 1544 together, all three LE-BO components - SORBES-SE, enabling mutation (MUTANG), and ROBOT
 1545 - contribute meaningful improvements. However, enabling mutation with MUTANG provides the
 1546 most substantial benefit: across all ablations, the LE-BO variants incorporating MUTANG consis-
 1547 tently outperform their counterparts without it.

Walk	Mutation	LL13	RC16	KI21	Difference from LE-BO
Euclidean	×	0.787 ± 0.408	1.141 ± 0.298	1.170 ± 0.283	1.033 ± 0.174
SORBES-SE	×	1.009 ± 0.408	1.128 ± 0.298	1.120 ± 0.283	0.709 ± 0.204
-	✓	1.343 ± 0.356	1.225 ± 0.275	0.803 ± 0.238	0.747 ± 0.125
Euclidean	✓	0.630 ± 0.190	0.751 ± 0.291	0.505 ± 0.199	0.253 ± 0.091
LE-BO	SORBES-SE	0.482 ± 0.226	0.458 ± 0.243	0.190 ± 0.191	0.000 ± 0.000

1554 **Table 8: Confirmation of improved performance of LE-BO compared to ablations for different seeds**
 1555 **than in Table 3.3**

1557 **K.3.3 ANALYSIS OF RESULTS FOR DIFFERENT ORACLES**
 1558

1559 To evaluate robustness across different oracles and to verify that additional peptide properties benefit
 1560 from geometry-aware exploration, we successfully applied LE-BO and show that geometry-aware
 1561 exploration improves prototype peptides for three tasks:

- 1562
 1. minimize MIC using DEEP-AMP Pandi et al. (2023) regressor other than APEX as oracle,
 2. minimize toxicity using ToxiPrep Guan et al. (2025) classification probabilities as oracle,
 3. maximize hydrophobicity computed in the Eisenberg scale Eisenberg et al. (1982) as oracle.

	KY14	KF16	KK16	FL14	mammuthusin-3	hydrodamin-2
log ₂ (MIC) (seed)	2.00	5.06	1.85	1.82	4.02	5.09
Euclidean LE-BO	-0.66 ± 0.50	0.04 ± 0.19	-0.56 ± 0.65	-0.55 ± 0.46	-0.48 ± 0.72	-0.15 ± 0.56
LE-BO	-0.80 ± 1.06	-0.47 ± 0.38	-1.20 ± 1.23	-2.35 ± 0.41	-0.77 ± 0.75	-0.93 ± 0.89

Table 9: Minimization of $\log_2(\text{MIC})$ values using LE-BO with DEEP-AMP Pandi et al. (2023) as oracle.

	KY14	KF16	KK16	FL14	mammuthusin-3	hydrodamin-2
Toxicity (seed)	0.8789	0.9283	0.9044	0.9932	0.5193	0.0522
Euclidean LE-BO	0.014 ± 0.002	0.016 ± 0.003	0.016 ± 0.005	0.015 ± 0.001	0.014 ± 0.002	0.016 ± 0.002
LE-BO	0.0115 ± 0.0007	0.0117 ± 0.0012	0.0120 ± 0.0022	0.0122 ± 0.0015	0.0126 ± 0.0013	0.0135 ± 0.0012

Table 10: Table Minimization of toxicity values using LE-BO with toxicity probabilities returned by ToxiPep Pandi et al. (2023) as oracle.

	KY14	KF16	KK16	FL14	mammuthusin-3	hydrodamin-2
Hydrophobicity (seed)	-0.34	-0.60	0.03	0.12	0.17	-0.33
Euclidean LE-BO	1.114 ± 0.091	1.208 ± 0.048	1.162 ± 0.112	1.193 ± 0.029	1.323 ± 0.031	0.963 ± 0.084
LE-BO	1.255 ± 0.013	1.270 ± 0.048	1.242 ± 0.067	1.265 ± 0.026	1.325 ± 0.022	0.953 ± 0.095

Table 11: Table Maximization of hydrophobicity using LE-BO with hydrophobicity values computed in Eisenberg scale Eisenberg et al. (1982) oracle.

K.4 ANALYSIS OF THE κ -STABLE DIMENSION W.R.T. TO OPTIMIZATION AND PHYSICOCHEMICAL PROPERTIES

As evaluated for two models, HydrAMP and PepCVAE, the effective rank is almost always much lower than the latent dimension (see distribution of the effective rank quantified in Fig. 4A, B). Moreover, it strongly positively correlates with peptide length (see Fig. 4C,D in Appendix B).

Our new analysis in Figure 10 in Appendix K shows that LE-BO with SORBES-SE and mutation enabled (Fig 10 A, E) identifies more candidates in Local Enumeration for higher κ -stable dimensions / longer peptides than for lower dimensions / shorter peptides. Euclidean-based search exhibits the opposite trend (Fig. 10, B,C,F,G). As our optimization method relies on dense populations of peptide neighborhoods, this explains weaker performance of Euclidean approaches. Moreover, LE-BO with SORBES-SE and mutation produces substantially more candidates on average (16.730 ± 14.786 ; Fig. 10 A, E) than the Euclidean search with mutation (442 ± 151 ; Fig. 10 B, F). Even without mutation, SORBES-SE yields higher candidate counts (17 ± 5 ; Fig. 10, Panels D, H) than its Euclidean counterpart (13 ± 5 ; Fig. 10, Panels C, G).

Moreover, the κ -stable dimension of the latent space is strongly negatively correlated with charge (Spearman $r = -0.55$) and aromaticity ($r = -0.57$) (Fig. 11). This indicates that peptides with high latent dimension tend to have higher solubility, but lower membrane affinity and antimicrobial activity. As LE-BO identifies substantially more candidates than Euclidean methods in these high-dimensional regions, increasing the likelihood that the unfavorable antimicrobial properties can be corrected.

1607
1608
1609
1610
1611
1612
1613
1614
1615
1616
1617
1618
1619

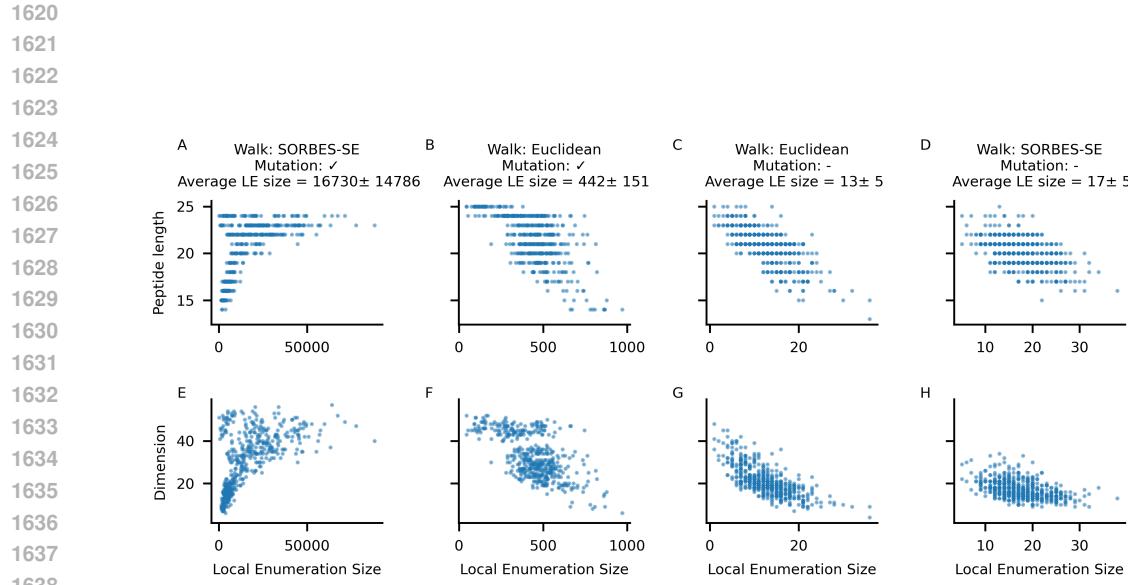


Figure 10: **Relationships between Local Enumeration (LE) result size and sequence properties under different ablations of LE-BO walk and mutation strategies.** Each panel (A–H) shows scatter plots comparing LE size with either peptide length (top row) or κ -stable dimension (bottom row) for four experimental LE-BO conditions: SORBES-SE walk with mutation (default LE-BO implementation) (A, E), Euclidean walk with mutation (B, F), Euclidean walk without mutation (C, G), and SORBES-SE walk without mutation (D, H). Titles report the mean \pm standard deviation of the LE size for each condition. Together, these comparisons highlight how walk geometry and mutation choice influence the size and structure of local neighborhoods explored during Local Enumeration.

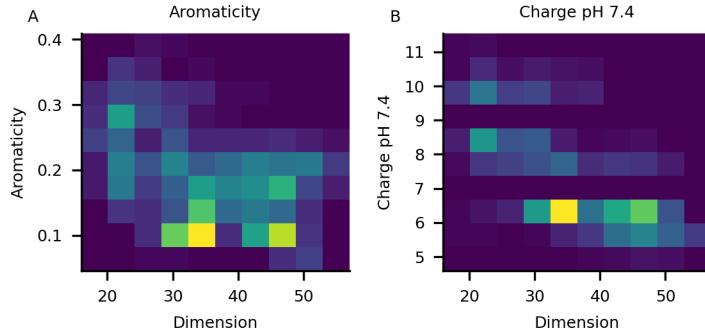


Figure 11: **Relationship between peptide physicochemical properties and latent dimension.** (A) Aromaticity versus embedding dimension, showing that peptides with lower aromatic residue content tend to occupy higher-dimensional regions of the embedding space. (B) Net charge at pH 7.4 versus embedding dimension, highlighting a similar trend in which peptides with lower cationic charge are more common at higher dimensions.

1674

1675

1676 Table 12: Bacterial strains used for experimental validation of antimicrobial peptide libraries.
1677 Strains marked with MDR are multidrug-resistant clinical isolates.
1678

ID	Bacterial Strain
AB1	<i>A. baumannii</i> ATCC 19606
AB2 _{MDR}	<i>A. baumannii</i> ATCC BAA-1605
EC1	<i>E. cloacae</i> ATCC 13047
EC2	<i>E. coli</i> ATCC 11775
EC3	<i>E. coli</i> AIC221
EC4 _{MDR}	<i>E. coli</i> AIC222
EC5 _{MDR}	<i>E. coli</i> ATCC BAA-3170
KP1	<i>K. pneumoniae</i> ATCC 13883
KP2 _{MDR}	<i>K. pneumoniae</i> ATCC BAA-2342
PA1	<i>P. aeruginosa</i> PAO1
PA2	<i>P. aeruginosa</i> PA14
PA3 _{MDR}	<i>P. aeruginosa</i> ATCC BAA-3197
SE1	<i>S. enterica</i> ATCC 9150
SE2	<i>S. enterica</i> Typhimurium ATCC 700720
BS1	<i>B. subtilis</i> ATCC 23857
SA1	<i>S. aureus</i> ATCC 12600
SA2 _{MDR}	<i>S. aureus</i> ATCC BAA-1556
EFS1 _{MDR}	<i>E. faecalis</i> ATCC 700802
EFU1 _{MDR}	<i>E. faecium</i> ATCC 700221

1701

1702

1703

1704

1705

1706 Table 13: Minimum inhibitory concentration (MIC, in $\mu\text{mol L}^{-1}$) values and peptide sequences
1707 for PepCompass-generated peptides tested against bacterial pathogen panel. ‘-’ indicates MIC
1708 $>64 \mu\text{mol L}^{-1}$. Strain IDs correspond to Table 12.

1709

1710

ID	Sequence	AB1	AB2	EC1	EC2	EC3	EC4	EC5	KP1	KP2	PA1	PA2	PA3	SE1	SE2	BS1	SA1	SA2	EFS1	EFU1		
Prototype-1-a	ILRWKKRKLWVKR	64	-	-	16	64	32	-	-	-	8	32	-	-	-	-	32	32	-	-		
Prototype-1-b	FLILRKRLPARVLL	8	-	-	64	8	16	-	-	-	-	-	-	-	-	-	16	8	32	4		
Seed-1	FLYKWWRIGRLKL	1	1	-	32	4	8	16	-	-	-	-	-	-	32	8	16	-	-	64		
LE-BO-1-1	RYAKINLRTAWRKRLKWLWIKVVMKKW	4	4	64	8	4	4	4	8	32	8	4	4	2	2	-	-	32	-	-		
LE-BO-1-2	RYAKINLRTAWRKRLKWLWIKVVMKKW	8	8	64	8	4	16	8	8	8	8	8	8	4	4	64	16	32	-	-		
LE-BO-1-3	RKANLKSRYAWLKLRLKLIKALIWK	2	4	-	8	4	4	4	4	8	8	4	4	2	2	-	32	64	-	-		
LE-BO-1-4	RKANLKSRYAWLKLRLKLIKALIWK	1	16	32	4	2	4	4	4	8	8	4	4	2	2	-	32	-	-	-		
LE-BO-1-5	RKANLKSRYAWLKLRLKLIKALIWK	4	8	-	16	8	8	4	16	8	8	8	4	2	-	-	-	-	-	-		
LE-BO-1-6	RKANLKSRYAWLKLRLKLIKALIWK	1	2	-	16	4	8	4	4	8	8	4	2	2	-	64	-	-	-	-		
LE-BO-1-7	RKANLKSRYAWLKLRLKLIKALIWK	1	2	-	16	4	8	4	4	8	8	4	2	2	-	-	16	-	-	-		
LE-BO-1-8	RKANLKSRYAWLKLRLKLIKALIWK	1	1	-	4	4	4	2	2	2	8	8	16	2	4	-	-	-	16	-		
LE-BO-1-9	RKANLKSRYAWLKLRLKLIKALIWK	4	-	-	8	4	4	2	-	-	16	-	-	-	-	-	-	-	-	4		
Prototype-2-a	KIWARGRKPKWKLQAIQLK	4	-	-	8	4	4	2	-	-	-	-	-	-	-	-	-	-	-	-		
Prototype-2-b	ILRWKKRKWVWLR	8	-	-	2	1	-	-	-	-	16	-	-	-	-	-	-	-	-	16		
Seed-2	KFRNRHBRKFKLIFRN	4	8	-	16	16	32	8	-	-	8	8	16	8	4	-	-	-	-	16		
LE-BO-2-1	NRRKYLRYWKLKLWIKLAKIAINW	8	8	-	16	4	16	16	4	8	8	16	8	4	4	-	-	-	-	-		
LE-BO-2-2	KARIKLYRBYWKLKLWIKLAKIAINW	8	8	-	16	4	16	8	8	4	8	16	64	8	32	32	-	-	-	-		
LE-BO-2-3	KARIKLYRBYWKLKLWIKLAKIAINW	1	1	64	2	1	2	2	1	1	2	1	4	1	1	-	-	-	-	-		
LE-BO-2-4	KARIKLYRBYWKLKLWIKLAKIAINW	2	1	-	2	2	2	4	4	8	4	2	4	4	4	-	-	-	-	-		
LE-BO-2-5	KARIKLYRBYWKLKLWIKLAKIAINW	32	64	-	64	32	64	16	16	64	16	64	-	-	-	-	-	-	-	-		
LE-BO-2-6	KARIKLYRBYWKLKLWIKLAKIAINW	4	16	-	16	8	16	16	16	-	16	16	16	64	-	-	-	-	-	-		
Prototype-3-a	ILRWKFRKVWVWLR	4	-	-	2	8	2	-	-	8	8	-	-	-	-	-	16	32	-	-		
Prototype-3-b	WRHKSLSWIKYKLNLALLA	0.78	-	-	3.12	0.78	1.56	-	25	-	6.25	-	-	-	-	-	-	50	25	-	3.12	
Seed-3	KKYWLWIKWLRWLWFLT	16	32	-	16	32	8	64	64	-	32	64	16	32	16	64	-	-	-	-	-	
LE-BO-3-1	KKARNLRKWAYLKYRLKLKILAINW	32	64	-	-	8	-	64	64	-	64	-	64	-	-	-	-	-	-	-	-	
LE-BO-3-2	KKARNLRKWAYLKYRLKLKILAINW	8	8	-	64	64	-	-	-	-	32	64	-	-	-	-	-	-	-	-	-	
LE-BO-3-3	KKRKLTKLKLKLLKLRL	2	1	64	2	2	4	4	64	8	4	1	4	4	8	-	-	-	-	-	-	
LE-BO-3-4	KKARNLRKWAYLKYRLKLKILAINW	32	32	-	-	4	-	16	64	-	16	32	64	-	-	-	-	-	-	-	-	
LE-BO-3-5	KLRLSKLARWLWIKMYLVWKVKAIIW	-	-	32	2	16	64	16	16	64	-	-	64	-	-	-	-	-	-	-	-	
LE-BO-3-6	KKRRLRKWTRWLWKLLELMMAAWFH	8	-	-	32	4	4	8	-	8	32	8	8	16	-	-	-	-	-	-	-	
Prototype-4-a	LRWKVWRVLAFRRL	4	-	-	4	8	8	-	16	-	16	-	-	-	-	-	-	64	32	1	-	
Prototype-4-b	KYCWRWFKLFLKKL	4	-	-	4	2	2	-	8	-	32	-	-	-	-	-	16	16	-	4	-	
Seed-4	KYCRFRFLWFLRFL	64	32	-	16	8	8	16	16	32	64	64	16	16	64	-	-	-	-	-	-	
LE-BO-4-1	KRARNYYRWKLWKLKLKILAINW	2	2	-	8	4	4	8	2	4	8	8	8	8	-	32	-	-	-	-	-	
LE-BO-4-2	KRIRKLRLIRLRTWKWVKLEMAAFAH	16	4	-	64	8	16	32	32	-	64	16	32	4	32	64	-	-	-	32	-	-
LE-BO-4-3	KRIRKLRLIRLRTWKWVKLEMAAFAH	16	16	-	32	16	16	16	-	64	32	16	2	8	-	-	-	64	-	-	-	
LE-BO-4-4	KRIRKLRLIRLRTWKWVKLEMAAFAH	16	8	-	64	16	4	8	16	-	64	16	4	8	64	-	-	-	-	-	-	
LE-BO-4-5	KRIRKLRLIRLRTWKWVKLEMAAFAH	32	16	-	-	8	16	16	64	-	16	16	4	16	-	-	-	-	-	-	-	
LE-BO-4-6	KRIRKLRLIRLRTWKWVKLEMAAFAH	4	4	-	16	8	4	16	16	-	-	16	4	-	-	-	-	-	-	64	-	-

1726

1727

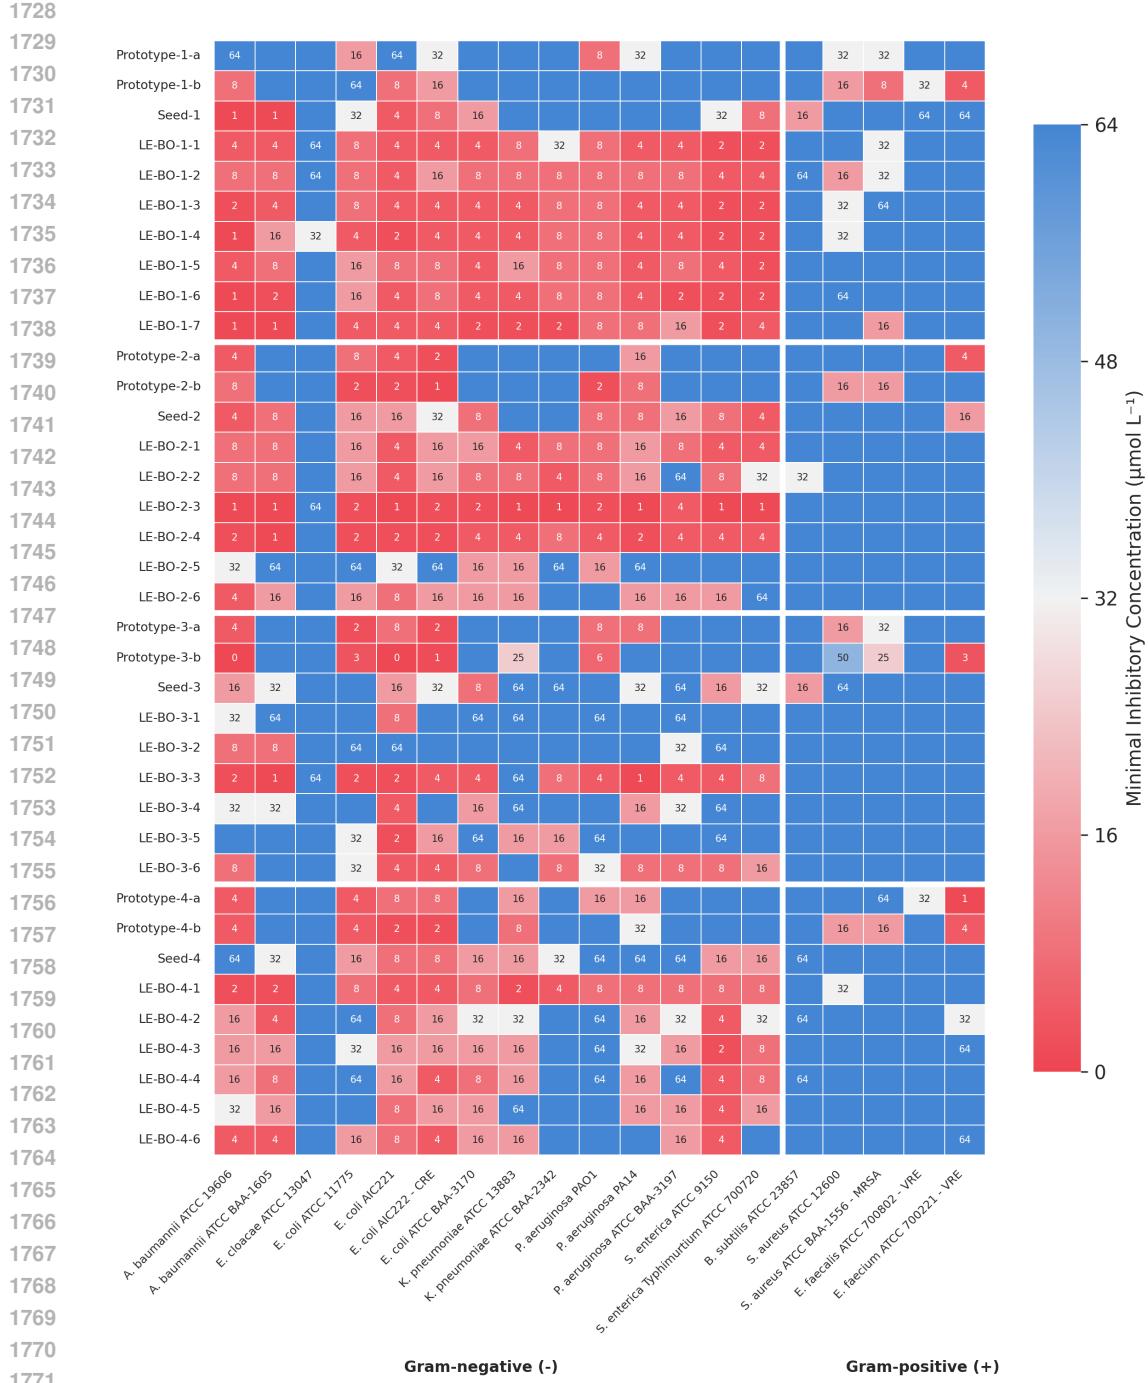


Figure 12: **Minimum inhibitory concentration profiles of antimicrobial peptide libraries against Gram-negative and Gram-positive bacterial pathogens.** MIC values (in $\mu\text{mol L}^{-1}$) for 37 peptide sequences evaluated against 19 bacterial strains. Peptides are stratified by seed family (seed-1 through seed-4), comprising parental prototype sequences and corresponding analogs generated via LE-BO. The bacterial panel encompasses 14 Gram-negative strains including carbapenem-resistant *Enterobacteriaceae* (CRE: *E. coli* AIC222 and ATCC BAA-3170), extended-spectrum β -lactamase-producing *K. pneumoniae* (ATCC BAA-2342), and fluoroquinolone-resistant *P. aeruginosa* (ATCC BAA-3197), alongside 5 Gram-positive strains including methicillin-resistant *S. aureus* (MRSA: ATCC BAA-1556) and vancomycin-resistant *Enterococcus* species (VRE: *E. faecalis* ATCC 700802, *E. faecium* ATCC 700221).

1782

1783

1784

1785

1786

1787

1788

1789

1790

1791

1792

1793

1794

1795

1796

1797

1798

1799

1800

1801

1802

1803

1804

1805

1806

1807

1808

1809

1810

1811

1812

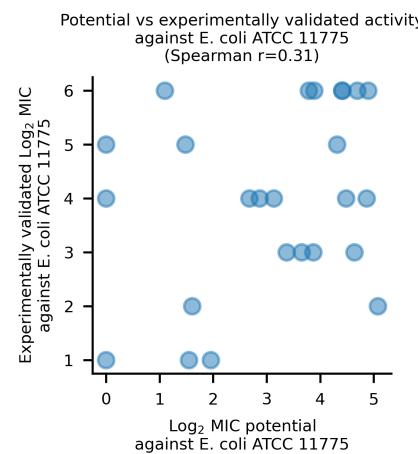
1813

1814

1815

1816

Figure 13: **Correlation between APEX potential and experimental antimicrobial activity for 29 validated peptides.** LE-BO-predicted Log_2 MIC potential shows a meaningful positive association with experimentally measured Log_2 MIC against *E. coli* ATCC 11775 (Spearman $r = 0.31$).



1835