
Token-Level Guided Discrete Diffusion for Membrane Protein Design

Shrey Goel,^{1,2} Peregrine Michael Schray,³ Yinuo Zhang,^{4,5} Sophia Vincoff,⁵
Huong T. Kratochvil,³ Pranam Chatterjee^{2,5,†}

¹Department of Computer Science, Duke University

²Department of Computer and Information Science, University of Pennsylvania

³Department of Chemistry, University of North Carolina at Chapel Hill

⁴Duke-NUS Medical School, Singapore

⁵Department of Bioengineering, University of Pennsylvania

†Corresponding author: pranam@seas.upenn.edu

Abstract

1 Reparameterized diffusion models (RDMs) have recently matched autoregressive
2 methods in protein generation, motivating their use for challenging design tasks
3 such as membrane proteins, which possess interleaved soluble and transmem-
4 brane (TM) regions. We introduce the *Membrane Diffusion Language Model*
5 (**MemDLM**), a fine-tuned RDM-based protein language model that enables con-
6 trollable membrane protein sequence design. MemDLM-generated sequences
7 recapitulate the TM residue density and structural features of natural proteins,
8 achieving comparable biological plausibility and outperforming state-of-the-art
9 diffusion baselines in motif scaffolding tasks by producing lower perplexity, higher
10 BLOSUM-62 scores, and improved pLDDT confidence. To enhance controllability,
11 we develop **PEr**-Token Guidance (PET), a novel classifier-guided sampling strategy
12 that selectively solubilizes residues while preserving conserved TM domains, yield-
13 ing sequences with reduced TM density but intact functional cores. Importantly,
14 MemDLM designs validated in TOXCAT β -lactamase growth assays demonstrate
15 successful TM insertion, distinguishing high-quality generated sequences from
16 poor ones. Together, our framework establishes the first experimentally validated
17 diffusion-based model for rational membrane protein generation, integrating *de*
18 *novo* design, motif scaffolding, and targeted property optimization.

19

1 Introduction

20 Membrane proteins play a crucial role in biological systems, regulating molecular transport, signal
21 transduction, and cellular communication [Jelokhani-Niaraki, 2022]. Their capacity to bind specific
22 ligands or undergo conformational changes renders them essential targets for drug development and
23 therapeutics for various diseases [Sanganna Gari et al., 2021]. Even more interestingly, *de novo*
24 design and engineering of membrane proteins offers a powerful therapeutic modality by enabling the
25 creation of highly-specific and stable proteins that can precisely modulate cell signaling pathways,
26 transport processes, and immune responses, making them ideal for targeting diseases such as cancer
27 and neurological disorders [Jelokhani-Niaraki, 2022]. Current methods for designing new protein
28 sequences or scaffolds rely on pre-trained structure prediction networks [Wang et al., 2022, Yin et al.,
29 2007, Elazar et al., 2022], which remains a particularly challenging prerequisite for membrane protein
30 targets. The scarcity of high-resolution structures hinders the training of high-fidelity deep learning

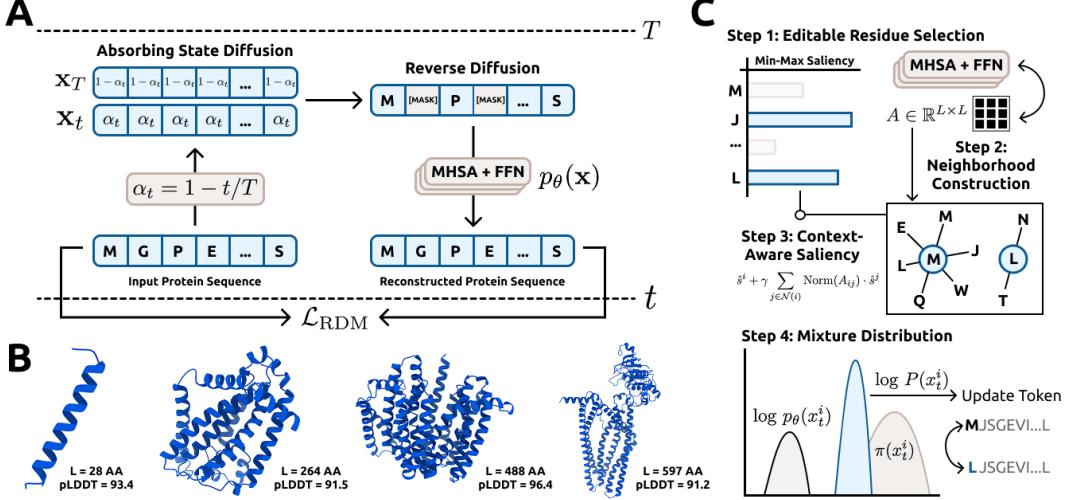


Figure 1: **MemDLM Schematic.** **A)** RDM-based model training diagram. **B)** AlphaFold3 visualizations of unconditional samples. **C)** Token-level classifier guided diffusion sampling with PET algorithm.

31 structure prediction models for membrane proteins: only $\sim 1\%$ of the current PDB structures are
 32 annotated as membrane proteins. Further, energy functions underlying physics-based computational
 33 models are suboptimal because they often require iterative optimizations to design analogs of
 34 membrane proteins [Vorobieva et al., 2021]. As a result, current methods in *de novo* membrane
 35 protein design are limited to simple helical barrel or beta-barrel folds with low sequence complexity.

36 While deep learning-based topology predictors (e.g., DeepLoc, AllesTM) aid in identifying helix
 37 regions and subcellular localization, they primarily analyze existing sequences and do not support
 38 *de novo* generation for function-specific design [Thumuluri et al., 2022] [Hönigschmid et al., 2020].
 39 Prior computational design efforts have achieved impressive results by designing zinc-transporting
 40 helices, yet they are often limited to fixed scaffolds, small proteins, or require extensive intervention
 41 [Joh et al., 2014]. What remains missing is a generative modeling framework that can autonomously
 42 produce membrane protein sequences with controllable structural features, including TM helices,
 43 soluble domains, and higher-order topologies, without relying on predetermined scaffolds or manual
 44 adjustments [Goverde et al., 2024].

45 In this work, we introduce **MemDLM**, a discrete diffusion protein language model for rational
 46 membrane protein design (Figure 1). At the core of our approach is **PER-Token Guidance** (PET), a
 47 novel classifier-guided sampling algorithm that combines attention scores and classifier rewards to
 48 optimize specific sequence tokens during inference. Unlike traditional classifier-guidance methods
 49 ([Gruver et al., 2024], [Li et al., 2024], [Vignac et al., 2022], [Dhariwal and Nichol, 2021], [Tang et al.,
 50 2025], [Chen et al., 2025]), PET ensures the retention of targeted tokens, an essential requirement
 51 in membrane protein design, where highly conserved transmembrane (TM) domains are critical to
 52 maintaining structural topology. We demonstrate that MemDLM generates biologically relevant
 53 proteins with structural features resembling membrane proteins (e.g. α -helices) and show that PET
 54 solubilizes natural membrane proteins while retaining key functional TM domains. Overall, our
 55 integrated pipeline serves as a versatile, end-to-end platform for designing and optimizing membrane
 56 protein sequences, with potential applications spanning therapeutics, drug delivery, and synthetic
 57 biology.

58 **Our key contributions are as follows:**

59 • We introduce **MemDLM**, a discrete diffusion protein language model specifically fine-tuned
 60 for *de novo* generation of membrane protein sequences with controllable structural features.

61 • We develop PET, a novel classifier-guided sampling algorithm to optimize specific sequence
 62 tokens during inference, ensuring the retention of targeted amino acid tokens like conserved
 63 TM domains.

64 • We demonstrate that MemDLM enables controllable sequence generation through token-
 65 level editing. In practice, we show MemDLM effectively solubilizes existing natural
 66 membrane protein sequences while preserving crucial functional TM regions.
 67 • We motivate MemDLM’s utility in real-world therapeutic design by showing it (i) outper-
 68 forms existing state-of-the-art models by achieving improved sequence-specific computa-
 69 tional benchmarks in *de novo* generation and sequence scaffolding tasks, and (ii) produces
 70 experimentally validated membrane protein designs that exhibit favorable growth curves
 71 under antibiotic selection.

72 **2 Methods**

73 **Language Modeling Preliminaries** Let $\mathbf{x} = (x^1, x^2, \dots, x^L) \in \{0, 1\}^{L \times |\mathcal{V}|}$ denote a discrete
 74 sequence of length L , where each token is represented as a one-hot vector over the vocabulary
 75 $\mathcal{V} = \{0, 1, \dots, 32\}$. The vocabulary includes 25 canonical and non-canonical amino acids, along
 76 with several special tokens [Lin et al., 2023]. *Language modeling* aims to estimate the underlying data
 77 distribution $\mathbf{x} \sim q(\mathbf{x})$ using a parameterized probabilistic model $p_\theta(\mathbf{x})$. Since the true distribution
 78 $q(\mathbf{x})$ is typically intractable, we approximate it using a neural network with parameters θ . In Sup-
 79 plementary A.1, we lay out the foundation for RDM-based protein language models by considering
 80 related modeling paradigms.

81 **2.1 MemDLM**

82 **Modeling** MemDLM is built on the Reparameterized Diffusion Model (RDM) framework [Zheng
 83 et al., 2023]. We define $\text{CAT}(\mathbf{x}; \mathbf{p})$ as the categorical distribution on the discrete sequence \mathbf{x}
 84 governed by the vector $\mathbf{p} \in \Delta^{|\mathcal{V}|-1}$, where $\Delta^{|\mathcal{V}|-1}$ denotes the $(|\mathcal{V}| - 1)$ -dimensional proba-
 85 bility simplex. Given a stationary noise distribution $\mathbf{q}_{\text{noise}}$, we define the unconditional prior as
 86 $q(\mathbf{x}_t) = \prod_{i=1}^L \text{CAT}(x_t^i; \mathbf{q}_{\text{noise}})$. We can then write the *forward* diffusion process as a transition kernel
 87 defined in closed-form as a convex combination of clean data and noise:

$$q(\mathbf{x}_t | \mathbf{x}_{t-1}) = \alpha_t \mathbf{x}_0 + (1 - \alpha_t) \mathbf{q}_{\text{noise}} \quad (1)$$

88 where $\alpha_t = \prod_{i=1}^t \beta_i = 1 - t/T$ is a linear noise schedule. This transition distribution in Eq. 1
 89 shows that the forward process is ultimately a convex combination of α_t , the probability of clean
 90 data \mathbf{x}_0 remaining unchanged, and $1 - \alpha_t$, the probability of \mathbf{x}_0 transitioning to the [MASK] token.
 91 By sampling $t \sim \mathcal{U}(0, T = 500)$, we can determine the identity of a token at the given timestep of
 92 the forward process:

$$x_t^i = \begin{cases} [\text{MASK}] & \text{if } u_i < \frac{t}{T}, \quad u_i \sim \text{Uniform}(0, 1) \\ x_0^i & \text{otherwise} \end{cases} \quad (2)$$

93 Importantly, the forward noising process is characterized by an *absorbing state*: $\lim_{t \rightarrow T} \alpha_t =$
 94 $\lim_{t \rightarrow T} (1 - t/T) = 0$, indicating all tokens are guaranteed to be replaced by noise. During inference,
 95 MemDLM $_\theta$ must *denoise* a fully masked sequence $\mathbf{x}_T = \{[\text{MASK}]\}_{i=1}^L$, rendering the absorbing
 96 state a necessary ingredient of the forward noising process. In Section 2.2, we formally outline a
 97 generalized denoising framework from [Peng et al., 2025] to obtain samples from masked diffusion
 98 models (e.g., RDMs).

$$q(\mathbf{x}_t) = \prod_{i=1}^L \text{CAT}(x_t^i; \mathbf{q}_{\text{noise}}) \quad (3)$$

$$q(\mathbf{x}_t | \mathbf{x}_{t-1}) = \alpha_t \mathbf{x}_0 + (1 - \alpha_t) \mathbf{q}_{\text{noise}}, \quad \alpha_t = \prod_{i=1}^t \beta_i = 1 - \frac{t}{T} \quad (4)$$

$$x_t^i = \begin{cases} [\text{MASK}], & \text{if } u_i < \frac{t}{T}, \\ x_0^i, & \text{otherwise} \end{cases} \quad u_i \sim \text{Uniform}(0, 1) \quad (5)$$

$$\lim_{t \rightarrow T} \alpha_t = \lim_{t \rightarrow T} \left(1 - \frac{t}{T} \right) = 0 \quad \Rightarrow \quad \mathbf{x}_T = \{[\text{MASK}]\}_{i=1}^L \quad (6)$$

99 **Loss Function** Following the proof in [Wang et al., 2024] (Appendix A), the RDM framework
100 simplifies the ELBO (Eq. 17) by breaking down the KL-divergence term to yield a simplified training
101 objective:

$$\begin{aligned} \mathcal{L}_{\text{RDM}} &= -\mathbb{E}_{q(\mathbf{x}_0)} \text{KL}[q(\mathbf{x}_{t-1} \mid \mathbf{x}_t, \mathbf{x}_0) \parallel p_{\theta}(\mathbf{x}_{t-1} \mid \mathbf{x}_t)] \\ &= \mathbb{E}_{q(\mathbf{x}_0)} \left[\lambda_t \sum_{i=1}^L b^i(t) \cdot \log p_{\theta}(x_0^i \mid \mathbf{x}_t) \right] \end{aligned} \quad (7)$$

$$\mathcal{L}_{\text{RDM}} = \mathbb{E}_{q(\mathbf{x}_0)} \left[\lambda_t \sum_{i=1}^L \mathbf{1}_{x_t^i \neq x_0^i} \cdot \log p_{\theta}(x_0^i \mid \mathbf{x}_t) \right] \quad (8)$$

102 where $\lambda_t := T - (t - 1)$ represents a linear, time-dependent coefficient and $b^i(t) = \mathbf{1}_{x_t^i \neq x_0^i}$. In
103 practice, \mathcal{L}_{RDM} can easily be computed using the cross-entropy loss between logits and sequence
104 labels. In Supplementary B.2, we detail the specific architectural and training schemes used to
105 construct MemDLM.

106 2.2 Path-Planning Sampling

107 To generate realistic membrane-like protein sequences from a trained MemDLM, we adopt the Path-
108 Planning (P2) paradigm introduced by [Peng et al., 2025], a novel sampling framework for masked
109 discrete diffusion language models. Notably, P2 breaks the assumption of uniform unmasking
110 probabilities and enhances generative quality compared to stochastic sampling from a Gumbel-
111 Softmax distribution or greedy decoding of softmax logits. We follow the *self-planner* variant of P2,
112 where the denoiser itself provides a planning signal used to identify and resample low-value tokens.
113 Here and in Algorithm 2, we outline the key steps of self-planning in P2 but direct the reader to [Peng
114 et al., 2025] for a complete background.

115 **Initial Token Sampling** Beginning with a fully masked sequence $\mathbf{x}_t = \{[\text{MASK}]\}_{i=0}^L$, MemDLM
116 predicts denoised logits $\mathbf{z}_{t-1} \in \mathbb{R}^{L \times |\mathcal{V}|}$ via $\mathbf{z}_{t-1} = p_{\theta}(\mathbf{x}_t)$ at each timestep. Candidate tokens are
117 sampled from the logits using Gumbel-softmax decoding with temperature parameter τ :

$$x_{t-1}^i = \arg \max_v \left(\log \text{softmax} \left(\frac{z_{t-1}^{i,v}}{\tau} + g^{i,v} \right) \right), \quad \mathbf{g}_i \sim \text{Gumbel}(0, 1) \quad (9)$$

118 **Self-Planning** An important requirement of self-planning is resampling low-value tokens using the
119 predictions of the denoising model. Accordingly, we use MemDLM’s log probabilities to compute
120 $s_t^i = \log p_{\theta}(x_t^i)$, a per-position score, and $\mathcal{R}_t = \mathbf{x}_{t-1}^{\setminus \mathcal{M}}$, the set of unmasked positions $\setminus \mathcal{M}$ eligible for
121 remasking. We select the top- K tokens from \mathcal{R}_t with the lowest log-probability scores s_t^i and remask
122 them. Specifically, we dynamically compute $K = \lfloor (1 - \kappa_t) \cdot |\mathcal{R}_t| \rfloor$ as a fixed proportion of unmasked
123 positions controlled by the monotonic scheduling function $\kappa_t = \kappa(i/N)$, where $i \in \{1, 2, \dots, N\}$
124 and $\kappa : [0, 1] \rightarrow [0, 1]$. This update forces the token predictions MemDLM was not confident about
125 ($low s_t^i$) to be remasked.

126 **Token Resampling** We sample new tokens at the remasked positions by copying the most recent
127 denoised tokens from the previous timestep \mathbf{x}_{t-1} into the current sequence \mathbf{x}_t at positions that were
128 masked but are no longer among the K lowest-scoring tokens. This step progressively commits
129 high-confidence tokens while leaving low-confidence regions available for further refinement in future

130 steps, a key advantage over ancestral and greedy sampling schemes. By following the self-planning
 131 scheme of P2, no additional model training or overhead is required, providing a lightweight inference
 132 mechanism for membrane protein design tasks.

133 2.3 Per-Token Classifier Guided Sampling

134 While generating arbitrary membrane proteins is valuable, it is insufficient for downstream applica-
 135 tions, as unconditional samples are unlikely to exhibit the functional properties required for their
 136 use as therapeutic modalities [Jelokhani-Niaraki, 2022]. *Classifier-guided sampling* has recently
 137 introduced controllability to deep generative models by following a gradient signal from a pre-trained
 138 classifier model [Gruver et al., 2024], [Li et al., 2024], [Vignac et al., 2022], [Dhariwal and Nichol,
 139 2021], [Tang et al., 2025], [Chen et al., 2025]. Although these methods bias the model’s sampling
 140 trajectory towards the desired class label, there is no guarantee that specific sequence tokens are
 141 preserved during inference.

142 To this end, we introduce *Per-Token Guidance* (PET), a novel classifier-guided sampling algorithm
 143 that selects and replaces specific sequence tokens with optimized analogues, moving the overall
 144 sequence towards the desired property (Figure 1C). In the case of membrane protein design, PET
 145 can readily be used to replace noncritical TM residues with soluble analogues to guarantee overall
 146 sequence solubility while maintaining biologically conserved TM domains. Solubilizing membrane
 147 proteins without disrupting these critical TM residues is essential for ensuring functional foldability
 148 and membrane localization, as TM residues often mediate key structural and biophysical interactions.
 149 Below, we carefully outline our PET algorithm and refer the reader to Supplementary A.2 for a
 150 background on discrete classifier guidance.

151 **Setup** Given a sequence consisting of only amino acid tokens, $\mathbf{x} = \{x_i \in \text{Canonical}\}_{i=1}^L$, PET first
 152 identifies a dynamic subset of editable positions $\mathcal{E} \subseteq \{1, \dots, L\}$ using existing residue annotations
 153 or a trained per-token solubility classifier $v_\phi : \mathbb{R}^{B \times L \times D} \rightarrow \mathbb{R}^{B \times L}$. This classifier operates over
 154 the hidden states h derived from the ESM-2-650M protein language model [Lin et al., 2023] and is
 155 trained on fully unmasked sequences. See Section B.3 for full training details regarding v_ϕ .

156 **Determining Editable Residues** PET first constructs a set of conserved, *non-editable* token indices
 157 \mathcal{C} based on solubility annotations or predictions:

- 158 1. If soluble residue annotations $\mathcal{S} \subseteq \{1, 2, \dots, L\}$ are provided (e.g. experimentally-derived
 159 labels for known membrane protein sequences), initialize $\mathcal{C} = \mathcal{S}$.
- 160 2. If no annotations are provided, initialize $\mathcal{C} = \{i \in \{1, \dots, L\} \mid v_\phi(h_t)_i \geq 0.5\}$. Inherently,
 161 it is assumed that some $v_\phi(h_t)_i < 0.5$.

162 Next, we consider low-value tokens, *i.e.*, insoluble amino acids with TM-like character. It is critical
 163 to maintain the most conserved TM regions during optimization to maintain the biological plausibility
 164 of the membrane protein. Thus, we guide the selection of *unimportant* TM residues under LaMBO-
 165 2’s (Supplementary A.2) definition of a token’s *saliency* $s^i(h)$, a score that quantifies a token’s
 166 importance relative to the classifier v_ϕ [Gruver et al., 2024]. Given a sequence’s latent representation,
 167 we construct a *saliency map* $\mathbf{s} = (s^1, s^2, \dots, s^L) \in \mathbb{R}^L$:

$$\mathbf{s}(h) := \max \left\{ \left(\sum_{d=1}^D |\nabla_h v_\phi(h)_d| \right)^{1/\tau}, \epsilon \right\}, \quad \hat{s}^i := \frac{s^i - \min \mathbf{s}}{\max \mathbf{s} - \min \mathbf{s} + \delta} \quad (10)$$

168 using temperature $\tau = 2.0$ and a ceiling $\epsilon = e^{-4}$ to stabilize gradient noise. Although LaMBO-2
 169 normalizes the saliency map to the probability distribution $P_{\text{edit}}(\mathbf{x}_t) = \mathbf{s} / \sum_i s_i$ ([Gruver et al., 2024],
 170 Eq. 5), PET opts for min-max scaling (Eq. 10) to prevent vanishing probabilities for large L . If v_θ is
 171 well-trained, high values of \mathbf{s} should correlate with low-value (TM-like) residues. To finalize \mathcal{C} , PET
 172 selects the top- K most salient tokens:

$$\mathcal{C} = \mathcal{C} \cup \text{top-}K(\hat{\mathbf{s}}, K = \max \{1, \frac{1}{10} \cdot (L - |\mathcal{C}|)\}), \quad \mathcal{E} = \{1, \dots, L\} \setminus \mathcal{C} \quad (11)$$

173 Together, these token selection strategies define \mathcal{E} , the set of editable token indices. This set excludes
 174 soluble and highly salient residues to preserve membrane protein character (TM-like residues) while
 175 optimizing for sequence solubility.

176 **Neighborhood Construction.** For each editable token $i \in \mathcal{E}$, PET constructs a context-aware
 177 neighborhood $\mathcal{N}(i)$ based on attention scores. Let $A \in \mathbb{R}^{L \times L}$ be the final-layer attention matrix ex-
 178 tracted from p_θ . The neighborhood $\mathcal{N}(i)$ is formed using top- p nucleus sampling over the normalized
 179 attention weights $\text{Norm}(A_{i,:} / \tau)$, excluding special tokens and the self-position i ; we set $\tau = 1/\log L$
 180 to ensure neighborhood selection is neither overly diffuse in long sequences nor overly narrow in
 181 short sequences. Thus, the final neighborhood contains all tokens j such that the cumulative attention
 182 probability $\sum_{j' \in \mathcal{N}(i)} A_{ij'}$ exceeds the threshold $p = 0.9$. The construction of an attention-informed
 183 neighborhood is necessary to propagate long-range residue information to avoid blindly modifying
 184 individual tokens.

185 **Context-Aware Saliency** PET then refines a token’s raw saliency score s_i with contributions from
 186 the token’s attention-weighted neighborhood $\mathcal{N}(i)$. The *context-aware saliency* score \tilde{s}^i is defined as:

$$\tilde{s}^i := \hat{s}^i + \gamma \sum_{j \in \mathcal{N}(i)} \frac{A_{ij}}{\sum_{j' \in \mathcal{N}(i)} A_{ij'}} \cdot \hat{s}^j \quad (12)$$

187 where $\gamma = 0.5$ controls the influence of the neighborhood saliency. Overall, the context-aware
 188 saliency blends both the intrinsic importance of the token x^i with the contributions of tokens it
 189 attends to most strongly, creating a holistic representation of an individual residue’s contribution to
 190 sequence-level solubility.

191 **Mixture Distribution** Let $\log p_\theta(x_t^i)$ be the log-probability distribution across the vocabulary for
 192 a singular token by the language model at timestep t , and let $\pi(x_t^i)$ be a prior token distribution in
 193 log-space. To update a token, PET defines a *mixture distribution* $\log P(x_t^i)$ for each editable position
 194 $i \in \mathcal{E}$:

$$\log P(x_t^i) = (1 - w^i) \cdot \log p_\theta(x_t^i) + w^i \cdot \pi(x_t^i) \quad (13)$$

195 By construction, $P(x_t^i)$ remains a valid probability distribution, as it is a convex combination of two
 196 normalized distributions. The mixture weight w^i can be computed as:

$$w_i = \sigma(\alpha \cdot \tilde{s}^i) \quad (14)$$

197 with $\sigma(\cdot)$ denoting the sigmoid function and $\alpha = 5.0$ controlling the sharpness of the transition. Eq.
 198 13 ensures that an updated token’s distribution is biased towards the prior when \tilde{s}^i is large since
 199 $s_i \rightarrow 1$ when $v_\theta(h_t^i) \rightarrow 0$. Biologically, this corresponds to a residue with high TM-like character
 200 that is thus conserved and should remain fixed. Conversely, when \tilde{s}^i is small, PET favors the model’s
 201 default prediction, allowing more flexibility in low-salency (non-critical) positions.

202 **Prior Distribution** In order to subtract the mixture distribution, we define a *temporal prior*
 203 $\pi(x_t^i) := \log p_\theta(x_{t-1}^i)$ in PET sampling that leverages the denoising model’s log probabilities from
 204 a previous diffusion timestep. This formulation maintains the likelihood of the original sequence
 205 while encouraging updates from the mixture weighting in Eq. 13.

206 **Token Sampling and Preservation.** A new token \hat{x}^i is sampled from $P(x^i)$ for each position
 207 $i \in \mathcal{E}$. By design, PET will not update positions $j \notin \mathcal{E}$, resulting in an optimized sequence that
 208 preserves soluble and conserved TM regions while refining low-salency, TM positions. To produce
 209 optimized amino acid tokens, we sample from a categorical distribution parameterized by the updated
 210 token probabilities at each position, $\hat{x}^i \sim \text{CAT}(\log P(x^i))$.

211 2.4 TOXCAT- β -Lactamase Growth Assay

212 The TOXCAT- β -lactamase assay was used to evaluate membrane insertion and TM association of
 213 MemDLM-generated sequences [Russ and Engelman, 1999, Lis and Blumenthal, 2006]. Candidate

214 designs were cloned between an N-terminal ToxR transcriptional activator and a C-terminal periplasmic β -lactamase in the pMAL_dst β L vector, and transformed into *E. coli* Cloni cells. Single colonies
 215 were used to inoculate LB cultures with spectinomycin, diluted to $OD_{600} = 0.05$, and normalized
 216 to 1.95×10^5 cells per well in 96-well plates. Cultures were grown in LB supplemented with
 217 spectinomycin (50 μ g/mL) and subjected to different selective pressures: carbenicillin (300 μ g/mL)
 218 to report on membrane insertion, or combined carbenicillin (100 μ g/mL) and chloramphenicol
 219 (100–120 μ g/mL) to report on TM-mediated oligomerization. Plates were incubated at 37°C with
 220 continuous shaking in a BioTek Synergy H1 plate reader, and growth was monitored by OD_{600} every
 221 10 minutes for 24 hours. Successful insertion positions β -lactamase in the periplasm to hydrolyze
 222 carbenicillin, while oligomerization activates the *ctx* promoter via ToxR dimerization, conferring
 223 chloramphenicol resistance.
 224

225 3 Results

226 3.1 *De Novo* Generation

227 Given the limited availability of experimentally verified membrane structures, we focused on
 228 sequence-based metrics (Supplementary B.4). Notably, we computed the TM Residue Density
 229 of the generated sequences by predicting TM and soluble residue regions with DeepTMHMM [Hall-
 230 gren et al., 2022]. To realize this comparison, we utilized all 1,098 sequences from the MemDLM
 231 model test set as the basis for our experiments, yielding a realistic evaluation of sequence plausibility
 232 and membrane character.

	PLDDT (\uparrow)	TM RESIDUE DENSITY	PPL (\downarrow)	ENTROPY (\uparrow)
Test Set	76.637	0.294	5.707	3.918
MemDLM	67.410	0.311	6.344	3.743

Table 1: Computational validation of generated and experimentally validated membrane proteins

233 Table 1 compares various metrics of experimentally annotated membrane proteins with *de novo*-
 234 generated sequences. The results show that MemDLM generates sequences with a soluble residue
 235 density closely matching that of experimentally verified membrane proteins, indicating that MemDLM
 236 has successfully learned their underlying distribution (Supplementary A1).

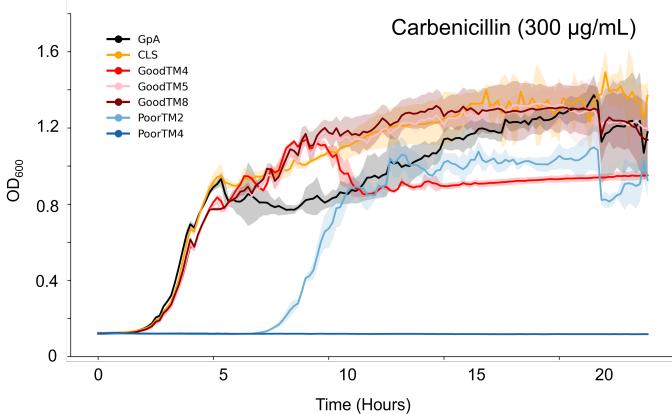


Figure 2: Growth curves of MemDLM-generated TM sequences under carbenicillin (300 μ g/mL).

252 and Mekalanos, 1995, Armstrong and Senes, 2016, Elazar et al., 2016) (Supplementary B.5, C.3).
 253 In these constructs, the design of interest is inserted between an N-terminal ToxR cytoplasmic do-
 254 main and a C-terminal periplasmic β -lactamase. *E. coli* survival under different antibiotic selection
 255 pressures then provides a direct functional readout: survival in carbenicillin indicates successful
 256 membrane insertion, which positions the β -lactamase in the periplasm to degrade the antibiotic, while
 257

258 To fully validate MemDLM's *de novo* generative capabilities, we selected three generated
 259 sequences considered to be *single-pass* membrane proteins ("GoodTM") from the top-100 and two from the bottom-22 ("PoorTM") set of MemDLM-generated sequences for experimental validation in *Escherichia coli* (*E. coli*) using TOXCAT- β -lactamase bacterial growth assays [Lis and Blumenthal, 2006], which employ a dual-reporter system for evaluating membrane insertion and oligomerization of single-pass peptides and proteins [Russ and Engelman, 1999, Lis and Blumenthal, 2006, Ottmann 2016].

261 growth in carbenicillin and chloramphenicol demonstrates TM-mediated oligomerization, where
262 multimerization of the ToxR transcription factors activates the downstream ctx promoter that confers
263 resistance to chloramphenicol.

264 Figure 2 shows the TOXCAT growth curves for poor and high-quality MemDLM sequences alongside
265 the positive insertion controls GpA and CLS (Supplementary A5). Under carbenicillin selection
266 (300 μ g/mL), GpA, CLS, GoodTM4, GoodTM5, and GoodTM8 all achieved similar growth kinet-
267 ics and reached the midpoint of log-phase growth at \sim 4 hours, demonstrating similar membrane
268 insertion efficiencies. PoorTM4 showed no growth in carbenicillin, much like our negative controls
269 (Supplementary A6), indicating that the sequence is not membrane-inserting. However, PoorTM2,
270 which contains six charged residues within the predicted TM span, also grew in carbenicillin but
271 with a noticeable delay, suggesting weaker membrane insertion propensity. The survival of GoodTM
272 designs under carbenicillin selection demonstrates that MemDLM can generate *de novo* TM-inserting
273 sequences and that filtering generated sequences with computational metrics effectively ranks TM-like
274 sequences. The poor survivability of PoorTM2 and PoorTM4, both ranked among the bottom 22
275 sequences by MemDLM, compared to the GoodTM designs further supports MemDLM’s ability to
276 distinguish TM-like sequences.

277 3.2 Motif Scaffolding

278 As a natural extension of *de novo* design, we scaffolded around TM and soluble motifs of experi-
279 mentally annotated membrane proteins. We take the entire test set, comprising 1,098 experimentally
280 verified membrane protein sequences with annotated TM and soluble motifs, and mask out all residues
281 except those in the TM or soluble motif(s). We use these partially masked sequences as input to the
282 models to assay their capability to generate scaffolds conditioned on known TM or soluble motifs.
283 We focused on these domains due to their distinct hydrophilic and hydrophobic regions that govern
284 the folding and thus function of the overall protein.

	PLDDT (\uparrow)		PPL (\downarrow)		BLOSUM-62 (\uparrow)		ENTROPY (\uparrow)	
	INSOL	SOL	INSOL	SOL	INSOL	SOL	INSOL	SOL
Test Set	76.637	76.637	5.707	5.707	–	–	3.918	3.918
EvoDiff	64.058	64.036	9.841	4.632	2.176	-0.188	3.841	3.841
MemDLM	62.762	70.112	8.748	3.242	2.964	0.512	3.876	3.803

Table 2: Reconstruction quality comparison of models scaffolding around TM and soluble motifs of 1,098 experimental membrane protein sequences that represent the MemDLM model test set.

285 Our results (Table 2, Supplementary A2, A3) show that MemDLM-inpainted sequences achieve lower
286 average pseudo-perplexities and higher pLDDT and BLOSUM-62 scores relative to EvoDiff-based
287 ([Alamdari et al., 2023]) scaffolds. These results suggest that MemDLM scaffolds functional motifs
288 with greater confidence while preserving biological relevance compared to SOTA diffusion models.

289 3.3 Solubilizing Targeted Residues

290 Finally, we apply PET to optimize specific residues spanning the insoluble regions of the test set
291 proteins, observing a decrease in TM Residue Density while still preserving critical TM domains
292 (Table 3, Supplementary A4).

	PLDDT (\uparrow)	TM RESIDUE DENSITY (\downarrow)	PPL (\downarrow)	BLOSUM-62 (\uparrow)	ENTROPY (\uparrow)
Test Set	76.637	0.294	5.707	–	3.918
MemDLM	62.979	0.181	8.472	0.495	3.870

Table 3: Computational validation of membrane proteins solubilized under the PET sampling strategy.

293 As a final validation, we visualize MemDLM-generated sequences with AlphaFold3 (Supplementary
294 D) and confirm the presence of hallmark membrane protein structures, including α -helical bundles
295 and distinct TM and soluble regions [Zhang et al., 2015].

296 **4 Discussion**

297 In this work, we introduce MemDLM, the first classifier-guided masked diffusion language model
298 designed specifically for *de novo* membrane protein generation. By leveraging the strengths of
299 masked diffusion over traditional structure-based models, MemDLM effectively captures long-range
300 dependencies critical to the structural and functional integrity of membrane proteins – an area where
301 structure-based models often fall short due to their reliance on pre-defined structural templates and
302 limited generation across diverse topologies. Furthermore, our integration of Per-Token Guidance
303 (PET) for classifier-guided sampling further enables property-guided optimization, enabling us to
304 generate soluble residues over existing TM domains while retaining an initial sequence scaffold.
305 MemDLM also outperforms existing models at demonstrating a robust capability in scaffolding
306 functional motifs, maintaining biological relevance, and achieving high similarity to natural proteins.
307 Moving forward, we aim to generate diverse membrane topologies, including β -barrel and higher-
308 order states and continue to experimentally characterize MemDLM-generated membrane proteins.
309 By evaluating the structural and functional properties of scaffolded TM domains and testing the
310 solubility and stability of membrane proteins generated through classifier-guided optimization, we
311 will validate MemDLM’s potential for advancing rational membrane protein design and expanding
312 its applications in therapeutic development.

313 **References**

314 Masoud Jelokhani-Niaraki. Membrane proteins: structure, function and motion. *International journal*
315 *of molecular sciences*, 24(1):468, 2022.

316 Raghavendar Reddy Sanganna Gari, Joel José Montalvo-Acosta, George R Heath, Yining Jiang,
317 Xiaolong Gao, Crina M Nimigean, Christophe Chipot, and Simon Scheuring. Correlation of
318 membrane protein conformational and functional dynamics. *Nature Communications*, 12(1):4363,
319 2021.

320 Jue Wang, Sidney Lisanza, David Juergens, Doug Tischer, Joseph L Watson, Karla M Castro, Robert
321 Ragotte, Amijai Saragovi, Lukas F Milles, Minkyung Baek, et al. Scaffolding protein functional
322 sites using deep learning. *Science*, 377(6604):387–394, 2022.

323 Hang Yin, Joanna S. Slusky, Bryan W. Berger, Robin S. Walters, Gaston Vilaire, Rustem I. Litvinov,
324 James D. Lear, Gregory A. Caputo, Joel S. Bennett, and William F. DeGrado. Computational
325 design of peptides that target transmembrane helices. *Science*, 315(5820):1817–1822, March
326 2007. ISSN 1095-9203. doi: 10.1126/science.1136782. URL <http://dx.doi.org/10.1126/science.1136782>.

328 Assaf Elazar, Nicholas J Chandler, Ashleigh S Davey, Jonathan Y Weinstein, Julie V Nguyen, Raphael
329 Trenker, Ryan S Cross, Misty R Jenkins, Melissa J Call, Matthew E Call, and Sarel J Fleishman.
330 De novo-designed transmembrane domains tune engineered receptor functions. *eLife*, 11, May
331 2022. ISSN 2050-084X. doi: 10.7554/elife.75660. URL <http://dx.doi.org/10.7554/eLife.75660>.

333 Anastassia A Vorobieva, Paul White, Binyong Liang, Jim E Horne, Asim K Bera, Cameron M
334 Chow, Stacey Gerben, Sinduja Marx, Alex Kang, Alyssa Q Stiving, et al. De novo design of
335 transmembrane β barrels. *Science*, 371(6531):eabc8182, 2021.

336 Vineet Thumuluri, José Juan Almagro Armenteros, Alexander Rosenberg Johansen, Henrik Nielsen,
337 and Ole Winther. Deeploc 2.0: multi-label subcellular localization prediction using protein
338 language models. *Nucleic acids research*, 50(W1):W228–W234, 2022.

339 Peter Hönigschmid, Stephan Breimann, Martina Weigl, and Dmitrij Frishman. Allestm: predicting
340 multiple structural features of transmembrane proteins. *BMC bioinformatics*, 21(1):242, 2020.

341 Nathan H Joh, Tuo Wang, Manasi P Bhate, Rudresh Acharya, Yibing Wu, Michael Grabe, Mei Hong,
342 Gevorg Grigoryan, and William F DeGrado. De novo design of a transmembrane zn^{2+} -transporting
343 four-helix bundle. *Science*, 346(6216):1520–1524, 2014.

344 Casper A Goverde, Martin Pacesa, Nicolas Goldbach, Lars J Dornfeld, Petra EM Balbi, Sandrine
345 Georgeon, Stéphane Rosset, Srajan Kapoor, Jagrity Choudhury, Justas Dauparas, et al. Computational
346 design of soluble and functional membrane protein analogues. *Nature*, 631(8020):449–458,
347 2024.

348 Nate Gruver, Samuel Stanton, Nathan Frey, Tim GJ Rudner, Isidro Hotzel, Julien Lafrance-Vanasse,
349 Arvind Rajpal, Kyunghyun Cho, and Andrew G Wilson. Protein design with guided discrete
350 diffusion. *Advances in neural information processing systems*, 36, 2024.

351 Xiner Li, Yulai Zhao, Chenyu Wang, Gabriele Scalia, Gokcen Eraslan, Surag Nair, Tommaso
352 Biancalani, Shuiwang Ji, Aviv Regev, Sergey Levine, et al. Derivative-free guidance in continuous
353 and discrete diffusion models with soft value-based decoding. *arXiv preprint arXiv:2408.08252*,
354 2024.

355 Clement Vignac, Igor Krawczuk, Antoine Siraudin, Bohan Wang, Volkan Cevher, and Pascal Frossard.
356 Digress: Discrete denoising diffusion for graph generation. *arXiv preprint arXiv:2209.14734*,
357 2022.

358 Prafulla Dhariwal and Alexander Nichol. Diffusion models beat gans on image synthesis. *Advances*
359 *in neural information processing systems*, 34:8780–8794, 2021.

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

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

364 Zeming Lin, Halil Akin, Roshan Rao, Brian Hie, Zhongkai Zhu, Wenting Lu, Nikita Smetanin,
365 Robert Verkuil, Ori Kabeli, Yaniv Shmueli, et al. Evolutionary-scale prediction of atomic-level
366 protein structure with a language model. *Science*, 379(6637):1123–1130, 2023.

367 Lin Zheng, Jianbo Yuan, Lei Yu, and Lingpeng Kong. A reparameterized discrete diffusion model for
368 text generation. *arXiv preprint arXiv:2302.05737*, 2023.

369 Fred Zhangzhi Peng, Zachary Bezemek, Sawan Patel, Jarrid Rector-Brooks, Sherwood Yao,
370 Avishek Joey Bose, Alexander Tong, and Pranam Chatterjee. Path planning for masked diffusion
371 model sampling. *arXiv preprint arXiv:2502.03540*, 2025.

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

374 William P Russ and Donald M Engelman. Toxcat: a measure of transmembrane helix association in
375 a biological membrane. *Proceedings of the National Academy of Sciences*, 96(3):863–868, 1999.

376 Maciej Lis and Kenneth Blumenthal. A modified, dual reporter toxcat system for monitoring
377 homodimerization of transmembrane segments of proteins. *Biochemical and biophysical research
378 communications*, 339(1):321–324, 2006.

379 Jeppe Hallgren, Konstantinos D Tsirigos, Mads Damgaard Pedersen, José Juan Almagro Armenteros,
380 Paolo Marcatili, Henrik Nielsen, Anders Krogh, and Ole Winther. Deeptmhmm predicts alpha and
381 beta transmembrane proteins using deep neural networks. *BioRxiv*, pages 2022–04, 2022.

382 Karen M Ottemann and John J Mekalanos. Analysis of vibrio cholerae toxr function by construction
383 of novel fusion proteins. *Molecular microbiology*, 15(4):719–731, 1995.

384 Claire R Armstrong and Alessandro Senes. Screening for transmembrane association in divisome
385 proteins using toxgreen, a high-throughput variant of the toxcat assay. *Biochimica et Biophysica
386 Acta (BBA)-Biomembranes*, 1858(11):2573–2583, 2016.

387 Assaf Elazar, Jonathan Weinstein, Ido Biran, Yearit Fridman, Eitan Bibi, and Sarel Jacob Fleishman.
388 Mutational scanning reveals the determinants of protein insertion and association energetics in the
389 plasma membrane. *Elife*, 5:e12125, 2016.

390 Sarah Alamdari, Nitya Thakkar, Rianne Van Den Berg, Neil Tenenholz, Bob Strome, Alan Moses,
391 Alex Xijie Lu, Nicolo Fusi, Ava Pardis Amini, and Kevin K Yang. Protein generation with
392 evolutionary diffusion: sequence is all you need. *BioRxiv*, pages 2023–09, 2023.

393 Shao-Qing Zhang, Daniel W Kulp, Chaim A Schramm, Marco Mravic, Ilan Samish, and William F
394 DeGrado. The membrane-and soluble-protein helix-helix interactome: similar geometry via
395 different interactions. *Structure*, 23(3):527–541, 2015.

396 Jacob Devlin. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv
397 preprint arXiv:1810.04805*, 2018.

398 Sophia Vincoff, Shrey Goel, Ksenia Kholina, Rishab Pulugurta, Pranay Vure, and Pranam Chatterjee.
399 Fuson-plm: a fusion oncoprotein-specific language model via adjusted rate masking. *Nature
400 Communications*, 16(1):1436, 2025.

401 Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. *Advances in
402 neural information processing systems*, 33:6840–6851, 2020.

403 Jascha Sohl-Dickstein, Eric Weiss, Niru Maheswaranathan, and Surya Ganguli. Deep unsupervised
404 learning using nonequilibrium thermodynamics. In *International conference on machine learning*,
405 pages 2256–2265. pmlr, 2015.

406 Subham Sekhar Sahoo, Marianne Arriola, Yair Schiff, Aaron Gokaslan, Edgar Marroquin, Justin T
407 Chiu, Alexander Rush, and Volodymyr Kuleshov. Simple and effective masked diffusion language
408 models. *arXiv preprint arXiv:2406.07524*, 2024.

409 Andrei L. Lomize, Spencer C. Todd, and Irina D. Pogozheva. Spatial arrangement of proteins in planar
410 and curved membranes by <scp>ppm</scp> 3.0. *Protein Science*, 31(1):209–220, November 2021.
411 ISSN 1469-896X. doi: 10.1002/pro.4219. URL <http://dx.doi.org/10.1002/pro.4219>.

412 Mikhail A Lomize, Andrei L Lomize, Irina D Pogozheva, and Henry I Mosberg. Opm: orientations
413 of proteins in membranes database. *Bioinformatics*, 22(5):623–625, 2006.

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

416 Marcin Wolny, Matthew Batchelor, Gail J Bartlett, Emily G Baker, Marta Kurzawa, Peter J Knight,
417 Lorna Dougan, Derek N Woolfson, Emanuele Paci, and Michelle Peckham. Characterization
418 of long and stable de novo single alpha-helix domains provides novel insight into their stability.
419 *Scientific reports*, 7(1):44341, 2017.

420 A Wettig, T Gao, Z Zhong, and D Chen. Should you mask 15% in masked language modeling? arxiv
421 2022. *arXiv preprint arXiv:2202.08005*.

422 Renhao Li, Roman Gorelik, Vikas Nanda, Peter B Law, James D Lear, William F DeGrado, and Joel S
423 Bennett. Dimerization of the transmembrane domain of integrin α ii β subunit in cell membranes.
424 *Journal of Biological Chemistry*, 279(25):26666–26673, 2004.

425 William P Russ and Donald M Engelman. The gxxxg motif: a framework for transmembrane
426 helix-helix association. *Journal of molecular biology*, 296(3):911–919, 2000.

427 **A Extended Background**

428 **A.1 Language Modeling**

429 **Masked Language Models** Masked Language Models (MLMs) employ Transformer-based archi-
 430 tectures to learn bi-directional sequence context, distant token relationships, and predict the identity
 431 of corrupted (masked) amino acid tokens. The model is trained under a sequence-recovery training
 432 objective:

$$\mathcal{L}_{\text{MLM}} = - \sum_{i \in \mathcal{M}} \log p_{\theta}(x^i | x^{\setminus \mathcal{M}}) \quad (15)$$

433 where the set of masked positions \mathcal{M} is a fraction of the sequence tokens. MLMs are strong
 434 representation-learners and excel at understanding both protein and natural languages. However,
 435 training these models to reconstruct only a minor fraction of tokens (15-40%) across a sequence
 436 makes complete *de novo* sequence generation difficult. [Devlin, 2018] [Lin et al., 2023] [Vincoff
 437 et al., 2025].

438 **Autoregression** AR language models apply the chain rule to obtain a sequential factorization.
 439 These models are trained to maximize the log-likelihood of the data:

$$\mathbb{E}_{q(\mathbf{x})} \log p_{\theta}(\mathbf{x}) = \mathbb{E}_{q(\mathbf{x})} \sum_{i=1}^L \log p_{\theta}(\mathbf{x}^i | \mathbf{x}^{1:L}) \quad (16)$$

440 New samples can be drawn ancestrally in L steps ($x^1 \sim p_{\theta}(x^1), \dots, x^L \sim p_{\theta}(x^L | x^{1:L-1})$) following
 441 a strictly left-to-right unidirectional protocol. These models are a viable choice for natural language
 442 modeling schemes where a linear relationship between past and present values is inherently assumed.
 443 However, in biological contexts, such as protein sequences, AR models are limited by their inability
 444 to capture non-linear and long-range dependencies. For example, multi-pass membrane proteins
 445 consist of interleaved TM and soluble regions that are spatially and functionally coupled but may be
 446 separated by long sequence distances.

447 **Denoising Diffusion Models** Diffusion models are a class of generative models defined by
 448 Markov processes [Ho et al., 2020] [Sohl-Dickstein et al., 2015]. The *forward* diffusion steps
 449 $q(\mathbf{x}_{1:T} | \mathbf{x}_0) = \prod_{t=1}^T q(\mathbf{x}_t | \mathbf{x}_{t-1})$ progressively corrupt an initial data sample $\mathbf{x}_0 \sim q(\mathbf{x}_0)$ into a noisy
 450 prior $\mathbf{x}_T \sim q_{\text{noise}}$ across T timesteps. The noise distribution q_{noise} typically corresponds to an
 451 isotropic Gaussian, $\mathcal{N}(0, I)$, in continuous latent spaces, or a uniform categorical distribution over
 452 the vocabulary, $\text{Cat}(|\mathcal{V}|)$, in the discrete case. During inference, the learned *backward* process
 453 $p_{\theta}(\mathbf{x}_{0:T}) = p(\mathbf{x}_t) \prod_{t=1}^T p_{\theta}(\mathbf{x}_{t-1} | \mathbf{x}_t)$ gradually denoises the corrupted data sample to obtain samples
 454 from the true data distribution. Diffusion models are trained to maximize the evidence lower bound
 455 (ELBO):

$$\begin{aligned} \mathbb{E}_{q(\mathbf{x}_0)} [\log p_{\theta}(\mathbf{x}_0)] &\geq \mathbb{E}_{q(\mathbf{x}_{0:T})} \left[\log \frac{p_{\theta}(\mathbf{x}_{0:T})}{q(\mathbf{x}_{1:T} | \mathbf{x}_0)} \right] \\ &= \mathbb{E}_{q(\mathbf{x}_0)} \left[\log p_{\theta}(\mathbf{x}_0 | \mathbf{x}_1) + \text{const.} - \sum_{t=2}^T \underbrace{\text{KL}(q(\mathbf{x}_{t-1} | \mathbf{x}_t, \mathbf{x}_0) \parallel p_{\theta}(\mathbf{x}_{t-1} | \mathbf{x}_t))}_{\mathcal{F}_t} \right] \end{aligned} \quad (17)$$

456 New data samples can be drawn by sampling from $q_{\text{noise}}(\mathbf{x}_T)$ and iteratively applying the learned
 457 denoising process $p_{\theta}(\mathbf{x}_{t-1}) = p_{\theta}(\mathbf{x}_{t-1} | \mathbf{x}_t)$. Various authors ([Sahoo et al., 2024], [Zheng et al.,
 458 2023]) have made simplifying assumptions about the reverse process to derive a computationally
 459 inexpensive loss function that reduces to a weighted negative log-likelihood, akin to a weighted form
 460 of Eq. 15.

461 **A.2 Classifier-Guided Sampling**

462 **Preliminaries** Given a property y , guided diffusion aims to maximize $q(y|\mathbf{x})$ by sampling from the
 463 joint distribution $\mathbf{x} \sim q(\mathbf{x}_0, y)$. Therefore, the reverse transition can be conditioned on the property
 464 value y and prior sequence samples. Using Bayes theorem, the conditional joint distribution can be
 465 decomposed:

$$q(\mathbf{x}_{t-1}|\mathbf{x}_t, y) = \frac{q(y|\mathbf{x}_{t-1}, \mathbf{x}_t)}{q(y|\mathbf{x}_t)} \quad (18)$$

466 In practice, the true distribution of $q(y|\mathbf{x}_t)$ is unknown and can be learned with a neural network
 467 $p_\phi(y|\mathbf{x}_t)$. To yield a tractable marginal reverse transition from Eq. 18, we can substitute the true
 468 distribution $q(\cdot)$ with our learned neural networks:

$$p_{\theta, \phi}(\mathbf{x}_{t-1}|\mathbf{x}_t, y) = \frac{p_\theta(y|\mathbf{x}_{t-1}, \mathbf{x}_t)}{p_\phi(y|\mathbf{x}_t)} \quad (19)$$

469 The normalization term in the denominator $p_\phi(y|\mathbf{x}_t)$ can be safely dropped since the model's parameters
 470 learn the normalized distribution. We can update the parameters θ, ϕ at each iteration in the
 471 direction given by the gradient

$$\nabla_{\mathbf{x}_{t-1}} \log p_{\theta, \phi}(\mathbf{x}_{t-1}|\mathbf{x}_t, y) = \nabla_{\mathbf{x}_{t-1}} \log p_\phi(y|\mathbf{x}_{t-1}) + \nabla_{\mathbf{x}_{t-1}} \log p_\theta(\mathbf{x}_{t-1}|\mathbf{x}_t) \quad (20)$$

472 With this formulation, we can steer the denoising trajectory of the unconditional diffusion model to
 473 maximize the target attribute y using gradients from an external classifier [Dhariwal and Nichol, 2021].
 474 Unlike classifier-free guidance, classifier-guidance prevents expensive retraining of existing denoising
 475 network on high-quality, task-specific labeled data and opens avenues for flexible, plug-and-play
 476 conditioning for various downstream applications.

477 **Discrete Classifier Guidance** While classifier guidance is well-formulated for diffusion models
 478 that operate over continuous data in Euclidean space [Dhariwal and Nichol, 2021], applying it to
 479 discrete spaces requires additional approximation. One common approach treats discrete tokens as
 480 continuous relaxations on the probability simplex and uses a first-order Taylor expansion around
 481 \mathbf{x}_t to approximate $\log p_\phi(y|\mathbf{x}_{t-1})$ by making $\nabla_{\mathbf{x}_t}(\cdot)$ a valid operator. However, this approximation
 482 can be inaccurate when the local linearization poorly captures the classifier's behavior over discrete
 483 transitions, especially in regions with sharp decision boundaries. To remedy this, several methods
 484 ([Li et al., 2024], [Vignac et al., 2022]) have been proposed to circumvent the lack of continuous
 485 representations in discrete gradient guidance; most relevant to our work is LaMBO-2 introduced by
 486 [Gruver et al., 2024].

487 **LaMBO-2** To realize classifier-guidance for discrete sequences, LaMBO-2 first conducts sequence
 488 optimization using a Langevin process over a property-informed latent space. We begin with the
 489 discrete Langevin dynamics used in score-based models:

$$\mathbf{x}'_t = \mathbf{x}_t - \eta \nabla_{\mathbf{x}} \log p_\theta(y | \mathbf{x}_t) + \sqrt{2\eta\tau} \epsilon, \quad \epsilon \sim \mathcal{N}(0, I), \quad (21)$$

490 and generalize this update to the continuous latent space $h'_t \in \mathbb{R}^{1 \times D}$ guided by a differentiable
 491 surrogate of the discrete generative model. The batch size dimension B is set to 1 for simplicity. The
 492 latent update step is defined as:

$$h'_t \leftarrow h'_t - \eta \nabla_{h'_t} [\lambda \text{KL}(p_\theta(\mathbf{x}_t|h'_t) \parallel p_\theta(\mathbf{x}_t|h_t)) - \sigma(v_\theta(h'_t)_d)] + \sqrt{2\eta\tau} \epsilon, \quad \epsilon \sim \mathcal{N}(0, I) \quad (22)$$

493 with step size η , temperature τ , and regularization strength λ , where the sigmoid operator $\sigma(\cdot)$ can
 494 be applied to produce a sequence-level binary class probability from the classifier's unnormalized
 495 logit. The explore-exploit loss $\mathcal{L}_{\text{EE}} := \lambda [\text{KL}(p_\theta(\mathbf{x}_t|h'_t) \parallel p_\theta(\mathbf{x}_t|h_t)) - \sigma(v_\theta(h'_t)_d)]$ guides the latent
 496 representation towards high values of the property with the gradient $\nabla_h \sigma(v_\theta(h))$, while the KL

497 term ensures the transition distribution maximizes the original sequence likelihood. Given a discrete
498 sequence \mathbf{x}_t and its corresponding latent representation h_t , one can take N Langevin steps of Eq. 22
499 to realize optimized sequence latent representations before using the language-modeling head of the
500 denoising network to project continuous embeddings to the discrete logit space ([Gruver et al., 2024],
501 Appendix B.2). However, this construction does not guarantee the retention of specific tokens during
502 inference because even if gradients are suppressed for particular positions, the subsequent projection
503 through the language modeling head back into discrete logits does not ensure that the tokens with
504 minimal gradient updates will be preserved.

505 **B Extended Methods**

506 **B.1 Dataset Curation**

507 **MemDLM** Bioassembly structures from X-ray scattering or electron microscopy with better than
508 3.5 Å resolution, annotated by PDBTM1, mpstruc2, OPM3, or MemProtMD4, were used to curate
509 membrane protein sequences for fine-tuning. *de novo* designed membrane proteins were added
510 manually to the database. The proteins were culled at 100% sequence identity and 30% sequence
511 identity to result in a non-redundant set and a sequence-diverse set, respectively. Integral membrane
512 residues, defined as residues with at least one atom within the bilayer, were parsed from the resulting
513 bioassembly structures using the membrane boundaries predicted by PPM 3.0 [Lomize et al., 2021].
514 From the dataset of integral membrane residues, only structures with at least one TM chain spanning
515 the entire membrane bilayer were included in the dataset. Additionally, chains without integral
516 membrane residues were removed from the structure. All peripheral membrane proteins, defined as
517 proteins with no TM chain, were filtered out. The TM protein sequences at the two sequence identity
518 cut-offs and the Python script that parses the sequences from the PPM predictions are included in
519 the SI. After these steps, 9,329 sequences with corresponding per-residue annotations remained. To
520 augment this set of sequences, we obtained 2,579 unique PDB IDs from the Orientations of Proteins
521 in Membranes (OPM) database with the provided "subunits" file [Lomize et al., 2006]. PDB IDs were
522 converted to corresponding protein sequences and per-residue labels (TM or soluble) were assigned
523 using the subunits file. The final set of 11,908 TM sequences were then split using the MMSeqs2
524 easy clustering module with a minimum sequence identity of 80% and a coverage threshold of 50%.
525 The resulting clusters were split to an 80-10-10 ratio into the training set (9,802 sequences, 82.31%),
526 the validation set (1,008 sequences, 8.47%), and the testing set (1,098 sequences, 9.22%).

527 **PET Sampling Classifier** We leveraged the same train/test/val set of 11,908 membrane sequences
528 from the MemDLM dataset to develop a binary classifier that predicts the solubility of each amino
529 acid within a protein sequence. Each sequence was annotated on a per-residue basis, with TM (class
530 1) and soluble (class 0) labels assigned according to the sequence's uppercase and lowercase residues,
531 respectively.

532 **B.2 Modeling MemDLM**

533 **Model Architecture** EvoFlow is a protein language model consisting of 33 Transformer-
534 encoder layers and a language modeling head that is capable of *de novo* generating protein se-
535 quences. More formally, it can denoise a protein sequence consisting of all [MASK] tokens,
536 making it a natural choice for a discrete diffusion-based protein language model. We use the
537 pre-trained EvoFlow protein language model checkpoint (<https://huggingface.co/fredzzp/EvoFlow-650M-context-3070>) as the basis of our neural network p_θ since EvoFlow was trained
538 under the RDM framework (forward process as defined by Eq. 1 and loss computation defined by Eq.
539 8). The Diffusion Protein Language Model (DPLM) was also trained under the RDM framework by
540 [Wang et al., 2024] and is thereby an alternative choice for p_θ . However, we opt for EvoFlow over
541 DPLM as the architecture for p_θ as DPLM is restricted by its shorter context length of 1,024 tokens,
542 compared to EvoFlow's extended context length of 3,070 tokens.

544 **Training** To achieve membrane protein-specific generation, we fine-tuned EvoFlow by selectively
545 updating a subset of the encoder's attention layers. Specifically, the final $N = 3$ Transformer encoder
546 layers $\{\mathcal{L}_{M-N+1}, \dots, \mathcal{L}_M\}$ are partially unfrozen, where $M = 33$ is the total number of encoder
547 layers. Within each layer, we enable gradient updates to only the key, query, and value projection

548 matrices (W_K , W_Q , and W_V) of the self-attention mechanism and keep all other weights frozen. With
 549 this training recipe, we bias the pre-existing EvoFlow latent space with physicochemical features of
 550 membrane proteins without overfitting on the new sequences. MemDLM was trained to minimize the
 551 objective in Eq. 8 on a 4xA6000 NVIDIA DGX server with 200 GB of shared VRAM for 3K steps
 552 using the AdamW optimizer (betas=($\beta_1 = 0.99$, $\beta_2 = 0.98$), weight decay $\lambda = 0.01$), a learning rate
 553 (LR) of 4×10^{-5} with a cosine schedule (150 warmup steps, LR minimum = 1×10^{-5}).

554 **B.3 Per-Token Solubility Classifier**

555 Let $v_\phi : \mathbb{R}^{B \times L \times D} \rightarrow \mathbb{R}^{B \times L}$ be a neural network trained to predict per-token solubility scores
 556 from continuous latent representations h_t . The model is trained using clean protein sequences \mathbf{x}
 557 with corresponding binary per-residue solubility labels $y \in \{0, 1\}^L$ (0 = insoluble, 1 = soluble).
 558 Each input sequence is first embedded using the pretrained ESM-2-650M protein language model
 559 checkpoint (https://huggingface.co/facebook/esm2_t33_650M_UR50D) [Lin et al., 2023].
 560 The resulting contextualized token embeddings are passed through a lightweight classifier v_ϕ with
 561 the following architecture: (i) trainable 2-layer Transformer encoder Transformer_ϕ ; (ii) LayerNorm
 562 and dropout ($p = 0.5$); and (iii) a trainable 2-layer projection head MLP_ϕ outputs a scalar logit for
 563 each token position. All parameters in ESM-2 are frozen, and only the transformer encoder and MLP
 564 layers are updated during training. The classifier is optimized using a per-token binary cross-entropy
 565 loss with logits:

$$\mathcal{L}_{\text{BCE}}(\phi) = -[y \cdot \log \sigma(z) + (1 - y) \cdot \log(1 - \sigma(z))] \quad (23)$$

566 where $\sigma(z)$ is the sigmoid activation function and $\mathbf{z} = v_\phi(h)$ is a vector of per-token logit predictions.
 567 The loss is computed without reduction to allow for masking padded positions and is averaged over
 568 all valid tokens in the batch. v_ϕ is trained on a 1xA6000 NVIDIA DGX server with 50 GB of shared
 569 VRAM for 50K steps using the AdamW optimizer (betas=($\beta_1 = 0.99$, $\beta_2 = 0.98$), weight decay
 570 $\lambda = 0.01$), a learning rate (LR) of $3e^{-5}$ with a cosine schedule (5000 warmup steps, LR minimum =
 571 $1e^{-5}$). The PET classifier was trained using the same train, test, and validation sequence splits as
 572 MemDLM pre-training.

573 **B.4 Computational Metrics**

574 Sequence generation quality was computationally verified using the following metrics:

575 **Pseudo Perplexity** The model’s generation quality was assessed using the ESM-2 [Lin et al.,
 576 2023] pseudo-perplexity metric. Typically, a lower pseudo-perplexity value indicates higher confi-
 577 dence. Specifically, the pseudo-perplexity is computed as the exponential of the negative pseudo-
 578 loglikelihood of a sequence. This metric yields a deterministic value for each sequence but necessitates
 579 L forward passes for computation, where L represents the input sequence length. It is formally defined
 580 as $\text{PPL}(\mathbf{x}) = \exp(-\frac{1}{L} \sum_{i=1}^L \log p(x^i | x^{\backslash i})$.

581 **pLDDT** The structural confidence of generated sequences was assessed using predicted Local
 582 Distance Difference Test (pLDDT) scores from ESMFold v1 with chunk size of 128 [Lin et al.,
 583 2023], a protein language model-based tool to predict protein structures from amino acid sequences
 584 alone. Higher pLDDT indicates ESMFold is more confident in the produced structure, suggesting the
 585 initial input sequence is biologically plausible.

586 **Shannon Entropy** To measure the diversity and uncertainty of the model’s token predictions, we
 587 compute the average Shannon entropy across the sequence. Let $p(x^i)$ denote the model’s probability
 588 distribution over the vocabulary \mathcal{V} at position i . Higher entropy values indicate greater diversity in
 589 the model’s predictions, while lower values suggest more repetitive distributions. The entropy is
 590 defined as: $\text{Entropy}(\mathbf{x}) = -\frac{1}{L} \sum_{i=1}^L \sum_{v \in \mathcal{V}} p(x^i = v) \cdot \log p(x^i = v)$.

591 **BLOSUM62 Substitution Score** The average BLOSUM62 score is a quantitative approach to
 592 determining whether an amino acid substitution is conservative or nonconservative. This value
 593 becomes an important computational metric for protein sequence infilling tasks (both unconditional
 594 and PET-based solubilization) to determine if the model is introducing non-conserved residue changes.
 595 For each aligned position between a generated sequence $\hat{\mathbf{x}}$ and reference sequence \mathbf{x} , we extract the
 596 substitution score $B(\hat{x}^i, x^i)$ from the BLOSUM62 matrix [Henikoff and Henikoff, 1992]. Higher

597 scores indicate greater biochemical similarity to the native sequence, while lower scores suggest
 598 more divergent or potentially deleterious substitutions. The final score is computed as the mean over
 599 all aligned residues $\text{BLOSUM}(\hat{\mathbf{x}}, \mathbf{x}) = \frac{1}{L} \sum_{i=1}^L B(\hat{x}^i, x^i)$.

600 **TM Residue Density** To estimate the membrane-localizing potential of generated sequences,
 601 we used DeepTMHMM v1.0 tool (<https://services.healthtech.dtu.dk/services/DeepTMHMM-1.0/>) [Hallgren et al., 2022] to produce per-residue topology annotations. Each residue
 602 is classified into one of six categories: signal peptide (S), inside cell/cytosol (I), alpha membrane
 603 (M), beta membrane (B), periplasm (P), or outside cell/lumen (O). For our analysis, we consider
 604 residues labeled as alpha membrane (M) to be “soluble” in the membrane context, and all other
 605 classes, including beta membrane (B), to be “insoluble.” We explicitly exclude B-labeled residues
 606 from the soluble category due to the structural and biophysical differences between beta-barrel and
 607 alpha-helical transmembrane domains, the latter being dominant in our training set. Using these
 608 annotations, we define the *TM Residue Density* of a sequence as the number of residues predicted to lie
 609 within alpha membrane (“M” predictions) regions divided by the sequence length as a normalization
 610 factor.

612 B.5 Wet-Lab Experiments

613 B.5.1 Cloning and Plasmid Construction

614 DNA sequences of our MemDLM-designed and control peptides were cloned. Target sequences
 615 derived from MemDLM were cloned into the pMAL_dst β L vector (Addgene plasmid #73805)
 616 between the genes encoding for ToxR and β -lactamase using blunt-end ligation. The resulting
 617 constructs were initially transformed into *E. coli* XL-10 Gold cells. Transformants were selected
 618 on Luria Broth (LB) agar plates containing spectinomycin and sequences were verified by Sanger
 619 sequencing. Confirmed plasmids were subsequently transformed into *E. coli* Cloni cells for the assay.

620 Cell lines:

REAGENT	CATALOG INFORMATION
E. Cloni 10G DUOs Chemically Competent Cells	Cat. No. 60107-1 (BioSearch Technologies)
XL 10-Gold Ultracompetent Cells	Cat. No. 200315 (Agilent)

Table 4: Competent cell reagents used in this study.

621 Genes inserted into the pMAL_dst β L plasmid vector:

- 622 • **Human CLS:**
 - 623 – Uniprot: UPI000007083D
 - 624 – Amino acid sequence: PLFIPVAVMVTAFSGLAFIIWLA
 - 625 – Gene: CCGCTGTTCATCCGGTTGCAGTTATGGTTACCGCTTTAGTGGATTG-
626 GCGTTTATCATCTGGCTGGCT
- 627 • **GpA-TM Region:**
 - 628 – Uniprot: UPI000012B75E
 - 629 – Amino acid sequence: LIIFGVMAGVIGTIL
 - 630 – Gene: TTAATTATTTCGGAGTGATGGCCGGAGTTATCGGCACAATTAAATC
- 631 • **ErbB2 TM Region:**
 - 632 – Uniprot: P04626-1
 - 633 – Amino acid sequence: SIISAVVGILLVVVLGVVFGIL
 - 634 – Gene: TCCATTATCTCCGCTGTCGTAGGAATCTGTTAGTTGTCGTC-
635 CTTGGGGTTGTGTTGGAATTAA
- 636 • **Qsox2 TM Region:**
 - 637 – Uniprot: Q6ZRP7
 - 638 – Amino acid sequence: SLCVVLYVASSLFMVMYFF

639 – Gene: AGTCTTGCGTCGTACTTACGTCGCATCTCACTGTTATGGTGATG-
 640 TATTCTTT

641 • **EK3 Water Soluble Helix** [Wolny et al., 2017]:

642 – Amino acid sequence: SAEEEKKAAEEKKKAEEEEEKKKAE

643 – Gene: TCCGCAGAGGAAGAAAAAGAAAAAGCTGAAGAAGAAAAGAAAAAG-
 644 GCAGAAGAAGAGAAAAAGGCAGAG

645 • **PoorTM2**

646 – MemDLM amino acid sequence: SSLLFSYQGAKKEERVFLDNF

647 – Gene: AGTTCTTGTTATTCAAGCTATCAGGGAGCCAAGAAAGAAGAA-
 648 GAACGTGTGTTCTGGATAACTTC

649 • **PoorTM4**

650 – MemDLM amino acid sequence: GTHAKDWRVTSWKRYGEIE

651 – Gene: GGAACACATGCTAAAGATTGGCGTGTGACATCTGGAAGCGTTACG-
 652 GCGAGATTGAA

653 • **GoodTM4**

654 – MemDLM amino acid sequence: DLSKWLGIVLLLLAILALLIR

655 – Gene: GATTTAAGCAAATGGCTGGGTATCGTACTGTTACTGTTACTGGC-
 656 TATTTGGCTTATTACTGATTGCT

657 • **GoodTM5**

658 – MemDLM amino acid sequence: SLRWLWSLVIGLLIVAFYLLR

659 – Gene: AGCCTGCGTGGTTGGTCTTAGTGATCGGCTTACTGCT-
 660 TATCGTTGCCTTACCTGCTGCTCGC

661 • **GoodTM8**

662 – MemDLM amino acid sequence: DFLRKAVIVLLVLVIVAGLLVIR

663 – Gene: GATTTCTGCGTAAGGCAGTGATTGTATTACTGCTTGGTTATTGTCG-
 664 GCGGGTCTGCTGGTTATTGCT

665 **B.5.2 TOXCAT- β -Lactamase Growth Assay**

666 Single colonies of plasmid-containing *E. coli* Cloni cells were used to inoculate 6-mL LB cultures
 667 supplemented with 50 μ g/mL spectinomycin. Glycerol stocks were made and used to inoculate new
 668 fresh LB culture tubes with 50 μ g/mL spectinomycin. Cultures were incubated for \sim 8 h or overnight
 669 at 37°C with shaking. Optical density at 600 nm (OD₆₀₀) was measured, and cultures were diluted
 670 with fresh LB + spectinomycin to an OD₆₀₀ of 0.05. Growth was continued until an OD₆₀₀ of \sim 0.1
 671 was reached.

672 To ensure consistent inoculation density across assays, the number of cells per well was normalized
 673 to 1.95×10^5 cells. This value was calculated using the relationship of 1 OD₆₀₀ \approx 8×10^8 cells/mL
 674 and adjusted for the measured absorbance at OD₆₀₀ of each culture. Growth under spectinomycin
 675 confirmed that the pMal_dsTBL plasmid was successfully introduced into *E. coli* Cloni cells across
 676 all conditions. All cultures grew equally under this condition, demonstrating comparable inoculation
 677 densities and consistent plasmid uptake.

678 Assays were performed in 96-well plates, with each well containing a final total volume of \sim 200
 679 μ L LB medium supplemented with the appropriate antibiotics in the following concentrations:
 680 Spectinomycin (50 μ g/mL), Carbenicillin (300 μ g/mL), Carbenicillin (100 μ g/mL) + chloramphenicol
 681 (100 μ g/mL), Carbenicillin (100 μ g/mL) + chloramphenicol (120 μ g/mL). Wells were inoculated
 682 with the calculated volume of diluted culture corresponding to 1.95×10^5 cells. Each antibiotic
 683 reporter was run in triplicate. Plates were incubated at 37°C in a pre-heated plate reader (BioTek
 684 Synergy H1). Bacterial growth was monitored by measuring absorbance at 600 nm for 24 hours with
 685 measurements taken every 10 minutes under continuous shaking.

686 **C Extended Results**

687 **C.1 Density Plots**

688 We visualize the density distribution of the various computational metrics to assess membrane protein
689 sequences. When using P2 Self-Planning to generate sequences, we set $\tau = 0.7$ to have a slight bias
690 towards deterministic model outputs.

691 **Unconditional Generation** We unconditionally generate 1,000 membrane protein sequences.
692 Lengths are randomly chosen from 50-250 residues.

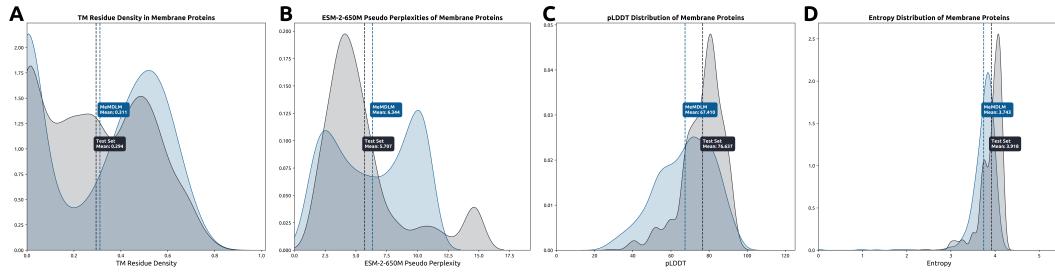


Figure A1: *De novo*-generated and natural membrane protein sequences.

693 **Motif Scaffolding** We mask out and infill both the insoluble and soluble regions of natural mem-
694 brane proteins derived from the model’s test set.

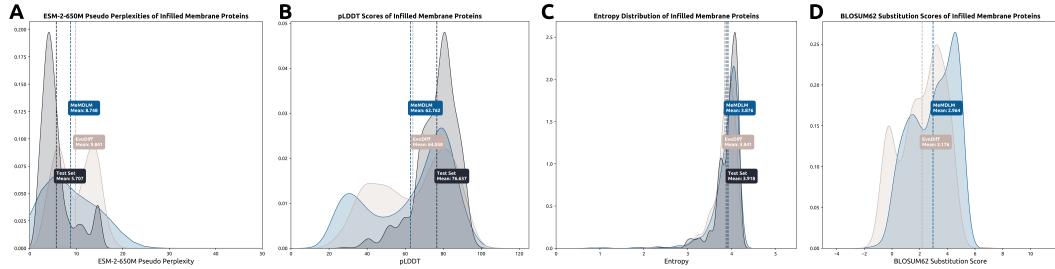


Figure A2: Infilling Insoluble Domain

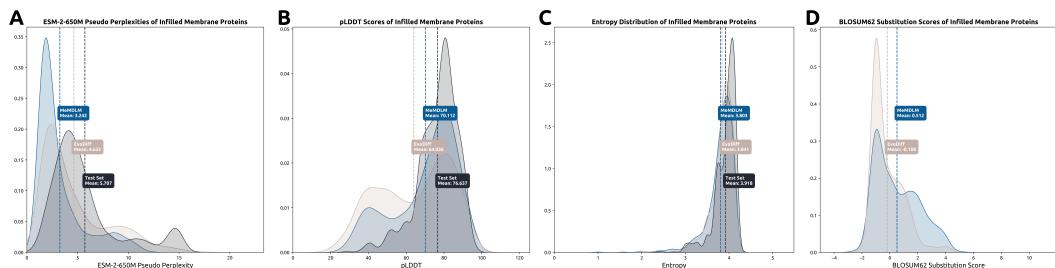


Figure A3: Infilling Soluble Domain

695 **Solubilization** We optimize the solubility of the proteins in the model’s test set by applying our
696 PET algorithm.

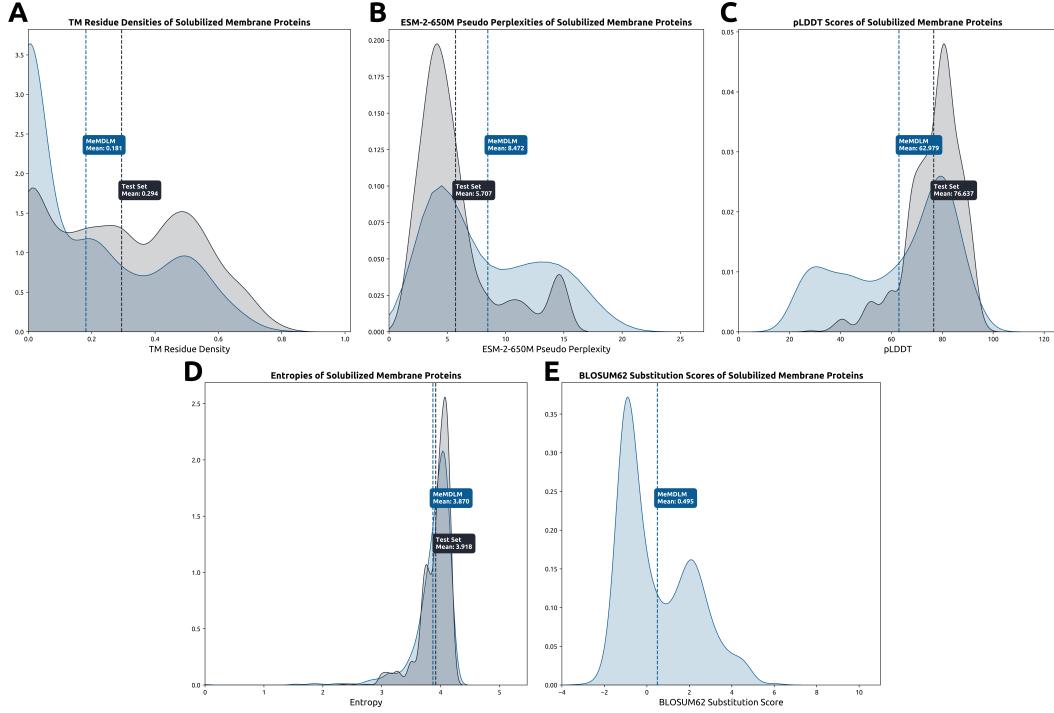


Figure A4: Solubilizing TM Domains

697 C.2 Physicochemical Property Prediction

698 As a surrogate task, we assessed if RDM training retains physicochemical information critical to
 699 membrane protein function by predicting per-residue solubility and membrane localization (Table 3).
 700 We use embeddings from three models—vanilla ESM-2-650M, ESM-2-650M fine-tuned on membrane
 701 protein sequences, and MemDLM—as inputs to a per-residue solubility and sequence-level membrane
 702 localization classifiers. We outline the dataset, training details, and evaluation results of these models
 703 in the following.

704 C.2.1 Datasets

705 **Solubility Prediction** We leveraged the same set of 11,908 membrane sequences from the
 706 MemDLM training dataset to develop a binary classifier that predicts the solubility of each amino
 707 acid within a protein sequence. Each sequence was annotated on a per-residue basis, with TM (class
 708 1) and soluble (class 0) labels assigned according to the sequence’s uppercase and lowercase residues,
 709 respectively. The same training, testing, and validation data splits used to train MemDLM were also
 710 utilized to train and evaluate this classifier.

711 **Membrane Localization** We collected 30,020 protein sequences from DeepLoc 2.0 thumuluri2022deeploc to build a binary classifier that predicts a protein sequence’s cellular localization. The
 712 authors of the dataset provided a multi-label label for each sequence indicating its localization(s). We
 713 used the authors’ provided data splits, with training sequences having 11 labels and testing sequences
 714 having 8 labels.

716 C.2.2 Models

717 **Solubility Prediction** We first predicted TM and soluble residues, a hallmark characteristic of
 718 membrane protein sequences. We utilized embeddings from each pLM’s latent space (ESM-2-150M,
 719 ESM-MLM, and MemDLM) as inputs to train a two-layer perceptron classifier that minimized the
 720 standard binary cross-entropy (BCE) loss to compute the probability that each residue in the sequence
 721 is either soluble (probability < 0.5, class 0) or TM (probability > 0.5, class 1).

722 **Membrane Localization Prediction** Proteins originating from the endomembrane system and
723 localizing in the plasma membrane differ in conformation and function from those in the cytosol and
724 other cellular organelles. We predicted the subcellular localization of protein sequences by utilizing
725 embeddings from each pLM’s latent space (ESM-2-150M, ESM-MLM, and MemDLM) to train a
726 XGBoost classifier that minimized the standard BCE loss to compute the probability that a protein
727 sequence localizes in the plasma membrane (probability > 0.5 , class 1) or in other regions (probability
728 < 0.5 , class 0).

729 **Fine-Tuning ESM-2** We fine-tune the ESM-2 pLM ([Lin et al., 2023]) to achieve an encoder that
730 produces membrane-aware protein sequence embedding used as a baseline comparison for the RDM
731 training task. We trained a MLM head on top of ESM-2-650M using membrane protein sequences to
732 force comprehension of membrane protein properties. We chose to randomly mask 40% of amino
733 acid tokens during training over the standard 15% to more closely resemble the dynamics of diffusion-
734 based (RDM) training; masking rates above 40% have been seen as detrimental during MLM training
735 tasks [Wettig et al.]. Corrupted sequences were passed into ESM-2-650M to retrieve their output
736 embeddings. During training, we unfroze the key, query, and value weights in the attention heads
737 of the final three encoder layers, similar to fine-tuning EvoFlow during MemDLM training. During
738 ESM-2 fine-tuning, the model performed a *masked-prediction* task over masked amino acid tokens
739 to minimize the NLL loss in Eq. (15). 2xH100 NVIDIA GPUs, learning rate of 5e-3, the Adam
740 optimizer, and a batch size of 8 over 10 epochs were used.

741 C.2.3 Results

742 We leveraged the trained solubility prediction and membrane localization classifiers to determine
743 if latent spaces from RDM-based generative models are aligned with relevant membrane protein
744 properties. Table 5 shows that MemDLM latent embeddings achieve predictive performance that
745 closely parallels SOTA pLM embeddings, which are designed specifically for delivering precise
746 representations.

MODEL	SOLUBILITY (\uparrow)	MEMBRANE LOCALIZATION (\uparrow)
ESM-2-650M	0.9383	0.6011
Fine-Tuned ESM-2	0.9375	0.6000
MemDLM	0.9375	0.5964

Table 5: Performance comparison (AUROC) of embeddings derived from various models in predicting physico-chemical properties of MemDLM test set sequences.

747 In total, these results demonstrate that MemDLM accurately captures the biological features under-
748 pinning functional membrane proteins despite being trained on a sequence generation task rather than
749 a masked-prediction task.

750 **C.3 Wet-Lab Experiments**

751 **C.3.1 TOXCAT Assay**

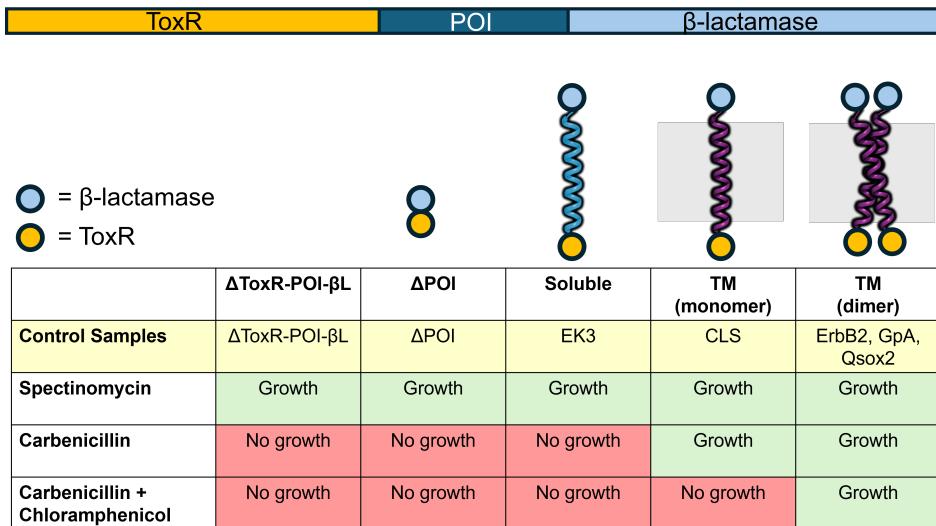


Figure A5: Summary of control constructs for the TOXCAT- β -lactamase assay and their expected growth responses to antibiotics.

752 Schematic showing gene ToxR-POI- β L, where POI is the peptide of interest and β L is β -lactamase.
 753 Periplasmic β -lactamase and cytoplasmic ToxR proteins are represented by blue and yellow dots,
 754 respectively. Expected growth phenotypes under spectinomycin and carbenicillin +/-chloramphenicol
 755 are indicated for each control. Negative controls Δ ToxR-POI- β L, Δ POI, and EK3 should not survive
 756 in carbenicillin because they lack a TM domain. Positive controls CLS, ErbB2, GpA, and Qsox2 all
 757 have TM domains and should survive in carbenicillin. Further, ErbB2, GpA, and Qsox2 are dimers.
 758 Expression of these controls should also confer resistance to chloramphenicol.

759 **C.3.2 TOXCAT Sequence Selection**

760 From 1,000 MemDLM-generated sequences, three sequences from the top 100 predicted performers
 761 ("GoodTM") and two sequences from the bottom 22 predicted performers ("PoorTM") were selected
 762 for screening in the TOXCAT assay. The following selection criteria was used:

CATEGORY	PLDDT	PPL	TM RESIDUE DENSITY	SEQUENCES SELECTED
GoodTM (Top 100)	> 60	< 10	Non-zero	3
PoorTM (Bottom 22)	< 60	< 15	Non-zero	2

Table 6: Selection criteria and sequence counts for MemDLM-generated sequences screened in the TOXCAT assay.

763 The top-ranked (GoodTM) sequences represented a diverse set of high-scoring designs. For example,
 764 GoodTM5 (SLRWLWLSLVIGLLIVAFYLLR, rank 57) contained a small-X₃-small motif known
 765 to promote TM helix association [Russ and Engelman, 1999] [Li et al., 2004] [Russ and Engelman,
 766 2000]. This further demonstrates that MemDLM generates plausible protein sequences with TM-like
 767 character.

768 **C.3.3 Growth Curves**

769 **Control Plasmids** Growth curves of *E. coli* Cloni cells containing control plasmids.

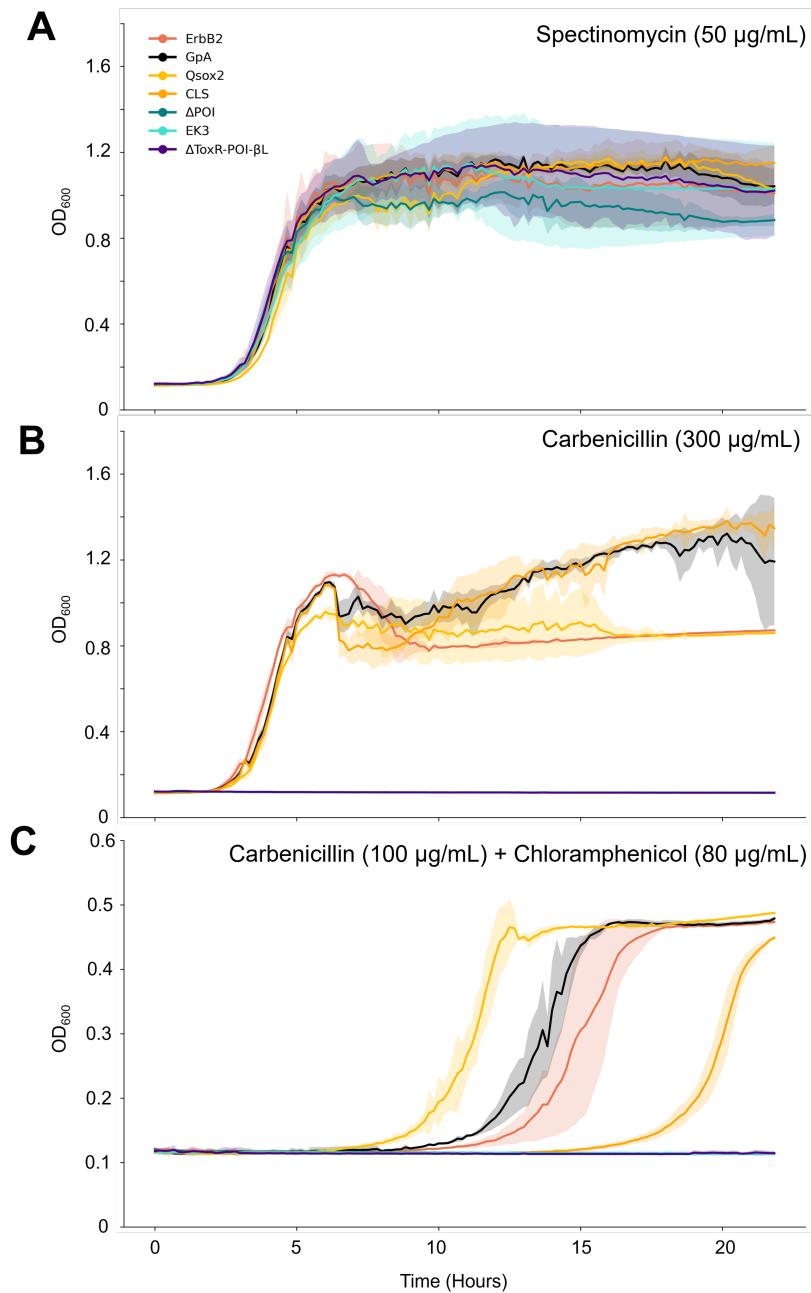


Figure A6: **A)** Survival in spectinomycin (50 µg/mL) confirmed plasmid uptake for all controls. **B)** Growth curves of control plasmids under carbenicillin (300 µg/mL) showed that control plasmids containing TM sequences survived selective pressure. **C)** Growth curves of control plasmids under combined carbenicillin (100 µg/mL) and chloramphenicol (80 µg/mL) selection, which tests both transmembrane insertion and association, show that the dimeric Qsox2, GpA, and ErbB2 controls begin growing in chloramphenicol earlier than the monomeric CLS control.

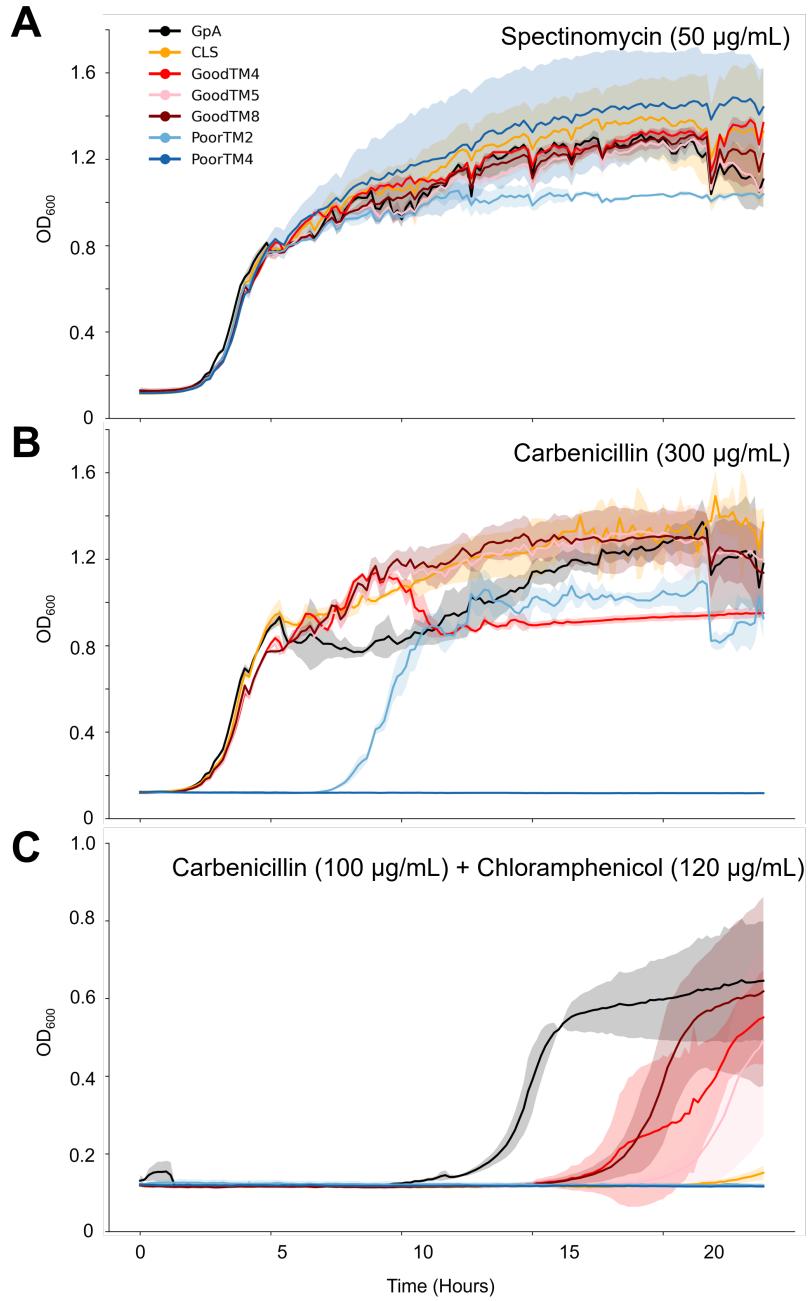


Figure A7: GpA is used as a positive control for insertion and TM association. CLS is the positive insertion and negative TM association control. **A)** Growth curve of *E. coli* Cloni cells containing *de novo* MemDLM TM sequences under spectinomycin (50 µg/mL) confirmed plasmid uptake. **B)** Growth curves of MemDLM peptides under carbenicillin (300 µg/mL) show GoodTM4, GoodTM5, and GoodTM8 growing as expected. PoorTM4 did not survive, indicating that it is not membrane inserting. PoorTM2 did not survive, indicating that it has lower membrane insertion propensity than the GoodTM constructs. **C)** Growth curves of MemDLM plasmids under combined carbenicillin (100 µg/mL) and chloramphenicol (120 µg/mL), used to select for both transmembrane insertion and transmembrane association, reveal that some of the TM designs may be oligomeric.

771 **D Visualizations**

772 AlphaFold3 visualizations of MemDLM-generated membrane protein sequences. TM Residue
 773 Density (TMRD) scores are derived from DeepTMHMM predictions. Structures and colors are from
 774 AlphaFold3 predictions, and pLDDT scores are from ESMFold.

775 **D.1 De novo Generation**

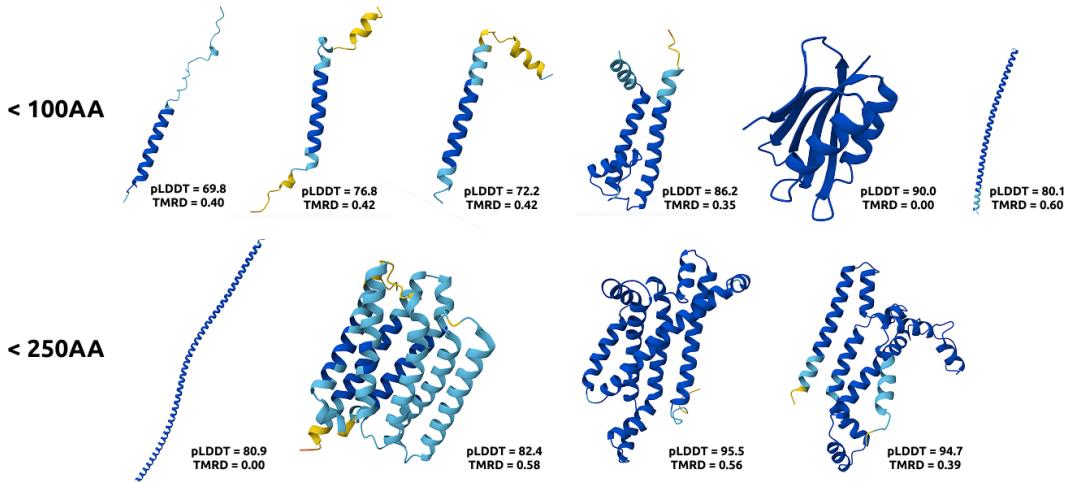


Figure A1: *De novo*-generated protein sequences from MemDLM across different lengths.

776 **D.2 Solubilization**

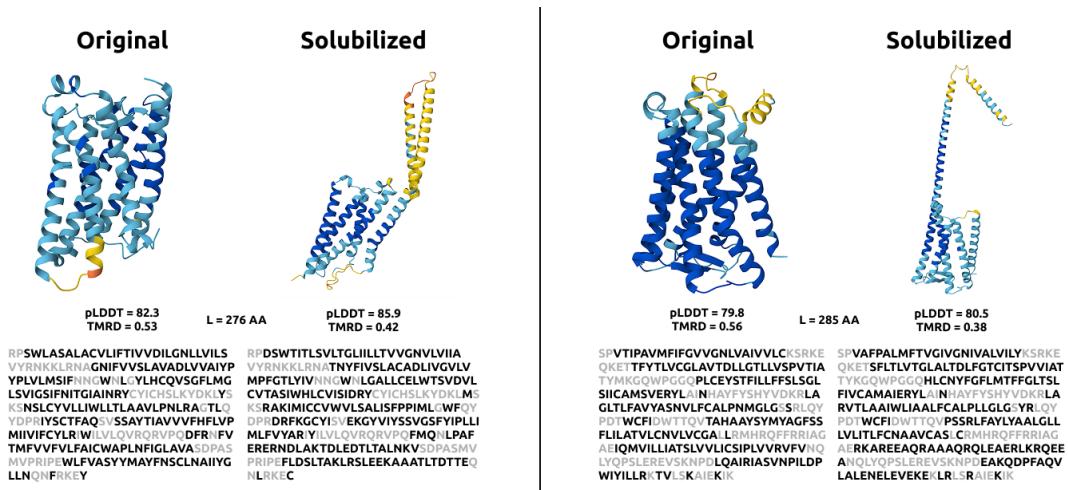


Figure A2: Original and solubilized versions of MemDLM test set protein sequences. Grey residues were annotated as soluble in the given sequence and were thus "fixed" during PET sampling.

777 **E Algorithm Pseudocode**

Algorithm 1 MemDLM Training

Require: Protein sequence dataset \mathcal{D} , diffusion model p_θ , number of diffusion timesteps T

```

1: while not converged do
2:   Sample batch  $\mathbf{x}_0 \sim \mathcal{D}$ 
3:   Sample timestep  $t \sim \mathcal{U}(1, T)$ 
4:   Corrupt sequence:  $\mathbf{x}_t \sim q(\mathbf{x}_t \mid \mathbf{x}_{t-1})$ 
5:   Compute RDM loss:  $\mathcal{L}_{\text{RDM}} = -\lambda_t \sum_{i=1}^L \log p_\theta(x_0^i \mid \mathbf{x}_t)$ 
6:   Take gradient descent step on:  $\nabla_\theta \mathcal{L}_{\text{RDM}}$ 
7: end while
8: return Trained MemDLM  $p_\theta$ 

```

Algorithm 2 MemDLM Sampling with P2 Self-Planning and Optional Sequence Refinement

Require: Fully masked sequence $\mathbf{x}_T = \{[\text{MASK}]\}_{i=1}^L$, trained MemDLM p_θ , number of denoising steps T

```

1: for  $t \in \{T, T-1, \dots, 0\}$  do
2:   Compute logits:  $\mathbf{z}_{t-1} = p_\theta(\mathbf{x}_t)$ 
3:   Sample candidate tokens:  $x_{t-1}^i = \arg \max_v \left( \frac{z_{t-1}^{i,v}}{\tau} + g^{i,v} \right)$ ,  $g^{i,v} \sim \text{Gumbel}(0, 1)$ 
4:   Compute per-token log-probabilities:  $s_t^i = \log p_\theta(x_t^i)$ 
5:   Identify unmasked positions:  $\mathcal{R}_t = \{i \mid x_{t-1}^i \neq [\text{MASK}]\}$ 
6:   Compute  $K = \lfloor (1 - \kappa_t) \cdot |\mathcal{R}_t| \rfloor$ 
7:   Select top- $K$  lowest scoring tokens from  $\mathcal{R}_t$  and remask them:  $x_t^i = [\text{MASK}]$  for  $i \in \text{top-}K(s_t^i)$ 
8:   Copy high-confidence predictions:  $x_{t-1}^i \leftarrow x_t^i$  for positions previously masked but not in top- $K$ 
9: end for
10: if PET Optimization then
11:   Perform Algorithm 3
12: end if
13: return Final decoded sequence  $\mathbf{x}_0$ 

```

Algorithm 3 PET-based MemDLM Sampling

Require: Candidate protein sequence \mathbf{x} , trained MemDLM p_θ , trained solubility classifier v_ϕ , pre-trained encoder Encoder_ϕ , number of optimization steps N

```

1: Produce sequence embeddings  $h = \text{Encoder}_\phi(\mathbf{x})$ 
2: Compute saliency map  $\mathbf{s}$  using gradients  $\nabla_h v_\phi(h)$ 
3: Normalize saliency map  $\hat{s}^i \leftarrow s_i$ 
4: Determine editable positions  $\mathcal{E}$  based on soluble residues and saliency scores
5: for each  $i \in \mathcal{E}$  do
6:   Define neighborhood  $\mathcal{N}(i)$ 
7:   Compute  $\tilde{s}^i = \hat{s}^i + \gamma \sum_{j \in \mathcal{N}(i)} \text{Norm}(A_{ij}) \cdot \hat{s}^j$ 
8:   Construct prior distribution  $\pi(x^i)$ 
9:   Compute guidance distribution:  $\log P(x^i) = (1 - \sigma(\alpha \tilde{s}^i)) \cdot \log p_\theta(x^i) + \sigma(\alpha \tilde{s}^i) \cdot \pi(x^i)$ 
10:  Sample token  $\hat{x}^i \sim \text{CAT}(\log P(x^i))$ 
11:  Update  $\mathbf{x}[i] \leftarrow \hat{x}^i$ 
12: end for
13: return Optimized sequence  $\hat{\mathbf{x}}$ 

```
