# PROTEIN GENERATION WITH EMBEDDING LEARNING FOR MOTIF DIVERSIFICATION

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

000

001

003 004

010 011

012

013

014

015

016

017

018

019

021

025

026027028

029

031

033

034

035

036

040

041

042

043

045

046

047

048

050 051

052

#### **ABSTRACT**

A fundamental challenge in protein design is the trade-off between generating structural diversity while preserving motif biological function. Current state-ofthe-art methods, such as partial diffusion in RFdiffusion, often fail to resolve this trade-off: small perturbations yield motifs nearly identical to the native structure, whereas larger perturbations violate the geometric constraints necessary for biological function. We introduce Protein Generation with Embedding Learning (PGEL), a general framework that learns high-dimensional embeddings encoding sequence and structural features of a target motif in the representation space of a diffusion model's frozen denoiser, and then enhances motif diversity by introducing controlled perturbations in the embedding space. PGEL is thus able to loosen geometric constraints while satisfying typical design metrics, leading to more diverse yet viable structures. We demonstrate PGEL on three representative cases: a monomer, a protein-protein interface, and a cancer-related transcription factor complex. In all cases, PGEL achieves greater structural diversity, better designability, and improved self-consistency, as compared to partial diffusion. Our results establish PGEL as a general strategy for embedding-driven protein generation allowing for systematic, viable diversification of functional motifs.

#### 1 Introduction

Designing proteins that achieve precise biological functions while allowing for structural diversity has long been a central goal in computational protein design. Recent advances in structure prediction models like AlphaFold (Jumper et al., 2021; Abramson et al., 2024), RoseTTAFold (Baek et al., 2021), ESMFold (Lin et al., 2023) and Boltz (Wohlwend et al., 2024; Passaro et al., 2025) have revolutionized protein generative models and paved the way for improved diffusion models in protein design. Among them, RFdiffusion (Watson et al., 2023), which results from fine-tuning RoseTTAFold, has shown strong performance in both unconditional and conditional generation.

Yet, targeted local modification remains a challenge. A common approach is *partial diffusion* in RFdiffusion, in which a native or designed structure undergoes only a few denoising steps to induce diversification (Watson et al., 2023; Vázquez Torres et al., 2024). However, this method faces a fundamental diversity-fidelity trade-off: small structural perturbations keep near-native conformations, but lack diversity, while larger perturbations induce excessive geometric drift that disrupts functional features (Lin et al., 2024). Overcoming this limitation requires rethinking how diffusion models can introduce controlled variation while still anchoring designs to essential geometric constraints.

A promising direction comes from recent advances in conditional image generation. Models such as Stable Diffusion and Latent Diffusion Models (LDMs) generate images from noise guided by text prompts (Ho et al., 2020; Rombach et al., 2022). Beyond standard prompting, textual inversion learns new prompt embeddings to represent unseen visual concepts (Gal et al., 2022; Jin et al., 2024). Once learned, these embeddings can be diversified to generate outputs that preserve the original concept while exploring novel variations. Here we adopt this embedding-centric view in the context of protein generation.

We present Protein Generation with Embedding Learning (PGEL), a general framework representing the first adaptation of textual inversion principles to protein diffusion models. PGEL introduces two key approaches with broad applicability: (1) learning high-dimensional embeddings that capture the sequence and structural characteristics of target protein regions of interest, thus shifting the

paradigm from coordinate-space to embedding-space perturbations, and (2) relaxing evolutionary and structural constraints by masking embeddings. Although we present our work here using RFdiffusion's representation space, our method is general and readily adaptable to other protein diffusion models, and can thus leverage the rich representational capacity of pre-trained diffusion models without expensive retraining or fine-tuning.

We focus on *motif diversification*. Here, a *motif* denotes a set of residues with a particular geometric arrangement, which may govern functional activity. The objective is to generate a set of backbones that keep a fixed scaffold in real space within tight bounds, and realize diverse yet functionally plausible conformations of a motif, while satisfying standard designability criteria so that downstream sequence design and structure prediction can recover the intended structures. Some existing approaches address related challenges, but differ in scope and implementation: structure inpainting methods (*e.g.*, masked region generation) fully marginalize a region by masking and regenerating it *de novo*, discarding the specific native geometry (Zhang et al., 2023), whereas flexible backbone loop remodeling in Rosetta (KIC/Next-Generation KIC) samples local conformations under explicit geometric and energetic restraints to achieve high-fidelity but relatively localized exploration (Mandell et al., 2009; Stein & Kortemme, 2013; Leman et al., 2020). Hence these tools do not explicitly target controlled exploration of a *neighborhood* around an existing functional motif while keeping a surrounding scaffold nearly fixed.

Thus, we compare chiefly to partial diffusion in RFdiffusion, the prevailing stochastic baseline for local variation which has been recently applied in therapeutically relevant design settings, including *de novo* creation of high-affinity peptide binders and venom toxin neutralizers (Vázquez Torres et al., 2024; 2025). Across three representative scenarios involving a monomeric protein (calmodulin), a protein–protein binding site (barstar-barnase), and a p53 binder within the p53-MDM2 complex, PGEL (1000 samples) produces more designable structures (motif pLDDT  $\geq$  70, scRMSD  $\leq$  1Å, mRMSD  $\leq$  2Å) than partial diffusion: 1000 vs 411 (monomer), 990 vs 331 (binding site), and 802 vs 252 (binder). PGEL also yields more structurally diverse TM-score clusters distinguishable from native, and shows better self-consistency after inverse folding and refolding (meeting mRMSD and pAE thresholds), while maintaining predicted binding affinities comparable to native and exceeding those obtained with partial diffusion. Our results support embedding learning combined with masking as a general, efficient strategy for systematic motif diversification.

#### 2 Background

**Functional motifs.** Conditional generation around functional residues, often framed as *motif scaf-folding*, has been a focal point for recent protein design methods. In that setting, the motif and scaffold are defined as disjoint subsets with the scaffold varied while the motif geometry is preserved. Approaches like RFdiffusion (where the motif coordinates are fixed), the Monte Carlobased Twisted Diffusion Sampler (Wu et al., 2023) applied to FrameDiff, and Genie2 (Lin et al., 2024) have made progress on this task, though performance remains task-dependent and can yield few or no backbones meeting success criteria in specific cases. In our motif diversification task, the scaffold is held fixed and the motif is diversified to explore multiple, function-preserving geometric realizations, enabling improvements in *e.g.*, affinity, specificity or stability, while maintaining the broader structural context.

Protein embeddings. The limited availability of structural data motivated the development of models that transform sequences into sequence embeddings that encode structural information. These embeddings have been employed for various tasks such as property prediction using a Gaussian Process regression model (Yang et al., 2018) and residue-residue contact prediction via a Bidirectional Long Short-Term Memory (BiLSTM) architecture (Bepler & Berger, 2019). Transformers (Vaswani et al., 2017) have been used in generating sequence embeddings, including for antibody-specific applications like paratope prediction (Leem et al., 2022). Transfer learning has also been shown to significantly improve performance across architectures by enabling the use of pre-trained embeddings that capture fundamental sequence-structure relationships (Detlefsen et al., 2022). Other approaches explicitly include structural information (Ali et al., 2024), such as contact maps-derived embeddings, and have shown enhanced performance in particular downstream tasks such as structure similarity assessment (Kandathil et al., 2025), structure searching (Greener & Jamali, 2024), property prediction (Blaabjerg et al., 2024; Danner et al., 2025), and domain classification (Lau

et al., 2024). Similarly, protein function annotation and local flexibility prediction have benefited from Graph Convolutional Networks, which combine structure-derived graphs to propagate contextual signals from protein sequence embeddings obtained with pre-trained models (Gligorijević et al., 2021; Michalewicz et al., 2025).

**Diffusion models for proteins.** Earlier works adapted Denoising Diffusion Probabilistic Models (DDPMs) to protein design by conditioning on local structural elements or coarse fold constraints (Wu et al., 2024; Anand & Achim, 2022; Trippe et al., 2023; Luo et al., 2022) yet, while encouraging, they produced few sequences that refolded to target backbones. RFdiffusion subsequently emerged as the diffusion approach that reliably yields designable structures and sequences that recover the intended geometry. In RFdiffusion, a highly accurate protein structure prediction method (RoseTTAFold (Baek et al., 2021)) is fine-tuned to undo random perturbations of atomic coordinates introduced via 3D Gaussian noise (*i.e.*, to denoise). RFdiffusion can be constrained to specific binding targets, or symmetry specifications, and once trained it can be viewed as a *frozen denoiser*. RoseTTAFold/AlphaFold-style models (including RFdiffusion) learn so-called *state* and *pair* embeddings (related to per-residue and residue-residue properties of the protein structure, respectively) and MSA embeddings related to multiple sequence alignment (Jumper et al., 2021).

**Textual inversion.** Gal et al. (2022) builds on LDMs (Rombach et al., 2022), a specific class of DDPMs, to perform textual inversion. In the context of text-to-image models, let x represent an image, s a text prompt,  $\epsilon_{\theta}$  a pre-trained denoising network, and  $\varepsilon$  an image encoder. LDMs aim to minimize the following loss:

$$\mathcal{L}_{\text{LDM}} := \mathbb{E}_{z \sim \varepsilon(x), s, \epsilon \sim \mathcal{N}(0, 1), t} \left[ \| \epsilon - \epsilon_{\theta}(z_t, t, c_{\theta}(s)) \|_2^2 \right]$$
 (1)

Here,  $c_{\theta}(s)$  represents a pre-trained text encoder that conditions the denoiser  $\epsilon_{\theta}$  based on the text prompt s, and  $z_t$  is a noised version of the image embedding z at timestep t. The goal of textual inversion is to learn a new text embedding  $v_*$  corresponding to a particular concept  $s_*$  such that it minimizes the LDM loss (equation 1). This means conditioning  $\epsilon_{\theta}$  on  $v_*$  so the generated image  $\tilde{x}$  closely resembles the original image x. Neutral prompts, such as "A photo of  $s_*$ " or "A portrait of  $s_*$ " are used while keeping  $\epsilon_{\theta}$  and  $c_{\theta}$  frozen. Multi-Concept Prompt Learning (Jin et al., 2024) extends this idea to handle multiple concepts by incorporating three regularization techniques: attention masking, bind adjective, and prompts contrastive loss.

#### 3 METHODS

We now present our method, *Protein Generation with Embedding Learning (PGEL)*, and describe how we learn the embedding representation of a motif in Section 3.1. In Section 3.2, we propose an approach to increase motif diversity, and Section 3.3 details the evaluation metrics.

#### 3.1 PROTEIN GENERATION WITH EMBEDDING LEARNING (PGEL)

We generalize the notion of textual inversion with LDMs to proteins, treating the structure as analogous to an image, and the sequence as analogous to a text prompt. Let  $R_*$  be a region of interest, or motif, defined as a continuous or discontinuous set of  $L_*$  amino acids within a protein. The motif has structure  $x_*$  and sequence  $s_*$ , where the coordinates of  $x_*$  are obtained from an experimental Protein Data Bank (PDB) entry, and the sequence  $s_*$  is masked when passed as an input to PGEL, i.e., the amino acid range of the motif is specified, but not its exact composition.

PGEL learns a representation of  $R_*$  in embedding space, which we denote as  $v_*$ . The remainder of the protein constitutes the *scaffold*, with structure  $x_c$  and sequence  $s_c$  of length  $L_c$ , from which the protein LDM frozen ENCODER computes an embedding representation  $v_c$ .

The procedure (see Figure 1 and Algorithm 1) starts by building a noised protein structure in which the scaffold coordinates are retained while the motif coordinates are subjected to T rounds of Gaussian noise injection, following Trippe et al. (2023). At each timestep t, the protein LDM frozen DENOISER predicts a denoised motif structure  $\hat{x}_*^{(0)}$ , conditioned jointly on the learnable motif embedding  $v_*$  and the fixed embedding  $v_c$ . These embeddings include state, pair and MSA embeddings. Then, by using structure  $x^{(t)}$  and the intermediate structure  $x^{(t)}$ , a reverse diffusion step REVERSESTEP, which does not contain any learnable parameters, yields  $x^{(t-1)}$  (see Algorithm 3). In

practice, we employ pre-trained building blocks of RFdiffusion for both the ENCODER and DE-NOISER, though alternative models could be substituted if desired.

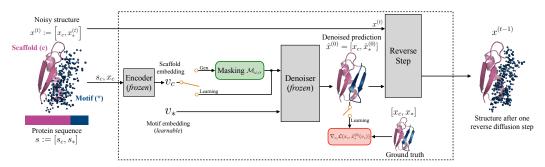


Figure 1: Outline of the PGEL learning and generation procedures during one reverse diffusion step.

**Embedding optimization.** The embedding  $v_*$  is learned by minimizing:

$$\mathcal{L} = \mathcal{L}_{MSE} + \lambda_{DM} \mathcal{L}_{DM} + \lambda_{torsion} \mathcal{L}_{torsion}$$
(2)

This loss function is composed of three terms, described hereafter, which compare different features of the ground truth structure  $x_*$  and the predicted structure  $\hat{x}_*^{(0)}(v_*)$  of the motif with the coefficients  $\lambda_{\rm DM}, \lambda_{\rm torsion} \in \mathbb{R}_{\geq 0}$  controlling the relative weight of the terms.

Data fidelity term (backbone atoms). For each motif residue  $i \in R_*$  we consider the A=4 backbone atoms (nitrogen N,  $\alpha$ -carbon  $C_{\alpha}$ , carbon C, oxygen O). Let  $\hat{x}_{i,a}^{(0)} \in \mathbb{R}^3$  denote the predicted position of atom a in residue i and  $x_{i,a} \in \mathbb{R}^3$  its ground truth counterpart. We then compute the mean squared error (MSE) between the backbone atoms of the ground truth and predicted motif:

$$\mathcal{L}_{\text{MSE}}(x_*, \hat{x}_*^{(0)}(v_*)) = \frac{1}{AL_*} \sum_{i \in R_*} \sum_{a \in \{\text{N,C}_\alpha, \text{C,O}\}} \left\| \hat{x}_{i,a}^{(0)}(v_*) - x_{i,a} \right\|^2$$
(3)

**Distance matrix between**  $\alpha$ -carbons. Let  $\hat{x}_{i,C_{\alpha}}^{(0)} \in \mathbb{R}^{3}$  denote the predicted position of the  $\alpha$ -carbon atom in residue i, and  $x_{i,C_{\alpha}} \in \mathbb{R}^{3}$  its ground truth counterpart. We define the following loss term based on  $\alpha$ -carbon Distance Matrices (DM), inspired by the distrogram notion (Senior et al., 2020):

$$\mathcal{L}_{\text{DM}}(x_*, \hat{x}_*^{(0)}(v_*)) = \frac{1}{L_*^2} \sum_{i \in R_*} \sum_{j \in R_*} \left( \left\| \hat{x}_{i, \mathsf{C}_\alpha}^{(0)}(v_*) - \hat{x}_{j, \mathsf{C}_\alpha}^{(0)}(v_*) \right\| - \|x_{i, \mathsf{C}_\alpha} - x_{j, \mathsf{C}_\alpha}\| \right)^2$$
(4)

In contrast to  $\mathcal{L}_{MSE}$ ,  $\mathcal{L}_{DM}$  is invariant under rigid motions (translations and rotations), thus encouraging global shape consistency.

**Backbone torsion angles.** Let  $\hat{\phi}_i$  and  $\hat{\psi}_i$  denote the predicted backbone torsion angles at residue i, computed from  $\hat{x}_*^{(0)}(v_*)$ , and let  $\phi_i$  and  $\psi_i$  be the corresponding ground truth values (Ramachandran et al., 1963). We impose a constraint on angular torsions through a cosine-based loss term akin to that of AlphaFold (Jumper et al., 2021):

$$\mathcal{L}_{\text{torsion}}(x_*, \hat{x}_*^{(0)}(v_*)) = \frac{1}{L_* - 2} \sum_{i=2}^{L_* - 1} \left[ 1 - \cos(\hat{\phi}_i(v_*) - \phi_i) + 1 - \cos(\hat{\psi}_i(v_*) - \psi_i) \right]$$
(5)

This term penalizes sterically implausible geometries, helping improve performance under the predicted local distance difference test (pLDDT). We do not include the third backbone angle  $\omega$  as it is typically considered fixed at 180 degrees (Cutello et al., 2006).

Once the embeddings are learned, we employ Algorithm 2 to generate novel proteins containing a diversification of the region of interest  $R_*$  (see Figure 1).

#### 3.2 Enhancing the diversity of generated motifs

MSA embeddings in sequence-to-structure predictors contain evolutionary covariation information about residues, thereby capturing geometric constraints such as residue proximity. In RFdiffusion,

#### Algorithm 1 PGEL – Embedding learning **Input:** region of interest/motif $R_*$ with masked sequence $s_*$ and structure $x_*$ , fixed scaffold with sequence $s_c$ and structure $x_c$ , pre-trained ENCODER and DENOISER. **Output:** learned embedding $v_*$ for region $R_*$ . initialize $v_*$ with zeros. while not converged do Build noised structure $x^{(T)} := [x_c, x_*^{(T)}]$ with associated sequence $s := [s_c, s_*]$ . $v_c = \text{Encoder}(s_c, x_c)$ for t = T down to 1 do $\hat{x}_*^{(0)} = \text{Denoiser}(v_c, v_*)$ $x^{(t-1)} = \text{REVERSESTEP}(x^{(t)}, [x_c, \hat{x}_*^{(0)}])$ Update $v_*$ by taking a gradient step $\nabla_{v_*} \mathcal{L}(x_*, \hat{x}_*^{(0)}(v_*))$ end for end while Return $v_*$

### Algorithm 2 PGEL – Generation with embedding masking

```
Input: region of interest/motif R_* with learned embeddings v_*, fixed scaffold with sequence s_c and structure x_c, pre-trained ENCODER and DENOISER. Output: generated structure. Build noised structure x^{(T)} := [x_c, x_*^{(T)}]. Draw at random the sample masking type \omega \sim \mathrm{Ber}(\frac{1}{2}) (row if 0, column if 1). Sample masking rate \alpha \sim \mathcal{U}[0,1]. Define \mathcal{M}_{\omega,\alpha}(\cdot) as a zero mask with type \omega and rate \alpha. v_c = \mathrm{ENCODER}(s_c, x_c) for t = T down to 1 do \hat{x}_*^{(0)} = \mathrm{DENOISER}(\mathcal{M}_{\omega,\alpha}(v_c), v_*) x^{(t-1)} = \mathrm{REVERSESTEP}(x^{(t)}, [x_c, \hat{x}_*^{(0)}]) end for Return x^{(0)}
```

however, such embeddings are derived solely from the input sequence  $s:=[s_c,s_*]$  rather than from a full stack of aligned sequences, and can be represented as a  $d_{\mathrm{MSA}} \times L$  matrix, where  $d_{\mathrm{MSA}}$  is the depth of the MSA embeddings and  $L:=L_*+L_c$  the total protein length. With PGEL, we show that the diversity of generated structures can be increased by applying perturbations to the scaffold MSA embeddings  $v_c \in \mathbb{R}^{d_{\mathrm{MSA}} \times L_c}$ . These embeddings couple through attention mechanisms with state and pair embeddings produced by an internal RFdiffusion encoder, and also interact with the learned motif embedding  $v_*$ , which provides an independent conditioning signal for the frozen denoiser.

Embedding masking. We studied the effect of applying zero masks, *i.e.* masks zeroing specific elements, to the scaffold MSA embeddings during generation (see Algorithm 2). Row masking corresponds to masking specific features for all residues, whereas column masking zeroes out all features of specific residues. Both strategies lift some constraints on inter-residue distances, and modulate which co-variation patterns remain accessible. We sample  $\omega \sim \text{Ber}(\frac{1}{2})$  to choose the masking mode ( $\omega=0$  for row masking and  $\omega=1$  for column masking) and  $\alpha \sim \mathcal{U}[0,1]$ , the masking rate, to set the fraction of rows or columns masked. This defines the operator  $\mathcal{M}_{\omega,\alpha}(\cdot)$ , which implements zero masking with type  $\omega$  and rate  $\alpha$ . As such, masking  $v_c$  relaxes the geometric constraints of the generated motif.

#### 3.3 EVALUATION METRICS

**Designability.** To quantify designability in the motif diversification task, we first require a motif  $pLDDT \geq 70$  as computed by an RFdiffusion internal block, following the threshold adopted for related tasks by Lin et al. (2024). We also require the scaffold RMSD to be  $scRMSD \leq 1\text{Å}$ , to ensure that the residues surrounding the motif remain fixed. Finally, we set the threshold for the motif RMSD to  $mRMSD \leq 2\text{Å}$ , to allow for structural diversification of the motif.

272

273

275

277 278

279

280

281

282

283

284 285 286

287

288

289

290

291

292 293

295

296

297

298

299

300

301

302

303

304

305

306

307

308

309 310

311 312

313

314

315

316

317

318 319

320 321

322

323

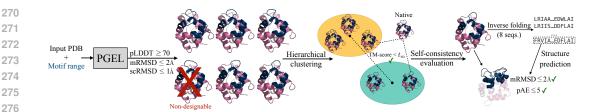


Figure 2: Summary of the evaluation metrics. PGEL takes as input a PDB entry and the amino acid range corresponding to the motif. From 1000 PGEL-generated backbones, designable candidates are filtered by root mean square deviation (RMSD) and pLDDT thresholds, and structural diversity is assessed via hierarchical clustering. Backbones are also required to be distinguishable from the native. Cluster representatives undergo self-consistency evaluation: sequences assigned to the designable backbones with ProteinMPNN are refolded, and at least one predicted structure must satisfy set mRMSD and predicted alignment error (pAE) conditions relative to the generated backbone.

**Diversity.** To quantify structural diversity among generated proteins (Figure 2), we compute the pairwise TM-scores (Zhang & Skolnick, 2004) across all designable candidates, and employ hierarchical clustering (Lin et al., 2024) with several linkage thresholds  $t_m$  to group similar backbones under this score. Diversity is measured by the number of clusters. We evaluate also the TM-score with respect to the native motif: a cluster is considered *distinguishable* from native if, for TM-score threshold  $t_m \in [0,1]$ , at least one cluster member exhibits lower similarity than  $t_m$  relative to the native. This analysis ensures that we capture the structural distinctiveness of the backbones.

**Self-consistency.** For the distinguishable backbones we use the procedure in Trippe et al. (2023) based on inverse folding to assess self-consistency between generated and predicted structures. Specifically, we use ProteinMPNN with default parameters (Dauparas et al., 2022) to assign 8 plausible sequences to each backbone, followed by a sequence-to-structure model, here AlphaFold3 (Abramson et al., 2024), to predict 8 full proteins. A designed backbone is deemed selfconsistent if it satisfies for at least one of the 8 predicted structures: mRMSD  $\leq 2$ Å and pAE  $\leq 5$ (Figure 2). Previous studies included this procedure under the designability assessment (Watson et al., 2023; Lin et al., 2024). However, this is computationally expensive and, when prioritizing diversity, often inefficient: many backbones either fail the initial scRMSD or mRMSD filters or exhibit negligible structural diversity. In motif diversification, diversity among generated proteins and distinguishability with respect to the native are decisive. We therefore invert the pipeline to enforce these criteria first and reserve the costly self-consistency evaluation only for diverse candidates.

Binding affinity. For protein-protein complexes, we run PRODIGY (Vangone & Bonvin, 2015; Xue et al., 2016) to estimate the binding affinity  $\Delta G$  expressed in kcal/mol, with larger  $|\Delta G|$ values indicating stronger binding. It is desirable that new designs present binding affinity values comparable to, or larger in magnitude than, those of the native complex. Note that learning is not optimized to enhance binding affinity; rather, this serves as an a posteriori assessment.

#### EXPERIMENTS

We focus on three representative test cases proposed in Watson et al. (2023) for different tasks: (1) Calmodulin, a monomer that plays a pivotal role in regulating the activity of nearly 100 diverse target enzymes and structural proteins (Fallon & Quiocho, 2003); (2) the barstar-barnase complex, in which the binding interface of barstar was diversified to probe its interaction with the extracellular ribonuclease barnase (Caro et al., 2023); (3) the cancer-related transcription factor p53 bound to its negative regulator MDM2 (Klein & Vassilev, 2004; Li et al., 2010).

#### 4.1 PROTOCOL

We established a protocol to systematically compare our method with RFdiffusion's partial diffusion using the metrics introduced in Section 3.3. For partial diffusion, we generated 1000 protein backbones by uniformly sampling the number of diffusion timesteps,  $T \sim \mathcal{U}\{2, 3, \dots, 49\}$ , as T = 50corresponds to the full diffusion process in RFdiffusion. In this way, we cover a spectrum of struc-

Table 1: Comparison of partial diffusion in RFdiffusion and PGEL. The number of self-consistent clusters (diversity) is computed at TM-score threshold  $t_m = 0.6$ .

	Ι	Designability		Diversity		
	(No. of viab	le structures out	of 1000)	(No. of self-consistent clusters)		
	Monomer	Binding site	Binder	Monomer	Binding site	Binder
Partial diffusion	411	331	252	0	6	1
PGEL	1000	990	802	2	10	6

tural perturbations ranging from near-native backbones to unrelated ones. For PGEL, we performed the learning of  $v_*$  with Stochastic Gradient Descent with learning rate  $l_r = 4 \times 10^{-4}$  and momentum p = 0.9,  $\lambda_{\rm DM} = 0.01$  and  $\lambda_{\rm torsion} = 0.05$  (Algorithm 1), and we then generated 1000 protein backbones (Algorithm 2).

For both sets of 1000 generated structures, we evaluated designability and, among those deemed designable, we computed TM-scores between all generated motifs and with respect to the native structure. We then plotted the number of clusters as a function of the TM-score. For structures that were designable and diverse according to a typical TM-score threshold  $t_m=0.6$  (Lin et al., 2024), we performed inverse folding through ProteinMPNN to generate compatible sequences, followed by AlphaFold3 inference to assess whether the predicted sequences refolded into the intended backbones, fulfilling the self-consistency requirement defined in Section 3.3.

#### 4.2 Example 1: Monomer

Calmodulin (PDB entry: 1PRW), a monomeric protein containing a double EF-hand motif spanning residues 16–35 and 52–71, was considered as the representative test case for single-chain proteins. All the 1000 backbones candidates generated by PGEL resulted to be designable, well exceeding the 411 obtained by partial diffusion (Table 1). All of the backbones generated by partial diffusion satisfied the pLDDT constraint, consistent with the fact that RFdiffusion's training favors high-confidence local structures, but 589 of them failed to meet the expected motif RMSD threshold. In these cases, the added noise during diffusion excessively perturbed the initial backbone, leading to conformations that no longer preserved the intended geometry of the EF-hand motif.

We then evaluated the structural diversity of the designable backbones, recording the number of clusters as a function of the TM-score threshold (Figure 3A). PGEL consistently produced a higher number of clusters across thresholds, demonstrating that embedding perturbations through masking introduce greater variability in backbone conformations. On the other hand, partial diffusion yielded structures too similar to the native backbone, and hence not distinguishable from it.

We carried out the self-consistency assessment at  $t_m=0.6$ , as per protocol. For PGEL, the two clusters had backbones that successfully refolded into the intended conformations after sequence design and AlphaFold3 inference. Figure 3B illustrates these two successful cases, along with examples of backbones generated by partial diffusion that either did not satisfy the mRMSD metric condition or the distinguishability from native.

#### 4.3 Example 2: BINDING SITE

Barstar is a small protein that binds the active site of barnase to prevent the latter from breaking RNA. This toxin-antitoxin pair (PDB entry: 7MRX) has been increasingly exploited in cancer therapy for targeted cytotoxicity (Kalinin et al., 2023). As target for motif diversification, we took the barstar's binding interface region comprising residues 25 to 46 (see Watson et al. (2023)).

Of the 1000 backbones generated with PGEL, 990 were classified as designable, nearly tripling the 331 obtained with partial diffusion. Correspondingly, Figure 4A demonstrates that PGEL consistently outperforms partial diffusion across the entire range of  $t_m \in [0,1]$ , with pronounced differences observed at  $t_m > 0.9$  and within  $0.45 < t_m < 0.55$ , around canonical TM-score thresholds. At  $t_m = 0.6$ , PGEL yielded 15 structural clusters compared to 13 for partial diffusion, which were reduced to 10 and 6, respectively, after self-consistency checks (Table 1).

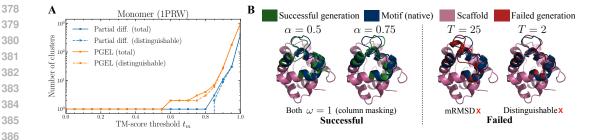


Figure 3: Results for example 1. (A) Number of clusters, both total and distinguishable from native, as a function of the TM-score threshold  $t_m$  for PGEL and partial diffusion. (B) Left: two successful PGEL designs at  $t_m=0.6$  using column masking with rates  $\alpha=0.5$  and  $\alpha=0.75$ . Right: two partial diffusion failed backbones at  $t_m=0.6$ , one obtained with T=25 timesteps that violates the motif RMSD constraint, and one with T=2 timesteps that is not distinguishable from the native.

In Figure 4A, we also display an overlay version of the native motif and 10 representative motifs derived from these clusters, highlighting the sequence variability both among generated barstar binding interfaces and relative to the native PDB structure. When predicting in silico the binding affinity of the generated complexes with PRODIGY, two of the generated structures exhibited binding affinities higher than the native complex (see Figure 5A and Table 3), while the remaining eight retained at least 80% of the original affinity  $\Delta G_{\rm native}$ . In contrast, only one complex generated by partial diffusion had a binding affinity value comparable to that of the native ( $\Delta G_{\rm design} > 0.9\Delta G_{\rm native}$ ).

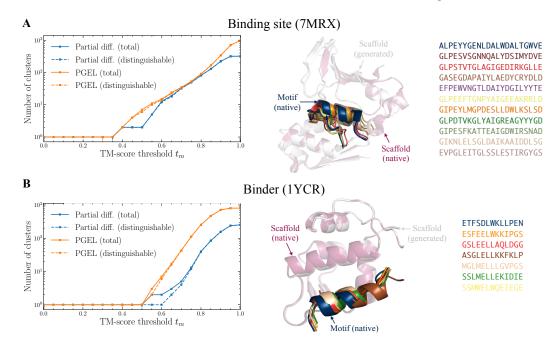


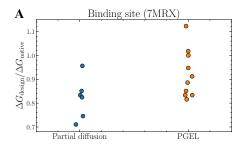
Figure 4: (A) Example 2: Left: Number of clusters identified PGEL and partial diffusion in the binding site example, both total and distinguishable from native, as a function of the TM-score threshold  $t_m$ . Right: generated binding site backbones (overlaid with native), alongside the native sequence and sequences that refold to self-consistent structures. (B) Same as A but for example 3.

#### 4.4 EXAMPLE 3: BINDER

The interaction between the transcription factor p53 and its negative regulator MDM2 is a key molecular process in cancer progression. Specifically, pharmacological disruption of the p53-MDM2 complex restores p53 activity and has been proven beneficial in cancer therapy (Hu et al., 2021).

Starting from PDB entry 1YCR, we addressed the motif diversification task by redesigning the complete p53 under the RMSD constraints described in Section 3.3. PGEL generated 802 designable backbones out of 1000 trials, with most non-designable cases attributable to low pLDDT confidence scores (Table 1). Partial diffusion produced only 252 designable backbones with considerably reduced structural diversity (a single cluster at TM-score threshold  $t_m=0.6$ , see Figure 4B). PGEL, by comparison, gave six clusters at  $t_m=0.6$ , all of which passed the self-consistency checks.

When assessing the binding affinity a posteriori, five out of six representatives of PGEL clusters exhibited lower affinity compared to the sole valid instance of partial diffusion (Figure 5B, Table 3). Notably, sequence SSMWELWQEIEGE (see Figure 4B), designed with PGEL in combination with ProteinMPNN, folded, as predicted by AlphaFold3, into a structure with a binding affinity comparable to that of the native structure, despite sharing only around 15% of sequence identity. This result highlights PGEL's ability to generate backbones that can accommodate sequences unrelated to the native while refolding into structures that preserve function.



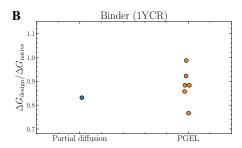


Figure 5: Ratios of the binding affinity predicted with PRODIGY of the native structure *vs* the AlphaFold3-predicted structures from backbones (one per cluster) generated by PGEL and partial diffusion for example 2 (A) and example 3 (B).

Computational remarks. Across the three case studies, column masking was more effective than row masking (see Table 2, with 80% of successful outcomes with  $\omega = 1$ ).

The time required per timestep during the generation process is nearly indistinguishable between PGEL and partial diffusion: average timestep  $0.7\,s$  for partial diffusion and  $0.71\,s$  for PGEL on a single NVIDIA GeForce RTX 3090 GPU with 24GB of memory.

#### 5 Limitations and future work

PGEL inherits the biases and limitations of the underlying frozen RFdiffusion denoiser, including its training data distribution and architectural constraints. Moreover, learning embeddings requires additional optimization time, which ranged in our examples from 2 minutes (example 2) to 2 hours (example 3) on a single GPU, with a trade-off between speed and improved results. Our experiments were limited to motifs of up to 40 residues, with practical limits of around 50 residues given available memory, though scaling to longer motifs should be feasible with larger hardware or engineering optimization. Beyond this, our evaluation is entirely *in silico* (pLDDT/RMSD/TM-score filtering, AlphaFold3 refolding, and PRODIGY  $\Delta G$ ) and thus predictive rather than experimental.

Future work will investigate alternative ways of perturbing embeddings, as this strategy for motif diversification remains largely unexplored, as well as different strategies for sampling the masking parameters  $\omega$  and  $\alpha$ . For instance, instead of sampling  $\alpha$  uniformly between 0 and 1, one could bias it toward smaller masking rates (e.g., using a Poisson distribution with rate  $\lambda$ , where  $\lambda$  tunes how conservative or aggressive the masking is), thus providing finer control over structural perturbations. A more systematic mapping between PGEL's  $\omega$  and  $\alpha$  and partial diffusion's T would also clarify the relationship between the diversity-fidelity trade-off in both methods. Finally, experimental validation will be pursued in follow-up work.

**Code availability.** Upon publication, we will release code and configurations to facilitate reproducibility.

#### REFERENCES

- Josh Abramson, Jonas Adler, Jack Dunger, Richard Evans, Tim Green, Alexander Pritzel, Olaf Ronneberger, Lindsay Willmore, Andrew J. Ballard, Joshua Bambrick, Sebastian W. Bodenstein, David A. Evans, Chia-Chun Hung, Michael O'Neill, David Reiman, Kathryn Tunyasuvunakool, Zachary Wu, Akvilė Žemgulytė, Eirini Arvaniti, Charles Beattie, Ottavia Bertolli, Alex Bridgland, Alexey Cherepanov, Miles Congreve, Alexander I. Cowen-Rivers, Andrew Cowie, Michael Figurnov, Fabian B. Fuchs, Hannah Gladman, Rishub Jain, Yousuf A. Khan, Caroline M. R. Low, Kuba Perlin, Anna Potapenko, Pascal Savy, Sukhdeep Singh, Adrian Stecula, Ashok Thillaisundaram, Catherine Tong, Sergei Yakneen, Ellen D. Zhong, Michal Zielinski, Augustin Žídek, Victor Bapst, Pushmeet Kohli, Max Jaderberg, Demis Hassabis, and John M. Jumper. Accurate structure prediction of biomolecular interactions with alphafold 3. *Nature*, 630(8016):493–500, Jun 2024. ISSN 1476-4687. doi: 10.1038/s41586-024-07487-w. URL https://doi.org/10.1038/s41586-024-07487-w.
- Sarwan Ali, Prakash Chourasia, and Murray Patterson. When protein structure embedding meets large language models. *Genes*, 15(1), 2024. ISSN 2073-4425. doi: 10.3390/genes15010025. URL https://www.mdpi.com/2073-4425/15/1/25.
- Namrata Anand and Tudor Achim. Protein structure and sequence generation with equivariant denoising diffusion probabilistic models. *arXiv* preprint arXiv:2205.15019, 2022.
- Minkyung Baek, Frank DiMaio, Ivan Anishchenko, Justas Dauparas, Sergey Ovchinnikov, Gyu Rie Lee, Jue Wang, Qian Cong, Lisa N. Kinch, R. Dustin Schaeffer, Claudia Millán, Hahnbeom Park, Carson Adams, Caleb R. Glassman, Andy DeGiovanni, Jose H. Pereira, Andria V. Rodrigues, Alberdina A. van Dijk, Ana C. Ebrecht, Diederik J. Opperman, Theo Sagmeister, Christoph Buhlheller, Tea Pavkov-Keller, Manoj K. Rathinaswamy, Udit Dalwadi, Calvin K. Yip, John E. Burke, K. Christopher Garcia, Nick V. Grishin, Paul D. Adams, Randy J. Read, and David Baker. Accurate prediction of protein structures and interactions using a three-track neural network. *Science*, 373(6557):871–876, 2021. doi: 10.1126/science.abj8754. URL https://www.science.org/doi/abs/10.1126/science.abj8754.
- Tristan Bepler and Bonnie Berger. Learning protein sequence embeddings using information from structure. In *International Conference on Learning Representations*, ICLR'19, 2019.
- Lasse M. Blaabjerg, Nicolas Jonsson, Wouter Boomsma, Amelie Stein, and Kresten Lindorff-Larsen. Ssemb: A joint embedding of protein sequence and structure enables robust variant effect predictions. *Nature Communications*, 15(1):9646, Nov 2024. ISSN 2041-1723. doi: 10.1038/s41467-024-53982-z. URL https://doi.org/10.1038/s41467-024-53982-z.
- José A. Caro, Kathleen G. Valentine, Taylor R. Cole, and A. Joshua Wand. Pressure, motion, and conformational entropy in molecular recognition by proteins. *Biophysical Reports*, 3(1), Mar 2023. ISSN 2667-0747. doi: 10.1016/j.bpr.2022.100098. URL https://doi.org/10.1016/j.bpr.2022.100098.
- Vincenzo Cutello, Giuseppe Narzisi, and Giuseppe Nicosia. A multi-objective evolutionary approach to the protein structure prediction problem. *Journal of The Royal Society Interface*, 3(6):139–151, 2006. doi: 10.1098/rsif.2005.0083. URL https://royalsocietypublishing.org/doi/abs/10.1098/rsif.2005.0083.
- Martin Danner, Matthias Begemann, Miriam Elbracht, Ingo Kurth, and Jeremias Krause. Utilizing protein structure graph embeddings to predict the pathogenicity of missense variants. *NAR Genomics and Bioinformatics*, 7(3):lqaf097, 07 2025. ISSN 2631-9268. doi: 10.1093/nargab/lqaf097. URL https://doi.org/10.1093/nargab/lqaf097.
- J. Dauparas, I. Anishchenko, N. Bennett, H. Bai, R. J. Ragotte, L. F. Milles, B. I. M. Wicky, A. Courbet, R. J. de Haas, N. Bethel, P. J. Y. Leung, T. F. Huddy, S. Pellock, D. Tischer, F. Chan, B. Koepnick, H. Nguyen, A. Kang, B. Sankaran, A. K. Bera, N. P. King, and D. Baker. Robust deep learning-based protein sequence design using proteinmpnn. *Science*, 378(6615):49-56, 2022. doi: 10.1126/science.add2187. URL https://www.science.org/doi/abs/10.1126/science.add2187.

Nicki Skafte Detlefsen, Søren Hauberg, and Wouter Boomsma. Learning meaningful representations of protein sequences. *Nature Communications*, 13(1):1914, Apr 2022. ISSN 2041-1723. doi: 10.1038/s41467-022-29443-w. URL https://doi.org/10.1038/s41467-022-29443-w.

Jennifer L Fallon and Florante A Quiocho. A closed compact structure of native ca2+-calmodulin. Structure, 11(10):1303-1307, 2003. ISSN 0969-2126. doi: https://doi.org/10.1016/j.str. 2003.09.004. URL https://www.sciencedirect.com/science/article/pii/S0969212603002053.

- Rinon Gal, Yuval Alaluf, Yuval Atzmon, Or Patashnik, Amit H. Bermano, Gal Chechik, and Daniel Cohen-Or. An image is worth one word: Personalizing text-to-image generation using textual inversion, 2022. URL https://arxiv.org/abs/2208.01618.
- Vladimir Gligorijević, P. Douglas Renfrew, Tomasz Kosciolek, Julia Koehler Leman, Daniel Berenberg, Tommi Vatanen, Chris Chandler, Bryn C. Taylor, Ian M. Fisk, Hera Vlamakis, Ramnik J. Xavier, Rob Knight, Kyunghyun Cho, and Richard Bonneau. Structure-based protein function prediction using graph convolutional networks. *Nature Communications*, 12(1):3168, May 2021. ISSN 2041-1723. doi: 10.1038/s41467-021-23303-9. URL https://doi.org/10.1038/s41467-021-23303-9.
- Joe G Greener and Kiarash Jamali. Fast protein structure searching using structure graph embeddings. *Bioinformatics Advances*, 5(1):vbaf042, 03 2024. ISSN 2635-0041. doi: 10.1093/bioadv/vbaf042. URL https://doi.org/10.1093/bioadv/vbaf042.
- Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. *Advances in neural information processing systems*, 33:6840–6851, 2020.
- Jiahao Hu, Jiasheng Cao, Win Topatana, Sarun Juengpanich, Shijie Li, Bin Zhang, Jiliang Shen, Liuxin Cai, Xiujun Cai, and Mingyu Chen. Targeting mutant p53 for cancer therapy: direct and indirect strategies. *Journal of Hematology & Oncology*, 14(1):157, Sep 2021. ISSN 1756-8722. doi: 10.1186/s13045-021-01169-0. URL https://doi.org/10.1186/s13045-021-01169-0.
- Chen Jin, Ryutaro Tanno, Amrutha Saseendran, Tom Diethe, and Philip Alexander Teare. An image is worth multiple words: Discovering object level concepts using multi-concept prompt learning. In *Forty-first International Conference on Machine Learning*, 2024. URL https://openreview.net/forum?id=F3x6uYILgL.
- John Jumper, Richard Evans, Alexander Pritzel, Tim Green, Michael Figurnov, Olaf Ronneberger, Kathryn Tunyasuvunakool, Russ Bates, Augustin Žídek, Anna Potapenko, Alex Bridgland, Clemens Meyer, Simon A. A. Kohl, Andrew J. Ballard, Andrew Cowie, Bernardino Romera-Paredes, Stanislav Nikolov, Rishub Jain, Jonas Adler, Trevor Back, Stig Petersen, David Reiman, Ellen Clancy, Michal Zielinski, Martin Steinegger, Michalina Pacholska, Tamas Berghammer, Sebastian Bodenstein, David Silver, Oriol Vinyals, Andrew W. Senior, Koray Kavukcuoglu, Pushmeet Kohli, and Demis Hassabis. Highly accurate protein structure prediction with alphafold. *Nature*, 596(7873):583–589, Aug 2021. ISSN 1476-4687. doi: 10.1038/s41586-021-03819-2. URL https://doi.org/10.1038/s41586-021-03819-2.
- R. S. Kalinin, V. O. Shipunova, Y. P. Rubtsov, V. M. Ukrainskay, A. Schulga, E. V. Konovalova, D. V. Volkov, I. A. Yaroshevich, A. M. Moysenovich, A. A. Belogurov, G. B. Telegin, A. S. Chernov, M. A. Maschan, S. S. Terekhov, V. D. Knorre, E. Khurs, N. V. Gnuchev, A. G. Gabibov, and S. M. Deyev. Barnase-barstar specific interaction regulates car-t cells cytotoxic activity toward malignancy. *Doklady Biochemistry and Biophysics*, 508(1):17–20, Feb 2023. ISSN 1608-3091. doi: 10. 1134/S1607672922700041. URL https://doi.org/10.1134/S1607672922700041.
- Shaun M Kandathil, Andy M Lau, Daniel W A Buchan, and David T Jones. Foldclass and merizosearch: scalable structural similarity search for single- and multi-domain proteins using geometric learning. *Bioinformatics*, 41(5):btaf277, 05 2025. ISSN 1367-4811. doi: 10.1093/bioinformatics/btaf277. URL https://doi.org/10.1093/bioinformatics/btaf277.

- C. Klein and L. T. Vassilev. Targeting the p53-mdm2 interaction to treat cancer. *British Journal of Cancer*, 91(8):1415-1419, Oct 2004. ISSN 1532-1827. doi: 10.1038/sj.bjc.6602164. URL https://doi.org/10.1038/sj.bjc.6602164.
  - Andy M. Lau, Nicola Bordin, Shaun M. Kandathil, Ian Sillitoe, Vaishali P. Waman, Jude Wells, Christine A. Orengo, and David T. Jones. Exploring structural diversity across the protein universe with the encyclopedia of domains. *Science*, 386(6721):eadq4946, 2024. doi: 10.1126/science.adq4946. URL https://www.science.org/doi/abs/10.1126/science.adq4946.
  - Jinwoo Leem, Laura S. Mitchell, James H.R. Farmery, Justin Barton, and Jacob D. Galson. Deciphering the language of antibodies using self-supervised learning. *Patterns*, 3(7):100513, 2022. ISSN 2666-3899. doi: https://doi.org/10.1016/j.patter.2022.100513. URL https://www.sciencedirect.com/science/article/pii/S2666389922001052.
  - Julia Koehler Leman, Brian D Weitzner, Steven M Lewis, Jared Adolf-Bryfogle, Nawsad Alam, Rebecca F Alford, Melanie Aprahamian, David Baker, Kyle A Barlow, Patrick Barth, et al. Macromolecular modeling and design in rosetta: recent methods and frameworks. *Nature methods*, 17 (7):665–680, 2020.
  - Chong Li, Marzena Pazgier, Changqing Li, Weirong Yuan, Min Liu, Gang Wei, Wei-Yue Lu, and Wuyuan Lu. Systematic mutational analysis of peptide inhibition of the p53-mdm2/mdmx interactions. *Journal of Molecular Biology*, 398(2):200-213, 2010. ISSN 0022-2836. doi: https://doi.org/10.1016/j.jmb.2010.03.005. URL https://www.sciencedirect.com/science/article/pii/S0022283610002433.
  - Yeqing Lin, Minji Lee, Zhao Zhang, and Mohammed AlQuraishi. Out of many, one: Designing and scaffolding proteins at the scale of the structural universe with genie 2, 2024. URL https://arxiv.org/abs/2405.15489.
  - Zeming Lin, Halil Akin, Roshan Rao, Brian Hie, Zhongkai Zhu, Wenting Lu, Nikita Smetanin, Robert Verkuil, Ori Kabeli, Yaniv Shmueli, Allan dos Santos Costa, Maryam Fazel-Zarandi, Tom Sercu, Salvatore Candido, and Alexander Rives. Evolutionary-scale prediction of atomic-level protein structure with a language model. *Science*, 379(6637):1123–1130, 2023. doi: 10.1126/science.ade2574. URL https://www.science.org/doi/abs/10.1126/science.ade2574.
  - Shitong Luo, Yufeng Su, Xingang Peng, Sheng Wang, Jian Peng, and Jianzhu Ma. Antigen-specific antibody design and optimization with diffusion-based generative models for protein structures. *Advances in Neural Information Processing Systems*, 35:9754–9767, 2022.
  - Daniel J. Mandell, Evangelos A. Coutsias, and Tanja Kortemme. Sub-angstrom accuracy in protein loop reconstruction by robotics-inspired conformational sampling. *Nature Methods*, 6:551–552, 2009. doi: 10.1038/nmeth0809-551.
  - Kevin Michalewicz, Mauricio Barahona, and Barbara Bravi. Integrating protein sequence embeddings with structure via graph-based deep learning for the prediction of single-residue properties, 2025. URL https://arxiv.org/abs/2502.17294.
  - Saro Passaro, Gabriele Corso, Jeremy Wohlwend, Mateo Reveiz, Stephan Thaler, Vignesh Ram Somnath, Noah Getz, Tally Portnoi, Julien Roy, Hannes Stark, David Kwabi-Addo, Dominique Beaini, Tommi Jaakkola, and Regina Barzilay. Boltz-2: Towards accurate and efficient binding affinity prediction. *bioRxiv*, 2025. doi: 10.1101/2025.06.14.659707.
  - G.N. Ramachandran, C. Ramakrishnan, and V. Sasisekharan. Stereochemistry of polypeptide chain configurations. *Journal of Molecular Biology*, 7(1):95-99, 1963. ISSN 0022-2836. doi: https://doi.org/10.1016/S0022-2836(63)80023-6. URL https://www.sciencedirect.com/science/article/pii/S0022283663800236.
  - Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High-resolution image synthesis with latent diffusion models. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pp. 10684–10695, 2022.

Andrew W. Senior, Richard Evans, John Jumper, James Kirkpatrick, Laurent Sifre, Tim Green, Chongli Qin, Augustin Žídek, Alexander W. R. Nelson, Alex Bridgland, Hugo Penedones, Stig Petersen, Karen Simonyan, Steve Crossan, Pushmeet Kohli, David T. Jones, David Silver, Koray Kavukcuoglu, and Demis Hassabis. Improved protein structure prediction using potentials from deep learning. *Nature*, 577(7792):706–710, Jan 2020. ISSN 1476-4687. doi: 10.1038/s41586-019-1923-7. URL https://doi.org/10.1038/s41586-019-1923-7.

- Amelie Stein and Tanja Kortemme. Improvements to robotics-inspired conformational sampling in rosetta. *PLoS ONE*, 8(5):e63090, 2013. doi: 10.1371/journal.pone.0063090.
- Brian L. Trippe, Jason Yim, Doug Tischer, David Baker, Tamara Broderick, Regina Barzilay, and Tommi S. Jaakkola. Diffusion probabilistic modeling of protein backbones in 3d for the motif-scaffolding problem. In *The Eleventh International Conference on Learning Representations*, 2023. URL https://openreview.net/forum?id=6TxBxqNME1Y.
- Anna Vangone and Alexandre MJJ Bonvin. Contacts-based prediction of binding affinity in protein-protein complexes. *eLife*, 4:e07454, jul 2015. ISSN 2050-084X. doi: 10.7554/eLife.07454. URL https://doi.org/10.7554/eLife.07454.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. In I. Guyon, U. Von Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett (eds.), Advances in Neural Information Processing Systems, volume 30. Curran Associates, Inc., 2017. URL https://proceedings.neurips.cc/paper\_files/paper/2017/file/3f5ee243547dee91fbd053c1c4a845aa-Paper.pdf.
- Susana Vázquez Torres, Philip J. Y. Leung, Preetham Venkatesh, Isaac D. Lutz, Fabian Hink, Huu-Hien Huynh, Jessica Becker, Andy Hsien-Wei Yeh, David Juergens, Nathaniel R. Bennett, Andrew N. Hoofnagle, Eric Huang, Michael J. MacCoss, Marc Expòsit, Gyu Rie Lee, Asim K. Bera, Alex Kang, Joshmyn De La Cruz, Paul M. Levine, Xinting Li, Mila Lamb, Stacey R. Gerben, Analisa Murray, Piper Heine, Elif Nihal Korkmaz, Jeff Nivala, Lance Stewart, Joseph L. Watson, Joseph M. Rogers, and David Baker. De novo design of high-affinity binders of bioactive helical peptides. *Nature*, 626(7998):435–442, Feb 2024. ISSN 1476-4687. doi: 10.1038/s41586-023-06953-1. URL https://doi.org/10.1038/s41586-023-06953-1.
- Susana Vázquez Torres, Melisa Benard Valle, Stephen P Mackessy, Stefanie K Menzies, Nicholas R Casewell, Shirin Ahmadi, Nick J Burlet, Edin Muratspahić, Isaac Sappington, Max D Overath, et al. De novo designed proteins neutralize lethal snake venom toxins. *Nature*, 639(8053):225–231, 2025.
- Joseph L. Watson, David Juergens, Nathaniel R. Bennett, Brian L. Trippe, Jason Yim, Helen E. Eisenach, Woody Ahern, Andrew J. Borst, Robert J. Ragotte, Lukas F. Milles, Basile I. M. Wicky, Nikita Hanikel, Samuel J. Pellock, Alexis Courbet, William Sheffler, Jue Wang, Preetham Venkatesh, Isaac Sappington, Susana Vázquez Torres, Anna Lauko, Valentin De Bortoli, Emile Mathieu, Sergey Ovchinnikov, Regina Barzilay, Tommi S. Jaakkola, Frank DiMaio, Minkyung Baek, and David Baker. De novo design of protein structure and function with rfdiffusion. *Nature*, 620(7976):1089–1100, Aug 2023. ISSN 1476-4687. doi: 10.1038/s41586-023-06415-8. URL https://doi.org/10.1038/s41586-023-06415-8.
- Jeremy Wohlwend, Gabriele Corso, Saro Passaro, Noah Getz, Mateo Reveiz, Ken Leidal, Wojtek Swiderski, Liam Atkinson, Tally Portnoi, Itamar Chinn, Jacob Silterra, Tommi Jaakkola, and Regina Barzilay. Boltz-1: Democratizing biomolecular interaction modeling. *bioRxiv*, 2024. doi: 10.1101/2024.11.19.624167.
- Kevin E Wu, Kevin K Yang, Rianne van den Berg, Sarah Alamdari, James Y Zou, Alex X Lu, and Ava P Amini. Protein structure generation via folding diffusion. *Nature communications*, 15(1): 1059, 2024.
- Luhuan Wu, Brian L Trippe, Christian A. Naesseth, David M Blei, and John P Cunningham. Practical and asymptotically exact conditional sampling in diffusion models. *arXiv preprint arXiv:2306.17775*, 2023.

- Li C. Xue, João Pglm Rodrigues, Panagiotis L. Kastritis, Alexandre Mjj Bonvin, and Anna Vangone. Prodigy: a web server for predicting the binding affinity of protein–protein complexes. *Bioinformatics*, 32(23):3676–3678, 08 2016. ISSN 1367-4803. doi: 10.1093/bioinformatics/btw514. URL https://doi.org/10.1093/bioinformatics/btw514.
- Kevin K Yang, Zachary Wu, Claire N Bedbrook, and Frances H Arnold. Learned protein embeddings for machine learning. *Bioinformatics*, 34(15):2642–2648, 03 2018. ISSN 1367-4803. doi: 10.1093/bioinformatics/bty178. URL https://doi.org/10.1093/bioinformatics/bty178.
- Cheng Zhang, Adam Leach, Thomas Makkink, Miguel Arbesú, Ibtissem Kadri, Daniel Luo, Liron Mizrahi, Sabrine Krichen, Maren Lang, Andrey Tovchigrechko, Nicolas Lopez Carranza, Uğur Şahin, Karim Beguir, Michael Rooney, and Yunguan Fu. FrameDiPT: SE(3) Diffusion Model for Protein Structure Inpainting. *bioRxiv*, 2023. doi: 10.1101/2023.11.21.568057. URL https://www.biorxiv.org/content/10.1101/2023.11.21.568057v2.
- Yang Zhang and Jeffrey Skolnick. Scoring function for automated assessment of protein structure template quality. *Proteins: Structure, Function, and Bioinformatics*, 57(4):702–710, 2004. doi: https://doi.org/10.1002/prot.20264. URL https://onlinelibrary.wiley.com/doi/abs/10.1002/prot.20264.

### A REVERSESTEP ALGORITHM

Let  $x^{(t)} = \{(r_l^{(t)}, u_l^{(t)})\}_{l=1}^L$  denote the noisy protein backbone structure at diffusion step t, where each residue l is represented by a rotation  $r_l^{(t)} \in SO(3)$ , with SO(3) the special orthogonal group in three dimensions, and a translation  $u_l^{(t)} \in \mathbb{R}^3$ . Let  $\hat{x}^{(0)} = \{(\hat{r}_l^{(0)}, \hat{u}_l^{(0)})\}_{l=1}^L$  denote the predicted denoised structure. Let  $\{\beta^{(t)}\}_{t=1}^T$  be a variance schedule with  $\gamma^{(t)} = 1 - \beta^{(t)}$  and  $\bar{\gamma}^{(t)} = \prod_{s=1}^t \gamma^{(s)}$ . For translations, let  $u_l^{(t-1)}$  be sampled from a Gaussian distribution with covariance  $\beta^{(t)}I_3$ . For rotations, let  $s_l$  denote the score approximation presented in Watson et al. (2023),  $\epsilon_{l,d}$  isotropic Gaussian perturbations and  $\{f_d\}_{d=1}^3$  a basis of the Lie algebra SO(3).

## Algorithm 3 REVERSESTEP function (Watson et al., 2023)

```
\begin{array}{l} \textbf{Input:} \ \text{noisy structure} \ x^{(t)}, \ \text{denoised prediction} \ \hat{x}^{(0)}. \\ \textbf{Output:} \ \text{updated structure} \ x^{(t-1)}. \\ \textbf{for} \ l = 1, \dots, L \ \textbf{do} \\ (r_l^{(t)}, u_l^{(t)}) = x_l^{(t)} \\ (\hat{r}_l^{(0)}, \hat{u}_l^{(0)}) = \hat{x}_l^{(0)} \\ u_l^{(t-1)} \sim \mathcal{N} \bigg( \frac{\sqrt{\bar{\gamma}^{(t-1)}} \beta^{(t)}}{1-\bar{\gamma}^{(t)}} \hat{u}_l^{(0)} + \frac{\sqrt{\bar{\gamma}^{(t)}} (1-\bar{\gamma}^{(t-1)})}{1-\bar{\gamma}^{(t)}} u_l^{(t)}, \ \beta^{(t)} I_3 \bigg) \\ \text{// Updating rotations below} \\ s_l = \text{ROTATIONSCOREAPPROXIMATION}(r_l^{(t)}, \hat{r}_l^{(0)}, \sigma_t^2) \\ \epsilon_{l,1}, \epsilon_{l,2}, \epsilon_{l,3} \stackrel{\text{iid}}{\sim} \mathcal{N}(0,1) \\ r_l^{(t-1)} = r_l^{(t)} \exp_{I_3} \bigg\{ (\sigma_t^2 - \sigma_{t-1}^2) r_l^{(t)\top} s_l + \sqrt{\sigma_t^2 - \sigma_{t-1}^2} \sum_{d=1}^3 \epsilon_{l,d} f_d \bigg\} \\ x_l^{(t-1)} = (r_l^{(t-1)}, u_l^{(t-1)}) \\ \textbf{end for} \\ \textbf{Return} \ x^{(t-1)} \end{array}
```

## B PGEL EVALUATION RESULTS

Table 2: Detailed results of PGEL successes for examples 1, 2 and 3.

	Table	כ ∠. ו		ills of FOE	L successes for examples 1, .	z anu 3.	
PDB & design ID	$\alpha$	$\omega$	mRMSD (Å)	Motif pLDDT	Sequence	mRMSD AF3 (Å)	pAE
1PRW							
17	0.75	1	1.78	79	FRVIAGGEDGLVTLEQLARY/VRRVAGRGGRLISFEDFLAI	1.54	4.43
860	0.5	1	1.96	81	ARWLDKGGSGAVFGEQLGEF/VAAALEGGKEAKLEEWFLNY	1.25	4.84
7MRX							
0	0.55	0	0.74	79	GLPESVSGNNQALYDSIMYDVE	0.89	3.17
31	0.4	0	1.35	78	GLPDTVKGLYAIGREAGYYYGD	0.83	3.30
100	0.95	1	1.26	73	GLPSTVTGLAGIGEDIRKGLLE	1.78	3.73
114	0.75	1	0.83	75	GASEGDAPAIYLAEDYCRYDLD	1.22	3.79
145	0.4	0	0.78	78	EFPEWVNGTLDAIYDGILYYTE	0.69	4.93
308	0.5	1	1.57	74	GIPESFKATTEAIGDWIRSNAD	1.25	4.97
352	0.85	1	1.53	75	GLPEEFTGNPYAIGEEAKRRLD	1.98	3.81
730	0.8	1	1.85	72	GIPEYLMGPDESLLDWLKSLSD	1.42	4.90
744	0.8	1	0.88	75	EVPGLEITGLSSLESTIRGYGS	1.96	2.42
814	0.3	1	0.77	78	GIKNLELSGLDAIKAAIDDLSG	1.13	3.79
1YCR							
14	0.3	1	1.36	75	GSLEELLAQLDGG	1.56	2.88
36	0.7	1	1.09	72	MGLMELLLGVPGS	1.25	3.19
285	0.3	1	1.59	72	ASGLELLKKFKLP	1.60	3.53
334	0.65	1	1.03	72	SSLMELLEKIDIE	1.25	2.88
619	0.2	1	0.59	87	ESFEELWKKIPGS	1.70	2.45
695	0.25	1	1.04	74	SSMWELWQEIEGE	0.86	2.42

## C BINDING AFFINITY RESULTS

Table 3: PRODIGY-predicted binding affinities for examples 2 and 3.

DIGY-predicted binding affinities for ex-						
PDB & design ID	Method	$\Delta G$ (kcal/mol)				
7MRX						
Native	_	-11.4				
0	PGEL	-9.3				
31	PGEL	-11.4				
100	PGEL	-10.8				
114	PGEL	-9.5				
145	PGEL	-10.1				
308	PGEL	-10.4				
352	PGEL	-12.8				
730	PGEL	-11.6				
744	PGEL	-9.5				
814	PGEL	-9.7				
1	Partial diff.	-9.5				
18	Partial diff.	-10.9				
307	Partial diff.	-8.5				
327	Partial diff.	-8.1				
513	Partial diff.	-9.7				
780	Partial diff.	-9.4				
1YCR						
Native	_	-7.7				
14	PGEL	-7.1				
36	PGEL	-5.9				
285	PGEL	-6.8				
334	PGEL	-6.8				
619	PGEL	-6.6				
695	PGEL	-7.6				
101	Partial diff.	-6.4				