

---

# Atom-level Protein Representation Learning Improves Protein Structure Prediction

---

Taewon Kim\* Hyosoon Jang\* Hyunjin Seo Seonghwan Seo Hyeongwoo Kim  
Wonho Zhung Mingyeong Shin Wooyoun Kim Sungsoo Ahn  
KAIST  
{maxkim139,hyosoon.jang,sungsoo.ahn}@kaist.ac.kr

## Abstract

Recent advances in generative modeling show that pretrained representations can improve generation as conditioning features or alignment targets. Motivated by this, we study protein representations for predicting structures beyond conventional function annotation. We propose TRIPROREP, a structure-aware pretraining method that jointly models three aligned residue-level views: amino-acid identity, backbone geometry, and local full-atom geometry, discretely encoded via VQ-VAE tokenizers. By pretraining to recover original tokens from generator-corrupted views, TRIPROREP learns to distinguish plausible but incorrect cross-view augmentations from the original protein. We further introduce REPSP, a benchmark for evaluating protein representations in structure-predictive settings. REPSP tests three uses of representations: homodimer co-folding from apo-chain representations, residue-level prediction of homodimer-derived interaction properties, and representation-aligned monomer structure prediction. Across these tasks, TRIPROREP improves over sequence-only and prior structure-aware representation models, while maintaining competitive performance on conventional benchmarks.

## 1 Introduction

Structure-aware protein representation learning aims to produce residue-level features that can be reused by downstream models for three-dimensional reasoning such as complex prediction, interface modeling, and structure prediction. Yet this goal is only indirectly tested by many common protein representation benchmarks. Enzyme commission (EC) and gene ontology (GO) prediction [1] measure broad biological utility, but they do not ask whether a representation exposes the geometric information needed to predict interfaces, assemble complexes, or supervise structure-prediction models. As a result, sequence-only representation models can remain competitive with explicitly structure-aware models on these benchmarks, which we also observe in Table 6, without resolving whether structural supervision has improved the representation for its intended use.

This mismatch is increasingly important as protein representation learning moves beyond sequence-only pretraining. SaProt [2] and ProstT5 [3] augment amino-acid sequences with Foldseek 3Di tokens [4]. ESM3 represents proteins with coupled sequence, structure, and function tracks [5]; and related methods condition sequence modeling on backbone structure or align sequence and structure representations [6, 7]. These approaches provide increasingly rich structural supervision, but their evaluation often remains centered on tasks that only weakly isolate residue-level geometry. A useful structure-aware representation should not merely improve functional annotation: it should provide features that downstream structure models can directly exploit.

---

\*Equal contribution. Author order was determined randomly.  
Code is available at <https://holymollyhao.github.io/TriProRep/>.

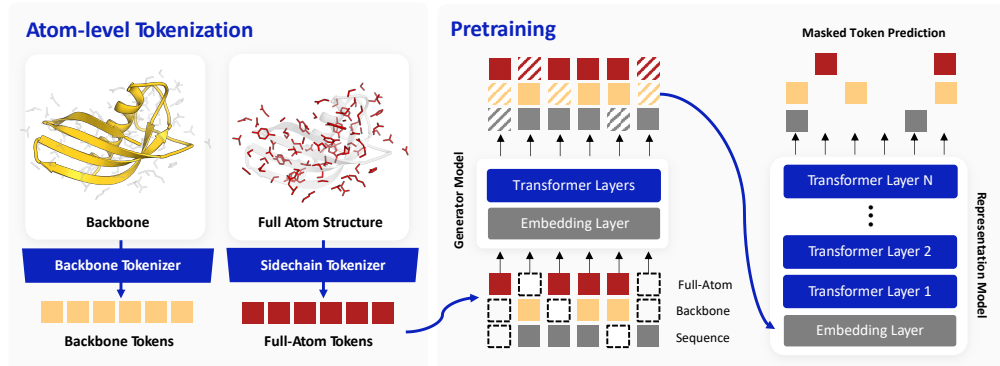


Figure 1: **TRIPROREP**. (a) Three-view tokenization. A protein is independently tokenized into amino-acid, backbone, and full-atom token sequences at the residue level. (b) ELECTRA-style discriminative pretraining. A small generator corrupts each of the three sequences, and a large discriminator predicts the original token at every position.

This perspective mirrors a broader trend in representation learning. In computer vision, pretrained representations are increasingly used not only as task features, but also to improve structured-output models through conditioning [8, 9, 10] and representation alignment-based distillation [11, 12]. We ask the analogous question for proteins:

**Question.** Can pretrained protein representations serve as useful geometric signals for structure-predictive modeling?

We study this question from both a methodological and evaluation perspective. Methodologically, we propose TRIPROREP, a structure-aware protein representation model that jointly encodes three aligned residue-level views: amino-acid identity, backbone geometry, and full-atom residue geometry. The two geometric views are discretized by VQ-VAE tokenizers [13] into backbone-geometry and full-atom-geometry tokens. Unlike prior tokenizations that mainly capture backbone-level geometry [4, 5, 14], our full-atom tokens preserve local atomic information, including side-chain arrangements that are important for residue environments and interfaces. To train these coupled views, we use corrective pretraining with generator-corrupted token sequences [15]: a lightweight generator replaces masked tokens in each view with plausible alternatives, and the representation model recovers the original tokens.

For evaluation, we introduce REPS, the first benchmark specifically designed for structure-predictive protein representation learning. Built on homodimer complexes from the recent AFDB multimer release [16], REPS provides a uniformly processed million-scale setting for controlled evaluation of binding-relevant geometry from apo-chain representations, while acknowledging that homodimers are simpler than general heteromeric complexes. REPS contains three complementary tasks: *homodimer co-folding* from frozen apo-chain representations, *homodimer-derived residue-level prediction* of binding sites, solvent-accessibility changes, interface regions, and interaction types, and *representation-aligned structure prediction*, where pretrained representations serve as dense alignment targets for structure prediction learning [11].

We pretrain and evaluate TRIPROREP at four scales, 35M, 150M, 650M, and 2.8B. On REPS, TRIPROREP improves over sequence-only and prior structure-aware representation models across homodimer co-folding, homodimer-derived residue-level prediction, and representation-aligned structure prediction. On the EC/GO benchmarks, TRIPROREP remains competitive with the strongest overall results, preserving broad biological representation quality.

Our contributions are summarized as follows:

- We propose TRIPROREP, a three-view structure-aware protein representation model trained with generator-corrupted token recovery.
- We introduce REPS, a benchmark that evaluates whether protein representations are useful for structure-predictive tasks.
- We show that TRIPROREP improves homodimer co-folding, homodimer-derived residue-level prediction, and representation-aligned structure prediction, while remaining competitive on EC/GO.

## 2 Related work

**Protein structure tokenization.** A line of work uses VQ-VAE-based tokenizers [13] to convert protein structures into discrete tokens for scalable structure-aware protein representation modeling. Foldseek defines 3Di tokens as a structural alphabet for residue-residue interaction geometry [4]. ESM3 also uses tokens that discretize local backbone geometry into reconstructable tokens [5], and AminoAseed improves tokenizer quality through codebook reparameterization and Pareto-optimal codebook configuration [14]. However, these tokenizations mainly focus on backbone-level geometry and omit full-atom geometry within residues, such as side-chain packing, and rotameric states.

**Structure-aware protein representation learning.** Recent protein representation models incorporate 3D structural information into protein language modeling. SaProt [2] pairs amino-acid tokens with Foldseek’s 3Di tokens [4], and ProstT5 [3] uses a T5 architecture to model bidirectional translation between sequence and 3Di tokens. ESM3 [5] introduces structure tokens alongside sequence and function tracks. In addition, MIF-ST [6] performs masked inverse folding with backbone-structure conditioning, while S-PLM [7] aligns sequence and structure representations through contrastive learning. These methods show that structural information can improve protein representations. We study whether such representations transfer to structure-predictive tasks.

**Protein structure and complex prediction.** Protein structure prediction has progressed from single-chain folding to protein-protein complex prediction. AlphaFold-Multimer [17], AlphaFold3 [18], and Boltz [19, 20] predict protein complexes using MSA-derived representations, while SimpleFold [21] revisits single-chain folding with flow matching and a general-purpose Transformer architecture using protein sequence representation. Our work uses structure prediction as an evaluation setting, rather than proposing a new structure prediction model. We evaluate whether structure-aware representations extracted from a single apo protein support flexible-docking prediction of the protein complex and serve as distillation targets for training single-protein folding models.

**Representation-based structure modeling.** In computer vision, structured-output models often use pretrained representations as conditioning signals [8, 9, 10] or distillation sources [11, 12]. Analogously, protein structure prediction models such as ESMFold [22] and SimpleFold [21] use sequence-only protein representations, particularly ESM2 representations [22]. In this paper, we study structure-aware protein representations in both roles: as input conditioning signals for flexible protein-protein docking, and as dense distillation sources for monomer structure prediction.

## 3 TRIPROREP: A three-view protein structure representation

In this section, we introduce TRIPROREP, a structure-aware protein representation model that jointly encodes three residue-level views of a protein: amino-acid identity, backbone geometry, and full-atom geometry. We first describe how these views are converted to discrete tokens (Section 3.1). We then describe a corrective pretraining objective over generator-corrupted token sequences (Section 3.2). The full pipeline is illustrated in Figure 1, with implementation details in Appendix A.

### 3.1 Three-view residue tokenization

Given a protein, we construct three aligned residue-level token sequences. For each residue, we define an amino-acid token, a backbone-geometry token, and a full-atom-geometry token. These tokens provide three complementary views of the same residue: sequence identity, local backbone geometry, and intra-residue full-atom geometry. The amino-acid token encodes residue identity, as in sequence-only protein language models such as ESM2 [22]. We pair this sequence-identity view with two structural views, each defined at the residue level.

**Backbone-geometry tokens.** The backbone-geometry token provides a residue-level structural view of the local protein backbone. We obtain these tokens using AminoAseed [14], a VQ-VAE-based tokenizer that maps local backbone substructures around each residue into discrete codes. Compared with Foldseek 3Di tokens used by a prior work [2] that also performs residue backbone tokenization, AminoAseed improves codebook utilization and token diversity, providing a more expressive backbone tokenization.

**Full-atom-geometry tokens.** We newly introduce the full-atom-geometry token to encode intra-residue full-atom geometry, complementing the backbone-geometry token with local atomic infor-

mation within each residue (as shown in Figure 5 of Appendix A.1). The full-atom-geometry token focuses on heavy-atom arrangements in a backbone-defined local frame, including side-chain rotamer geometry. This complementary tokenization provides information that is largely inaccessible from backbone geometry alone, allowing the model to incorporate residue-level atomic details that are useful for downstream protein-geometry understanding. We obtain this token by training a residue-level VQ-VAE tokenizer that takes heavy-atom coordinates expressed in an SE(3)-invariant local frame defined by the  $N$ ,  $C\alpha$ , and  $C$  atoms, together with dihedral-angle features, and assigns each residue a categorical code. We provide details in Appendix A.1.

### 3.2 Corrective pretraining with generator-corrupted views

We pretrain TRIPROREP using a corrective token-recovery objective over the three aligned token sequences. The objective is inspired by ELECTRA-style pretraining [15], but differs in that the representation model predicts the original token rather than a binary replacement label. In our three-view setting, generator-based replacement serves as token-level augmentation. Since each residue is represented by amino-acid, backbone-geometry, and full-atom-geometry tokens, independent replacement can produce inputs that are plausible within each view but mutually inconsistent across views. Recovering the original tokens encourages the representation model to model consistency among sequence identity, backbone geometry, and residue full-atom geometry.

At each step, we independently sample masked positions for each token type. We feed the partially masked amino-acid, backbone-geometry, and full-atom-geometry token sequences to a small generator. The generator is trained with masked-token cross-entropy to predict the original token distribution at each masked position. We then sample replacement tokens from the generator distributions and construct generator-corrupted token sequences.

The representation model takes the corrupted token sequences as input and predicts the original token at every position for all three token types, using separate token-recovery heads for the three views. We use the corrective objective [23], which recovers the original token rather than predicting a binary replaced-or-unchanged label. This provides dense supervision at every residue position and exposes the model to a large set of generator-produced cross-view corruptions. After pretraining, we discard both the generator and the token-recovery heads, and use the hidden states of the representation model as the transferable protein representation. We provide details in Appendix A.2.

**Datasets.** We pretrain TRIPROREP on 83.7M predicted protein structures from the AlphaFold Protein Structure Database (AFDB) [24] and ESMAtlas [22]. AFDB provides AlphaFold2-predicted structures over proteome-scale sequences in UniProt [25], and ESMAtlas provides ESMFold predictions over metagenomic sequences in MGnify [26] augmenting the AFDB structures. For AFDB, we use the representative protein from each UniRef30 cluster and retain structures with pLDDT at least 50, yielding 18.2M representative proteins. For ESMAtlas, we retain structures with pLDDT greater than 70, yielding 65.5M additional predicted structures.

## 4 A benchmark for structure-predictive protein representation learning

We introduce REPSP, which evaluates whether protein representations improve structure prediction across three tasks: (1) inferring homodimer complex structures, (2) predicting residue-level labels derived from homodimer structures, and (3) accelerating monomer protein structure prediction via distillation [11]. To the best of our knowledge, our work is the first to use structure-aware representations for flexible-docking prediction and distillation. The full pipeline is illustrated in Figure 2, with implementation details in Appendix B.

**Dataset curation and split.** This benchmark uses 1.8M homodimer complexes released in the March 2026 AlphaFold Protein Structure Database (AFDB) [16]. We focus on homodimers because they are currently the largest uniformly processed complex dataset in AFDB, providing a controlled setting for evaluation at scale: the binding partner is identical to the query chain, thereby removing variation along the partner-pairing axis while still requiring interface-relevant geometric information.<sup>2</sup>

To be specific, we pair each holo homodimer complex with its corresponding monomer protein prediction in the AFDB [24, 27, 28]. We then split the dataset by sequence clusters at 50% sequence

<sup>2</sup>We also evaluate representations in heterodimer settings, as shown in Table 5 of Section 5.4.

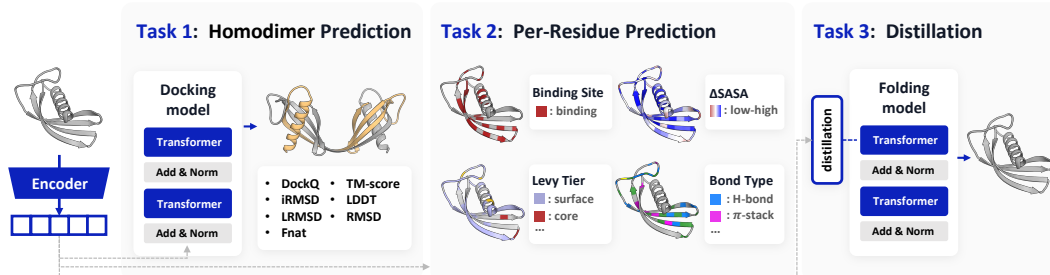


Figure 2: **REPSP**. We define three structure-generative tasks that use protein representations as input: **(task 1)** homodimer structure prediction, **(task 2)** per-residue homodimer binding-property prediction via MLP probing, and **(task 3)** distillation into a monomer structure prediction model.

identity. From the resulting cluster representatives, we select 400 validation and 1,000 test sequences, and use the remaining representatives for training. We provide details in Appendix B.1.

**Homodimer structure prediction (Section 5.1).** We evaluate whether protein representations are predictive of the corresponding homodimer structure. Strong performance indicates that the protein representation contains structural information useful for protein complex structure inference. Specifically, a 100M flexible-docking model takes monomer protein representations as input and learns to predict the homodimer structures. We modify SimpleFold [21], a Transformer-based single-protein folding model [22], into a homodimer flexible-docking model, with details provided in Appendix B.2. We evaluate the predicted structures using interface-quality metrics, including DockQ, iRMSD, LRMSD, and Fnat, and overall quality metrics, including TM-score, LDDT, and RMSD.

**Per-residue homodimer binding property prediction (Section 5.2).** We next evaluate whether monomer protein representations are predictive of residue-level binding properties of the corresponding homodimer complex. Strong probing performance indicates that the representation captures binding-relevant local signals, including binding residues, exposure changes, interface regions, and interaction types. Given residue-wise monomer representations, we use MLP probing with four per-residue targets: binary binding site, continuous change in solvent-accessible surface area between the monomer and homodimer structures ( $\Delta$ SASA) [29, 30], Levy tier with five classes, and multi-label bond type over five interaction classes [31, 32]. We describe details in Appendix B.3.

**Distillation into monomer structure prediction (Section 5.3).** We evaluate whether protein representations provide useful distillation targets for monomer structure prediction. This experiment is motivated by recent work showing that high-quality representations can improve generative model training through distillation [11]. Following this, we distill monomer protein representations into the SimpleFold-100M model [21] by maximizing the cosine similarity between monomer protein representations and the model’s hidden states for structure prediction. Implementation details are provided in Appendix B.4. We assess representation quality by the resulting improvement in structure prediction, using TM-score, GDT-TS, GDT-HA, LDDT, backbone-based LDDT<sub>bb</sub>, and RMSD.

## 5 Results

We pretrain TRIPROREP at four model scales: 35M, 150M, 650M, and 2.8B. Then, we compare against sequence-only ESM2 [22], structure-aware representation models, including SaProt [2], S-PLM [7], MIF-ST [6], ESM3 [5], and ProstT5 [3]. We report on REPSP, a new benchmark that evaluates the structure-generative capacity of protein representations (Sections 5.1 to 5.3), and on conventional protein representation benchmarks (Section 5.4).

### 5.1 Homodimer structure prediction

Table 1 reports homodimer flexible-docking performance on REPSP, with qualitative examples shown in Figure 3. Across scales from 35M to 2.8B parameters, TRIPROREP consistently improves both interface quality and overall structural accuracy, indicating favorable scaling with model size. Notably, the 650M TRIPROREP model already outperforms the huge baselines on all reported metrics,

Table 1: **Homodimer flexible-docking performance on REPS**. Monomer protein representations predict homodimer structures. **Bold** denotes the best within each size class. Overall, flexible-docking with representations from TRIPROREP outperforms the baselines across all parameter scales.

Models	Params	Interface quality				Overall quality		
		DockQ $\uparrow$	iRMSD $\downarrow$	LRMSD $\downarrow$	Fnat $\uparrow$	TM-score $\uparrow$	LDDT $\uparrow$	RMSD $\downarrow$
<i>Small models</i>								
ESM2	35M	0.278	11.549	25.942	0.417	0.531	0.603	16.176
SaProt	35M	0.296	10.380	25.308	0.414	0.567	0.712	14.660
<b>Ours</b>	35M	<b>0.371</b>	<b>8.915</b>	<b>21.964</b>	<b>0.486</b>	<b>0.633</b>	<b>0.771</b>	<b>12.778</b>
<i>Medium models</i>								
ESM2	150M	0.310	10.743	25.328	0.450	0.563	0.667	15.170
<b>Ours</b>	150M	<b>0.419</b>	<b>8.170</b>	<b>20.292</b>	<b>0.512</b>	<b>0.666</b>	<b>0.812</b>	<b>11.874</b>
<i>Large models</i>								
ESM2	650M	0.374	9.517	22.803	0.533	0.613	0.721	13.913
SaProt	650M	0.443	7.449	18.878	0.548	0.690	0.807	11.196
S-PLM	704M	0.366	9.350	22.678	0.507	0.619	0.708	13.591
MIF-ST	643M	0.299	10.598	25.420	0.384	0.563	0.818	14.521
<b>Ours</b>	650M	<b>0.477</b>	<b>6.883</b>	<b>18.243</b>	<b>0.588</b>	<b>0.705</b>	<b>0.829</b>	<b>10.554</b>
<i>Huge models</i>								
ESM2	3B	0.387	8.679	21.455	0.547	0.635	0.732	13.025
ESM3	1.4B	0.448	7.060	18.554	0.579	0.695	<b>0.835</b>	10.719
ProstT5	1.2B	0.409	7.790	20.403	0.541	0.661	0.782	11.843
<b>Ours</b>	2.8B	<b>0.499</b>	<b>6.370</b>	<b>17.087</b>	<b>0.612</b>	<b>0.724</b>	0.834	<b>9.875</b>

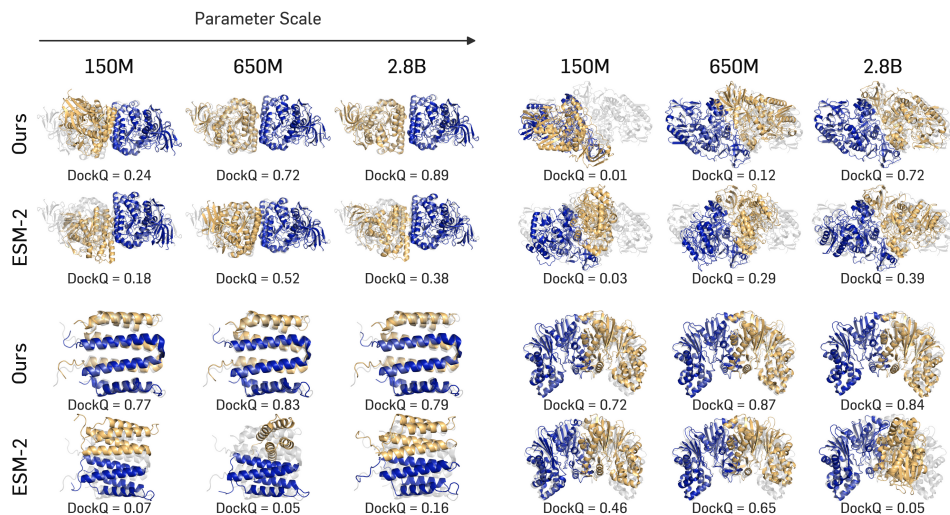


Figure 3: **Scaling of flexible-docking**. Predicted homodimer structures (chain A blue, chain B gold) overlaid on ground truth (gray) across encoder sizes (150M, 650M, 3B) for four test records.

while the huge TRIPROREP model achieves the strongest performance on nearly all metrics, with ESM3 only marginally higher in LDDT.

The gains are most pronounced on interface-level metrics, which depend not only on accurate monomer structures but also on their relative placement and contact geometry. The qualitative examples in Figure 3 show the same trend: larger TRIPROREP models reliably recover the correct monomer orientation, whereas baseline predictions often misplace the interface. These results suggest that TRIPROREP provides representations that better support complex-level geometric inference.

Our REPS benchmark also highlights the importance of structural information beyond parameter count. Structure-aware models, including SaProt-650M, ESM3-1.4B, and ProstT5-1.2B, outperform the larger sequence-only ESM2-3B on most metrics. This pattern suggests that homodimer flexible docking is a structure-sensitive task, making REPS a useful evaluation setting for structure-dependent representation quality.

Table 2: **Per-residue homodimer binding properties probing on REPSP**. MLP probes on monomer protein representations predict properties on the homodimer. **Bold** indicates the best within each size class. Results are averaged over three runs. Our method shows strongest performance.

Models	Params	Binding site		$\Delta$ SASA		Levy tier		Bond type	
		AUPRC	AUROC	Pearson	Spearman	Macro F1	Acc.	AUPRC	AUROC
<i>Small</i>									
ESM2	35M	0.6635	0.8862	0.5955	0.3445	0.5386	0.7706	0.2936	0.9338
SaProt	35M	0.6291	0.8794	0.6218	0.3923	0.5248	0.8036	0.2663	0.9388
<b>Ours</b>	35M	<b>0.7614</b>	<b>0.9243</b>	<b>0.7047</b>	<b>0.4167</b>	<b>0.6260</b>	<b>0.8342</b>	<b>0.3566</b>	<b>0.9518</b>
<i>Medium</i>									
ESM2	150M	0.7087	0.9062	0.6379	0.3703	0.5665	0.7969	0.3227	0.9425
<b>Ours</b>	150M	<b>0.8161</b>	<b>0.9433</b>	<b>0.7801</b>	<b>0.4671</b>	<b>0.6843</b>	<b>0.8668</b>	<b>0.4204</b>	<b>0.9642</b>
<i>Large</i>									
ESM2	650M	0.7915	0.9369	0.7146	0.4113	0.6442	0.8281	0.3902	0.9566
SaProt	650M	0.8273	0.9508	0.7881	0.4715	0.6946	0.8693	0.4305	<b>0.9692</b>
S-PLM	704M	0.7789	0.9324	0.6967	0.4036	0.6215	0.8166	0.3602	0.9523
MIF-ST	643M	0.5846	0.8671	0.6726	0.4338	0.5261	0.8435	0.2623	0.9430
<b>Ours</b>	650M	<b>0.8524</b>	<b>0.9561</b>	<b>0.8088</b>	<b>0.4813</b>	<b>0.7290</b>	<b>0.8780</b>	<b>0.4425</b>	0.9673
<i>Huge</i>									
ESM2	3B	0.8247	0.9483	0.7446	0.4281	0.6851	0.8426	0.4106	0.9614
ESM3	1.4B	0.8470	0.9558	0.8121	<b>0.4878</b>	0.7285	<b>0.8898</b>	0.4229	0.9673
ProstT5	1.2B	0.8291	0.9480	0.7574	0.4449	0.6848	0.8538	0.4077	0.9594
<b>Ours</b>	2.8B	<b>0.8650</b>	<b>0.9608</b>	<b>0.8152</b>	0.4854	<b>0.7408</b>	0.8811	<b>0.4516</b>	<b>0.9691</b>




Figure 4: **Acceleration of monomer structure prediction via representation alignment on REPSP**. We compare the no-REPA baseline against ESM2, SaProt, S-PLM, MIF-ST, and TRIPROREP as the alignment target. TRIPROREP provides strongest alignment target for structure prediction model.

## 5.2 Per-residue homodimer binding property prediction

Table 2 reports per-residue probing results on REPSP. Across all model scales, TRIPROREP improves over the baselines on most tasks and metrics, with the strongest overall performance attained at the huge scale. The consistency of these gains indicates that the benefits of the proposed representation are not confined to a particular model size, but scale reliably across the model family.

These probing results offer a residue-level explanation for the co-folding improvements in Section 5.1. Since the probes are trained on frozen monomer representations, performance gains reflect that information is already encoded in the representation rather than capacity added by downstream fine-tuning. The consistent improvements across all four prediction targets therefore suggest that TRIPROREP captures homodimer binding signals at residue-level resolution, providing mechanistic evidence for the complex-level gains observed in Table 1.

Table 3: **Monomer structure prediction with distillation on REPSP.** We report test performance of a monomer structure prediction model. **Bold** denotes the best performance. TRIPROREP provides the strongest alignment target among the evaluated large-scale representation models.

REPA target	Params	TM-score $\uparrow$	GDT-TS $\uparrow$	GDT-HA $\uparrow$	LDDT $\uparrow$	LDDT <sub>bb</sub> $\uparrow$	RMSD ( $\text{\AA}$ ) $\downarrow$
Baseline	–	0.5422	0.3432	0.1990	0.6208	0.6746	12.98
SaProt	650M	0.5659	0.3677	0.2192	0.6410	0.6947	12.26
S-PLM	704M	0.5569	0.3559	0.2081	0.6301	0.6838	12.40
MIF-ST	643M	0.5491	0.3520	0.2075	0.6310	0.6827	12.73
<b>Ours</b>	650M	<b>0.5822</b>	<b>0.3858</b>	<b>0.2344</b>	<b>0.6513</b>	<b>0.7051</b>	<b>11.91</b>

Table 4: **Ablations on per-residue homodimer binding properties probing.** MLP probes on monomer protein representations predict properties derived from the holodimer. Ours<sub>AF</sub> and Ours<sub>MLM</sub> denote variants trained only on AlphaFold structures and with masked language modeling, respectively. Overall, our full model yields strongest performance compared to variants.

Models	Params	Binding site		$\Delta$ SASA		Levy tier		Bond type	
		AUPRC	AUROC	Pearson	Spearman	Macro F1	Acc.	AUPRC	AUROC
ESM2	150M	0.7104	0.9066	0.6376	0.3701	0.5701	0.7972	0.3534	0.9542
Ours <sub>AF</sub>	150M	0.7759	0.9314	0.7285	0.4344	0.6458	0.8470	0.3818	0.9576
Ours <sub>MLM</sub>	150M	0.7677	0.9297	0.7309	0.4380	0.6424	0.8471	0.3748	0.9588
<b>Ours</b>	150M	<b>0.8156</b>	<b>0.9432</b>	<b>0.7797</b>	<b>0.4662</b>	<b>0.6819</b>	<b>0.8668</b>	<b>0.4184</b>	<b>0.9644</b>

### 5.3 Distillation into monomer structure prediction

Unlike the co-folding and probing evaluations, this experiment uses pretrained representations not as frozen input features, but as supervisory targets for the structure prediction model. It therefore tests a complementary property of a representation: whether its structural signal is organized in a form that can guide the learning of a generative structure predictor.

Figure 4 and Tab. 3 report monomer structure prediction results under this representation-alignment setting. TRIPROREP achieves the best performance across all reported metrics, indicating that it provides the most effective supervisory signal for structure prediction training. This suggests that the structural information captured by TRIPROREP is not only decodable by downstream models, but also transferable as a training target for structure-generative learning.

We further evaluate huge-scale baselines in Table 10 of Appendix C.2. At this scale, performance saturates, with ProstT5 and TRIPROREP achieving comparable results. Although ProstT5 is pretrained for sequence-to-structure translation, TRIPROREP matches its performance without this objective, while also performing well in flexible docking-related evaluations where ProstT5 gives only limited gains. These results indicate that TRIPROREP learns rich representations that support flexible docking and serve as effective alignment targets for monomer structure prediction.

### 5.4 Ablation studies

#### Ablation on pretraining data and objective.

Table 4 ablates two design choices in TRIPROREP at the 150M scale: pretraining data and pretraining objective. First, Ours<sub>AF</sub> is trained only on AlphaFold structures, removing ESMAtlas from the pretraining corpus. Next, Ours<sub>MLM</sub> replaces the generator-corrupted corrective objective with standard masked language modeling. Both variants outperform ESM2-150M, showing that three-view structure-aware pretraining already provides a strong representation, but the full model achieves the best performance. This indicates that both the augmented data and the generator-based corruption improve representation training.

Table 5: **Binding-site probing for binder screening.** Probes to predict binding site and contact count, conditioned on binder representation. TRIPROREP shows strongest performance.

Model	Params	Binding site		Contact count	
		AUPRC	AUROC	Pearson	Spearman
ESM2	3B	0.7646	0.7743	0.4162	0.3987
SaProt	650M	0.7978	0.8112	0.4996	0.4757
Ours	650M	0.8184	0.8271	0.4830	0.4700
S-PLM	704M	0.7541	0.7640	0.4006	0.3903
ProstT5	1.2B	0.7608	0.7688	0.4392	0.4244
ESM3	1.4B	0.7472	0.7719	0.0543	0.0551
<b>Ours</b>	2.8B	<b>0.8332</b>	<b>0.8431</b>	<b>0.5169</b>	<b>0.5020</b>

Table 7: **Homodimer structure prediction on RCSB structures deposited after June 1, 2023.** We evaluate the co-folding models in Table 1 on real-world homodimers. **Bold** denotes the best performance. The relative gains of TRIPROREP are preserved in predicting real-world structures.

Models	Params	Max seq. id. < 90% (N=252)			Max seq. id. < 40% (N=86)		
		DockQ $\uparrow$	Medium% $\uparrow$	High% $\uparrow$	DockQ $\uparrow$	Medium% $\uparrow$	High% $\uparrow$
S-PLM	704M	0.275	29.8	4.0	0.185	17.4	3.5
ProstT5	1.2B	0.297	33.3	5.6	0.193	18.6	2.3
ESM2	3B	0.296	32.9	6.0	0.197	18.6	2.3
SaProt	650M	0.320	36.5	7.1	0.210	19.8	4.7
ESM3	1.4B	0.322	37.7	8.7	0.195	19.8	3.5
<b>Ours</b>	2.8B	<b>0.342</b>	<b>39.7</b>	<b>13.5</b>	<b>0.228</b>	<b>20.9</b>	<b>8.1</b>

**Target-conditioned binder-interface probing.** Table 5 evaluates whether protein representations identify target residues involved in heteromeric binding. This setting differs from homodimer probing because the interface is formed with a non-identical binder, and therefore tests whether the representation transfers to asymmetric cross-chain interactions. TRIPROREP shows competitive or stronger performance on both binding-site classification and contact-count regression. These results suggest that TRIPROREP encodes residue-level interface information that transfers beyond homodimeric interactions, supporting its use for target-conditioned binder screening. See the full results and details in Table 11 of Appendix C.3.

**Conventional benchmarks.** We also evaluate TRIPROREP on conventional benchmarks, namely Enzyme Commission and Gene Ontology [33]. We use a two-layer MLP probing setup. We describe the detailed hyperparameters in Appendix C.4. As shown in Table 6, TRIPROREP remains competitive with the strongest baselines on these benchmarks, while sequence-only ESM2-3B performs close to structure-aware models.

**Performance on real-world protein structures.** We further evaluate whether the co-folding results transfer beyond the AFDB-derived benchmark. Table 7 reports homodimer structure prediction performance on RCSB structures deposited after June 1, 2023 [34]. Since our pretraining and co-folding training use predicted AFDB structures rather than experimentally resolved structures, we mainly use loss-based filtering and additionally report results under the conventional cutoff. The relative trend observed on REPSP is preserved on these real-world structures: TRIPROREP consistently outperforms the baselines in average DockQ score and in medium- and high-quality success rates (DockQ $\geq$  0.49 and DockQ $\geq$  0.8).

Table 6: **Per-protein function probing on EC and GO.** MLP probing results for EC and GO benchmarks on mean-pooled representations. TRIPROREP shows competitive performance.

Models	Params	EC	GO		
			MF	BP	CC
<i>Small models</i>					
ESM2	35M	0.888	0.618	0.456	0.521
SaProt	35M	0.882	0.614	0.449	<b>0.539</b>
<b>Ours</b>	35M	<b>0.914</b>	<b>0.640</b>	<b>0.474</b>	0.504
<i>Medium models</i>					
ESM2	150M	0.900	0.633	0.463	<b>0.529</b>
<b>Ours</b>	150M	<b>0.906</b>	<b>0.642</b>	<b>0.467</b>	0.528
<i>Large models</i>					
ESM2	650M	0.906	0.646	0.481	0.533
SaProt	650M	0.897	0.637	0.472	<b>0.555</b>
S-PLM	704M	<b>0.915</b>	<b>0.660</b>	0.480	0.531
<b>Ours</b>	650M	<b>0.915</b>	0.655	<b>0.488</b>	0.511
<i>Huge models</i>					
ESM2	3B	0.920	0.654	<b>0.490</b>	<b>0.532</b>
ESM3	1.4B	<b>0.926</b>	0.648	0.487	0.516
<b>Ours</b>	2.8B	0.923	<b>0.660</b>	<b>0.490</b>	0.508

## 6 Conclusion

In this paper, we presented TRIPROREP, a structure-aware representation model that incorporates amino-acid identity, backbone geometry, and full-atom residue geometry, together with REPSP, a benchmark for evaluating whether pretrained protein representations provide useful geometric signals for structure-predictive modeling. On REPSP, structure-aware representations outperform sequence-only protein representations, and TRIPROREP further improves over prior structure-aware baselines. These results highlight the utility of structure-aware protein representations for structure-predictive modeling, with finer geometric detail further improving their utility in these tasks.

**Limitations.** Following prior structure-aware representation learning studies that use predicted structures [2, 3, 7], TRIPROREP and REPSP use high-confidence AFDB structures for million-scale, uniformly processed benchmarking, while broader validation on experimentally resolved structures remains future work despite similar initial trends on real PDB complexes (Table 7).

## References

- [1] Vladimir Gligorijević, P. Douglas Renfrew, Tomasz Kosciolatek, 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, 2021.
- [2] Jin Su, Chenchen Han, Yuyang Zhou, Junjie Shan, Xibin Zhou, and Fajie Yuan. Saprot: Protein language modeling with structure-aware vocabulary. In *The Twelfth International Conference on Learning Representations*, 2024.
- [3] Michael Heinzinger, Konstantin Weissenow, Joaquin Gomez Sanchez, Adrian Henkel, Milot Mirdita, Martin Steinegger, and Burkhard Rost. Bilingual language model for protein sequence and structure. *NAR Genomics and Bioinformatics*, 6(4):lqae150, 12 2024.
- [4] Michel Van Kempen, Stephanie S Kim, Charlotte Tumescheit, Milot Mirdita, Jeongjae Lee, Cameron LM Gilchrist, Johannes Söding, and Martin Steinegger. Fast and accurate protein structure search with foldseek. *Nature biotechnology*, 42(2):243–246, 2024.
- [5] Thomas Hayes, Roshan Rao, Halil Akin, Nicholas J Sofroniew, Deniz Oktay, Zeming Lin, Robert Verkuil, Vincent Q Tran, Jonathan Deaton, Marius Wiggert, et al. Simulating 500 million years of evolution with a language model. *Science*, 387(6736):850–858, 2025.
- [6] Kevin K Yang, Niccolò Zanichelli, and Hugh Yeh. Masked inverse folding with sequence transfer for protein representation learning. *Protein Engineering, Design and Selection*, 36:gzad015, 2023.
- [7] Duolin Wang, Mahdi Pourmirzaei, Usman L Abbas, Shuai Zeng, Negin Manshour, Farzaneh Esmaili, Biplab Poudel, Yuexu Jiang, Qing Shao, Jin Chen, et al. S-plm: structure-aware protein language model via contrastive learning between sequence and structure. *Advanced Science*, 12(5):2404212, 2025.
- [8] Tianhong Li, Dina Katabi, and Kaiming He. Return of unconditional generation: A self-supervised representation generation method. *Advances in Neural Information Processing Systems*, 37:125441–125468, 2024.
- [9] Hu Ye, Jun Zhang, Sibol Liu, Xiao Han, and Wei Yang. Ip-adapter: Text compatible image prompt adapter for text-to-image diffusion models. 2023.
- [10] Nicolas Sereyjol-Garros, Ellington Kirby, Victor Letzelter, Victor Besnier, and Nermin Samet. Test-time conditioning with representation-aligned visual features. *arXiv preprint arXiv:2602.03753*, 2026.
- [11] Sihyun Yu, Sangkyung Kwak, Huiwon Jang, Jongheon Jeong, Jonathan Huang, Jinwoo Shin, and Saining Xie. Representation alignment for generation: Training diffusion transformers is easier than you think. In *The Thirteenth International Conference on Learning Representations*, 2025.
- [12] Xingjian Leng, Jaskirat Singh, Yunzhong Hou, Zhenchang Xing, Saining Xie, and Liang Zheng. Repa-e: Unlocking vae for end-to-end tuning of latent diffusion transformers. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 18262–18272, 2025.
- [13] Aaron Van Den Oord, Oriol Vinyals, et al. Neural discrete representation learning. *Advances in neural information processing systems*, 30, 2017.
- [14] Xinyu Yuan, Zichen Wang, Marcus D Collins, and Huzefa Rangwala. Protein structure tokenization: Benchmarking and new recipe. In *International Conference on Machine Learning*, pages 73645–73670. PMLR, 2025.
- [15] Kevin Clark, Minh-Thang Luong, Quoc V. Le, and Christopher D. Manning. ELECTRA: Pre-training text encoders as discriminators rather than generators. In *ICLR*, 2020.

- [16] Yewon Han, Maxim I Tsenkov, Niccolo AE Venanzi, Damian Bertoni, Sooyoung Cha, Alejandro Chacon, Nick Dietrich, Boris Fomitchev, Yonathan Goldtzvik, Darren Hsu, et al. Alphafold database expands to proteome-scale quaternary structures. *bioRxiv*, pages 2026–03, 2026.
- [17] Richard Evans, Michael O’neill, Alexander Pritzel, Natasha Antropova, Andrew Senior, Tim Green, Augustin Žídek, Russ Bates, Sam Blackwell, Jason Yim, et al. Protein complex prediction with alphafold-multimer. *bioRxiv*, pages 2021–10, 2021.
- [18] 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*, 630(8016):493–500, 2024.
- [19] Jeremy Wohlwend, Gabriele Corso, Saro Passaro, Noah Getz, Mateo Reveiz, Ken Leidal, Wojtek Swiderski, Liam Atkinson, Tally Portnoi, Itamar Chinn, et al. Boltz-1 democratizing biomolecular interaction modeling. *BioRxiv*, pages 2024–11, 2025.
- [20] Saro Passaro, Gabriele Corso, Jeremy Wohlwend, Mateo Reveiz, Stephan Thaler, Vignesh Ram Somnath, Noah Getz, Tally Portnoi, Julien Roy, Hannes Stark, et al. Boltz-2: Towards accurate and efficient binding affinity prediction. *BioRxiv*, 2025.
- [21] Yuyang Wang, Jiarui Lu, Navdeep Jaitly, Joshua M. Susskind, and Miguel Ángel Bautista. Simplefold: Folding proteins is simpler than you think. In *The Fourteenth International Conference on Learning Representations*, 2026.
- [22] 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.
- [23] Zhenhui Xu, Linyuan Gong, Guolin Ke, Di He, Shuxin Zheng, Liwei Wang, Jiang Bian, and Tie-Yan Liu. Mc-bert: Efficient language pre-training via a meta controller. *arXiv preprint arXiv:2006.05744*, 2020.
- [24] Mihaly Varadi, Stephen Anyango, Mandar Deshpande, Sreenath Nair, Cindy Natassia, Galabina Yordanova, David Yuan, Oana Stroe, Gemma Wood, Agata Laydon, et al. Alphafold protein structure database: massively expanding the structural coverage of protein-sequence space with high-accuracy models. *Nucleic acids research*, 50(D1):D439–D444, 2022.
- [25] The UniProt Consortium. Uniprot: the universal protein knowledgebase in 2023, 2023.
- [26] Alex L Mitchell, Alexandre Almeida, Martin Beracochea, Miguel Boland, Josephine Burgin, Guy Cochrane, Michael R Crusoe, Varsha Kale, Simon C Potter, Lorna J Richardson, et al. Mgnify: the microbiome analysis resource in 2020. *Nucleic acids research*, 48(D1):D570–D578, 2020.
- [27] Mihaly Varadi, Damian Bertoni, Paulyna Magana, Urmila Paramval, Ivanna Pidruchna, Malarvizhi Radhakrishnan, Maxim Tsenkov, Sreenath Nair, Milot Mirdita, Jingsi Yeo, Oleg Kovalevskiy, Kathryn Tunyasuvunakool, Agata Laydon, Augustin Žídek, Hamish Tomlinson, Dhavanthi Hariharan, Josh Abrahamson, Tim Green, John Jumper, Ewan Birney, Martin Steinegger, Demis Hassabis, and Sameer Velankar. Alphafold protein structure database in 2024: providing structure coverage for over 214 million protein sequences. *Nucleic Acids Research*, 52(D1):D368–D375, 01 2024.
- [28] John Jumper, Richard Evans, Alexander Pritzel, Tim Green, Michael Figurnov, Olaf Ronneberger, Kathryn Tunyasuvunakool, Russ Bates, Augustin Žídek, Anna Potapenko, et al. Highly accurate protein structure prediction with alphafold. *nature*, 596(7873):583–589, 2021.
- [29] Byungkook Lee and Frederic M Richards. The interpretation of protein structures: estimation of static accessibility. *Journal of molecular biology*, 55(3):379–IN4, 1971.
- [30] Andrew Shrake and John A Rupley. Environment and exposure to solvent of protein atoms. lysozyme and insulin. *Journal of molecular biology*, 79(2):351–371, 1973.

- [31] Sebastian Salentin, Sven Schreiber, V Joachim Haupt, Melissa F Adasme, and Michael Schroeder. Plip: fully automated protein–ligand interaction profiler. *Nucleic acids research*, 43(W1):W443–W447, 2015.
- [32] Harry C Jubb, Alicia P Higuero, Bernardo Ochoa-Montaño, Will R Pitt, David B Ascher, and Tom L Blundell. Arpeggio: a web server for calculating and visualising interatomic interactions in protein structures. *Journal of molecular biology*, 429(3):365–371, 2017.
- [33] Vladimir Gligorijević, P Douglas Renfrew, Tomasz Kosciolk, Julia Koehler Leman, Daniel Berenberg, Tommi Vatanen, Chris Chandler, Bryn C Taylor, Ian M Fisk, Hera Vlamakis, et al. Structure-based protein function prediction using graph convolutional networks. *Nature communications*, 12(1):3168, 2021.
- [34] Stephen K Burley, Charmi Bhikadiya, Chunxiao Bi, Sebastian Bittrich, Li Chen, Gregg V Crichlow, Cole H Christie, Kenneth Dalenberg, Luigi Di Costanzo, Jose M Duarte, et al. Rcsb protein data bank: powerful new tools for exploring 3d structures of biological macromolecules for basic and applied research and education in fundamental biology, biomedicine, biotechnology, bioengineering and energy sciences. *Nucleic acids research*, 49(D1):D437–D451, 2021.
- [35] Martin Steinegger and Johannes Söding. Mmseqs2 enables sensitive protein sequence searching for the analysis of massive data sets. *Nature biotechnology*, 35(11):1026–1028, 2017.
- [36] William Peebles and Saining Xie. Scalable diffusion models with transformers. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 4195–4205, 2023.
- [37] Peter J. A. Cock, Tiago Antao, Jeffrey T. Chang, Brad A. Chapman, Cymon J. Cox, Andrew Dalke, Iddo Friedberg, Thomas Hamelryck, Frank Kauff, Bartek Wilczynski, and Michiel J. L. de Hoon. Biopython: freely available python tools for computational molecular biology and bioinformatics. *Bioinformatics*, 25(11):1422–1423, 2009.
- [38] Emmanuel D Levy. A simple definition of structural regions in proteins and its use in analyzing interface evolution. *Journal of molecular biology*, 403(4):660–670, 2010.
- [39] Anton Bushuiev, Roman Bushuiev, Petr Kouba, Anatolii Filkin, Marketa Gabrielova, Michal Gabriel, Jiri Sedlar, Tomas Pluskal, Jiri Damborsky, Stanislav Mazurenko, and Josef Sivic. Learning to design protein-protein interactions with enhanced generalization. In *The Twelfth International Conference on Learning Representations*, 2024.
- [40] Wolfgang Kabsch and Christian Sander. Dictionary of protein secondary structure: pattern recognition of hydrogen-bonded and geometrical features. *Biopolymers*, 22(12):2577–2637, 1983.

## A Details in TRIPROREP

### A.1 Full-atom tokenization

For each residue, the tokenizer computes heavy-atom geometry features from Atom37 coordinates. These features include (i) atomic displacements relative to  $C_\alpha$  in a per-residue canonical frame, which gives SE(3) invariance, (ii) distances from  $C_\alpha$  to each heavy atom, (iii) binned backbone bond angles, and (iv) binned side-chain torsions  $\chi_1, \dots, \chi_4$ . These features initialize a per-atom single representation and a  $37 \times 37$  inter-atomic pair representation. The two representations are jointly refined by  $N$  stacked Pairformer-style layers. We pool the resulting 37 atom embeddings within each residue to obtain a continuous embedding  $z$ , and assign  $z$  to its nearest codebook entry  $q_k$  from a learned codebook of size  $V$ .

**Training objective.** We pretrain the tokenizer on AlphaFoldDB monomer structures with

$$\mathcal{L} = \mathcal{L}_{\text{VQ-VAE}} + \mathcal{L}_{\text{pair}} + \mathcal{L}_\chi, \quad (1)$$

where  $\mathcal{L}_{\text{VQ-VAE}}$  is the standard VQ-VAE objective,  $\mathcal{L}_{\text{pair}}$  is a clamped pairwise inter-atomic distance loss, and  $\mathcal{L}_\chi$  is a per-bin cross-entropy loss for side-chain torsion prediction.

**Hyperparameters.** The tokenizer uses single width 256, pair width 128, and  $N = 6$  Pairformer-style layers with 8 attention heads. The output embedding dimension is 256. The codebook contains  $V = 512$  entries, uses EMA updates with decay 0.99, and uses entropy regularization with weight 0.1. Backbone angles and side-chain torsions are discretized into 21 bins over  $[-\pi, \pi]$ . The decoder is a 4-layer Transformer with hidden size 256 and 8 attention heads. All loss terms have weight 1.0. We train with AdamW using learning rate  $5 \times 10^{-4}$ ,  $\beta = (0.9, 0.99)$ , and weight decay 0.01.

**Ablations.** We compare whether the backbone and full-atom tokenizers capture side-chain information by measuring their sensitivity to side-chain rotamers (Figure 5). For each token, we compute its distribution over the three canonical  $\chi_1$  rotamers: **Trans** ( $\approx 180^\circ$ ), **Gauche<sup>+</sup>** ( $\approx +60^\circ$ ), and **Gauche<sup>-</sup>** ( $\approx -60^\circ$ ). We then plot each token in a triangle whose vertices correspond to the three rotamer states. Tokens near a vertex are associated with a specific rotamer, while tokens near the center mix multiple rotamers.

Backbone tokens are broadly spread near the center, indicating that they do not separate side-chain rotamer states well. In contrast, full-atom tokens are concentrated near the vertices, showing that they capture side-chain rotamer geometry. This result confirms that full-atom tokens provide side-chain information that is largely absent from backbone tokens, making the two tokenizers complementary.

### A.2 Corrective pretraining with generator-corrupted views

Let  $\mathbf{x} = (\mathbf{x}^{\text{aa}}, \mathbf{x}^{\text{bb}}, \mathbf{x}^{\text{fa}})$  denote the token sequences for a protein of length  $L$ , where  $\mathbf{x}^{\text{aa}}$ ,  $\mathbf{x}^{\text{bb}}$ , and  $\mathbf{x}^{\text{fa}}$  are the amino-acid, backbone, and full-atom token sequences, respectively. We pretrain TRIPROREP with a corrective token-recovery objective. The objective follows the generator-discriminator structure of ELECTRA [15]. A lightweight generator  $\mathcal{G}$  replaces selected tokens with plausible alternatives, and the representation model  $\mathcal{E}$  recovers the original tokens from the corrupted input.

**Generator-based tri-token augmentation.** We use generator-based replacement as an augmentation scheme for tri-token protein pretraining. Instead of masking tokens and asking the model to reconstruct them from explicit mask symbols, we first use  $\mathcal{G}$  to replace selected tokens with plausible in-distribution alternatives. We sample replacement positions independently for the three views.

Let  $\mathbf{m}^v \subset \{1, \dots, L\}$  denote the positions selected for replacement. We first form a masked input  $\mathbf{x}^{\text{masked}}$  by replacing  $x_i^v$  with [MASK] for all  $i \in \mathbf{m}^v$ . The generator predicts a token distribution




Figure 5: **Tokens vs. sidechain rotamer.** Density of codes in the  $\chi_1$  simplex. (a) Backbone tokens. (b) Full-atom tokens.

Table 8: ELECTRA model configuration across the four scales.

	35M	150M	650M	2.8B
embedding dim	480	640	1280	2560
encoder depth	10	28	30	33
decoder depth	2	2	3	3
attention heads	20	20	20	40

$p_G^v(\cdot | \mathbf{x}^{\text{masked}})_i$  at each selected residue  $i$  and samples a replacement token:

$$\hat{x}_i^v \sim p_G^v(\cdot | \mathbf{x}^{\text{masked}})_i, \quad i \in \mathbf{m}^v. \quad (2)$$

We then replace the original tokens at the selected positions with the sampled tokens to obtain the corrupted tri-token input  $\mathbf{x}^{\text{corrupt}}$ .

**Training objective.** The generator is trained with masked-token cross-entropy:

$$\mathcal{L}_{\text{MLM}} = \mathbb{E} \left[ \sum_v \sum_{i \in \mathbf{m}^v} \mathcal{L}_{\text{CE}}(p_G^v(\mathbf{x}^{\text{masked}})_i, x_i^v) \right], \quad (3)$$

where  $\mathcal{L}_{\text{CE}}$  denotes the cross-entropy loss over the predicted token distribution.

The representation model is trained to recover the original tokens from the corrupted tri-token input:

$$\mathcal{L}_{\text{rec}} = \mathbb{E} \left[ \sum_v \sum_{i=1}^L \mathcal{L}_{\text{CE}}(p_E^v(\mathbf{x}^{\text{corrupt}})_i, x_i^v) \right], \quad (4)$$

where  $p_E^v(\cdot | \mathbf{x}^{\text{corrupt}})_i$  is its predicted token distribution at residue  $i$  and view  $v$ . Unlike the generator loss, the recovery loss is applied to all residue positions and all views.

**Architecture.** The generator  $\mathcal{G}$  and the representation model  $\mathcal{E}$  are Transformer-based sequence models over residue-level tri-token inputs. For each residue  $i$ , the amino-acid, backbone, and full-atom tokens ( $x_i^{\text{aa}}, x_i^{\text{bb}}, x_i^{\text{fa}}$ ) are embedded with view-specific embedding tables  $e^{\text{aa}}, e^{\text{bb}},$  and  $e^{\text{fa}}$ . The three embeddings are concatenated along the feature dimension and projected to the Transformer hidden dimension with a fuse layer:

$$h_i^{(0)} = W_{\text{fuse}} [e^{\text{aa}}(x_i^{\text{aa}}); e^{\text{bb}}(x_i^{\text{bb}}); e^{\text{fa}}(x_i^{\text{fa}})]. \quad (5)$$

The sequence of fused residue representations ( $h_1^{(0)}, \dots, h_L^{(0)}$ ) is then processed by a Transformer stack with rotary positional information. Note that the generator is smaller than the representation model. After pretraining, we discard the generator and all token-recovery heads. We use the hidden states from the Transformer stack of  $\mathcal{E}$  as transferable protein representations.

**Hyperparameters.** We pretrain four model scales, 35M, 150M, 650M, and 2.8B, using the configurations in Table 8. Across all scales, the generator  $\mathcal{G}$  uses an 8-layer Transformer encoder and a 2-layer Transformer decoder, and is discarded after pretraining. The representation model  $\mathcal{E}$  follows the scale-specific Transformer configuration in Table 8.

All models are pretrained for 500K steps with a crop size of 512 residues. We use NVIDIA B200 nodes with 8 GPUs per node. The 35M, 150M, 650M, and 2.8B models are trained on 1, 2, 4, and 4 nodes, respectively, with per-GPU batch sizes of 128, 64, 32, and 32. This gives an effective batch size of 1024 protein sequences for all scales. We independently mask each token view with probability 0.6. Training uses AdamW with peak learning rate  $4 \times 10^{-4}$ , weight decay 0.01,  $\beta_1 = 0.9$ ,  $\beta_2 = 0.95$ , gradient clipping at 1.0, and bf16-mixed precision.

## B Details in REPS

### B.1 Dataset curation

We use 1.8M high-confidence homodimer complexes released in the March 2026 AlphaFold Protein Structure Database (AFDB) complex expansion [17, 27, 16]. For each dimer, we obtain the corresponding single-chain apo monomer prediction from AFDB, keyed by UniProt accession. Samples without an AFDB monomer prediction are removed, leaving 1,719,497 apo-holo pairs. We split the benchmark by sequence clusters using MMseqs2 [35] `easy-cluster` with a 50% sequence-identity threshold. This yields 404,961 cluster representatives, from which we select 400 validation and 1,000 test sequences. We construct the initial training split from the remaining representatives, then remove training sequences that share more than 50% sequence identity with any validation or test sequence. After this filtering step, the training split contains 390,861 samples.

### B.2 Homodimer structure prediction

To construct a co-folding model for benchmark purposes, we modify SimpleFold [21], an open Transformer-based single-protein folding model, into a homodimer co-folding model. Our model has a per-atom representation module, a per-residue trunk that fuses the monomer features at its entrance, and a per-atom decoder that outputs the denosing predictions.

**Tokenization and monomer representation extraction.** For each protein, we convert its structure into per-residue and per-atom token tables using the Boltz tokenizer [20, 19], which standardizes residue indexing, atom typing, and bond information across the training pool. We then compute per-residue features for the monomer protein with the frozen representation model  $h_{\text{monomer}}$ . Across runs, we change only the monomer representation model. The atom-processing layers, trunk, and atom decoder are shared across all runs.

**Atom representation module and representation fusion.** The atom representation module is a windowed Transformer with depth 1, hidden size 256, and 4 attention heads. This processes per-atom features, e.g., atom and bond types, together with per-atom coordinates of the current diffusion state  $x_t$ , and produces per-atom latents  $a \in \mathbb{R}^{N_{\text{atom}} \times 256}$ . We aggregate  $a$  to per-residue latents via mean pooling along an atom-to-token assignment, and project up to the trunk hidden size of 768.

Then, we construct the monomer input by concatenating two copies of  $h$  along the residue dimension and adding learned chain-identity embeddings  $E_{\text{chain}}$  that distinguish chain-A and chain-B positions:

$$\tilde{h} = \text{concat}(h_{\text{monomer}}, h_{\text{monomer}}) + E_{\text{chain}} \in \mathbb{R}^{2L \times D_{\text{monomer}}}. \quad (6)$$

We then fuse  $\tilde{h}$  into the per-residue latent at the trunk entrance via concatenation followed by a learned linear layer  $W_{\text{cat}} : \mathbb{R}^{768 + D_{\text{monomer}}} \rightarrow \mathbb{R}^{768}$ :

$$z = W_{\text{cat}}(\text{concat}(\text{token-latent}, \tilde{h})) \in \mathbb{R}^{2L \times 768}. \quad (7)$$

**Trunk.** The trunk is an 8-layer DiT [36] stack with hidden size 768, 12 attention heads (head dim 64), SwiGLU MLPs of ratio 4, and QK-norm along (chain, residue, atom, sub-atom) axes. The trunk operates only on the residue-level latents and outputs  $z' \in \mathbb{R}^{2L \times 768}$ .

**Per-atom decoder.** We expand  $z'$  back to per-atom latents through the atom-to-token assignment, residual-add the per-atom representation output  $a$ , project down to the decoder hidden size 256, and pass through a windowed Transformer decoder of depth 1, followed by a final AdaLN layer that produces the predicted velocity  $v_{\text{pred}} \in \mathbb{R}^{2L \times N_{\text{atom}} \times 3}$ .

**Flow-matching training.** We train the model via flow matching. Given clean homodimer coordinates  $x_1$  and noise  $x_0 \sim \mathcal{N}(0, I)$ , training algorithm samples  $t \sim U(0, 1)$  and computes the interpolant  $x_t = (1 - t)x_0 + tx_1$ . Then, the model is trained to minimize:

$$\mathcal{L} = \mathbb{E}_{t, x_0, x_1} |v_{\theta}(x_t, t, \tilde{h}) - (x_1 - x_0)|^2 + \lambda_{\text{IDDT}} \mathcal{L}_{\text{LDDT}}^{\text{smooth}}, \quad (8)$$

where  $\mathcal{L}_{\text{LDDT}}^{\text{smooth}}$  is a smooth IDDT auxiliary on the denoised coordinates. At inference, we sample  $x_0 \sim \mathcal{N}(0, I)$  and integrate  $v_{\theta}$  from  $t = 0$  to  $t = 1$  with an Euler-Maruyama integrator to obtain the predicted homodimer coordinates.

**Hyperparameters and experimental settings.** For all experiments, we use four NVIDIA B200 GPUs. We train the model for 100,000 steps with AdamW using a learning rate of  $2 \times 10^{-4}$  and weight decay of 0.01. We use gradient clipping with norm 2.0 and use norm 0.5 for 2.8B representations to prevent exploding gradients. We use bf16-mixed precision. Following Boltz [19, 20], we apply a multiplicity factor of 8 to each sample, which replicates each sample with 8 independent random SE(3) augmentations within a step. With 4 samples per GPU and DDP across 4 GPUs, the effective batch size is 128 samples per step. During training, we crop each protein to 512 residue tokens using the Boltz interface-aware cropper and pad to a static shape for `torch.compile`. The auxiliary smooth-IDDT loss weight is  $\lambda_{\text{IDDT}} = 1.0$ . At inference, we use 500 time steps and noise scale  $\tau = 0.3$  to obtain the predicted homodimer coordinates.

### B.3 Per-residue homodimer binding property prediction

**Label construction.** We first compute binding-relevant properties from the homodimer structures. These properties cannot be directly obtained from the monomer alone because they require the corresponding homodimer structure. To be specific, we derive four targets for residue  $i$  as follows.

*Binding-site prediction (binary classification).* A residue is labeled positive if its  $C_\alpha$  atom lies within  $8 \text{ \AA}$  of any  $C_\alpha$  atom on the partner chain.

*$\Delta$ SASA prediction (regression).* For the monomer and the homodimer, we compute the per-residue solvent-accessible surface area using the Shrake-Rupley algorithm [30] as implemented in BioPython [37], with probe radius  $1.4 \text{ \AA}$  and 960 sphere points and get labels:

$$\Delta\text{SASA}_i = \text{SASA}_i^{\text{mono}} - \text{SASA}_i^{\text{dimer}}.$$

*Levy tier (multi-class classification).* We assign each residue to one of five structural regions following the support-core-rim definition [38]. This computes the relative SASA  $r_i = \text{SASA}_i / \text{SASA}_{aa(i)}^{\text{max}}$ , where  $\text{SASA}_{aa(i)}^{\text{max}}$  is the maximum SASA for residue type  $aa(i)$ . Then, this computes  $r_i$  for both the monomer and the homodimer. Residues without inter-chain contacts are assigned to *surface* if  $r_i^{\text{mono}} > 0.25$  and to *interior* if  $r_i^{\text{mono}} \leq 0.25$ . Contacting residues are assigned to *support* if they are already buried in the monomer ( $r_i^{\text{mono}} \leq 0.25$ ), to *rim* if they remain exposed in both states ( $r_i^{\text{mono}} > 0.25$  and  $r_i^{\text{dimer}} > 0.25$ ), and to *core* if they are exposed in the monomer but become buried in the dimer ( $r_i^{\text{mono}} > 0.25$  and  $r_i^{\text{dimer}} \leq 0.25$ ).

*Bond type (multi-label classification).* For each inter-chain residue pair, we test five non-covalent interaction types: hydrogen bonds, salt bridges, hydrophobic contacts,  $\pi$ - $\pi$  stacking, and cation- $\pi$  interactions. We compute interaction labels using PLIP [31]. Residue-level labels are obtained by taking the union of interaction types over all inter-chain partners of each residue.

**Aggregation across chain instances.** Each residue position appears twice in the homodimer, once in each chain. Since the representation model takes the monomer as input and is chain-blind, we merge the two chain-specific labels into a single per-position label: OR for binding site and bond type, mean for  $\Delta$ SASA, and max-rank aggregation for Levy tier.

**Hyperparameters and experimental settings.** We use a deterministic 10% subset of the training split for probe training to better evaluate generalization capability of representations, while using the full validation (400) and test (1,000) sets. For each monomer protein, we extract residue-wise monomer representations and train a two-hidden-layer MLP with hidden width 1280, GELU activations, and dropout 0.1. We train each probe for 10 epochs with a residue-wise batch size of 16,824, which corresponds to approximately 40 to 50 proteins per batch.

### B.4 Monomer structure prediction through distillation

**Architecture.** Our folding model follows the atom-token-atom design of SimpleFold [21]. A per-atom encoder maps atomic features together with the current diffusion state  $x_t$  into atom latents, which are pooled into per-residue tokens. The tokens pass through a residue-level Transformer trunk, and a per-atom decoder reads the trunk output to predict the flow-matching velocity.

The trunk also incorporates per-residue embeddings from a frozen ESM2-650M sequence backbone [22]. These embeddings are added to the residue tokens at the trunk entrance. The trunk is an 8-layer DiT [36] stack with hidden size 768, 12 attention heads, head dimension 64, SwiGLU MLPs

with ratio 4, and QK-norm. The atom encoder and decoder are windowed Transformers with depth 1, hidden size 256, and 4 attention heads, identical to SimpleFold.

**Representation alignment.** We evaluate whether frozen monomer representations provide useful auxiliary supervision for monomer structure prediction. For each representation model, we first extract residue-wise monomer protein representations as representation alignment targets  $h_{\text{tgt}}$  and cache them. During training, we align the intermediate representation of the structure prediction model with the frozen monomer representation. Specifically, we use the output of the 4-th DiT block’s representation  $z$ , and align it to the cached target representation  $h_{\text{tgt}}$ . A two-layer MLP maps the source representation to the target representation dimension:

$$\hat{h} = g(z) \tag{9}$$

where the hidden size of  $g$  is  $\max(768, D_{h_{\text{tgt}}})$ . We then minimize a residue-wise negative cosine similarity between the projected source representation and the cached target representation:

$$\mathcal{L}_{\text{REPA}} = -\lambda_{\text{REPA}} \frac{1}{|\mathcal{V}|} \sum_{i \in \mathcal{V}} \frac{\langle \hat{h}_i, h_{\text{tgt},i} \rangle}{\|\hat{h}_i\|_2 \|h_{\text{tgt},i}\|_2}, \tag{10}$$

where  $\mathcal{V}$  denotes the valid residue indices in the cropped sample. Gradients from  $\mathcal{L}_{\text{REPA}}$  are applied only to the DiT trunk and the projection head  $g$ .

**Hyperparameters and experimental settings.** We train on four NVIDIA B200 GPUs for 30,000 steps with AdamW (learning rate  $4 \times 10^{-4}$ , weight decay 0.01), gradient-norm clipping at 2.0, and bf16-mixed precision. We crop with the Boltz cropper to `max_tokens = 256`, padded to static shapes for `torch.compile`. We use a multiplicity factor of 8, a per-GPU batch of 4, and DDP across 4 GPUs, giving an effective batch size of 128 samples per step. The smooth-IDDT weight is  $\lambda_{\text{IDDT}} = 1.0$  and the REPA weight is  $\lambda_{\text{REPA}} = 2.0$ . At inference, we use 500 time steps and noise scale  $\tau = 0.3$  to obtain the predicted monomer coordinates.

## C Additional experiments

### C.1 Additional per-residue property prediction

Table 9: **Per-residue homodimer binding properties probing on REPSP.** MLP probes on monomer protein representations predict properties derived from the holodimer. **Bold** indicates the best within each size class. We use three random seeds.

Models	Params	Binding site		$\Delta$ SASA		Levy tier		Bond type	
		AUPRC	AUROC	Pearson	Spearman	Macro F1	Acc.	AUPRC	AUROC
<i>Small</i>									
ESM2	35M	0.6635 $\pm$ 0.0009	0.8862 $\pm$ 0.0005	0.5955 $\pm$ 0.0015	0.3445 $\pm$ 0.0015	0.5386 $\pm$ 0.0027	0.7706 $\pm$ 0.0002	0.2936 $\pm$ 0.0004	0.9338 $\pm$ 0.0003
SaProt	35M	0.6291 $\pm$ 0.0009	0.8794 $\pm$ 0.0005	0.6218 $\pm$ 0.0007	0.3923 $\pm$ 0.0010	0.5248 $\pm$ 0.0041	0.8036 $\pm$ 0.0003	0.2663 $\pm$ 0.0014	0.9388 $\pm$ 0.0001
Ours	35M	<b>0.7614</b> $\pm$ 0.0003	<b>0.9243</b> $\pm$ 0.0002	<b>0.7047</b> $\pm$ 0.0003	<b>0.4167</b> $\pm$ 0.0006	<b>0.6260</b> $\pm$ 0.0037	<b>0.8342</b> $\pm$ 0.0003	<b>0.3566</b> $\pm$ 0.0002	<b>0.9518</b> $\pm$ 0.0002
<i>Medium</i>									
ESM2	150M	0.7087 $\pm$ 0.0025	0.9062 $\pm$ 0.0009	0.6379 $\pm$ 0.0006	0.3703 $\pm$ 0.0014	0.5665 $\pm$ 0.0031	0.7969 $\pm$ 0.0003	0.3227 $\pm$ 0.0034	0.9425 $\pm$ 0.0003
Ours	150M	<b>0.8161</b> $\pm$ 0.0004	<b>0.9433</b> $\pm$ 0.0002	<b>0.7801</b> $\pm$ 0.0003	<b>0.4671</b> $\pm$ 0.0012	<b>0.6843</b> $\pm$ 0.0021	<b>0.8668</b> $\pm$ 0.0001	<b>0.4204</b> $\pm$ 0.0024	<b>0.9642</b> $\pm$ 0.0002
<i>Large</i>									
ESM2	650M	0.7915 $\pm$ 0.0006	0.9369 $\pm$ 0.0006	0.7146 $\pm$ 0.0007	0.4113 $\pm$ 0.0007	0.6442 $\pm$ 0.0033	0.8281 $\pm$ 0.0010	0.3902 $\pm$ 0.0020	0.9566 $\pm$ 0.0001
SaProt	650M	0.8273 $\pm$ 0.0008	0.9508 $\pm$ 0.0003	0.7881 $\pm$ 0.0002	0.4715 $\pm$ 0.0012	0.6946 $\pm$ 0.0014	0.8693 $\pm$ 0.0004	0.4305 $\pm$ 0.0016	<b>0.9692</b> $\pm$ 0.0002
S-PLM	704M	0.7789 $\pm$ 0.0017	0.9324 $\pm$ 0.0006	0.6967 $\pm$ 0.0002	0.4036 $\pm$ 0.0013	0.6215 $\pm$ 0.0021	0.8166 $\pm$ 0.0001	0.3602 $\pm$ 0.0024	0.9523 $\pm$ 0.0001
MIF-ST	643M	0.5846 $\pm$ 0.0016	0.8671 $\pm$ 0.0007	0.6726 $\pm$ 0.0008	0.4338 $\pm$ 0.0011	0.5261 $\pm$ 0.0022	0.8435 $\pm$ 0.0002	0.2623 $\pm$ 0.0023	0.9430 $\pm$ 0.0003
Ours	650M	<b>0.8524</b> $\pm$ 0.0011	<b>0.9561</b> $\pm$ 0.0002	<b>0.8088</b> $\pm$ 0.0001	<b>0.4813</b> $\pm$ 0.0005	<b>0.7290</b> $\pm$ 0.0009	<b>0.8780</b> $\pm$ 0.0001	<b>0.4425</b> $\pm$ 0.0050	0.9673 $\pm$ 0.0008
<i>Huge</i>									
ESM2	3B	0.8247 $\pm$ 0.0008	0.9483 $\pm$ 0.0002	0.7446 $\pm$ 0.0009	0.4281 $\pm$ 0.0012	0.6851 $\pm$ 0.0054	0.8426 $\pm$ 0.0004	0.4106 $\pm$ 0.0027	0.9614 $\pm$ 0.0003
ESM3	1.4B	0.8470 $\pm$ 0.0016	0.9558 $\pm$ 0.0005	0.8121 $\pm$ 0.0009	<b>0.4878</b> $\pm$ 0.0010	0.7285 $\pm$ 0.0048	<b>0.8898</b> $\pm$ 0.0002	0.4229 $\pm$ 0.0056	0.9673 $\pm$ 0.0002
ProstT5	1.2B	0.8291 $\pm$ 0.0013	0.9480 $\pm$ 0.0004	0.7574 $\pm$ 0.0011	0.4449 $\pm$ 0.0004	0.6848 $\pm$ 0.0025	0.8538 $\pm$ 0.0005	0.4077 $\pm$ 0.0040	0.9594 $\pm$ 0.0002
Ours	2.8B	<b>0.8650</b> $\pm$ 0.0007	<b>0.9608</b> $\pm$ 0.0005	<b>0.8152</b> $\pm$ 0.0009	0.4854 $\pm$ 0.0020	<b>0.7408</b> $\pm$ 0.0016	0.8811 $\pm$ 0.0003	<b>0.4516</b> $\pm$ 0.0017	<b>0.9691</b> $\pm$ 0.0011

### C.2 Additional results on monomer structure prediction

Table 10: **Additional results on monomer structure prediction with disillation on REPSP.** **Bold** denotes the best performance.

REPA target	Params	TM-score $\uparrow$	GDT-TS $\uparrow$	GDT-HA $\uparrow$	LDDT $\uparrow$	LDDT <sub>bb</sub> $\uparrow$	RMSD ( $\text{\AA}$ ) $\downarrow$
Baseline	–	0.5422	0.3432	0.1990	0.6208	0.6746	12.98
SaProt	650M	0.5659	0.3677	0.2192	0.6410	0.6947	12.26
S-PLM	704M	0.5569	0.3559	0.2081	0.6301	0.6838	12.40
MIF-ST	643M	0.5491	0.3520	0.2075	0.6310	0.6827	12.73
ESM3	1.4B	0.5684	0.3714	0.2223	0.6366	0.6900	12.25
ProstT5	1.2B	<b>0.5833</b>	0.3854	0.2328	<b>0.6535</b>	<b>0.7119</b>	<b>11.76</b>
Ours	650M	0.5822	<b>0.3858</b>	<b>0.2344</b>	0.6513	0.7051	11.91

### C.3 Per-residue heterodimer binding-property probing

**Dataset curation.** We probe per-target-residue binding properties on heterodimers from PPIRef50K [39], restricting to entries where the binder and target chains have non-identical sequences ( $n = 15,913$  complexes; sequence-identical homodimers are excluded so that labels strictly reflect cross-chain interactions between distinct proteins). Each complex consists of two distinct chains (binder, target), so labels are not aggregated across chain instances. The representation model takes the target monomer as input, and we evaluate two complementary target-conditioned labels: a binary *binding-site* indicator and a continuous *contact count*.

**Binding-site prediction (binary classification).** Per-residue solvent-accessible surface area is computed for both the target monomer and the binder–target complex with DSSP [40], using the Sander–Rost relative-ASA normalization. Let  $r_i^{\text{mono}}$  and  $r_i^{\text{cpx}}$  denote the relative ASA of target residue  $i$  in the monomer and the complex, respectively, and define the burial change  $\Delta r_i = r_i^{\text{mono}} - r_i^{\text{cpx}}$ . A residue is labeled positive if  $\Delta r_i \geq 0.05$  and both  $r_i^{\text{mono}}$ ,  $r_i^{\text{cpx}}$  are finite, i.e., its surface is measurably occluded by the binder. Residues with non-finite DSSP output (missing atoms, alternative locations) are dropped from training and evaluation.

**Contact-count prediction (regression).** For each target residue  $i$  we count the number of partner  $C_\alpha$  atoms within  $8 \text{\AA}$  of  $i$ ’s  $C_\alpha$ :

$$c_i = |\{j \in \text{binder} : \|C_{\alpha,i}^t - C_{\alpha,j}^b\| \leq 8 \text{\AA}\}|. \quad (11)$$

Compared with the binary binding-site label,  $c_i$  retains a graded signal of how deeply each interface residue is packed against the partner, capturing local contact density and hotspot strength. Targets are stored as  $\log(1 + c_i)$  and the probe is trained with mean-squared error on the transformed target; we report Pearson  $r$  and Spearman  $\rho$  on the original count scale.

**Class-balanced subsampling.** Across heterodimer targets, only  $\sim 10.9\%$  of residues are interface residues, which biases probes toward majority-class collapse for the binary task and zero-inflation bias for the regression task (a trivial  $\hat{c}_i = 0$  predictor would attain a low MSE while learning nothing). To control prevalence we construct a single deterministic class-balanced subsample that we apply uniformly to both tasks: we keep *all* interface residues ( $\Delta r_i \geq 0.05$ ) and globally Bernoulli-sample non-interface residues with probability  $p = n_{\text{interface}}/n_{\text{non}}$ , yielding  $\approx 50/50$  prevalence ( $\sim 620,000$  residues kept). The same per-PPI keep-mask is applied identically to the training, validation, and test splits (the split is PDB-level, so a target’s residue indices retain their split tag while the mask determines residue-level inclusion). The mask is generated once with a fixed seed (default 42) so that all PLMs see exactly the same residues across the binding-site and contact-count tasks. Because train and test share the same balancing rule, chance baselines shift accordingly: AUPRC chance is 0.5 for binding-site, and the variance of  $\log(1 + c_i)$  on the balanced subset replaces the zero-inflated baseline for regression.

**Hyperparameters and experimental settings.** For each target monomer, we extract residue-wise representations from the frozen PLM and feed them into a probe head. The probe head is a two-hidden-layer MLP with hidden width 1,280, GELU activations, and dropout 0.1, optimized with AdamW (weight decay 0.01). We train each probe for 10 epochs with a residue-wise batch size of 16,824 and a learning rate of  $5 \times 10^{-4}$ . Binding-site uses cross-entropy with uniform class weights (the subsample already removes majority-class dominance, so reweighting would collapse to identity); contact-count uses mean-squared error on  $\log(1 + c_i)$ . We evaluate on the full validation and test splits (residues filtered by the same balanced mask). Binding-site is reported with AUPRC and AUROC; contact-count with Pearson  $r$  and Spearman  $\rho$ .

Table 11: **Per-residue heterodimer binding properties probing on REPS.** MLP probes on monomer protein representations predict properties derived from the heterodimer complex. Both labels are target-conditioned and evaluated on a class-balanced subset. **Bold** indicates the best performance within each model-size group. Our model achieves the strongest performance.

Models	Params	Binding site		Contact count	
		AUPRC	AUROC	Pearson	Spearman
<i>Small</i>					
ESM-2	35M	0.717 $\pm$ 0.001	0.656 $\pm$ 0.002	0.339 $\pm$ 0.002	0.322 $\pm$ 0.001
Ours	35M	<b>0.755<math>\pm</math>0.001</b>	<b>0.693<math>\pm</math>0.001</b>	<b>0.411<math>\pm</math>0.002</b>	<b>0.401<math>\pm</math>0.002</b>
<i>Medium</i>					
ESM-2	150M	0.741 $\pm$ 0.001	0.676 $\pm$ 0.001	0.376 $\pm$ 0.003	0.362 $\pm$ 0.003
Ours	150M	<b>0.796<math>\pm</math>0.001</b>	<b>0.727<math>\pm</math>0.002</b>	<b>0.460<math>\pm</math>0.004</b>	<b>0.445<math>\pm</math>0.004</b>
<i>Large</i>					
ESM-2	650M	0.759 $\pm$ 0.002	0.692 $\pm$ 0.001	0.407 $\pm$ 0.002	0.391 $\pm$ 0.002
SaProt	650M	0.797 $\pm$ 0.001	0.730 $\pm$ 0.001	<b>0.500<math>\pm</math>0.002</b>	<b>0.476<math>\pm</math>0.002</b>
S-PLM	704M	0.755 $\pm$ 0.001	0.690 $\pm$ 0.001	0.402 $\pm$ 0.002	0.390 $\pm$ 0.001
MIF-ST	643M	0.757 $\pm$ 0.000	0.695 $\pm$ 0.001	0.448 $\pm$ 0.002	0.435 $\pm$ 0.003
Ours	650M	<b>0.817<math>\pm</math>0.001</b>	<b>0.744<math>\pm</math>0.000</b>	0.482 $\pm$ 0.001	0.468 $\pm$ 0.001
<i>Huge</i>					
ESM-2	3B	0.764 $\pm$ 0.000	0.697 $\pm$ 0.000	0.415 $\pm$ 0.002	0.396 $\pm$ 0.002
ESM3	1.4B	0.746 $\pm$ 0.007	0.696 $\pm$ 0.004	0.047 $\pm$ 0.007	0.045 $\pm$ 0.010
ProSST	1.2B	0.761 $\pm$ 0.000	0.694 $\pm$ 0.000	0.442 $\pm$ 0.002	0.426 $\pm$ 0.002
Ours	2.8B	<b>0.834<math>\pm</math>0.001</b>	<b>0.761<math>\pm</math>0.000</b>	<b>0.517<math>\pm</math>0.000</b>	<b>0.502<math>\pm</math>0.000</b>

#### C.4 Conventional benchmark

Following SaProt [2], instead of native crystal structures we use AlphaFold predicted structures from the AFDB corresponding to each protein sequence. For each protein, we extract last-layer residue-wise representations from the pretrained encoder and mean-pool over non-special tokens to

obtain a single per-protein embedding. We then train a two-hidden-layer MLP probe with hidden width 1280, GELU activations, and dropout 0.1, using binary cross-entropy. Each probe is trained for 100 epochs with Adam (learning rate  $5 \times 10^{-4}$ ) and a per-protein batch size of 128. We select the epoch with the highest validation  $F_{\max}$  and report the corresponding test  $F_{\max}$ .