SLAE: STRICTLY LOCAL ALL-ATOM ENVIRONMENT FOR PROTEIN REPRESENTATION

Anonymous authors

Paper under double-blind review

ABSTRACT

Building physically grounded protein representations is central to computational biology, yet most existing approaches rely on sequence-pretrained language models or backbone-only graphs that overlook side-chain geometry and chemical detail. We present SLAE, a unified all-atom framework for learning protein representations from each residue's local atomic neighborhood using only atom types and interatomic geometries. To encourage expressive feature extraction, we introduce a novel multi-task autoencoder objective that combines coordinate reconstruction, sequence recovery, and energy regression. SLAE reconstructs all-atom structures with high fidelity from latent residue environments and achieves state-of-the-art performance across diverse downstream tasks via transfer learning. SLAE's latent space is chemically informative and environmentally sensitive, enabling quantitative assessment of structural qualities and smooth interpolation between conformations at all-atom resolution.

1 Introduction

Proteins are the fundamental machinery of life, carrying out processes from catalysis and signaling to structural organization. Their remarkable functional diversity arises not only from their amino acid sequences but from the intricate three-dimensional structures into which those sequences fold.

Within protein structures, the backbone and side chain atoms act as an intricately coupled system that establishes local atomic environments through hydrophobic packing, hydrogen-bonding networks, and electrostatic interactions. These residue-level environments mediate conformational preferences and side chain dynamics, linking the global fold to the specific interactions that underlie protein function. Representing these interactions in a concise, learnable form is therefore essential for generalizable and physically grounded models of protein structure and function.

Despite the importance of atomistic detail, many existing structure representations remain coarse-grained, as reasoning over thousands of atoms per protein is often computationally intractable. While all-atom graph encoders have been proposed using convolution and hierarchical pooling (Torng & Altman, 2017; Hermosilla et al., 2021; Wang et al., 2023), they still incur a considerable computational burden making them impractical for large-scale pretraining with previously proposed denoising objectives (Zaidi et al., 2022; Jamasb et al., 2024). In contrast, most current representation learning efforts rely on protein language models (PLMs) pretrained on massive amino acid sequence corpora (Meier et al., 2021; Lin et al., 2023), which introduce inductive biases that may not align with physical reality (Weinstein et al., 2023; Ding & Steinhardt, 2024). More recently, hybrid approaches that augment sequence representations with structural representations have been explored (Su et al., 2023b; Gao et al., 2024; Hayes et al., 2024), but the incorporated structural information is typically restricted to backbone coordinates, omitting side-chain detail important for protein function. As a result, the field remains limited by the absence of a general-purpose pretraining framework that extracts, compresses, and transfers knowledge of all-atom structure across proteins and downstream applications.

We propose **SLAE** (Strictly Local All-atom Environment autoencoder), a new framework for protein representation learning. We model a protein as a collection of residue-centric chemical environments. Fully connected local atom graphs capture interactions between a residue and its nearby atoms, and are computationally tractable during pretraining. We show that these local representations contain sufficient information to recover all-atom Cartesian coordinates at high fidelity.

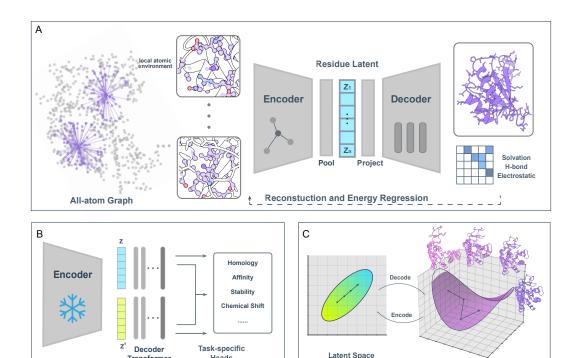


Figure 1: **Overview of the SLAE framework. A. Pretraining** A graph encoder maps local atomic neighborhoods to residue embeddings, which a transformer decoder uses to reconstruct all-atom coordinates and predict inter-residue energy terms. **B. Transfer learning** The pretrained embeddings are fed to lightweight heads for diverse downstream tasks. **C. Latent geometry** Linear interpolations on latent space decode to physically coherent structures that follow changes on the underlying chemical-environment manifold.

Chemical Environment Space

We design an all-atom autoencoder architecture that separates local and global reasoning across the encoding and decoding stages. An SE(3)-equivariant graph encoder maps each local environment to a rotation/translation-invariant residue token. A Transformer decoder with self-attention then aggregates these tokens to model long-range couplings and reconstruct coherent global geometry. This residue-level bottleneck forces the encoder to distill the packing signals such as covalent bonds, hydrogen-bond motifs, and steric/electrostatic cues that the global decoder requires to reconstruct long-range geometry, facilitating transfer across tasks. We introduce a physics-augmented pretraining objective that couples self-supervised (i) all-atom coordinate reconstruction, (ii) sequence recovery, and supervised (iii) Rosetta-derived inter-residue energies. These complementary signals act as a multi-view regularizer, aligning the latent space with atomistic structure, biochemical signal and energetics, yielding embeddings that vary smoothly with conformation and are interpretable along axes of side-chain chemistry, solvent exposure, and secondary structure.

SLAE supports multiscale readouts: atom and residue embeddings for fine-grained local characterization, and pooled protein-level features for global structure. This flexibility allows downstream task heads to focus on single residues, interfaces, or entire folds using a single pretrained representation. We demonstrate that pretraining directly on all-atom protein structures yields features that transfer effectively. Across benchmarks on multiple resolution scale tasks including fold classification, protein–protein binding affinity, single-point mutation stability, and NMR chemical shifts, SLAE achieves state-of-the-art or on-par performance.

Main contributions: With the SLAE framework, we (i) propose a residue-centered, local atomgraph protein representation, and show it is sufficient for high-fidelity all-atom reconstruction. (ii) design a self- and weakly supervised pretraining objective coupling all-atom coordinate reconstruction, amino-acid sequence recovery, and inter-residue energy prediction. (iii) pretrain an all-atom autoencoder that pairs local graph encoder with global transformer decoder to encourage a compact, transferable latent space. (iv) achieve state-of-the-art on diverse downstream tasks with transfer learning. (v) show that the above design allow an interpretable latent space.

RELATED WORK

110 111 112

113

114

115

116

117

118

119

120

121

122

123

124

125 126

127

128

129

130

131

132

133

134

135

108

109

Protein Representation Pretraining Protein representation learning has followed two main tracks. Sequence pretraining with protein language models (PLMs) on massive corpora captures evolutionary constraints but lacks explicit structure information (Meier et al., 2021; Lin et al., 2023). In parallel, graph denoising objectives noises sequence or structural features and train graph models to recover them (Zaidi et al., 2022; Jamasb et al., 2024), capturing global context while abstracting away side-chain geometry. Neither paradigm learns atomistic features as the primary signal. SLAE departs by pretraining directly on all-atom coordinates reconstruction and showing that features learned from atomistic geometry are sufficient for high-fidelity coordinate reconstruction and downstream transfer.

Sequence-structure co-embedding approaches pair PLM embeddings with structural features to inject geometry into sequence representations, improving downstream performance without learning at all-atom resolution. Representative methods include SaProt (Su et al., 2023b), FoldToken (Gao et al., 2024), ProSST (Li et al., 2024), and ESM3 (Hayes et al., 2024). Most hybrid models augment sequence tokens with backbone-level descriptors, and the learned tokens remain sequence-anchored. SLAE instead learns structure and energetics-anchored residue tokens, reducing sequence-only bias while increasing structure representation resolution.

Geometric GNNs for Atomistic Systems Representing atomistic systems as geometric graphs is natural. Equivariant GNNs such as DimeNet (Gasteiger et al., 2022), NequIP (Batzner et al., 2022) and MACE (Batatia et al., 2023) excel at small-molecule property prediction and interatomic potentials. For scalability, many adopt low-order interactions with truncated neighborhoods, closely related to Atomic Cluster Expansion (ACE) formulations (Drautz, 2019). Works which extend atomistic modeling to proteins are emerging (Pengmei et al., 2024; Bojan et al., 2025), but existing approaches typically pretrain on small-molecule datasets, reuse features from pretrained potential models or are trained in a task-specific manner. There remains a gap in methods amenable to large-scale, all-atom pretraining on proteins. SLAE addresses this by modeling two-body local interactions over cutoff graphs and pretrain a physics-informed autoencoder that yields a general, task-agnostic latent space at protein scale: thousands of atoms per system compared to tens of atoms.

136 137 138

THE SLAE FRAMEWORK

139 140 141

142

143

144

145

We introduce the SLAE autoencoder and its end-to-end pretraining objectives (Fig. 1A). SLAE solves a deliberately difficult two-part problem: the geometric graph encoder projects interatomic interactions within each atom's local neighborhood into compact residue tokens, while the decoder learns a global prior over how these local environments compose into coherent macromolecular structures. This residue-level bottleneck over all-atom inputs makes large-scale pretraining tractable and learns meaningful embeddings.

146 147

3.1 STRUCTURE REPRESENTATION

148 149 150

154

155

156

157 158

159

160

161

Given a protein structure, we construct a directed graph $\mathcal{G} = (\mathcal{V}, \mathcal{E})$, where:

151 152 153

Nodes Each node $v_i \in \mathcal{V}$ represents heavy atom a_i . The node feature is a one-hot encoding of the atom's chemical type. **Edges** For each pair of atoms a_i, a_j with $\|a_j - a_i\|_2 \le 8$ Å, we define a directed edge $e_{j\to i} \in \mathcal{E}$

with features $h_{i \to i}^{(e)}$ that is a concatenation of: (i) the scaler interatomic distance $\|a_j - a_i\|_2$ in terms of Bessel radial basis functions $\phi_r(a_i, a_j)$ and (ii) the unit vector interatomic direction projected onto spherical harmonics $Y_{\ell m} \phi_a(\boldsymbol{a}_i, \boldsymbol{a}_i)$.

Design Motivation This representation is *minimal yet physically complete*: it encodes interatomic distances and orientations without relying on torsion angles, amino acids, or residue indices. As such, it enables generalization to arbitrary biomolecular complexes, which we leave for future work. Bond connectivity and hydrogen patterns are learned implicitly through the autoencoder objective detailed in Section 3.4.

3.2 ENCODER

The encoder maps each atom's local chemical environment into residue-level latent embeddings $\{z_1,\ldots,z_n\}, \quad z_i \in \mathbb{R}^{128}$.

Equivariant Neighborhood Embedding We employ a SE(3)-equivariant neural network, inspired by Musaelian et al. (2023), that operates on each heavy atom and its neighbors through learned edge embeddings. Each layer L maintains coupled latent spaces: a scalar space x_{ij}^L (invariant) and a tensor space \mathbf{V}_{ij}^L (equivariant). An equivariant tensor product incorporates interactions between the current equivariant state of the center-neighbor pair (i,j) and all other neighbors $k \in \mathcal{N}(i)$: $\mathbf{V}_{ij}^L = \mathbf{V}_{ij}^{L-1} \otimes (\sum_{k \in \mathcal{N}(i)} w_{ik}^L \phi(\mathbf{r}_{ik}))$, where $\phi(\mathbf{r}_{ik})$ is a geometric embedding of the neighbor direction and w_{ik}^L are learned weights derived from scalar features of edges (i,k). This can be viewed as a weighted projection of the atomic density around atom i, enabling equivariant interactions between the pair (i,j) and the environment of i.

Following the tensor product, scalar outputs are reintroduced into the scalar latent space with $x_{ij}^L = \mathrm{MLP}_{\mathrm{latent}}^L(x_{ij}^{L-1} \parallel \mathbf{V}_{ij}^L) \cdot u(r_{ij})$, where $u(r_{ij})$ is a smooth cutoff envelope. This step completes the coupling of scalar and equivariant latent spaces: scalars distilled from tensor products inject directional information back into x_{ij}^L , allowing the invariant channel to carry geometric cues that were previously only available to the equivariant representation.

Residue Environment Pooling After the final layer, we obtain scalar pair features x_{ij}^L . We first pool to atoms by mean-aggregating incoming edges, and then pool atom embeddings to residues: $s_i = \frac{1}{|\mathcal{N}(i)|} \sum_{j \in \mathcal{N}(i)} x_{ij}^L$, $z_r = \frac{1}{|\mathcal{A}(r)|} \sum_{i \in \mathcal{A}(r)} s_i$. This yields compact residue-level representations while retaining strictly local chemical information.

Design Motivation The encoder updates edge embeddings dynamically by incorporating information from neighboring edges. This paradigm originally developed for interatomic potentials in small-molecule graphs naturally extends to large protein graphs. This allows SLAE to capture strictly local but physically meaningful chemical environments. Pooling representations to the residue level serves as an efficient and natural information bottleneck for protein structure.

3.3 Decoder

Having distilled each residue's local chemistry and geometry into embeddings $z \in \mathbb{R}^{128}$, the decoder assembles these local descriptors into a single, coherent macromolecule that respects long-range couplings.

Architecture We first project each latent embedding to a model dimension of \mathbb{R}^{1024} . On top of these expanded embeddings, we employ a Transformer architecture with global self-attention and Rotary Positional Embeddings (RoPE) (Su et al., 2023a) to capture long-range residue interactions with a stack of multi-head self-attention layers.

The Transformer outputs are passed into three parallel MLP heads for structure reconstruction, sequence recovery, and energy prediction:

- 1. Reconstructs the 3D coordinates of up to 37 heavy and side-chain atoms per residue ($\hat{x} \in \mathbb{R}^{n \times 37 \times 3}$).
- 2. Recovers the amino acid identity at each residue position ($\hat{s} \in \mathbb{R}^{n \times 20}$).
- 3. Approximates inter-residue physical interactions using Rosetta scores, including hydrogen bonding, electrostatics, and solvation energies ($\hat{r} \in \mathbb{R}^{n \times n \times 3}$).

Design Motivation The decoder is designed to complement the encoder's strictly local representation by modeling *global* dependencies across residues. Global self-attention allows residue embeddings to exchange information across the entire protein, enabling the reconstruction of coherent backbone and side-chain geometries. The addition of energy prediction task guides the decoder toward physically meaningful structures, ensuring that the latent space encodes not only geometric detail but also the energetic constraints that govern protein stability and interactions.

3.4 Pretraining

We pretrain SLAE end-to-end on full atomic structures with three complementary objectives:

223

218

1. All-atom Structure Recovery To obtain the predicted structure \hat{x} , we mask the atom-37 template coordinates while providing the ground-truth residue identities to train the decoder to recover ground truth coordinates. We supervise this reconstruction with a combination of all-atom local distance difference test loss (SmoothLDDT) (Jumper et al., 2021) and frame-aligned point error (FAPE) (Anishchenko et al., 2024): $\mathcal{L}_{\text{struct}} = \alpha \text{LDDT}(\boldsymbol{x}, \hat{\boldsymbol{x}}) + \beta \text{FAPE}(\boldsymbol{x}, \hat{\boldsymbol{x}})$, where \boldsymbol{x} and $\hat{\boldsymbol{x}}$ denote the ground-truth and predicted all-atom coordinates.

224 225 226

2. Sequence Recovery We additionally recover the residue sequence from the latent space: $\mathcal{L}_{seq} =$ CrossEntropy (s, \hat{s}) , where s is the ground-truth amino-acid identity and \hat{s} are the predicted logits over 20 amino acid classes.

227 228 229

3. Energy Prediction To inject physically grounded supervision, we predict inter-residue energies approximated by Rosetta scores, including hydrogen bonding, electrostatics, and solvation: $\mathcal{L}_{\text{energy}} = \| \boldsymbol{r} - \hat{\boldsymbol{r}} \|_2^2$, where \boldsymbol{r} and $\hat{\boldsymbol{r}}$ are ground-truth and predicted energy terms.

231 232

230

The combined loss integrates all three components:

233 234

$$\mathcal{L} = w_{\text{coord}} \left(\alpha \, \text{LDDT} + \beta \text{FAPE} \right) + w_{\text{seq}} \, \text{CrossEntropy} + w_{\text{energy}} \, \text{MSE}$$
 (1)

235

with weights $w_{\text{struct}}, w_{\text{seq}}, w_{\text{energy}} \geq 0$ as tunable hyperparameters (Appendix C.1).

236 237

238

239

240

241

242

243

244

Implicit Latent Space Regularization By jointly optimizing geometry, identity, and energetics, SLAE's pretraining objective provides complementary constraints on the latent space: (i) Geometry losses depend smoothly on atomic coordinates, promoting continuous and physically plausible reconstructions. (ii) Sequence recovery encourages embeddings to encode amino acid identity, preserving biochemical interpretability and avoiding collapse. (iii) Energy prediction provides a physics-based signal, guiding embeddings toward inter-residue interactions such as hydrogen bonding, solvation, and electrostatics. These losses shape a latent manifold that maps cleanly onto valid, physically coherent protein conformations. The result is a structurally consistent, chemically informative, and energetically grounded representation without relying on explicit regularizers.

245 246

3.5 RESULTS AND ABLATIONS

247 248 249

250 251

Graph	Discretization	Codebook	Training	Seq. Acc.	RMSD	RMSD
Radius(Å)	Method	Size	Obj.	(%)	< 128 (Å)	< 512 (Å)
8	LFQ	32768	all	75.2	2.50	3.74
8	kNN	4096	all	97.5	2.96	4.03
8	kNN	32768	all	99.4	1.60	2.31
8	_	_	w/o. FAPE	97.2	3.89	5.22
8	_	_	w/o. Energy	98.0	3.26	5.17
6	_	_	all	99.9	1.24	2.55
8	-	_	all	99.9	1.12	1.92

257

258

Table 1: Reconstruction performance of SLAE and ablations. We report sequence recovery accuracy (%) and reconstruction RMSD (Å) on test structures. All further experiments use the highlighted best SLAE model.

259 260 261

262

263

We pretrain SLAE on a sequence-augmented CATH(Ingraham et al., 2019)-derived dataset (Lu et al., 2025b)(Appendix D). On the held-out test set with no family overlap, the autoencoder achieves 99.9% sequence recovery and all-atom RMSD of 1.1Å for structures shorter than 128 residues and 1.9Å across all lengths up to 512 residues.

264 265 266

267

268

269

We study the effect of model and pretraining design choices on pretraining performance (Table 6). For encoder locality, we swept cutoff radii and find an 8 Å neighborhood yields the best results (Appendix F). For discretization, we compare end-to-end VQ (van den Oord et al., 2018) and LFQ (Yu et al., 2023) against post-hoc kNN codebooks built on frozen encoder embeddings. End-toend quantization trades off sequence and structure accuracy, whereas reconstruction from post-hoc kNN-codebook quantized embeddings approaches continuous resolution as the codebook grows.

Ablation experiments (Table 6, Appendix F) further highlight the importance of both the FAPE loss and Rosetta-derived energy supervision, confirming the effectiveness of our multitask pretraining framework. These results validate the design choices and permit downstream evaluation on a faithful representation of protein structures.

4 DOWNSTREAM TASKS

We next demonstrate that SLAE embeddings pretrained on all-atom reconstruction and energetics objectives transfer effectively to diverse downstream tasks(Figure 1B). Across all four benchmarks spanning complementary biological scales, SLAE achieves better or on-par performance with state-of-the-art methods, underscoring the generality and flexibility of the SLAE framework.

Fold Classification Protein fold classification is a cornerstone of structural biology, linking structure to evolutionary relationships and functional annotation. Using the SCOPe 1.75 dataset Fox et al. (2014) and following Hou et al. (2018), we evaluate generalization under three test sets: Family, Superfamily, and Fold. An MLP is trained on pooled residue embeddings. SLAE achieves on-par or superior accuracy compared to prior state-of-the-art models across all splits (Table 2), demonstrating that global fold information can be recovered even from strictly local all-atom embeddings.

Method	Fold (%)	Superfamily (%)	Family (%)
GVP-GNN(Jing et al., 2021)	16.0	22.5	83.8
IEConv(Hermosilla et al., 2021)	45.0	69.7	98.9
GearNet-Edge-IEConv (Zhang et al., 2023)	48.3	70.3	99.5
ProNet-SCHull (Wang et al., 2024)	56.1	74.6	99.4
SLAE-finetuned	55.1	77.1	99.1

Table 2: Fold classification accuracy (%) on SCOPe 1.75 under three test splits

Protein-Protein Binding Affinity Prediction Protein-protein interactions underlie nearly all cellular processes, and accurate prediction of binding affinity is critical for understanding signaling pathways, complex assembly, and therapeutic design. We evaluate SLAE on the PPB-Affinity dataset (Liu et al., 2024), a recently curated large-scale benchmark that aggregates 12,062 experimental binding $\Delta\Delta G$ values from multiple sources and aligns them with high-quality structural complexes.

Complex structures are embedded chain-wise and interface-wise with the SLAE encoder, and pooled residue embeddings are passed into an MLP for regression. In 5-fold cross-validation, SLAE achieves lower RMSE and higher Pearson correlation than PLM-based baselines (Table 3). Despite being pretrained only on single-chain data, SLAE generalizes seamlessly to multi-chain contexts, thanks to its atomistic representation that does not rely on residue or chain indices.

Method	RMSE (kcal/mol)	Pearson Correlation
PPB-Affinity Baseline (Liu et al., 2024)	2.08	0.70
PPLM-Affinity (Liu et al., 2025)	1.89	0.76
SLAE-finetuned (w/o. interface)	2.01	0.73
SLAE-finetuned (with interface)	1.86	0.77

Table 3: Protein-protein binding affinity prediction on the PPB-Affinity dataset

Single-Point Mutation Thermostability Prediction Protein stability is fundamental to function, and predicting the impact of point mutations on thermostability $(\Delta \Delta G)$ is a central challenge for protein engineering, drug resistance modeling, and disease variant interpretation. We benchmark SLAE on the Megascale mutation dataset (Tsuboyama et al., 2023), filtered according to ThermoMPNN protocol with 272,712 mutations across 298 proteins Dieckhaus et al. (2024).

Pairs of wild-type and mutant structures are embedded with residue-level differences extracted at the mutation site. An MLP head predicts $\Delta\Delta G$. SLAE achieves 0.68 RMSE and 0.76 Pearson correlation (Table 4) on the test set, outperforming prior methods. Ablation experiments show that

removing mutation-site differencing degrades performance, highlighting the importance of local residue environment modeling for physical property prediction in the SLAE framework.

Method	RMSE (kcal/mol)	Pearson Correlation
Rosetta (Pancotti et al., 2022)	5.18	0.53
RaSP (Blaabjerg et al., 2023)	1.08	0.71
ThermoMPNN (Dieckhaus et al., 2024)	0.71	0.75
SLAE-finetuned (w/o. mutated site)	0.73	0.70
SLAE-finetuned (with mutated site)	0.68	0.76

Table 4: Single-point mutation thermostability prediction on the Megascale dataset test split

Chemical Shift Prediction NMR chemical shifts are among the most direct experimental probes of local atomic environments, among them the backbone nitrogen are notoriously difficult to predict accurately due to its large variance and contributions from ring currents, electrostatics, and subtle side-chain conformations. We benchmark on stringently filtered BMRB (Hoch et al., 2023) which contains 2,532 training and 594 validation chemical shift records and their corresponding Alphafold2 predicted structures. PLM-CS framework is adopted as baseline model architecture, which trains a lightweight predictor on top of pretrained representations (Zhu et al., 2025).

We report validation set performance of finetuned SLAE along with PLM-CS results using multiple protein residue embeddings, including ESM2, AlphaFold2, ProSST and SLAE. ¹ Finetuned SLAE achieves the lowest RMSE and highest correlation, substantially outperforming retrained PLM-CS baselines (Table 5). This demonstrates that SLAE embeddings capture fine-grained atomistic features essential for NMR observables.

Method	RMSE (ppm)	Pearson Correlation
PLMCS-AF2	2.94	0.82
PLMCS-ESM2	2.74	0.84
PLMCS-ProSST	2.53	0.87
PLMCS-SLAE	2.53	0.87
SLAE-finetuned	1.88	0.93

Table 5: Backbone nitrogen chemical shift prediction on BMRB

5 INTERPRETING THE LATENT SPACE

SLAE's downstream performance stems from a structured, interpretable latent space. We show that residue embeddings are organized along biochemically meaningful axes, are sensitive to local environment changes, and admit linear paths that decode to geometrically coherent structures (Figure 1C).

5.1 EMBEDDING VARIABILITY REFLECTS CHEMICAL ENVIRONMENT CHANGE

To probe what SLAE embeddings captures at the residue level, we analyze how they organize across local chemical environments. Dimensionality reduction of kNN centroids from CATH (Section 3.5, Appendix F) shows that residue latents cluster by side chain chemistry and broader structural context. The latent space also stratifies along gradients of solvent accessibility and separates by secondary structure, with helices, sheets, and coils occupying distinct submanifolds (Figure 3, App. Fig 6and 7). This indicates that SLAE representation is sensitive to both chemical identity and structural environment.

We then quantify this sensitivity using the mdCATH dataset (Mirarchi et al., 2024). Across 5,398 proteins, per-residue latent displacement between conformers correlates with physical measures of environment variability: changes in contact maps and solvent exposure explain over half of the

¹Fair comparison with open-sourced methods is not possible due to non-overlapping dataset splits (some entries from PLM-CS datasets do not pass the filter standard). We therefore re-trained a PLM-CS baseline on our splits and evaluate all embeddings under an identical protocol.

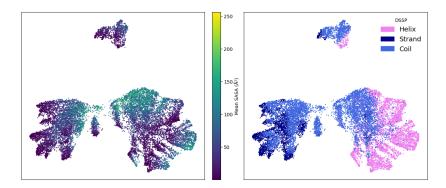


Figure 2: **SLAE latent organization.** UMAP visualization of kNN centroids shows clustering by solvent accessibility (left) and secondary structure (right).

variance in embedding similarity ($R^2 = 0.55$, $\rho \approx 0.74$; Appendix F). Thus, SLAE embeddings consistently track how residues respond to burial, packing, and secondary-structure transitions.

5.2 DISCRIMINATIVE POWER OVER NATIVE-DECOY RESIDUE ENVIRONMENTS

We show that SLAE residue latent capture local environments contain signal that zero-shot distinguishes native structures from decoys and provide a practical embedding space for evaluating backbone–sequence co-design.

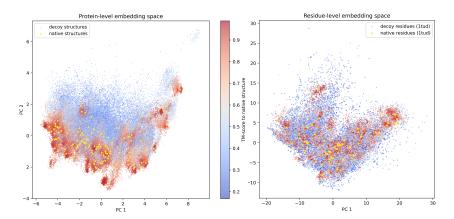


Figure 3: **SLAE** embedding comparison between native and decoy structures. (native in yellow, decoys colored by TM-score to their native; warmer = more native-like) **Left** protein-level PCA. Each point is a protein. **Right** residue-level PCA for 1TUD and its decoys. Decoy residues are colored by their parent decoy's TM-score. In both panels, SLAE embeddings organize along gradients of nativeness, revealing coherent neighborhoods that align with structural quality.

On the Rosetta decoy dataset (Park et al., 2016) containing 133 native protein structures with thousands of decoys each, native–decoy cosine margin is 0.136 across residues. We further fit a leave-protein-out logistic regression by training on all proteins except one and tested on the held-out protein's residues and report AUROC = 0.659 (Appendix F, indicating a moderate, generalizable linear signal at the residue level. Motivated by this discriminative signal, we use the SLAE embedding space to quantify the distributional coverage of generative models, extending prior metrics (Lu et al., 2025a) to all-atom resolution and residue granularity. As a proof of concept, we compute per-residue type Fréchet Protein Distance (FPD) between SLAE embeddings of the generated structures and the native CATH distribution for models such as Chroma (Ingraham et al., 2023), Protpardelle-1c (Lu et al., 2025b) and La-Proteina (Geffner et al., 2025). SLAE-based FPD identifies subtle model coverage differences in local amino acid environments that are averaged out in global fold level assessments (App. Fig 8)

5.3 SMOOTH LATENT INTERPOLATION CAPTURES CONFORMATIONAL TRANSITIONS

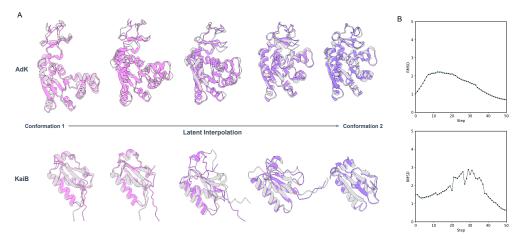


Figure 4: Latent space interpolation between two conformations A. Structures sampled by linear interpolation (purple) overlaid with MD simulation frames (grey) B. Alignment RMSD to MD simulation trajectories

Latent space smoothness is relevant for evaluating whether a representation supports continuous sampling of protein conformations. Unlike variational autoencoders that encourage smoothness via KL regularization to a simple prior, the SLAE autoencoder relies solely on physics-augmented pretraining objectives. We examine the smoothness of SLAE latent by linear interpolation between two conformation states $\mathbf{Z}^{(A)}$ and $\mathbf{Z}^{(B)}$. For each residue i and interpolation scale $t \in [0,1]$, the interpolated residue embeddings are given by $\mathbf{z}_i^{(t)} = (1-t)\mathbf{z}_i^{(A)} + t\mathbf{z}_i^{(B)}$. The interpolated set $\mathbf{Z}^{(t)} = \{\mathbf{z}_1^{(t)}, \dots, \mathbf{z}_n^{(t)}\}$ is then decoded into an all-atom structure with the pretrained SLAE decoder (Figure 1C).

For two proteins with known conformational changes, adenylate kinase (AdK) and KaiB, we linearly interpolate between the SLAE embedding of the two experimentally determined states(AdK: 1AKE, 4AKE; KaiB: 2QKE, 5JYT). We sample intermediate structures from 50 evenly spaced values of t and align their backbone coordinates to frames in MD simulation of the transitions (Seyler et al., 2015; Zhang et al., 2024). For AdK, the interpolated structures closely track the MD intermediates, as evidenced by smooth trajectories with low RMSD (Figure 4), and they agree better than interpolations from the generative model (App. Fig. 10). Notably, these interpolations are *unguided by any energy function or model likelihood*; they arise solely from linear paths in SLAE latent space anchored in pretraining with physics-based task. KaiB shows higher RMSD between steps 20 and 30 (Figure 4). Closer examination of the interpolated structures (App. Fig 9) reveals disagreement in the C-terminus, which is known to unfold during transition (Wayment-Steele et al., 2023). This degradation is expected as SLAE is pretrained on folded structures and thus treats unfolded segments as out-of-distribution, where local environment cues under-constrain reconstruction.

Within the folded structure regime, SLAE's latent space is sufficiently regular that simple linear paths often decode to geometrically coherent intermediates aligned with MD trajectories. These results support the view that SLAE embeddings approximate a continuous, chemically grounded manifold of protein structures. The latent space reflects local environmental variation while accommodating large-scale transitions, make it useful for downstream analysis and generative applications.

6 CONCLUSION

We introduced SLAE, a framework tailored to learning general-purpose representations of proteins at all-atom resolution. SLAE applies a strictly local graph neural network over atomic environments, using computationally simple layers to perform expressive geometric reasoning on atom-type and interatomic distance features. Pretraining is driven by a novel objective that combines full atomic coordinate reconstruction with energy score regression, yielding embeddings that are structurally faithful, chemically grounded, and energetically informed.

REPRODUCIBILITY STATEMENT

The model architecture, training objectives, and hyperparameters are specified in Appendix B. All datasets used in this work are described in detail in Appendix D, including preprocessing pipelines, filtering thresholds, and dataset splits. We fixed random seeds where applicable and followed standard evaluation protocols to minimize nondeterminism and ensure fair comparison to baselines. All training and experiments can be reproduced on a single NVIDIA A100 GPU or equivalent hardware. Evaluation metrics and procedures are fully described in the main text and appendix. We plan to release code and pretrained checkpoints upon archiving of this work to enable full reproduction of our results.

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, June 2024. ISSN 1476-4687. doi: 10.1038/s41586-024-07487-w.
- Ivan Anishchenko, Yakov Kipnis, Indrek Kalvet, Guangfeng Zhou, Rohith Krishna, Samuel J. Pellock, Anna Lauko, Gyu Rie Lee, Linna An, Justas Dauparas, Frank DiMaio, and David Baker. Modeling protein-small molecule conformational ensembles with ChemNet, September 2024.
- Ilyes Batatia, Dávid Péter Kovács, Gregor N. C. Simm, Christoph Ortner, and Gábor Csányi. MACE: Higher Order Equivariant Message Passing Neural Networks for Fast and Accurate Force Fields, January 2023.
- Simon Batzner, Albert Musaelian, Lixin Sun, Mario Geiger, Jonathan P. Mailoa, Mordechai Kornbluth, Nicola Molinari, Tess E. Smidt, and Boris Kozinsky. E(3)-equivariant graph neural networks for data-efficient and accurate interatomic potentials. *Nature Communications*, 13(1):2453, May 2022. ISSN 2041-1723. doi: 10.1038/s41467-022-29939-5.
- Lasse M Blaabjerg, Maher M Kassem, Lydia L Good, Nicolas Jonsson, Matteo Cagiada, Kristoffer E Johansson, Wouter Boomsma, Amelie Stein, and Kresten Lindorff-Larsen. Rapid protein stability prediction using deep learning representations. *eLife*, 12:e82593, May 2023. ISSN 2050-084X. doi: 10.7554/eLife.82593.
- Meital Bojan, Sanketh Vedula, Advaith Maddipatla, Nadav Bojan Sellam, Federico Napoli, Paul Schanda, and Alex M. Bronstein. Representing local protein environments with atomistic foundation models, June 2025.
- Alexander E Chu, Jinho Kim, Lucy Cheng, Gina El Nesr, Minkai Xu, Richard W Shuai, and Po-Ssu Huang. An all-atom protein generative model. *Proceedings of the National Academy of Sciences*, 121(27):e2311500121, 2024.
- Henry Dieckhaus, Michael Brocidiacono, Nicholas Z. Randolph, and Brian Kuhlman. Transfer learning to leverage larger datasets for improved prediction of protein stability changes. *Proceedings of the National Academy of Sciences*, 121(6):e2314853121, February 2024. doi: 10.1073/pnas.2314853121.
- Frances Ding and Jacob Steinhardt. Protein language models are biased by unequal sequence sampling across the tree of life, March 2024.
- Ralf Drautz. Atomic cluster expansion for accurate and transferable interatomic potentials. *Physical Review B*, 99(1), 2019. doi: 10.1103/PhysRevB.99.014104.

Naomi K. Fox, Steven E. Brenner, and John-Marc Chandonia. SCOPe: Structural Classification of Proteins—extended, integrating SCOP and ASTRAL data and classification of new structures. *Nucleic Acids Research*, 42(Database issue):D304–309, January 2014. ISSN 1362-4962. doi: 10.1093/nar/gkt1240.

Zhangyang Gao, Cheng Tan, Jue Wang, Yufei Huang, Lirong Wu, and Stan Z. Li. FoldToken: Learning Protein Language via Vector Quantization and Beyond, March 2024.

- Johannes Gasteiger, Janek Groß, and Stephan Günnemann. Directional Message Passing for Molecular Graphs, April 2022.
- Tomas Geffner, Kieran Didi, Zhonglin Cao, Danny Reidenbach, Zuobai Zhang, Christian Dallago, Emine Kucukbenli, Karsten Kreis, and Arash Vahdat. La-Proteina: Atomistic Protein Generation via Partially Latent Flow Matching, July 2025.
- Thomas Hayes, Roshan Rao, Halil Akin, Nicholas J. Sofroniew, Deniz Oktay, Zeming Lin, Robert Verkuil, Vincent Q. Tran, Jonathan Deaton, Marius Wiggert, Rohil Badkundri, Irhum Shafkat, Jun Gong, Alexander Derry, Raul S. Molina, Neil Thomas, Yousuf A. Khan, Chetan Mishra, Carolyn Kim, Liam J. Bartie, Matthew Nemeth, Patrick D. Hsu, Tom Sercu, Salvatore Candido, and Alexander Rives. Simulating 500 million years of evolution with a language model, December 2024.
- Pedro Hermosilla, Marco Schäfer, Matěj Lang, Gloria Fackelmann, Pere Pau Vázquez, Barbora Kozlíková, Michael Krone, Tobias Ritschel, and Timo Ropinski. Intrinsic-Extrinsic Convolution and Pooling for Learning on 3D Protein Structures, April 2021.
- Jeffrey C Hoch, Kumaran Baskaran, Harrison Burr, John Chin, Hamid R Eghbalnia, Toshimichi Fujiwara, Michael R Gryk, Takeshi Iwata, Chojiro Kojima, Genji Kurisu, Dmitri Maziuk, Yohei Miyanoiri, Jonathan R Wedell, Colin Wilburn, Hongyang Yao, and Masashi Yokochi. Biological Magnetic Resonance Data Bank. *Nucleic Acids Research*, 51(D1):D368–D376, January 2023. ISSN 0305-1048. doi: 10.1093/nar/gkac1050.
- Jie Hou, Badri Adhikari, and Jianlin Cheng. DeepSF: Deep convolutional neural network for mapping protein sequences to folds. *Bioinformatics*, 34(8):1295–1303, April 2018. ISSN 1367-4803. doi: 10.1093/bioinformatics/btx780.
- John Ingraham, Vikas Garg, Regina Barzilay, and Tommi Jaakkola. Generative Models for Graph-Based Protein Design. In *Advances in Neural Information Processing Systems*, volume 32. Curran Associates, Inc., 2019.
- John B. Ingraham, Max Baranov, Zak Costello, Karl W. Barber, Wujie Wang, Ahmed Ismail, Vincent Frappier, Dana M. Lord, Christopher Ng-Thow-Hing, Erik R. Van Vlack, Shan Tie, Vincent Xue, Sarah C. Cowles, Alan Leung, João V. Rodrigues, Claudio L. Morales-Perez, Alex M. Ayoub, Robin Green, Katherine Puentes, Frank Oplinger, Nishant V. Panwar, Fritz Obermeyer, Adam R. Root, Andrew L. Beam, Frank J. Poelwijk, and Gevorg Grigoryan. Illuminating protein space with a programmable generative model. *Nature*, 623(7989):1070–1078, November 2023. ISSN 1476-4687. doi: 10.1038/s41586-023-06728-8.
- Arian R. Jamasb, Alex Morehead, Chaitanya K. Joshi, Zuobai Zhang, Kieran Didi, Simon V. Mathis, Charles Harris, Jian Tang, Jianlin Cheng, Pietro Lio, and Tom L. Blundell. Evaluating representation learning on the protein structure universe, June 2024.
- Bowen Jing, Stephan Eismann, Patricia Suriana, Raphael J. L. Townshend, and Ron Dror. Learning from Protein Structure with Geometric Vector Perceptrons, May 2021.
- 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, August 2021. ISSN 1476-4687. doi: 10.1038/s41586-021-03819-2.

- Mingchen Li, Pan Tan, Xinzhu Ma, Bozitao Zhong, Huiqun Yu, Ziyi Zhou, Wanli Ouyang, Bingxin
 Zhou, Liang Hong, and Yang Tan. ProSST: Protein Language Modeling with Quantized Structure
 and Disentangled Attention, May 2024.
 - 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, March 2023. doi: 10.1126/science.ade2574.
 - Huaqing Liu, Peiyi Chen, Xiaochen Zhai, Ku-Geng Huo, Shuxian Zhou, Lanqing Han, and Guoxin Fan. PPB-Affinity: Protein-Protein Binding Affinity dataset for AI-based protein drug discovery. *Scientific Data*, 11(1):1316, December 2024. ISSN 2052-4463. doi: 10.1038/ s41597-024-03997-4.
 - Jun Liu, Hungyu Chen, and Yang Zhang. A Corporative Language Model for Protein–Protein Interaction, Binding Affinity, and Interface Contact Prediction, July 2025. ISSN 2692-8205.
 - Amy X. Lu, Wilson Yan, Kevin K. Yang, Vladimir Gligorijevic, Kyunghyun Cho, Pieter Abbeel, Richard Bonneau, and Nathan Frey. Tokenized and Continuous Embedding Compressions of Protein Sequence and Structure, August 2024.
 - Tianyu Lu, Melissa Liu, Yilin Chen, Jinho Kim, and Po-Ssu Huang. Assessing generative model coverage of protein structures with SHAPES. *Cell Systems*, 16(8):101347, August 2025a. ISSN 2405-4712. doi: 10.1016/j.cels.2025.101347.
 - Tianyu Lu, Richard Shuai, Petr Kouba, Zhaoyang Li, Yilin Chen, Akio Shirali, Jinho Kim, and Po-Ssu Huang. Conditional Protein Structure Generation with Protpardelle-1C, August 2025b. ISSN 2692-8205.
 - Joshua Meier, Roshan Rao, Robert Verkuil, Jason Liu, Tom Sercu, and Alexander Rives. Language models enable zero-shot prediction of the effects of mutations on protein function. In *Proceedings of the 35th International Conference on Neural Information Processing Systems*, NIPS '21, pp. 29287–29303, Red Hook, NY, USA, December 2021. Curran Associates Inc. ISBN 978-1-7138-4539-3.
 - Antonio Mirarchi, Toni Giorgino, and Gianni De Fabritiis. mdCATH: A Large-Scale MD Dataset for Data-Driven Computational Biophysics. *Scientific Data*, 11(1):1299, November 2024. ISSN 2052-4463. doi: 10.1038/s41597-024-04140-z.
 - Albert Musaelian, Simon Batzner, Anders Johansson, Lixin Sun, Cameron J. Owen, Mordechai Kornbluth, and Boris Kozinsky. Learning local equivariant representations for large-scale atomistic dynamics. *Nature Communications*, 14(1):579, February 2023. ISSN 2041-1723. doi: 10.1038/s41467-023-36329-y.
 - Corrado Pancotti, Silvia Benevenuta, Giovanni Birolo, Virginia Alberini, Valeria Repetto, Tiziana Sanavia, Emidio Capriotti, and Piero Fariselli. Predicting protein stability changes upon single-point mutation: A thorough comparison of the available tools on a new dataset. *Briefings in Bioinformatics*, 23(2):bbab555, March 2022. ISSN 1477-4054. doi: 10.1093/bib/bbab555.
 - Hahnbeom Park, Philip Bradley, Per Jr. Greisen, Yuan Liu, Vikram Khipple Mulligan, David E. Kim, David Baker, and Frank DiMaio. Simultaneous Optimization of Biomolecular Energy Functions on Features from Small Molecules and Macromolecules. *Journal of Chemical Theory and Computation*, 12(12):6201–6212, December 2016. ISSN 1549-9618. doi: 10.1021/acs.jctc.6b00819.
 - Zihan Pengmei, Zhengyuan Shen, Zichen Wang, Marcus Collins, and Huzefa Rangwala. Pushing the Limits of All-Atom Geometric Graph Neural Networks: Pre-Training, Scaling and Zero-Shot Transfer, October 2024.
 - Sean L. Seyler, Avishek Kumar, M. F. Thorpe, and Oliver Beckstein. Path Similarity Analysis: A Method for Quantifying Macromolecular Pathways. *PLOS Computational Biology*, 11(10): e1004568, October 2015. ISSN 1553-7358. doi: 10.1371/journal.pcbi.1004568.

- Jianlin Su, Yu Lu, Shengfeng Pan, Ahmed Murtadha, Bo Wen, and Yunfeng Liu. RoFormer: Enhanced Transformer with Rotary Position Embedding, November 2023a.
 - Jin Su, Chenchen Han, Yuyang Zhou, Junjie Shan, Xibin Zhou, and Fajie Yuan. SaProt: Protein Language Modeling with Structure-aware Vocabulary, October 2023b.
 - Wen Torng and Russ B. Altman. 3D deep convolutional neural networks for amino acid environment similarity analysis. *BMC Bioinformatics*, 18(1):302, June 2017. ISSN 1471-2105. doi: 10.1186/s12859-017-1702-0.
 - Kotaro Tsuboyama, Justas Dauparas, Jonathan Chen, Elodie Laine, Yasser Mohseni Behbahani, Jonathan J. Weinstein, Niall M. Mangan, Sergey Ovchinnikov, and Gabriel J. Rocklin. Megascale experimental analysis of protein folding stability in biology and design. *Nature*, 620(7973): 434–444, August 2023. ISSN 1476-4687. doi: 10.1038/s41586-023-06328-6.
 - Aaron van den Oord, Oriol Vinyals, and Koray Kavukcuoglu. Neural Discrete Representation Learning, May 2018.
 - Limei Wang, Haoran Liu, Yi Liu, Jerry Kurtin, and Shuiwang Ji. Learning Hierarchical Protein Representations via Complete 3D Graph Networks, March 2023.
 - Shih-Hsin Wang, Yuhao Huang, Justin M. Baker, Yuan-En Sun, Qi Tang, and Bao Wang. A Theoretically-Principled Sparse, Connected, and Rigid Graph Representation of Molecules. In *The Thirteenth International Conference on Learning Representations*, October 2024.
 - Hannah K. Wayment-Steele, Adedolapo Ojoawo, Renee Otten, Julia M. Apitz, Warintra Pitsawong, Marc Hömberger, Sergey Ovchinnikov, Lucy Colwell, and Dorothee Kern. Predicting multiple conformations via sequence clustering and AlphaFold2. *Nature*, pp. 1–3, November 2023. ISSN 1476-4687. doi: 10.1038/s41586-023-06832-9.
 - Eli N. Weinstein, Alan N. Amin, Jonathan Frazer, and Debora S. Marks. Non-identifiability and the Blessings of Misspecification in Models of Molecular Fitness, February 2023.
 - Lijun Yu, Jose Lezama, Nitesh Bharadwaj Gundavarapu, Luca Versari, Kihyuk Sohn, David Minnen, Yong Cheng, Agrim Gupta, Xiuye Gu, Alexander G. Hauptmann, Boqing Gong, Ming-Hsuan Yang, Irfan Essa, David A. Ross, and Lu Jiang. Language Model Beats Diffusion Tokenizer is key to visual generation. In *The Twelfth International Conference on Learning Representations*, October 2023.
 - Sheheryar Zaidi, Michael Schaarschmidt, James Martens, Hyunjik Kim, Yee Whye Teh, Alvaro Sanchez-Gonzalez, Peter Battaglia, Razvan Pascanu, and Jonathan Godwin. Pre-training via Denoising for Molecular Property Prediction, October 2022.
 - Ning Zhang, Damini Sood, Spencer C. Guo, Nanhao Chen, Adam Antoszewski, Tegan Marianchuk, Supratim Dey, Yunxian Xiao, Lu Hong, Xiangda Peng, Michael Baxa, Carrie Partch, Lee-Ping Wang, Tobin R. Sosnick, Aaron R. Dinner, and Andy LiWang. Temperature-dependent fold-switching mechanism of the circadian clock protein KaiB. *Proceedings of the National Academy of Sciences*, 121(51):e2412327121, December 2024. doi: 10.1073/pnas.2412327121.
 - Zuobai Zhang, Minghao Xu, Arian Jamasb, Vijil Chenthamarakshan, Aurelie Lozano, Payel Das, and Jian Tang. Protein Representation Learning by Geometric Structure Pretraining, January 2023.
 - He Zhu, Lingyue Hu, Yu Yang, and Zhong Chen. A novel approach to protein chemical shift prediction from sequences using a protein language model. *Digital Discovery*, 4(2):331–337, 2025. doi: 10.1039/D4DD00367E.

ADDITIONAL RELATED WORKS

All-atom protein generative models which simultaneously generate backbone and sidechain coordinates can also have an all-atom representation of protein structure. Protpardelle can be cast as a continuous normalizing flow to generate deterministic latent encodings of all-atom protein structures (Chu et al. (2024)). A joint embedding space of sequence and all-atom structure was proposed in CHEAP (Lu et al., 2024), in which the embeddings reconstruct all-atom protein structures and recover sequence. However, interpolation between two conformations of the same protein sequence is not possible as identical sequence would map to the same CHEAP embedding. Representations can also be derived from protein structure prediction models such as AlphaFold3 (Abramson et al., 2024), but the information is distributed across layers and in both single and pairwise representations.

Model В

AUTOENCODER PSEUDOCODE

The end-to-end SLAE autoencoder can be summarized as follows:

Algorithm 1 SLAE Autoencoder. $h^{(e)}$: edge features, \mathcal{E} : SE(3)-equivariant update, \mathcal{P} : pooling, \mathcal{D}_{Tr} : Transformer decoder. Outputs: $\hat{\mathbf{x}}$ (coordinates), $\hat{\mathbf{s}}$ (sequence), $\hat{\mathbf{r}}$ (energies).

```
1: Input: heavy-atom coordinates \{a_i\}_{i=1}^N
```

2: Build G = (V, E) with cutoff r_c

3: Init $\boldsymbol{h}^{(e)} \leftarrow (\phi_r, \phi_a)$

4: for
$$L=1$$
 to 2 do
5: $\{x_{ij}^L, \mathbf{V}_{ij}^L\} \leftarrow \mathcal{E}(\{x_{ij}^{L-1}, \mathbf{V}_{ij}^{L-1}\})$
6: end for

7: $\{\boldsymbol{z}_r\} \leftarrow \mathcal{P}(\{x_{ij}^L\})$

8: $(\hat{\boldsymbol{x}}, \hat{\boldsymbol{s}}, \hat{\boldsymbol{r}}) \leftarrow \mathcal{D}_{\mathrm{Tr}}(\{\boldsymbol{z}_r\})$

B.2 ENCODER ARCHITECTURE

Notation Let $a_i \in \mathbb{R}^3$ be the coordinate of atom $i, r_{ij} = a_j - a_i, r_{ij} = \|r_{ij}\|, \widehat{r}_{ij} = r_{ij}/r_{ij}$. The neighbor set of i is $\mathcal{N}(i) = \{j \mid r_{ij} \leq r_c\}$. Each directed edge (i,j) maintains invariant scalars $x_{ij}^L \in \mathbb{R}^{d_{\text{sc}}}$ and equivariant tensors \mathbf{V}_{ij}^L .

Two-body initialization Edge features are initialized with radial and angular bases:

$$x_{ij}^{0} = u(r_{ij}) \cdot \text{MLP}_{2b}(\text{onehot}(Z_i) \parallel \text{onehot}(Z_j) \parallel \phi_r(r_{ij})), \qquad (2)$$

$$\mathbf{V}_{ij}^{0} = \omega_{ij} \cdot \phi_a(\widehat{\mathbf{r}}_{ij}), \quad \omega_{ij} = \mathrm{MLP}_{\omega}(x_{ij}^0), \tag{3}$$

where ϕ_T are Bessel radial basis functions, ϕ_a are angular embeddings (e.g., spherical harmonics), and $u(r_{ij})$ is a smooth cutoff envelope.

Tensor product update At layer L, equivariant features of edge (i, j) interact with the embedded environment of atom i:

$$\mathbf{V}_{ij}^{L} = \mathbf{V}_{ij}^{L-1} \otimes \left(\sum_{k \in \mathcal{N}(i)} w_{ik}^{L} \phi(\mathbf{r}_{ik}) \right), \tag{4}$$

where $\phi(\mathbf{r}_{ik})$ encodes neighbor geometry and $w_{ik}^L = \mathrm{MLP_{embed}^L}(x_{ik}^{L-1})$ are learned weights. This corresponds to a weighted projection of the atomic density around atom i.

Latent scalar update. Scalar channels are updated with tensor product scalars:

$$x_{ij}^L = \text{MLP}_{\text{latent}}^L \left(x_{ij}^{L-1} \parallel \mathbf{V}_{ij}^L \right) \cdot u(r_{ij}),$$
 (5)

injecting geometric information from V_{ij}^L back into x_{ij}^L .

Hierarchical pooling Final edge scalars are aggregated:

$$s_i = \frac{1}{|\mathcal{N}(i)|} \sum_{j \in \mathcal{N}(i)} x_{ij}^L, \tag{6}$$

$$z_r = \frac{1}{|\mathcal{A}(r)|} \sum_{i \in \mathcal{A}(r)} s_i, \quad z_r \in \mathbb{R}^{128},$$
 (7)

producing residue-level embeddings $\{z_r\}$.

B.3 DECODER ARCHITECTURE

Transformer backbone We employ a standard pre-norm Transformer encoder with Rotary Positional Embeddings (RoPE) with $L_{\rm Tr}=8$ layers, h=16 heads, model width $d_{\rm model}=1024$. Each layer consists of:

- Multi-head self-attention with RoPE (pre-norm): MHA_{RoPE}(LayerNorm(·)).
- Residual connection.
- Feed-forward network with hidden dimension $d_{\rm ff}$ and SwiGLU, applied as FFN(LayerNorm(·)).
- No dropout.

Formally:

$$\mathbf{H}^{(\ell)} = \mathrm{MHA}_{\mathrm{RoPE}} \left(\mathrm{LayerNorm}(\mathbf{H}^{(\ell-1)}) \right) + \mathbf{H}^{(\ell-1)}, \tag{8}$$

$$\mathbf{H}^{(\ell)} = FFN\left(LayerNorm(\mathbf{H}^{(\ell)})\right) + \mathbf{H}^{(\ell)},\tag{9}$$

for $\ell = 1, \dots, L_{\text{Tr}}$, with $\mathbf{H}^{(0)} = [\boldsymbol{z}_1, \dots, \boldsymbol{z}_n]$.

Prediction heads From final hidden states $\mathbf{H} \in \mathbb{R}^{n \times d_{\text{model}}}$ ($d_{\text{model}} = 1024$), we apply three parallel heads:

(i) 3D coordinates (linear head) LayerNorm + Linear maps per-residue embeddings to all Atom37 coordinates:

$$\hat{\boldsymbol{x}} = \text{Unflatten}(\text{Linear}(\text{LN}(\mathbf{H})), 37 \times 3) \in \mathbb{R}^{n \times 37 \times 3}.$$

(Atoms 1–4 are N, C_{α} , C, and O; atoms 5–37 are side chain. Masking is applied via the Atom37 mask.)

(ii) Sequence logits on valid tokens An MLP head operates only on valid tokens (mask-compacted), then is re-padded for loss:

$$\hat{s} = \text{MLP}_{\text{seq}}(\mathbf{H}_{\text{valid}}) \in \mathbb{R}^{n_{\text{valid}} \times 20}$$
.

(iii) Pairwise energies A pairwise feature head first down-projects H, lifts to 2D by pairwise product/difference, applies a small MLP, then per-type linear heads with magnitude clamp to 1e-3:

$$\hat{\boldsymbol{r}} = \left[\hat{\boldsymbol{r}}^{\text{hbond}}, \, \hat{\boldsymbol{r}}^{\text{sol}}, \, \hat{\boldsymbol{r}}^{\text{elec}}\right] \in \mathbb{R}^{n \times n \times 3}.$$

B.4 TASK-SPECIFIC HEADS

Trainable decoder backbone. We expose a lightweight wrapper over the $\tt DecoderBackbone$ to enable fine-tuning the last N Transformer blocks while freezing the rest. Take the single site mutation stability task as an example, we document the layout of downstream task-specific finetuning here.

Contrastive and site-aware head A Siamese head takes two or more structure embeddings (e.g., wild-type and mutant), runs them through the shared DecoderBackbone, and regresses a scalar target (e.g., $\Delta\Delta G$). Beyond global contrastive pooling, it can extract *site-specific* residue representations, enabling residue-level tasks.

Backbone embeddings. Given masked inputs $(\mathbf{X}^{\text{wt}}, \mathbf{M}^{\text{wt}})$ and $(\mathbf{X}^{\text{mut}}, \mathbf{M}^{\text{mut}})$,

$$\mathbf{H}^{\mathrm{wt}} = \mathtt{DecoderBackbone}(\mathbf{X}^{\mathrm{wt}}, \mathbf{M}^{\mathrm{wt}}), \qquad \mathbf{H}^{\mathrm{mut}} = \mathtt{DecoderBackbone}(\mathbf{X}^{\mathrm{mut}}, \mathbf{M}^{\mathrm{mut}}).$$

Mask-aware pooling and site features. Let the mean-pooling operator be

$$\operatorname{Pool}(\mathbf{H}, \mathbf{M}) = \frac{\sum_{t=1}^{L} \mathbf{H}_{:,t,:} \mathbf{M}_{:,t}}{\sum_{t=1}^{L} \mathbf{M}_{:,t} + \varepsilon} \in \mathbb{R}^{B \times 1024}.$$

We form global embeddings $\mathbf{z}^{\mathrm{wt}} = \mathrm{Pool}(\mathbf{H}^{\mathrm{wt}}, \mathbf{M}^{\mathrm{wt}})$ and $\mathbf{z}^{\mathrm{mut}} = \mathrm{Pool}(\mathbf{H}^{\mathrm{mut}}, \mathbf{M}^{\mathrm{mut}})$. Given mutation indices $\iota \in \{1, \ldots, L\}^B$, we also extract site embeddings

$$\mathbf{s}^{\text{wt}} = \mathbf{H}^{\text{wt}}[\text{range}(B), \boldsymbol{\iota}], \quad \mathbf{s}^{\text{mut}} = \mathbf{H}^{\text{mut}}[\text{range}(B), \boldsymbol{\iota}] \in \mathbb{R}^{B \times 1024}$$

Contrastive feature and MLP regressor. We concatenate global and site representations together with their difference:

$$\mathbf{u} = \left[\left. \mathbf{z}^{\text{wt}}, \, \left. \mathbf{z}^{\text{mut}}, \, \mathbf{s}^{\text{wt}}, \, \mathbf{s}^{\text{mut}}, \, \mathbf{s}^{\text{mut}} - \mathbf{s}^{\text{wt}} \, \right] \right. \in \, \mathbb{R}^{B \times (5 \cdot 1024)}$$

A small MLP head predicts a scalar per pair:

$$\hat{\mathbf{y}} = \text{MLP}(\mathbf{u}) = \text{Linear} \circ \text{GELU} \circ \text{LayerNorm} \circ \text{Linear} \circ \text{GELU}(\mathbf{u}) \in \mathbb{R}^{B \times 1}.$$

General usage The same interface supports other pairwise or single-input tasks by (i) choosing one or multiple passes through <code>DecoderBackbone</code>, (ii) selecting global vs. site-wise features, and (iii) swapping the final MLP for the appropriate output dimensionality/loss. For atom-level tasks, the <code>DecoderBackbone</code> can be reinitialized with the smaller attention window.

C TRAINING

C.1 Losses

All-atom FAPE (Frame-Aligned Point Error) All-atom FAPE is computed by aligning the predicted and reference structures on every triplet of bonded atoms (i, j, k) (with the exception of symmetric side chain atoms) and then measuring per-atom positional deviations between the aligned structures. For each frame f(i, j, k) (with j as the origin), define an orthonormal basis for predicted/true coordinates via a deterministic map $\Phi: (\mathbb{R}^3)^3 \to SO(3)$:

$$\mathbf{U}_f^{\mathrm{pred}} = \Phi(\hat{\mathbf{x}}_i, \hat{\mathbf{x}}_j, \hat{\mathbf{x}}_k), \qquad \mathbf{U}_f^{\mathrm{true}} = \Phi(\mathbf{x}_i^{\star}, \mathbf{x}_j^{\star}, \mathbf{x}_k^{\star}),$$

where Φ constructs column vectors from the two edge directions at i,

$$\mathbf{v}_{j\to i} = \mathbf{x}_i - \mathbf{x}_j, \quad \mathbf{v}_{j\to k} = \mathbf{x}_k - \mathbf{x}_j,$$

then

$$\mathbf{e}_0 = (\mathbf{v}_{j \to k}) \times (\mathbf{v}_{j \to i}), \quad \mathbf{e}_2 = \mathbf{v}_{j \to i} - \mathbf{v}_{j \to k}, \quad \mathbf{e}_1 = \mathbf{e}_2 \times \mathbf{e}_0,$$

and column-normalizes $[e_0, e_1, e_2]$ to obtain a right-handed 3×3 matrix.

For any atom a in the same protein as f, rotate origin-subtracted positions into the local frames:

$$\mathbf{r}_{f,a}^{\mathrm{pred}} = \mathbf{U}_f^{\mathrm{pred}} \big(\hat{\mathbf{x}}_a - \hat{\mathbf{x}}_j \big), \qquad \mathbf{r}_{f,a}^{\mathrm{true}} = \mathbf{U}_f^{\mathrm{true}} \big(\mathbf{x}_a^{\star} - \mathbf{x}_j^{\star} \big).$$

Define $d_{f,a} = \|\mathbf{r}_{f,a}^{\text{pred}} - \mathbf{r}_{f,a}^{\text{true}}\|_2$, clamped at $c = 10 \,\text{Å}$ as $\tilde{d}_{f,a} = \min(d_{f,a}, c)$, and apply a Huber penalty with $\delta = 1.0$:

$$\rho_{\delta}(\tilde{d}) = \begin{cases} \frac{1}{2}\tilde{d}^2, & \tilde{d} \leq \delta, \\ \delta \tilde{d} - \frac{1}{2}\delta^2, & \tilde{d} > \delta. \end{cases}$$

We average first over frames and then over atoms, yielding an atom-weighted mean:

$$\mathcal{L}_{\text{FAPE}} = \frac{1}{B} \sum_{b=1}^{B} \frac{1}{|\mathcal{A}_b|} \sum_{a \in \mathcal{A}_b} \left(\frac{1}{|\mathcal{F}_b|} \sum_{f \in \mathcal{F}_b} \rho_{\delta}(\tilde{d}_{f,a}) \right).$$

All-atom smooth LDDT We use a differentiable, all-atom version of LDDT that compares pairwise distances within a cutoff. Let $\mathcal{P} = \{(i,a),(j,b)\}$ be all heavy-atom pairs with $\|\mathbf{x}_{i,a}^{\star} - \mathbf{x}_{j,b}^{\star}\| \leq R_{\max}$ and not in the same residue. Define ground-truth and predicted distances $d_{iabj}^{\star} = \|\mathbf{x}_{i,a}^{\star} - \mathbf{x}_{j,b}^{\star}\|$ and $\hat{d}_{iabj} = \|\hat{\mathbf{x}}_{i,a} - \hat{\mathbf{x}}_{j,b}\|$, and the absolute error $\Delta_{iabj} = |\hat{d}_{iabj} - d_{iabj}^{\star}|$. Using standard IDDT thresholds $\tau \in \{0.5, 1.0, 2.0, 4.0\}$ Å with smooth indicators $s_{\tau}(\Delta) = \sigma(\alpha(\tau - \Delta))$ (sigmoid, α controls sharpness),

$$\mathrm{sLDDT}_i = \frac{1}{|\mathcal{P}_i|} \sum_{(i,a),(j,b) \in \mathcal{P}_i} \ \frac{1}{|\mathrm{T}|} \sum_{\tau \in \mathrm{T}} s_\tau(\Delta_{iabj}), \qquad \mathcal{L}_{\mathrm{sLDDT}} = 1 - \frac{1}{N_{\mathrm{res}}} \sum_i \mathrm{sLDDT}_i.$$

Mean-squared error (MSE) Used for continuous targets regression:

$$\mathcal{L}_{\text{MSE}} = \frac{1}{|\Omega|} \sum_{u \in \Omega} \|\hat{y}_u - y_u^{\star}\|_2^2.$$

Huber loss. Used for continuous targets regression with $\delta = 1.35$:

$$\mathcal{L}_{\text{Huber}} = \frac{1}{|\Omega|} \sum_{u \in \Omega} \begin{cases} \frac{1}{2} \left(\hat{y}_u - y_u^{\star} \right)^2, & |\hat{y}_u - y_u^{\star}| \leq \delta, \\ \delta \left| \hat{y}_u - y_u^{\star} \right| - \frac{1}{2} \delta^2, & \text{otherwise.} \end{cases}$$

C.2 TRAINING SPECIFICS

The autoencoder is trained on a single NVIDIA A100 or H100 GPU using batch size 16. We train for 30 epochs with early stopping on validation loss not decreasing after 5 epochs. The learning rate schedule is linear warmup for 1,000 steps followed by cosine decay. Optimization uses AdamW with maximum learning rate 1×10^{-4} and standard β_1 =0.9, β_2 =0.999 (weight decay as in AdamW defaults). Unless noted otherwise, downstream task-specific fine-tuning uses the same batch size and maximum learning rate 1×10^{-5} .

D DATASETS

Pretraining Structure We train SLAE on an sequence-augmented CATH set (Lu et al., 2025b) by redesigning each domain with 32 ProteinMPNN sequences and predicting structures with ESMFold; we retain only high-confidence, self-consistent structure models ($pLDDT \geq 80, scRMSD \leq 2.0 \text{\AA}$), yielding 337936 structures, with 271 test structures from holdout CATH domains. We evaluate SLAE latent space on protein conformational ensembles sampled from the dataset of molecular dynamics (MD) simulations mdCATH (Mirarchi et al., 2024). We subsample 32 frames per protein across MD trajectory ensembles for each of the 5398 structures.

Pretraining Rosetta Score We use PyRosetta to compute residue pairwise energy scores for all pretraining structures under its default full-atom energy terms. For each pair of residues we compute (1) *fa_sol*: Lazaridis-Karplus solvation energy (2) *fa_elec*: Coulombic electrostatic potential with a distance-dependent dielectric (3) *hbond*: Sum of all hydrogen bonding terms for backbone and sidechain.

Fold Classification We obtain the dataset from Hermosilla et al. (2021), which consolidated 16,712 proteins with 1195 different folds from the SCOPe 1.75 database (Fox et al., 2014). Three test sets are used: (1) Family, which allows proteins from the same family to appear in both training and test; (2) Superfamily, which excludes proteins sharing family membership with the training set; and (3) Fold, which further excludes proteins from the same superfamily as those in training. All structures are obtained from the SCOPe 1.75 archive.

Stability We obtain the dataset curated by Dieckhaus et al. (2024) on Tsuboyama et al. (2023), composed of 272,712 single point mutations and their experimental $\Delta\Delta G$. The proteins were clustered using MMseqs2 with sequence identity cutoff of 25% to yield 239 training, 31 validation and 29 validation proteins. For wild type sequences we predict their structures with AlphaFold2. For all mutated structures we model the mutation with PyRosetta and relax within 8Å radius to obtain training structures.

Binding Affinity We use the PPB-AFfinity (Liu et al., 2024) which integrates experimental protein-protein binding affinity data from several source databases: SKEMPI v2.0, SAbDab, PDB-bind v2020, Affinity Benchmark v5.5, and ATLAS. This dataset contains 12062 unique binding complexes consisting of 3032 unique PDB codes and point mutations. We use the structures curated in the dataset and define interface residues as those within 5Å distance from other atoms of the neighboring chains. For all mutations we mutate the sidechain with PyRosetta and relax within 8Å radius to obtain training structures.

NMR Chemical Shift We retrieve the BMRB totaling 17,028 entries (2025-07-02) (Hoch et al., 2023). The entries were filtered and processed based on NMR experiment type, backbone chemical shift coverage, sequence consistency, basic experimental condition boundary plus any other routine re-referencing requirements. 3623 entries were retained and split into 2532 training and 594 validation entries at a 50% pairwise sequence-identity threshold after filtering entries without any nitrogen chemical shifts. Alphafold2 was used to generate all structures used in training.

MD Simulation For adenylate kinase (AdK), we use conformational ensembles generated using the Framework Rigidity Optimized Dynamics Algorithm (FRODA), yielding 200 trajectories (Seyler et al., 2015). For KaiB, we use the temperature-dependent fold-switching simulation from Zhang et al. (2024), subsampling every 10 frames out of the 4 successful fold-switching trajectories from the fold-switched state to ground state.

Rosetta Decoy To assess local residue environment embeddings distribution between native and decoy structure, we use structure dataset by Park et al. (2016), where each of the 133 native structures are accompanied with large numbers (≥ 1000 cluster centers) of alternative conformations (decoys).

E METRICS

Structure comparison We report RMSD after optimal Kabsch rigid alignment for $C\alpha$, backbone and all-atom. Given reference $\mathbf{X}^{\star} \in \mathbb{R}^{n \times 3}$ and prediction $\hat{\mathbf{X}}$, align $\hat{\mathbf{X}}$ to \mathbf{X}^{\star} then compute

$$\text{RMSD} = \sqrt{\frac{1}{n} \sum_{j=1}^{n} \left\| \hat{\mathbf{x}}_{j}^{\text{align}} - \mathbf{x}_{j}^{\star} \right\|_{2}^{2}}.$$

Numeric regression Given targets $\{y_i\}_{i=1}^N$ and predictions $\{\hat{y}_i\}_{i=1}^N$, we report

RMSE =
$$\sqrt{\frac{1}{N} \sum_{i=1}^{N} (\hat{y}_i - y_i)^2}$$
, $r = \frac{\sum_{i=1}^{N} (\hat{y}_i - \bar{\hat{y}})(y_i - \bar{y})}{\sqrt{\sum_{i=1}^{N} (\hat{y}_i - \bar{\hat{y}})^2} \sqrt{\sum_{i=1}^{N} (y_i - \bar{y})^2}}$.

Distribution comparison We compute Fréchet Protein Distance (FPD) following Lu et al. (2025a). Given N data points from a reference distribution $p_{\text{data}}(\mathbf{x})$, here the sequence-augmented CATH dataset, and M samples from a generative model $p_{\text{sample}}(\mathbf{x})$, we computed per-residue SLAE embeddings $\{\mathbf{z}_{\text{data}}^{(i)}\}_{i=1}^{N}$ and $\{\mathbf{z}_{\text{sample}}^{(j)}\}_{j=1}^{M}$ and then compute

$$FPD = \|\boldsymbol{\mu}_{data} - \boldsymbol{\mu}_{sample}\|_{2}^{2} + Tr\left(\boldsymbol{\Sigma}_{data} + \boldsymbol{\Sigma}_{sample} - 2\left(\boldsymbol{\Sigma}_{data}\boldsymbol{\Sigma}_{sample}\right)^{\frac{1}{2}}\right)$$
(10)

where μ_{data} and μ_{sample} are the means of the reference embeddings and the sample embeddings respectively, and Σ_{data} and Σ_{sample} are the covariance matrices of the reference embeddings and the sample embeddings respectively. We compute FPD using a smaller subset of 2000 samples as SHAPES showed that this is sufficient for an accurate FPD estimate Lu et al. (2025a).

F ADDITIONAL EXPERIMENTS AND RESULTS

F.1 PRETRAINING

We report in Table 6 additional results on the pretraining performance of the SLAE autoencoder. We note that encoders with 10Å graph radius cutoff is infeasible to train with a single GPU due to the number of edges.

Graph	Discretization	Codebook	Training	Seq. Acc.	RMSD	RMSD
Radius(Å)	Method	Size	Obj.	(%)	< 128 (Å)	< 512 (Å)
8	LFQ	16384	all	69.5	4.12	5.79
8	LFQ	32768	all	75.2	2.50	3.74
8	VQ	16384	all	65.7	5.02	5.88
8	VQ	32768	all	70.4	4.30	6.02
8	kNN	4096	all	97.5	2.96	4.03
8	kNN	16384	all	98.6	1.71	2.57
8	kNN	32768	all	99.4	1.60	2.31
8	_	_	w/o. FAPE	97.2	3.89	5.22
8	_	_	w/o. Energy	98.0	3.26	5.17
4	_	_	all	99.8	2.57	3.86
6	_	_	all	99.9	1.24	2.55
8	_	_	all	99.9	1.12	1.92

Table 6: Complete results of SLAE autoencoder ablation experiments.

F.2 LATENT SPACE CHARACTERIZATION

F.2.1 KNN CLUSTERING

We examine the CATH-kNN-quantized latent space, the k-means codebook of k=16,384 centroids. We assign each centroid the majority amino-acid identity among its members; the commitment loss is the L2 distance from an embedding to its assigned centroid. The commitment loss histogram is tightly concentrated around 3-5 L2 units (Figure 5), which is modest relative to the embedding norm (15 ± 4) , indicating that quantization preserves most geometric signal.

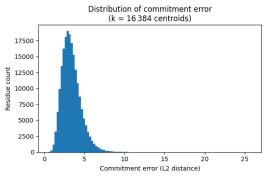


Figure 5: Commitment loss distribution during post-hoc quantization

We observe clear residue type mixing in the clusters. Although many centroids are quite pure (median majority fraction 0.96), the distribution is broad (mean 0.89 ± 0.15 ; entropy mean 0.52), with a substantial tail of mixed-composition clusters (10th-percentile majority 0.67). Along with the modest commitment error, this suggests that the observed mixing reflects genuinely overlapping local chemistries. Consistently, residue-conditioned intra-cluster distances show that some types form diffuse, mixed neighborhoods (A, G, S, C with ratios ≥ 1), while others are tighter and more typespecific (W, Y, R with ratios ≤ 1). These observations suggest that the kNN partitioning of residue embedding space yields chemically meaningful clusters but does not enforce one-residue exclusivity and captures real cross-type similarity in local environments.

Residue	# Centroids	Mean intra-cluster distance	Ratio to global distance
		mean \pm std	
A	1601	19.95 ± 6.50	[1.20]
C	215	17.67 ± 5.72	[1.06]
D	962	13.43 ± 7.47	[0.81]
E	1076	11.58 ± 3.97	[0.70]
F	641	11.44 ± 3.53	[0.69]
G	1192	18.38 ± 6.21	[1.11]
H	387	11.56 ± 3.53	[0.70]
I	899	14.33 ± 4.60	[0.86]
K	947	11.48 ± 3.76	[0.69]
L	1565	14.26 ± 4.63	[0.86]
M	272	13.69 ± 4.33	[0.82]
N	729	13.28 ± 4.16	[0.80]
P	737	13.83 ± 4.49	[0.83]
Q	492	11.48 ± 3.89	[0.69]
R	720	9.95 ± 3.14	[0.60]
S	1032	17.31 ± 5.58	[1.04]
T	920	15.81 ± 4.97	[0.95]
V	1253	15.96 ± 5.13	[0.96]
W	202	8.86 ± 2.76	[0.53]
Y	542	10.49 ± 3.16	[0.63]

Table 7: Residue-wise clustering statistics: number of centroids that each residue type dominates, mean intra-cluster distance (\pm standard deviation), and ratio relative to the global mean.

F.2.2 RESIDUE EMBEDDING VISUALIZATION

We project the 16,384-entry codebook (centroid) embeddings into three dimensions using UMAP and analyze how local chemical environments are organized in this latent space (Figs. 6–7). Each CATH residue is assigned to its nearest codebook entry, and for every centroid we aggregate properties across its assigned residues. We compute the mean SASA and the majority secondary-structure label. This yields a coarse-grained landscape in which centroids arrange along solvent-exposure gradients and segregate by secondary-structure preferences.

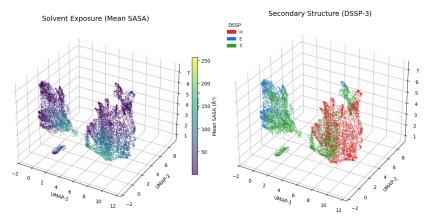


Figure 6: 3D UMAP projection of CATH residue embeddings colored by solvent accessibility and secondary structure

F.2.3 STRUCTURE ENSEMBLE ANALYSIS

Subsampled mdCATH For each residue, we measure how much its embedding changes across the ensemble by averaging pairwise differences between frames. For a given residue and set of frames, we compute two physical descriptors: *Contact-map change*: we form a binary contact row per frame (contact if residues are within a chosen distance threshold) and measure, on average, what fraction of those contacts differ between frames. *Solvent-exposure change*: we compute solvent-

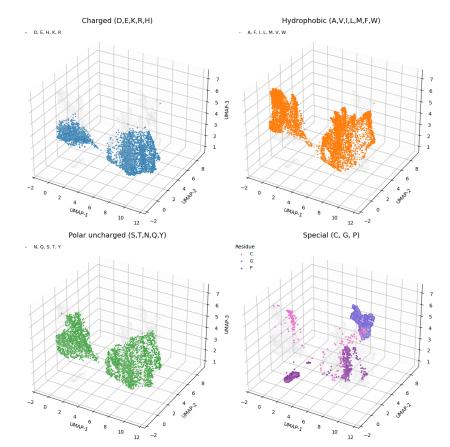


Figure 7: 3D UMAP projection of CATH residue embeddings colored by amino acid type

accessible surface area (SASA), convert to residue-type-normalized relative SASA, and take the average absolute change between frames. We fit a simple linear model that predicts per-residue embedding change from the two descriptors. We aggregate performance on held-out residues and report: (i) the proportion of variance explained and (ii) the Spearman rank correlation between observed and predicted embedding change.

Rosetta Decoys For each native protein we have a residue–embedding matrix and a set of its decoy matrices, aligned by residue index. We apply row-wise L2 normalization so that inner products equal cosine similarity. For a given protein, we compute the mean residue-wise cosine similarity between each decoy and its native, then take the average over decoys. The *native-decoy cosine margin* is defined as the difference between the native's self-similarity (equal to 1.0 after normalization) and this mean decoy similarity.

To test linear separability at the residue level and generalization to unseen proteins, we train a logistic-regression classifier on residue embeddings with leave-protein-out grouped cross-validation: each residue embedding is a sample (label 0=native, 1=decoy) and carries its protein ID for grouped CV. We split with GroupKFold so all residues from a held-out protein appear only in the test set, and train an L2-regularized LogisticRegression. On each test fold we report AUROC; metrics are aggregated as mean \pm sd across folds.

F.3 PER-RESIDUE GENERATIVE MODEL ASSESSMENT

We compare distribution coverage of all-atom chemical environments sampled by generative models, stratified by residue type. For each residue type, we extracted the SLAE embeddings of 2000 random examples from the sequence-augmented CATH dataset and from a collection of 20,000 unconditional samples of all-atom protein structures from La-Proteina, Protpardelle-1c, and Chroma.

The FPD metrics reveal subtle differences in the coverage of local amino acid environments by different generative models (Figure 8). For example, biased sampling is evident in La-Proteina samples for serine, threonine, and valine relative to Protpardelle-1c and Chroma. Using SLAE embeddings provides a more sensitive view on coverage of all-atom local environments which are ignored in backbone-only FPD and which may be averaged out on the global protein fold level as in previous assessments of generative model coverage of protein structures (Lu et al. (2025a)).

F.4 LATENT SPACE INTERPOLATION

In Figure 9 A and B we show 20 out of 50 interpolated structures for AdK and KaiB. In addition, we compare linearly interpolated AdK structures from the SLAE latent space to those from the all-atom generative model Protpardelle-1c (Figure 10) and show that SLAE interpolation is better matched to simulated intermediate structures.

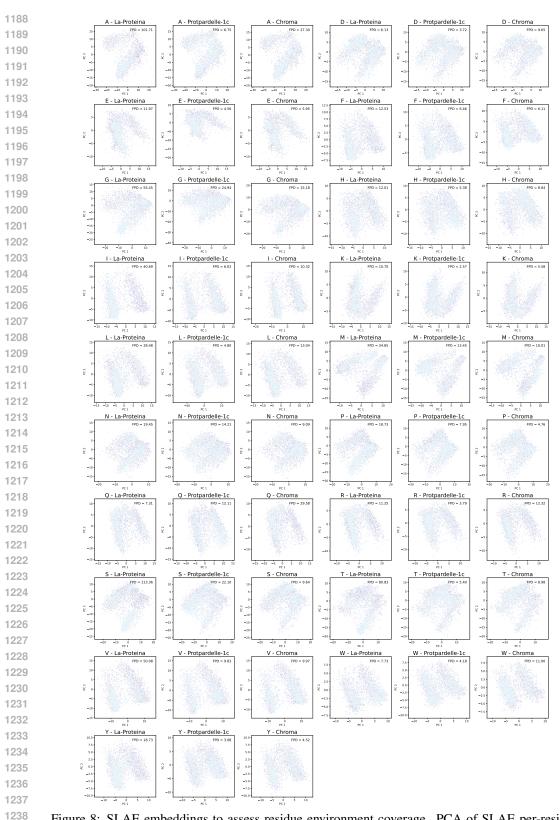


Figure 8: SLAE embeddings to assess residue environment coverage. PCA of SLAE per-residue embeddings of *de novo* structure samples (light blue) compared to the reference CATH distribution (purple) stratified by amino acid type given in the title. The two modes in each amino acid type correspond to residues belonging to a beta sheet or alpha helix.

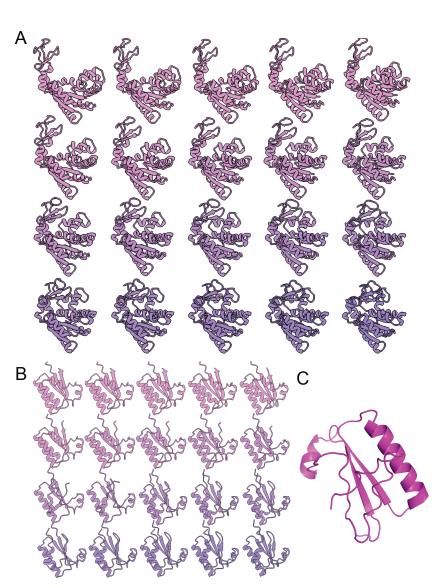


Figure 9: Structures decoded from SLAE latent interpolation. **A.** AdK **B.** KaiB **C.** Step 23 KaiB intermediate structure with under-characterized C-terminus showing disordered backbone collapsing onto itself.

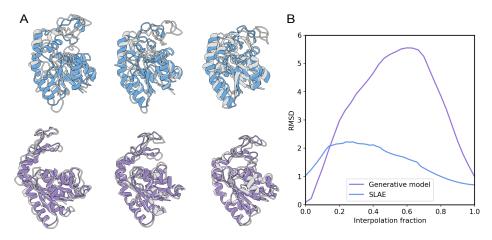


Figure 10: Comparison of SLAE and generative model (Protpardelle-1c) latent interpolation. **A.** Three representative steps from interpolation fraction 0.3 to 0.7. Top: Protpardelle-1c linear interpolation (blue) and best MD frame matches (grey). Bottom: SLAE linear interpolation (purple) and best MD frame matches (grey). **B.** RMSD of interpolation trajectories to their closest-match MD frames