

Appendix

A Diffusion Processes

A.1 Posteriors

Posterior of Amino Acid Types The generative diffusion kernel for amino acid types $p(s_j^{t-1} | \mathcal{R}^t, \mathcal{C})$ (Eq.3) should align to the posterior $q(s_j^{t-1} | s_j^t, s_j^0)$. It can be derived from Eq.1 and Eq.2 [Hoogeboom et al., 2021]:

$$q(s_j^{t-1} | s_j^t, s_j^0) = \text{Multinomial} \left(\left[\alpha_{\text{type}}^t \cdot \text{onehot}(s_j^t) + (1 - \alpha_{\text{type}}^t) \cdot \frac{1}{20} \cdot \mathbf{1} \right] \odot \left[\bar{\alpha}_{\text{type}}^{t-1} \cdot \text{onehot}(s_j^0) + (1 - \bar{\alpha}_{\text{type}}^{t-1}) \cdot \frac{1}{20} \cdot \mathbf{1} \right] \right). \quad (17)$$

The vector inside $\text{Multinomial}(\cdot)$ might not sum to one. In this case, the probability of a class is the ratio of the value in the sum of the vector.

Posterior of C_α Coordinates The generative diffusion kernel $p(\mathbf{x}_j^{t-1} | \mathcal{R}^t, \mathcal{C})$ (Eq.7) should align to the posterior obtained from Eq.5 and Eq.6 [Ho et al., 2020]:

$$q(\mathbf{x}_j^{t-1} | \mathbf{x}_j^t, \mathbf{x}_j^0) = \mathcal{N} \left(\mathbf{x}_j^{t-1} \mid \boldsymbol{\mu}_q(\mathbf{x}_j^t, \mathbf{x}_j^0), \frac{(1 - \bar{\alpha}_{\text{pos}}^{t-1})\beta_{\text{pos}}^t}{1 - \bar{\alpha}_{\text{pos}}^t} \mathbf{I} \right), \quad (18)$$

$$\text{where } \boldsymbol{\mu}_q(\dots) = \frac{\sqrt{\bar{\alpha}_{\text{pos}}^{t-1}}\beta_{\text{pos}}^t}{1 - \bar{\alpha}_{\text{pos}}^{t-1}} \mathbf{x}_j^0 + \frac{\sqrt{\alpha_{\text{pos}}^t}(1 - \bar{\alpha}_{\text{pos}}^{t-1})}{1 - \bar{\alpha}_{\text{pos}}^t} \mathbf{x}_j^t. \quad (19)$$

A.2 Amino Acid C_α Position Normalization

As amino acid C_α positions could be arbitrary in the 3D space. We need to normalize them to use the standard normal distribution with zero mean and unit variance as the prior distribution. First, we need to derive the statistics of CDR positions. For each CDR in the SAbDab dataset, we shift the overall structure such that the center point of the two CDR anchors is located in origin. Then, we aggregate C_α positions in the shifted CDRs. Finally, we calculate the mean and standard deviation of them. Before training and inference, we shift the whole structure according to their CDR anchors and further shift and scale the structure according to the pre-calculated mean and standard deviation to obtain the normalized coordinates.

B Distributions on SO(3)

B.1 Preliminary: Axis-Angle Representation of Rotations

Conventionally, a rotation is usually represented by 3 Euler angles (α, β, γ) , which can be interpreted as the composition of counter-clockwise rotations by α, β, γ about x, y, z axes. However, the Euler representation is unsuitable for defining useful operations and distributions w.r.t. rotations considered in this work. Alternatively, we introduce another rotation representation called *axis-angle representations*. This representation parameterized a rotation with an rotational axis \mathbf{u} ($\|\mathbf{u}\|_2 = 1$) and an angle θ ($\theta \in \mathbb{R}$).

B.2 Logarithm of Rotation Matrices and Exponential of Skew-Symmetric Matrices

Logarithm of Rotation Matrices Derived from the definition of matrix logarithm, the logarithm of a rotation matrix \mathbf{R} is a **skew-symmetric matrix** [Gallier and Xu, 2003], which can be represented as:

$$\mathbf{S} := \log \mathbf{R} = \begin{bmatrix} 0 & -v_z & v_y \\ v_z & 0 & -v_x \\ -v_y & v_x & 0 \end{bmatrix}. \quad (20)$$

It can be proven that $\mathbf{v} = [v_x, v_y, v_z]$ is the rotational axis of \mathbf{R} , and $\|\mathbf{v}\|_2$ is the rotational angle. For brevity, we can use the vector notation \mathbf{v} to represent a rotation in the logarithm space. The space is also known as $so(3)$ (different from the rotation group $SO(3)$, the symbol is in lowercase).

To efficiently compute the logarithm of a rotation matrix without computing matrix logarithm or solving rotational axis-angle, we can use the following formula [Gallier and Xu, 2003]:

$$\log \mathbf{R} = \frac{\theta}{2 \sin \theta} (\mathbf{R} - \mathbf{R}^\top), \quad (21)$$

where θ can be obtained from $\theta = \cos^{-1} \left(\frac{\text{Tr}(\mathbf{R}) - 1}{2} \right)$ by the fact that $\text{Tr}(\mathbf{R}) = 1 + 2 \cos \theta$. Specially, when $\theta = 0$ (or $\mathbf{R} = \mathbf{I}$), $\log \mathbf{R} = [\mathbf{0}, \mathbf{0}, \mathbf{0}]$.

Exponential of Skew-Symmetric Matrices The inversion of the rotation matrix logarithm is the exponential of skew-symmetric matrices. Derived from the definition of matrix exponential, the conversion formula is [Gallier and Xu, 2003]:

$$\exp \mathbf{S} = \mathbf{I} + \frac{\sin \|\mathbf{v}\|_2}{\|\mathbf{v}\|_2} \mathbf{S} + \frac{1 - \cos \|\mathbf{v}\|_2}{\|\mathbf{v}\|_2^2} \mathbf{S}^2, \quad (22)$$

where \mathbf{S} is a skew-symmetric matrix parameterized by three values $\mathbf{v} = [v_x, v_y, v_z]$, identical to the definition in Eq.20.

Remarks The logarithm and exponential defined above provide an easy way to create and manipulate rotations in the axis-angle parameterization space. For example, when we would like to create a rotation matrix with an axis and an angle, we can first create a vector \mathbf{v} whose direction is the same as the given axis and whose length equals the angle. Then, we rewrite the vector \mathbf{v} into a skew-symmetric matrix \mathbf{S} , and finally convert it to a rotation matrix by Eq.22. We can also manipulate a rotation matrix, for example, changing its rotational angle, by mapping it to the logarithm space, modifying the skew-symmetric matrix, and finally converting it back to a rotation matrix using the exponential formula.

B.3 ScaleRot: Rotation Scaling Function

When we parameterize a rotation matrix with an axis and an angle, it is natural to define the rotation scaling function ScaleRot as scaling the rotational angle. Formally, the definition is:

$$\text{ScaleRot}(k, \mathbf{R}) := \exp(k \log \mathbf{R}), \quad (23)$$

where k is the scaling factor and \mathbf{R} is a rotation matrix. Specially, $\text{ScaleRot}(0, \mathbf{R}) = \mathbf{I}$ for all rotation matrix \mathbf{R} . Intuitively, scaling a rotation matrix by 0 cancels its effect, leading to identity transformation.

B.4 $\mathcal{IG}_{SO(3)}$: Isotropic Gaussian Distribution on $SO(3)$

The isotropic Gaussian distribution on $SO(3)$, denoted as $\mathcal{IG}_{SO(3)}$, is defined on the axis-angle space of rotation: $\mathbb{S}^2 \times [0, \pi]$, where $\mathbb{S}^2 = \{\|\mathbf{x}\|_2 = 1 | \mathbf{x} \in \mathbb{R}^3\}$ is the unit sphere in \mathbb{R}^3 . $\mathcal{IG}_{SO(3)}$ is parameterized by a mean rotation \mathbf{M} and a scalar variance σ^2 . Let $\mathbf{u} \in \mathbb{S}^2$ and θ denote the rotational axis and angle random variables respectively. We first consider $\mathcal{IG}_{SO(3)}$ with the identity matrix as its mean: $\mathcal{IG}_{SO(3)}(\mathbf{u}, \theta | \mathbf{I}, \sigma^2)$. Its p.d.f. is defined by the product of the uniform distribution on \mathbb{S}^2 and a special angular distribution [Matthies et al., 1970, Nikolayev and Savyolov, 1970, Leach et al., 2022]:

$$p_{\mathcal{IG}_{SO(3)}}(\mathbf{u}, \theta | \mathbf{I}, \sigma^2) = p_{\text{uniform}(\mathbb{S}^2)}(\mathbf{u}) p_{\text{angular}}(\theta | \sigma^2), \quad (24)$$

$$\text{where } p_{\text{uniform}(\mathbb{S}^2)}(\mathbf{u}) = \frac{1}{4\pi} \delta(\|\mathbf{u}\|_2 - 1), \quad (\mathbf{u} \in \mathbb{S}^2) \quad (25)$$

$$\text{and } p_{\text{angular}}(\theta | \sigma^2) = \frac{1 - \cos \theta}{\pi} \sum_{l=0}^{\infty} (2l + 1) e^{-l(l+1)\sigma^2} \frac{\sin\left(\left(l + \frac{1}{2}\right)\theta\right)}{\sin\left(\frac{\theta}{2}\right)}. \quad (\theta \in [0, \pi]) \quad (26)$$

When the mean is other than \mathbf{I} , to sample from the distribution, we can first sample an rotation \mathbf{E} from $\mathcal{IG}_{SO(3)}(\mathbf{u}, \theta | \mathbf{I}, \sigma^2)$. Then, we left-multiply \mathbf{R} to \mathbf{E} to obtain the desired random value $\mathbf{R}\mathbf{E}$.

Sampling The algorithm for drawing samples from $\mathcal{IG}_{\text{SO}(3)}(\mathbf{I}, \sigma^2)$ (here, the mean rotation is identity) can be broken down into two steps.

The first step is to draw a unit vector \mathbf{u} from the uniform distribution on \mathbb{S}^2 , $p_{\text{uniform}(\mathbb{S}^2)}(\mathbf{u})$. This can be done efficiently by sampling from the 3D standard Gaussian distribution and then normalizing the sampled vector to unit length.

The second part is drawing samples from $p_{\text{angular}}(\theta|\sigma^2)$, which could be more tricky. We empirically use two different proximate sampling strategies depending on the variance σ^2 . When σ is larger than 0.1, the series (Eq.26) converges fast. In such cases, we use histograms to approximate the distribution. In specific, we evenly partition $[0, \pi]$ into 8192 bins, and use the probability density $p_{\text{angular}}(\theta|\sigma^2)$ at the center of each bin as the bin weight. We randomly select a bin according to the weights to draw samples from the discretized distribution. Then, we sample from the uniform distribution spanning from the lower bound to the upper bound of the bin. The discretization process is time-consuming. However, since the variances in the diffusion processes are predetermined, we pre-compute and cache the bins and weights to draw samples efficiently. When σ is smaller than 0.1, we approximate the distribution using the truncated Gaussian distribution whose mean is 2σ and the standard deviation is σ . Empirically, we find that the above proximate sampling algorithm is sufficient for training and sampling from our diffusion model.

To sample from $\mathcal{IG}_{\text{SO}(3)}$ with an arbitrary mean rotation \mathbf{R} , we first draw a rotation from $\mathcal{IG}_{\text{SO}(3)}(\mathbf{I}, \sigma^2)$, denoted as \mathbf{E} . Finally, we left-multiply \mathbf{R} to \mathbf{E} to get the desired sample.

B.5 Uniform Distribution on SO(3)

The uniform distribution on $\text{SO}(3)$ is equivalent to the uniform distribution of normalized quaternions on \mathbb{S}^3 [Shoemake, 1992]. To sample a random rotation uniformly, we first sample a random vector from the 4D standard normal distribution. Next, we normalize the vector and treat it as a quaternion. Finally, we convert the quaternion to a rotation matrix which can be regarded as a sample from the uniform distribution on $\text{SO}(3)$.

C Neural Network Parameterization

C.1 Computing Residue Orientations

The orientation of a residue is determined by the coordinate of its three backbone atoms: C_α , C, and N. Let \mathbf{x}_i^α , \mathbf{x}_i^C , and \mathbf{x}_i^N denote the 3D coordinates of the three backbone atoms of the i -th residue respectively. The orientation of the residue, denoted by \mathbf{O}_i , can be constructed using the following Gram-Schmidt-based algorithm:

$$\mathbf{v}_1 \leftarrow \mathbf{x}_i^C - \mathbf{x}_i^\alpha, \quad (27)$$

$$\mathbf{e}_1 \leftarrow \frac{\mathbf{v}_1}{\|\mathbf{v}_1\|}, \quad (28)$$

$$\mathbf{v}_2 \leftarrow \mathbf{x}_i^N - \mathbf{x}_i^\alpha, \quad (29)$$

$$\mathbf{u}_2 \leftarrow \mathbf{v}_2 - \langle \mathbf{e}_1, \mathbf{v}_2 \rangle \mathbf{e}_1, \quad (30)$$

$$\mathbf{e}_2 \leftarrow \frac{\mathbf{u}_2}{\|\mathbf{u}_2\|}, \quad (31)$$

$$\mathbf{e}_3 \leftarrow \mathbf{e}_1 \times \mathbf{e}_2, \quad (32)$$

$$\mathbf{O}_i \leftarrow [\mathbf{e}_1, \mathbf{e}_2, \mathbf{e}_3]. \quad (33)$$

C.2 Architectures

Amino Acid Embedding Layer The embedding layer for each amino acid takes into account the following information:

- **Amino acid type:** Each of the 20 amino acid types is represented by an embedding vector denoted by $\mathbf{e}_i^{\text{type}}$.
- **Heavy atom local coordinates:** The coordinate of each heavy atom in an amino acid is projected to the local coordinate frame using the rule $\mathbf{x}_i^{\text{local}} = \mathbf{O}_i^T (\mathbf{x}_i^{\text{atom}} - \mathbf{x}_i^\alpha)$. All

of the local coordinates are concatenated into a single vector denoted by e_i^{coord} . If some heavy atoms are missing, their local coordinates are filled with zeros. Note that the local coordinates are invariant to global rotation and translation thanks to the projection rule.

- **Backbone dihedral angles:** The backbone dihedrals of amino acid, including ϕ , ψ , and ω [Liljas et al., 2016, Ingraham et al., 2019], are transformed using a series of sine and cosine functions with different frequencies, which are then concatenated into a single vector e_i^{dihed} .
- **CDR flags and anchor flags:** Amino acids on the CDR or by the two ends of the CDR (anchors) are differentiated from other amino acids by special 0-1 flags denoted as e_i^{flag} .

All of the vectors above are concatenated and fed to an MLP to produce the final embedding vector for each residue.

Pairwise Embedding Layer Pairwise embeddings include information about the relationship between two residues. The pairwise embedding for residue i and j involves the following information:

- **Amino acid types of both amino acids:** There are $20 \times 20 = 400$ combinations of two amino acid types. We represent each of them using an embedding vector denoted by z_{ij}^{type} .
- **Sequential relative position:** If two residues are on the same chain and their distance on the sequence is less than or equal to 32 ($d_{ij}^{\text{seq}} \in \{-32 \dots 32\}$), the distance is represented by an embedding vector z_{ij}^{seq} . Otherwise, the distance embedding is filled with zeros.
- **Pairwise distances:** The distances between all pairs of atoms are flattened into a vector and transformed by $e^{-cd_{ij}}$ (c is a learnable coefficient) into the spatial distance embedding z_{ij}^{dist} . Missing pairs are filled with zeros.
- **Pairwise backbone dihedrals:** The backbone dihedrals between any two amino acids i and j are defined as $\phi_{ij} = \text{Dihedral}(\mathbf{x}_i^{\text{C}}, \mathbf{x}_j^{\text{N}}, \mathbf{x}_j^{\text{O}}, \mathbf{x}_i^{\text{C}})$ and $\psi_{ij} = \text{Dihedral}(\mathbf{x}_i^{\text{N}}, \mathbf{x}_i^{\text{O}}, \mathbf{x}_j^{\text{C}}, \mathbf{x}_j^{\text{N}})$. These two dihedrals are transformed by a series of sine and cosine functions into pairwise dihedral embeddings z_{ij}^{dihed} .

We concatenate the above vectors and feed them into an MLP to get the final pairwise embeddings for each pair of amino acids z_{ij} .

Encoder The encoder for encoding the current diffusion state consists of a stack of orientation-aware invariant 3D attention layers. Its aim is to capture relationships between amino acids and provide high-level representations for each residue to denoise.

Let \mathbf{h}_i^ℓ denote the hidden representation output from the last layer (when $\ell = 0$, the representation is the initial residue embedding). The formulas for computing the logit of attention weight between residue i (query) and j (key) is:

$$a_{ij} = \langle \mathbf{q}(\mathbf{h}_i^\ell), \mathbf{k}(\mathbf{h}_j^\ell) \rangle + f(z_{ij}) + g\left(\left\{\mathbf{O}_i^\top(\mathbf{x}_j^{\text{atom}} - \mathbf{x}_i^\alpha)\right\}_{\text{atom}}\right), \quad (34)$$

where $\mathbf{q}(\cdot)$, $\mathbf{k}(\cdot)$, $f(\cdot)$, and $g(\cdot)$ are MLP subnetworks. The attention weights can be obtained by taking softmax: $w_{ij} = \text{softmax}_{j=1}^N(a_{ij})$. Note that, for simplicity, we do not consider attention heads in the formula, but in practice we use multiple attention heads and different heads can be combined easily via concatenation.

The formula for computing the value passed from residue j to i is:

$$\mathbf{v}_{ij} = \mathbf{v}\left(\mathbf{h}_j^\ell, z_{ij}, \left\{\mathbf{O}_i^\top(\mathbf{x}_j^{\text{atom}} - \mathbf{x}_i^\alpha)\right\}_{\text{atom}}\right), \quad (35)$$

where $\mathbf{v}(\cdot)$ is a network consisting of MLPs. Finally, the values along with attention weights are used to update the amino acid representations with residual connection and layer normalization, same as the standard transformer [Vaswani et al., 2017].

C.3 Notes on the Notations of the Denoising Networks F , G , and H

The notations F , G , and H do *not only* denote the MLPs following the encoder that outputs denoising results. It refers to the embedding layers, the encoder, and the specific output MLP (for example, F includes the MLP for denoising amino acid types). Therefore, the input to F , G , and H is the diffusion state (sequence and structure) rather than hidden representations. Treating the three sections as a whole allows us to neatly express the equivariance property of the model.

C.4 Proof of Equivariance

Lemma 1. *The Euclidean distance function between two points is invariant to rotations and translations, i.e. $d(\mathbf{R}\mathbf{x}_1 + \mathbf{r}, \mathbf{R}\mathbf{x}_2 + \mathbf{r}) = d(\mathbf{x}_1, \mathbf{x}_2)$, $\forall \mathbf{R} \in \text{SO}(3), \mathbf{r} \in \mathbb{R}^3$.*

Proof.

$$\begin{aligned} d(\mathbf{R}\mathbf{x}_1 + \mathbf{r}, \mathbf{R}\mathbf{x}_2 + \mathbf{r}) &= \|(\mathbf{R}\mathbf{x}_1 + \mathbf{r}) - (\mathbf{R}\mathbf{x}_2 + \mathbf{r})\|_2 \\ &= \|\mathbf{R}(\mathbf{x}_1 - \mathbf{x}_2)\|_2 \\ &= (\mathbf{x}_1 - \mathbf{x}_2)^\top \mathbf{R}^\top \mathbf{R} (\mathbf{x}_1 - \mathbf{x}_2) \\ &= \|\mathbf{x}_1 - \mathbf{x}_2\|_2 \\ &= d(\mathbf{x}_1, \mathbf{x}_2). \quad \square \end{aligned}$$

Lemma 2. *The dihedral function for four points is invariant to rotations and translations, i.e. $\text{Dihedral}(\mathbf{R}\mathbf{x}_1 + \mathbf{r}, \mathbf{R}\mathbf{x}_2 + \mathbf{r}, \mathbf{R}\mathbf{x}_3 + \mathbf{r}, \mathbf{R}\mathbf{x}_4 + \mathbf{r}) = \text{Dihedral}(\mathbf{x}_1, \mathbf{x}_2, \mathbf{x}_3, \mathbf{x}_4)$, $\forall \mathbf{R} \in \text{SO}(3), \mathbf{r} \in \mathbb{R}^3$. Here, $\text{Dihedral}(\dots)$ is defined as:*

$$\text{Dihedral}(\mathbf{x}_1 \dots \mathbf{x}_4) = \text{atan2}(\mathbf{v}_2 \cdot ((\mathbf{v}_1 \times \mathbf{v}_2) \times (\mathbf{v}_2 \times \mathbf{v}_3)), \|\mathbf{v}_2\|(\mathbf{v}_1 \times \mathbf{v}_2) \cdot (\mathbf{v}_2 \times \mathbf{v}_3)), \quad (36)$$

where $\mathbf{v}_i = \mathbf{x}_{i+1} - \mathbf{x}_i$ ($i = 1, 2, 3$).

Proof. First, we note that:

$$(\mathbf{R}\mathbf{x}_{i+1} + \mathbf{r}) - (\mathbf{R}\mathbf{x}_i + \mathbf{r}) = \mathbf{R}(\mathbf{x}_{i+1} - \mathbf{x}_i) = \mathbf{R}\mathbf{v}_i.$$

By the equivariance of cross product ($\mathbf{R}\mathbf{a} \times \mathbf{R}\mathbf{b} = \mathbf{R}(\mathbf{a} \times \mathbf{b})$) and the invariance of inner product ($\mathbf{R}\mathbf{a} \cdot \mathbf{R}\mathbf{b} = \mathbf{a} \cdot \mathbf{b}$), we have:

$$\begin{aligned} \text{Dihedral}(\mathbf{R}\mathbf{x}_i + \mathbf{r} | i = 1 \dots 4) &= \text{atan2}(\mathbf{R}\mathbf{v}_2 \cdot (\mathbf{R}(\mathbf{v}_1 \times \mathbf{v}_2) \times \mathbf{R}(\mathbf{v}_2 \times \mathbf{v}_3)), \\ &\quad \|\mathbf{R}\mathbf{v}_2\| \mathbf{R}(\mathbf{v}_1 \times \mathbf{v}_2) \cdot \mathbf{R}(\mathbf{v}_2 \times \mathbf{v}_3)) \\ &= \text{atan2}(\mathbf{R}\mathbf{v}_2 \cdot \mathbf{R}((\mathbf{v}_1 \times \mathbf{v}_2) \times (\mathbf{v}_2 \times \mathbf{v}_3)), \\ &\quad \|\mathbf{v}_2\|(\mathbf{v}_1 \times \mathbf{v}_2) \cdot (\mathbf{v}_2 \times \mathbf{v}_3)) \\ &= \text{atan2}(\mathbf{v}_2 \cdot ((\mathbf{v}_1 \times \mathbf{v}_2) \times (\mathbf{v}_2 \times \mathbf{v}_3)), \\ &\quad \|\mathbf{v}_2\|(\mathbf{v}_1 \times \mathbf{v}_2) \cdot (\mathbf{v}_2 \times \mathbf{v}_3)) \\ &= \text{Dihedral}(\mathbf{x}_i | i = 1 \dots 4) \quad \square \end{aligned}$$

Lemma 3. *The per-amino-acid orientation \mathbf{O}_i is equivariant to rotations and translations, i.e., $\mathbf{O}(\mathbf{R}\mathbf{x}_i^\alpha + \mathbf{r}, \mathbf{R}\mathbf{x}_i^C + \mathbf{r}, \mathbf{R}\mathbf{x}_i^N + \mathbf{r}) = \mathbf{R}\mathbf{O}(\mathbf{x}_i^\alpha, \mathbf{x}_i^C, \mathbf{x}_i^N)$*

Proof. First, we show that the first two basis vectors \mathbf{e}_1 and \mathbf{e}_2 are equivariant:

$$\begin{aligned} \mathbf{e}_1(\mathbf{R}\mathbf{x}_i^\alpha + \mathbf{r}, \mathbf{R}\mathbf{x}_i^C + \mathbf{r}) &= \frac{(\mathbf{R}\mathbf{x}_i^C + \mathbf{r}) - (\mathbf{R}\mathbf{x}_i^\alpha + \mathbf{r})}{\|(\mathbf{R}\mathbf{x}_i^C + \mathbf{r}) - (\mathbf{R}\mathbf{x}_i^\alpha + \mathbf{r})\|} \\ &= \mathbf{R} \frac{\mathbf{x}_i^C - \mathbf{x}_i^\alpha}{\|\mathbf{x}_i^C - \mathbf{x}_i^\alpha\|} \\ &= \mathbf{R}\mathbf{e}_1(\mathbf{x}_i^\alpha, \mathbf{x}_i^C). \end{aligned}$$

Let $\mathbf{v}_2 = \mathbf{x}_i^N - \mathbf{x}_i^\alpha$. We have $(\mathbf{R}\mathbf{x}_i^N + \mathbf{r}) - (\mathbf{R}\mathbf{x}_i^\alpha + \mathbf{r}) = \mathbf{R}\mathbf{v}_2$. Then, we can prove the equivariance of \mathbf{e}_2 :

$$\begin{aligned} \mathbf{e}_2(\mathbf{R}\mathbf{x}_i^\alpha + \mathbf{r}, \mathbf{R}\mathbf{x}_i^C + \mathbf{r}, \mathbf{R}\mathbf{x}_i^N + \mathbf{r}) &= \mathbf{R}\mathbf{v}_2 - \langle \mathbf{R}\mathbf{e}_1, \mathbf{R}\mathbf{v}_2 \rangle \mathbf{R}\mathbf{e}_1 \\ &= \mathbf{R}\mathbf{v}_2 - \langle \mathbf{e}_1, \mathbf{v}_2 \rangle \mathbf{R}\mathbf{e}_1 \\ &= \mathbf{R}(\mathbf{v}_2 - \langle \mathbf{e}_1, \mathbf{v}_2 \rangle \mathbf{e}_1) \\ &= \mathbf{R}\mathbf{e}_2(\mathbf{x}_i^\alpha, \mathbf{x}_i^C, \mathbf{x}_i^N) \end{aligned}$$

By the equivariance of cross product, it is straightforward to show that \mathbf{e}_3 is also equivariant. Finally, combining the equivariance of \mathbf{e}_1 , \mathbf{e}_2 , and \mathbf{e}_3 , we prove the equivariance of the orientation matrix:

$$\begin{aligned} \mathbf{O}(\mathbf{R}\mathbf{x}_i^\alpha + \mathbf{r}, \mathbf{R}\mathbf{x}_i^C + \mathbf{r}, \mathbf{R}\mathbf{x}_i^N + \mathbf{r}) &= [\mathbf{R}\mathbf{e}_1, \mathbf{R}\mathbf{e}_2, \mathbf{R}\mathbf{e}_3] \\ &= \mathbf{R}\mathbf{O}(\mathbf{x}_i^\alpha, \mathbf{x}_i^C, \mathbf{x}_i^N). \quad \square \end{aligned}$$

Lemma 4. *The per-amino-acid and pairwise embedding layers are invariant to rotations and translations of the input structure. i.e.*

$$\begin{aligned} e(s_i, \{\mathbf{x}_i^{\text{atom}}\}_{\text{atom}}, \phi_i, \psi_i, \omega_i, \mathbf{e}_i^{\text{flag}}) &= e(s_i, \{\mathbf{R}\mathbf{x}_i^{\text{atom}} + \mathbf{r}\}_{\text{atom}}, \phi_i, \psi_i, \omega_i, \mathbf{e}_i^{\text{flag}}), \quad \text{and} \\ \mathbf{z}(\{d(\mathbf{x}_i^{\text{atom}1}, \mathbf{x}_j^{\text{atom}2})\}_{\text{atom}1, \text{atom}2, \dots}) &= \mathbf{z}(\{d(\mathbf{R}\mathbf{x}_i^{\text{atom}1} + \mathbf{r}, \mathbf{R}\mathbf{x}_j^{\text{atom}2} + \mathbf{r})\}_{\text{atom}1, \text{atom}2, \dots}). \end{aligned}$$

Proof. Before embedding atom positions for an amino acid, the network first projects the positions using the orientation by the rule:

$$\mathbf{x}_i^{\text{local}} = \mathbf{O}_i^\top (\mathbf{x}_i^{\text{atom}} - \mathbf{x}_i^\alpha)$$

The projection operation is invariant to rotations and translations, using Lemma 3:

$$\begin{aligned} \mathbf{x}_i^{\text{local}}(\mathbf{R}\mathbf{x}_i^{\text{atom}} + \mathbf{r}, \mathbf{R}\mathbf{x}_i^\alpha + \mathbf{r}) &= (\mathbf{R}\mathbf{O}_i)^\top ((\mathbf{R}\mathbf{x}_i^{\text{atom}} + \mathbf{r}) - (\mathbf{R}\mathbf{x}_i^\alpha + \mathbf{r})) \\ &= \mathbf{O}_i^\top \mathbf{R}^\top \mathbf{R} (\mathbf{x}_i^{\text{atom}} - \mathbf{x}_i^\alpha) \\ &= \mathbf{x}_i^{\text{local}}(\mathbf{x}_i^{\text{atom}}, \mathbf{x}_i^\alpha). \end{aligned}$$

The formulas for computing dihedral angles $(\phi_i, \psi_i, \omega_i)$ are also invariant by Lemma 2 Other variables (amino acid types and CDR flags) are independent of the 3D structure and hence they are invariant.

So far, we have shown that all the components of embedding layers are invariant to rotations and translations of the overall 3D structure. Therefore, the embedding layer is invariant.

Pairwise embedding layers involve distances between residues, which are invariant by Lemma 2. Other variables are irrelevant to 3D structures. Hence, the pairwise embedding layer is invariant. \square

Lemma 5. *The orientation-aware attention layer is invariant to rotations and translations if the input hidden representations $\mathbf{h}_i, \mathbf{z}_{ij}(i, j = 1 \dots N)$ come from invariant functions.*

Proof. First, we show that projecting atoms on the j -th amino acid to the orientation of the i -th amino acid is invariant to rotations and translations by Lemma 3:

$$(\mathbf{R}\mathbf{O}_i)^\top ((\mathbf{R}\mathbf{x}_j^{\text{atom}} + \mathbf{r}) - (\mathbf{R}\mathbf{x}_i^\alpha + \mathbf{r})) = \mathbf{O}_i^\top \mathbf{R}^\top \mathbf{R} (\mathbf{x}_j^{\text{atom}} - \mathbf{x}_i^\alpha).$$

As other inputs to the attention layer $(\mathbf{h}_i, \mathbf{z}_{ij}(i, j = 1 \dots N))$ are invariant to rigid transforms on the structure, the networks for computing attention weights and values are invariant. Hence, the attention layer is invariant.

In the case where we stack multiple attention layers, each layer outputs invariant representations for its next layer. Therefore, the network consisting of multiple attention layers is invariant. \square

Proposition 2. *For any proper rotation matrix $\mathbf{R} \in \text{SO}(3)$ and any 3D vector $\mathbf{r} \in \mathbb{R}^3$ (rigid transformation $(\mathbf{R}, \mathbf{r}) \in \text{SE}(3)$), F, G and H satisfy the following equivariance properties:*

$$F(\mathbf{R}\mathcal{R}^t + \mathbf{r}, \mathbf{R}\mathcal{C} + \mathbf{r}) = F(\mathcal{R}^t, \mathcal{C}), \quad (37)$$

$$G(\mathbf{R}\mathcal{R}^t + \mathbf{r}, \mathbf{R}\mathcal{C} + \mathbf{r}) = \mathbf{R}G(\mathcal{R}^t, \mathcal{C}), \quad (38)$$

$$H(\mathbf{R}\mathcal{R}^t + \mathbf{r}, \mathbf{R}\mathcal{C} + \mathbf{r}) = \mathbf{R}H(\mathcal{R}^t, \mathcal{C}), \quad (39)$$

where $\mathbf{R}\mathcal{R}^t + \mathbf{r} := \{\mathbf{s}_j^t, \mathbf{x}_j^t + \mathbf{r}, \mathbf{R}\mathbf{O}_j^t\}_{j=l+1}^{l+m}$ and $\mathbf{R}\mathcal{C} + \mathbf{r} := \{s_i, \mathbf{x}_i + \mathbf{r}, \mathbf{R}\mathbf{O}_i\}_{i \in \{1 \dots N\} \setminus \{l+1, \dots, l+m\}}$ denote the transformed structure.

Proof. By Lemma 5, we know that the encoder network produces invariant representations. Therefore, the MLP for predicting amino acid types that transforms the invariant representations into a probability over 20 categories is invariant, so F is invariant.

The MLP for predicting local coordinate changes $\text{MLP}_G(\mathbf{h}_i)$ is invariant. The local coordinate change is converted to the global coordinate change using the following rule:

$$\hat{\mathbf{e}}_j = \mathbf{O}_j^t \text{MLP}_G(\mathbf{h}_j).$$

By Lemma 3, the above rule is equivariant to rotations, and hence G is equivariant to rotations.

Similarly, the MLP for predicting changes in orientation $\text{MLP}_H(\mathbf{h}_i)$ is invariant. The changes is applied to the original orientation by:

$$\hat{\mathbf{O}}_j^{t-1} = \mathbf{O}_j^t M_j,$$

which is equivariant to rotations according to Lemma 3. Therefore, H is equivariant to rotations. \square

D Sampling Algorithms

D.1 Backbone Atoms and Sidechain C_β Construction

The coordinates of backbone atoms (N, C_α , C, O) and sidechain C_β can be determined by the orientation and the C_α position of an amino acid because the geometry of these atoms is inflexible [Liljas et al., 2016]. To construct the position of N, C_α , C, and C_β for the i -th amino acid, we use the following formula:

$$\mathbf{x}_i^{\text{atom}} = \mathbf{O}_i \mathbf{c}^{\text{atom}} + \mathbf{x}_i, \quad (\text{atom} \in \{\text{N}, C_\alpha, \text{C}, C_\beta\}) \quad (40)$$

where \mathbf{O}_i and \mathbf{x}_i is the model-predicted amino acid orientation and C_α position. \mathbf{c}^{atom} is the local coordinate derived from experimental data relative to the orientation and the C_α position, as shown in the following table.

Atom	c_x	c_y	c_z
N	-0.526	1.361	0.000
C_α	0.000	0.000	0.000
C	1.525	0.000	0.000
C_β	-0.500	-0.733	-1.154

The position of O depends on the ψ angle of the amino acid, which relies on the next amino acid in the sequence. Therefore, after constructing backbone atoms, we calculate the ψ angle for each amino acid ($\psi_i = \text{Dihedral}(\text{N}^i, C_\alpha^i, C^i, \text{N}^{i+1})$), and use the following rule to construct O coordinates:

$$\mathbf{x}_i^{\text{O}} = \mathbf{O}_i \mathbf{c}^{\text{O}}(\psi_i) + \mathbf{x}_i, \quad (41)$$

where

$$\mathbf{c}^{\text{O}}(\psi_i) = \begin{bmatrix} 1 & 0 & 0 \\ 0 & \cos \psi_i & -\sin \psi_i \\ 0 & \sin \psi_i & \cos \psi_i \end{bmatrix} \begin{bmatrix} 2.151 \\ -1.062 \\ 0.000 \end{bmatrix}. \quad (42)$$

D.2 Sidechain Packing and Full Atom Refinement

We use PackRotamersMover in PyRosetta [Chaudhury et al., 2010] to pack sidechains only for amino acids on the generated CDR. The packing program is based on the Dunbrack 2010 rotamer library [Shapovalov and Dunbrack Jr, 2011] and the REF2015 energy function [Alford et al., 2017].

After packing sidechains, we refine the structure with OpenMM [Eastman et al., 2017]. Specifically, we first use PDBFixer to prepare the structure for refinement. We minimize the potential energy of the structure. The potential energy is AMBER99SB force field plus quadratic constraint terms that restrain the position of atoms outside the generated CDR.

E Source Code

Code and data are available at <https://github.com/luost26/diffab>

Checklist

1. For all authors...
 - (a) Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope? [Yes]
 - (b) Did you describe the limitations of your work? [Yes]
 - (c) Did you discuss any potential negative societal impacts of your work? [Yes] Our method might be used to design malicious proteins.
 - (d) Have you read the ethics review guidelines and ensured that your paper conforms to them? [Yes]
2. If you are including theoretical results...
 - (a) Did you state the full set of assumptions of all theoretical results? [N/A]
 - (b) Did you include complete proofs of all theoretical results? [N/A]
3. If you ran experiments...
 - (a) Did you include the code, data, and instructions needed to reproduce the main experimental results (either in the supplemental material or as a URL)? [Yes] Appendix E
 - (b) Did you specify all the training details (e.g., data splits, hyperparameters, how they were chosen)? [Yes]
 - (c) Did you report error bars (e.g., with respect to the random seed after running experiments multiple times)? [Yes]
 - (d) Did you include the total amount of compute and the type of resources used (e.g., type of GPUs, internal cluster, or cloud provider)? [Yes] Single A100 GPU.
4. If you are using existing assets (e.g., code, data, models) or curating/releasing new assets...
 - (a) If your work uses existing assets, did you cite the creators? [Yes]
 - (b) Did you mention the license of the assets? [N/A]
 - (c) Did you include any new assets either in the supplemental material or as a URL? [No]
 - (d) Did you discuss whether and how consent was obtained from people whose data you're using/curating? [N/A]
 - (e) Did you discuss whether the data you are using/curating contains personally identifiable information or offensive content? [N/A]
5. If you used crowdsourcing or conducted research with human subjects...
 - (a) Did you include the full text of instructions given to participants and screenshots, if applicable? [N/A]
 - (b) Did you describe any potential participant risks, with links to Institutional Review Board (IRB) approvals, if applicable? [N/A]
 - (c) Did you include the estimated hourly wage paid to participants and the total amount spent on participant compensation? [N/A]