Karlsruhe Institute of Technology

KCDS Summer School 2026

4. Applications I

Jun.-Prof. Dr. Jan Stühmer

Outline

- Molecular Property Prediction
  - Local frames: Dimenet [Gasteiger & Günnemann 2020]
  - Equivariant representations: SEGNN [Brandstetter et al. 2022]

- Protein Design
  - FrameDiff and FrameFlow [Yim et al. 2023 and 2024]
  - GAFL [Wagner & Seute et al. 2024]

Molecular Property Prediction

Molecular Property Prediction

- In Molecular Property prediction, the task is to predict different properties of a molecule, given its atom types and atom coordinates
- A common benchmark is QM9, which contains data of about 130,000 molecules with 19 regression targets
- The target values are for example internal energy, free energy, or heat capacity
- Input features commonly are the atom type, the 3D coordinate of an atom, and vectors to neighboring atoms

Dimenet [Gasteiger & Günnemann 2020]

- Directional Message Passing

$$\boldsymbol{m}_{ji}^{(l+1)} = f_{\text{update}}(\boldsymbol{m}_{ji}^{(l)}, \sum_{k \in \mathcal{N}_j \setminus \{i\}} f_{\text{int}}(\boldsymbol{m}_{kj}^{(l)}, \boldsymbol{e}_{\text{RBF}}^{(ji)}, \boldsymbol{a}_{\text{SBF}}^{(kj,ji)}))$$

where e_{RBF} is the radial basis function representing the interatomic distance and a_{SBF} is a steerable 2D spherical Fourier-Bessel basis.

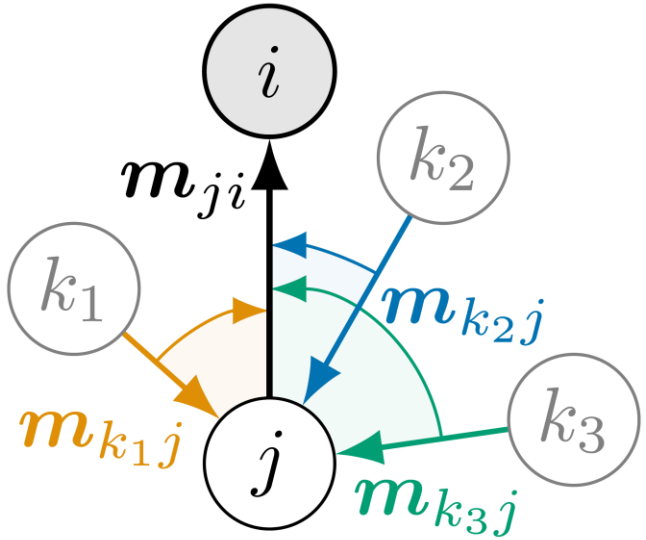

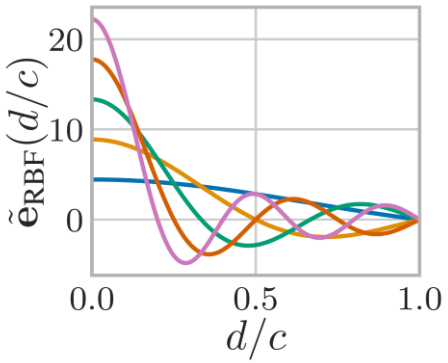

e_{RBF}

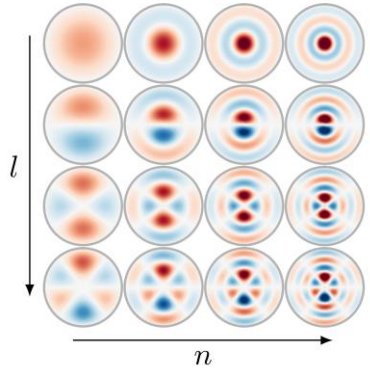

a_{SBF}

Dimenet [Gasteiger & Günnemann 2020]

- Directional Message Passing defines a local frame relative to the orientation of the edge for which a message is computed
- Local frames are a mechanism to make models equivariant, think of them as local gauges that transform with the group transformation
- However, the local frame does not define a second axis, thus the representation is invariant towards rotations around the tangent, resulting in 2D spherical filters
- Examples of learnt filters:

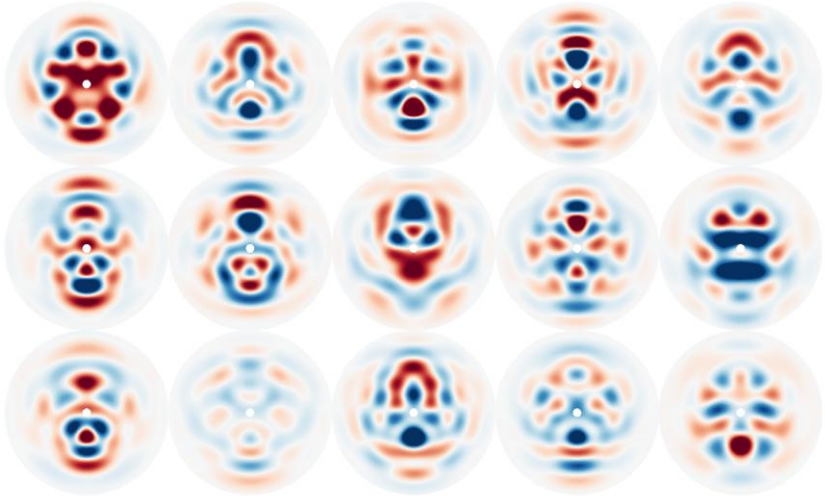

Dimenet - Results

Target	Unit	PPGN	SchNet	PhysNet	MEGNet-s	Cormorant	**DimeNet**
μ	D	0.047	0.033	0.0529	0.05	0.13	**0.0286**
α	${a_0}^3$	0.131	0.235	0.0615	0.081	0.092	**0.0469**
ϵ_{HOMO}	meV	40.3	41	32.9	43	36	**27.8**
ϵ_{LUMO}	meV	32.7	34	24.7	44	36	**19.7**
$\Delta\epsilon$	meV	60.0	63	42.5	66	60	**34.8**
$\langle R^2 \rangle$	${a_0}^2$	0.592	**0.073**	0.765	0.302	0.673	0.331
ZPVE	meV	3.12	1.7	1.39	1.43	1.98	**1.29**
U_0	meV	36.8	14	8.15	12	28	**8.02**
U	meV	36.8	19	8.34	13	-	**7.89**
H	meV	36.3	14	8.42	12	-	**8.11**
G	meV	36.4	14	9.40	12	-	**8.98**
c_{v}	$\frac{\text{cal}}{\text{mol K}}$	0.055	0.033	0.0280	0.029	0.031	**0.0249**
std. MAE	%	1.84	1.76	1.37	1.80	2.14	**1.05**
logMAE	-	−4.64	−5.17	−5.35	−5.17	−4.75	**−5.57**

SEGNN [Brandstetter et al. 2022]

- Brandstetter et al. introduce SE-GNN, an equivariant model for molecular property prediction
- They use spherical harmonics as steerable basis for SE(3)-equivariant message passing

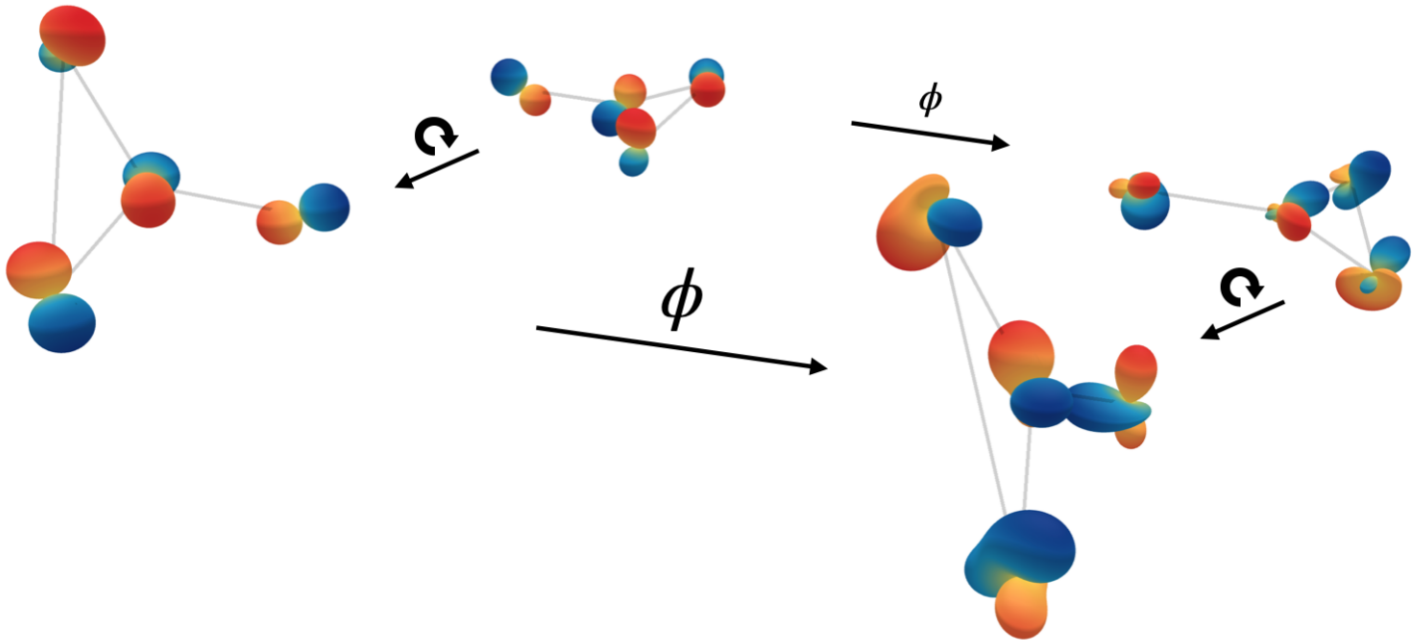

SEGNN [Brandstetter et al. 2022]

- The architecture is a message passing graph neural network with message function and node update functions:

$$\mathbf{m}_{ij} = \phi_m\left(\mathbf{f}_i, \mathbf{f}_j, \mathbf{a}_{ij}\right) \ ,$$

$$\mathbf{f}_i' = \phi_f\left(\mathbf{f}_i, \sum_{j \in \mathcal{N}(i)} \mathbf{m}_{ij}\right)$$

- Vectors to neighboring atoms are embedded as spherical harmonics

$$\tilde{\mathbf{a}}^{(l)} = \left(Y_m^{(l)}\left(\frac{\mathbf{x}}{\|\mathbf{x}\|}\right)\right)_{m=-l,-l+1,\ldots,l}^{T}$$

where the spherical harmonic functions $Y_m^{(l)}$ are functions on the sphere.

Spherical Harmonics

- Spherical harmonics originate from solving Laplace's equation on the sphere
- Thus a similar concept to how we derived the spectrum of the graph Laplacian
- They form a steerable basis for representing functions on the sphere:

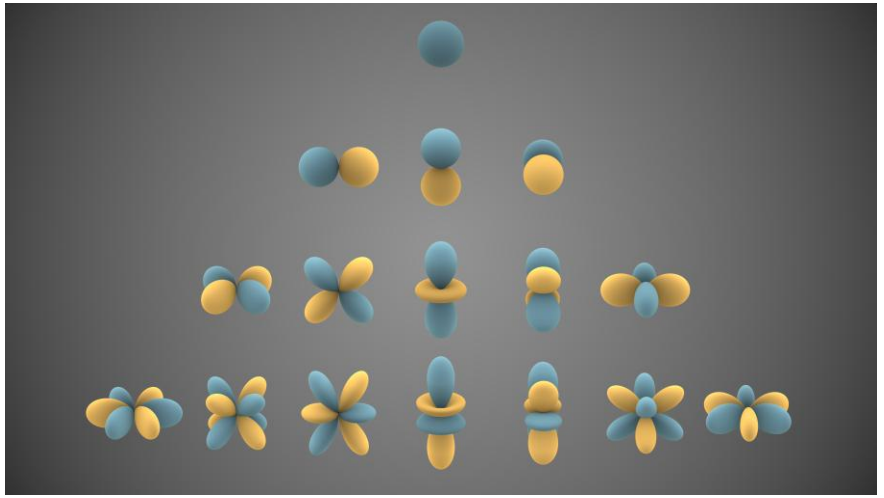

[Wikipedia]

- The vector space in organized into vectors of different grades (rows in above visualization)

Spherical Harmonics – Steerable Basis

- They form a steerable basis, and vectors transform by *Wigner-D matrices*, the representations of SE(3) transformations
- The transformation $D^l(\mathfrak{g})$ acts separately on each subspace V_l

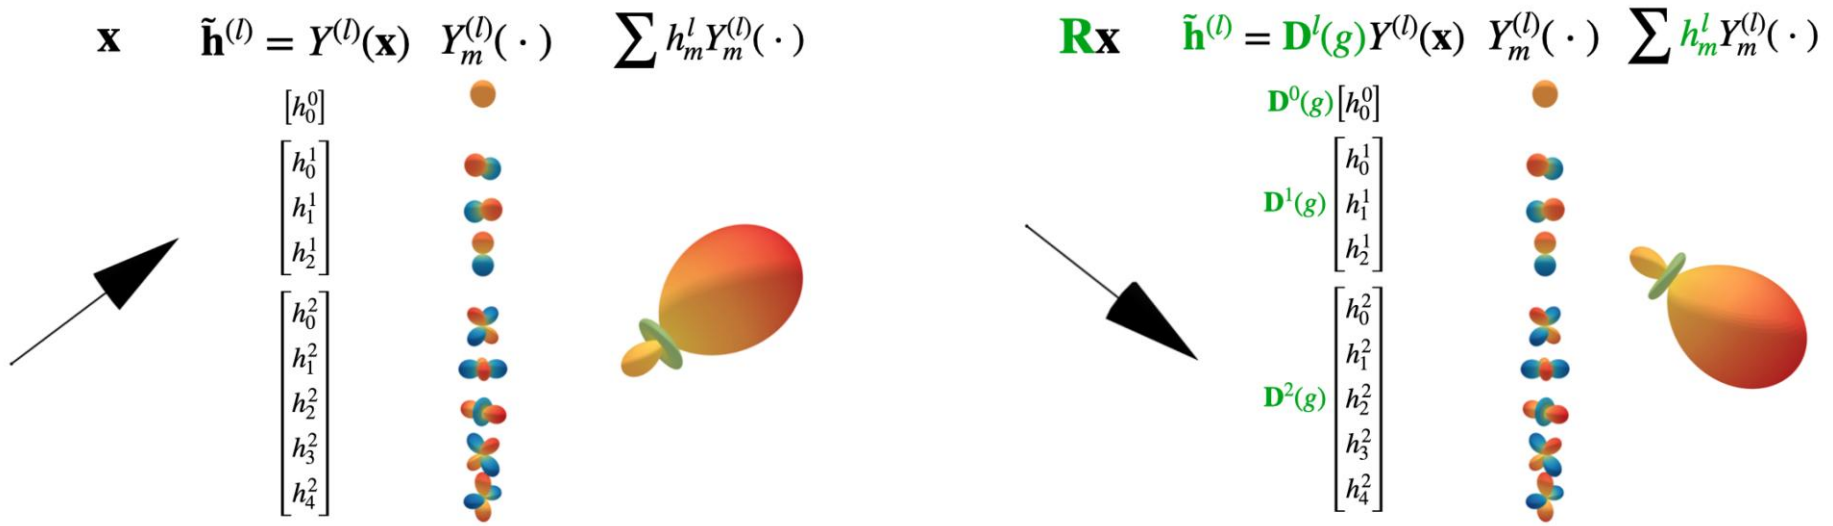

Steerable Tensor Product

- Interactions between feature vectors are described with the Clebsch-Gordon tensor product

$$(\tilde{\mathbf{h}}^{(l_1)} \otimes_{cg}^{w} \tilde{\mathbf{h}}^{(l_2)})_m^{(l)} = w \sum_{m_1=-l_1}^{l_1} \sum_{m_2=-l_2}^{l_2} C_{(l_1,m_1)(l_2,m_2)}^{(l,m)} h_{m_1}^{(l_1)} h_{m_2}^{(l_2)}$$

where the Clebsch-Gordon coefficients $C_{(l_1,m_1)(l_2,m_2)}^{(l,m)}$ (either 1 or 0) ensure that the resulting vector is steerable (only certain grades of vectors may be multiplied with each other)

- For each Clebsch-Gordon coefficient that equals 1, a trainable weight is associated
- This results in a *steerable group convolution*

SEGNN Results

	Energy MAE [eV] ↓				EwT ↑			
Model	ID	OOD Ads	OODCat	OOD Both	ID	OOD Ads	OOD Cat	OOD Both
Median baseline	1.7499	1.8793	1.7090	1.6636	0.71%	0.72%	0.89%	0.74%
CGCNN	0.6149	0.9155	0.6219	0.8511	3.40%	1.93%	3.10%	2.00%
SchNet	0.6387	0.7342	0.6616	0.7037	2.96%	2.33%	2.94%	2.21%
EdgeUpdateNet	0.5839	0.7252	0.6016	0.6862	3.48%	2.35%	3.30%	2.57%
EnergyNet	0.6366	0.717	0.6387	0.6626	3.30%	2.20%	3.07%	2.34%
DimeNet++	0.5620	0.7252	0.5756	0.6613	4.25%	2.07%	4.10%	2.41%
SphereNet	0.5630	0.7030	0.5710	**0.6380**	4.47%	2.29%	4.09%	2.41%
SEGNN (Ours)	**0.5327**	**0.6921**	**0.5369**	0.6790	**5.37%**	**2.46%**	**4.91%**	**2.63%**

Karlsruhe Institute of Technology

Approximate Equivariance

with Torben Berndt

Drawbacks of Strict Equivariance

Equivariant models can be …

Computationally expensive

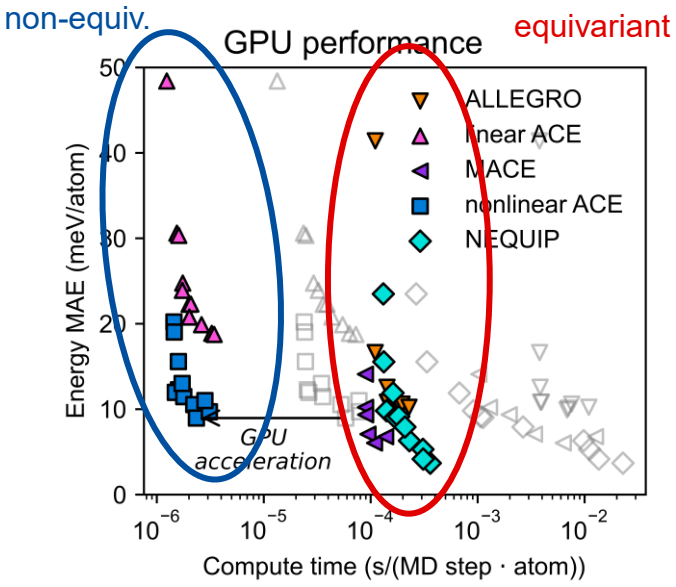

[Leimeroth et al., MSMSE, 2025]

Too constraining

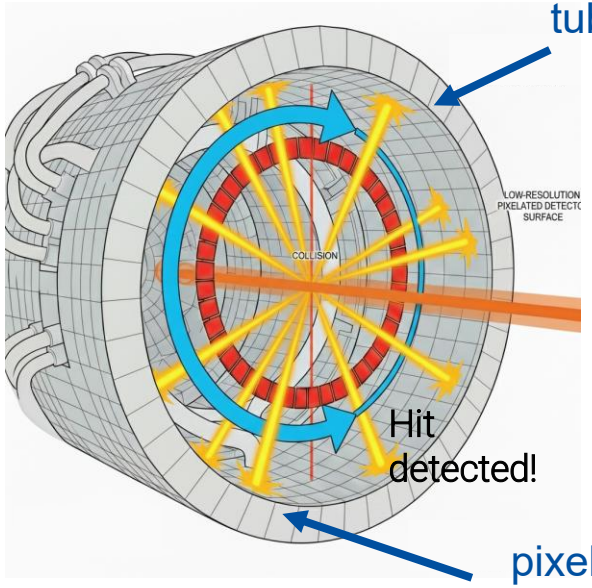

[T. Berndt]

More difficult to optimize

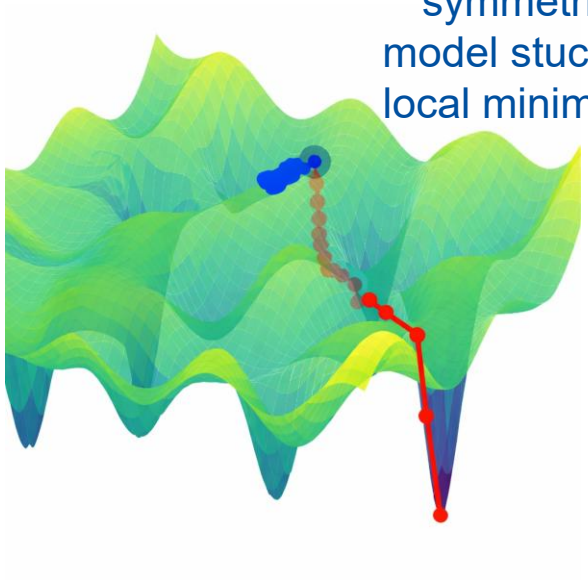

[Manolache, NeurIPS 2025]

Drawbacks of Strict Equivariance

Equivariant models can be …

Computationally expensive

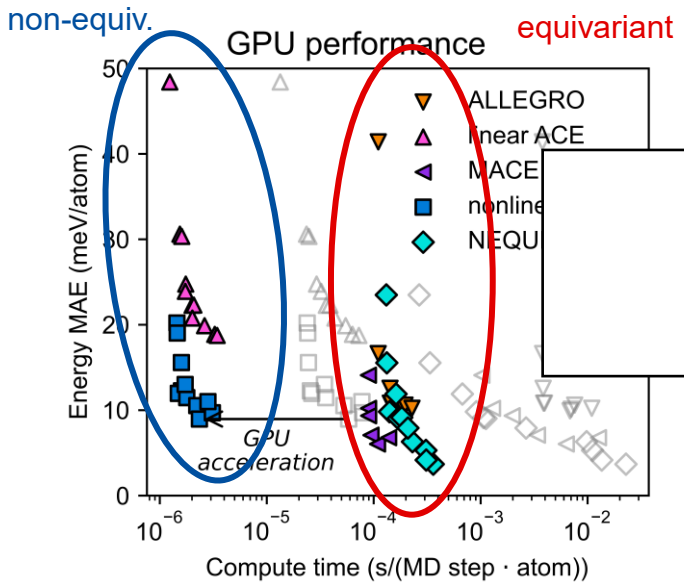

[Leimeroth et al., MSMSE, 2025]

Too constraining

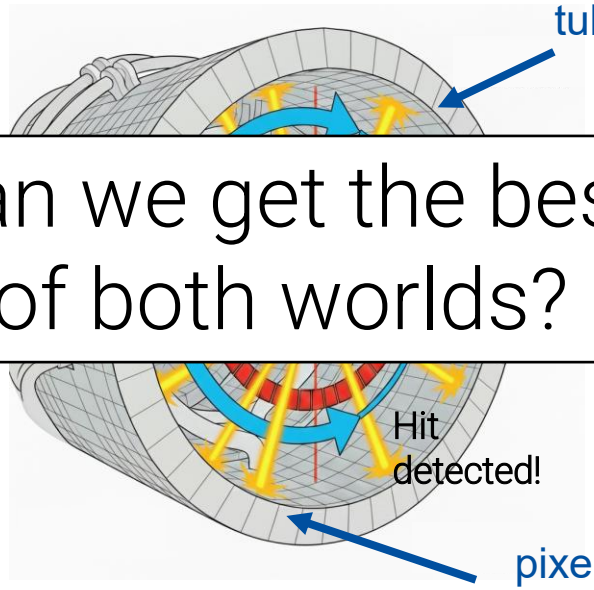

[T. Berndt]

More difficult to optimize

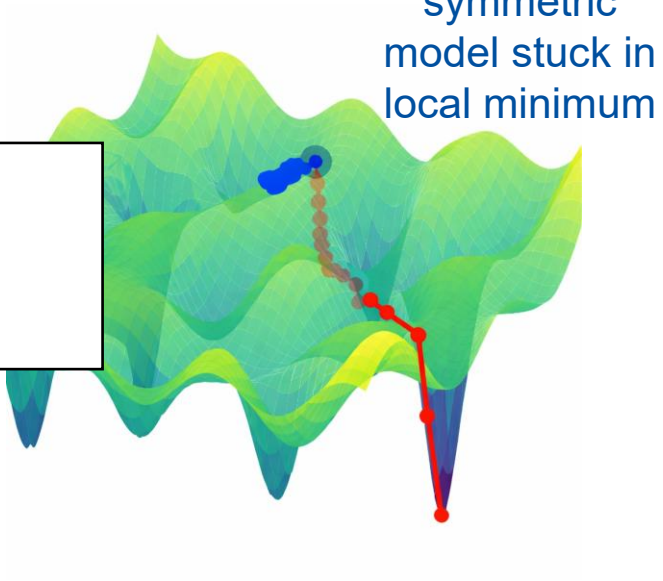

[Manolache, NeurIPS 2025]

Can we get the best of both worlds?

Equivariance Regularization

Idea: Use unconstrained f and find a measure $\mathrm{c}(f)$ which is minimized when f is equivariant

$$\hat{f} = \operatorname{argmin}_{f \in \mathcal{F}} \hat{\mathcal{R}}(f) + \lambda\, \mathrm{c}(f)$$

Why?

- 1. Regularizers only need to be calculated during training.
  → Inference compute cost is unaffected!

Equivariance Regularization

Idea: Use unconstrained f and find a measure $\mathrm{c}(f)$ which is minimized when f is equivariant

$$\hat{f} = \operatorname*{argmin}_{f \in \mathcal{F}} \hat{\mathcal{R}}(f) + \lambda\, \mathrm{c}(f)$$

Why?

- 2. Model itself is unconstrained

  → avoids strict loss barriers

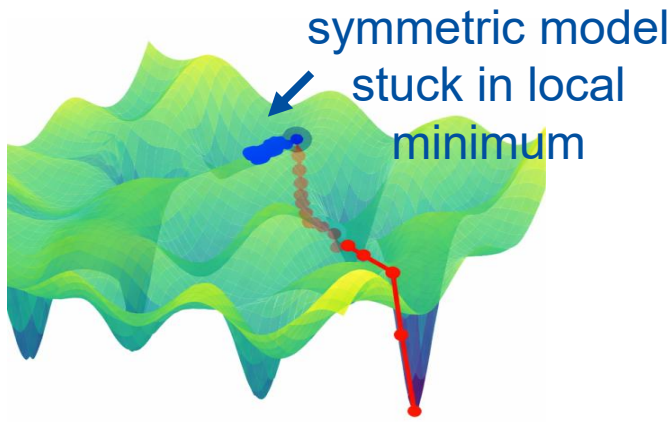

Figure: Loss landscape with training trajectory of equivariant (blue) and non-equivariant model (red) (Manolache, NeurIPS 2025)

Equivariance Regularization

Idea: Use unconstrained f and find a measure $\mathrm{c}(f)$ which is minimized when f is equivariant

$$\hat{f} = \operatorname{argmin}_{f \in \mathcal{F}} \hat{\mathcal{R}}(f) + \lambda\, \mathrm{c}(f)$$

Why?

- 3. Compare with standard regularization with complexity measure γ :

  Constrained form:

$$\hat{f}_\delta = \operatorname{argmin}_{f \in \mathcal{F}} \hat{\mathcal{R}}(f) \quad s.t. \;\; \gamma(f) \leq \delta$$

  Regularized form:

$$\hat{f} = \operatorname{argmin}_{f \in \mathcal{F}} \hat{\mathcal{R}}(f) + \lambda\, \gamma(f)$$

Group Averaging
(How to make any function Equivariant)

- Define the group averaging operator $P_{\mathfrak{G}}$ as

$$P_{\mathfrak{G}} f(u) \stackrel{\text{def}}{=} \int \rho(\mathfrak{g})^{-1} f \, \rho(\mathfrak{g}) \, \mathrm{d}\lambda(\mathfrak{g})$$

- Compare to $\mathfrak{G}$-*smoothing operator* (of a discrete, finite group):

$$S_{\mathfrak{G}} f \stackrel{\text{def}}{=} \frac{1}{|\mathfrak{G}|} \sum_{\mathfrak{g} \in \mathfrak{G}} f \circ \mathfrak{g}$$

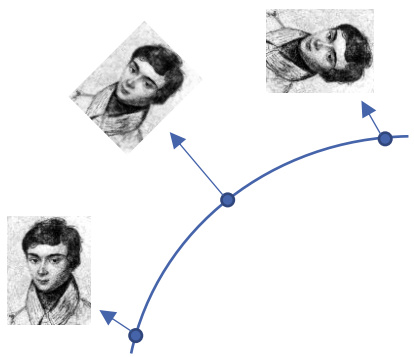

Group orbit $\mathfrak{G}.x = \{\mathfrak{g}.x; \mathfrak{g} \in \mathfrak{G}\}$

Group Averaging
(How to make any function Equivariant)

- Define the group averaging operator $P_{\mathfrak{G}}$ as

$$P_{\mathfrak{G}} f \stackrel{\text{def}}{=} \int \rho(\mathfrak{g})^{-1} f \, \rho(\mathfrak{g}) \, \mathrm{d}\lambda(\mathfrak{g})$$

- Observe that f is equivariant if and only if

$$P_{\mathfrak{G}} f = f \iff \|f - P_{\mathfrak{G}} f\|^2 = 0$$

- → We can use $\|f - P_{\mathfrak{G}} f\|^2$ to measure how equivariant f is

Equivariance Regularization

Idea: Orthogonal splitting into equivariant and non-equivariant part:

$$\hat{f} = \operatorname*{argmin}_{f \in \mathcal{F}} \hat{\mathcal{R}}(f) + \lambda_{\mathrm{G}} \| P_{\mathfrak{G}} f \|^2 + \lambda_{\perp} \| f - P_{\mathfrak{G}} f \|^2$$

Then we could use λ_{G} and $\lambda_{\perp}$ to "tune" a model towards equivariance

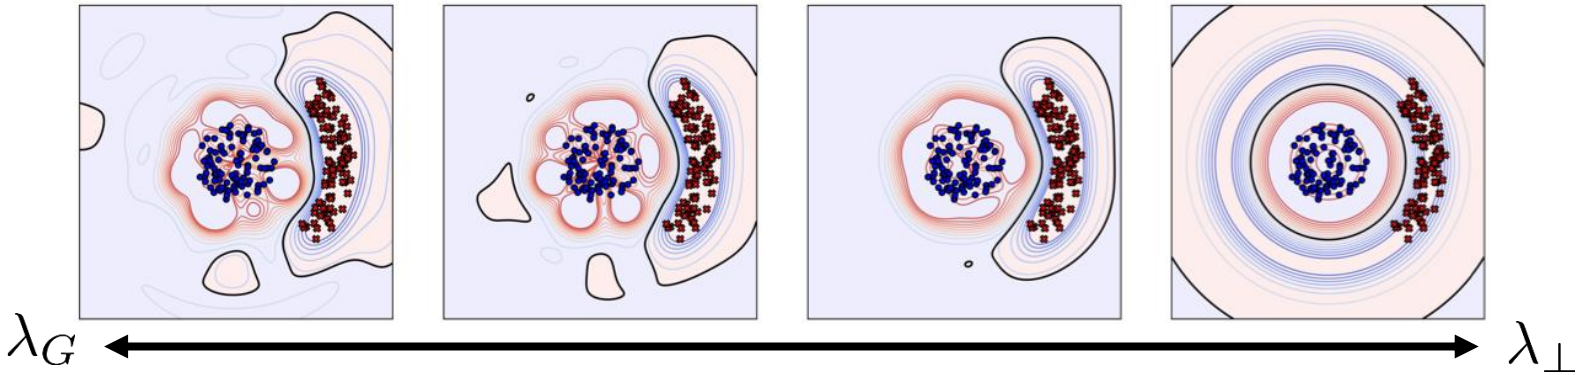

Figure: Controlling the degree of learned $SO(2)$ invariance by varying the values of λ_G and $\lambda_\perp$.

Equivariance Regularization

- Perform an orthogonal decomposition to regularize the equivariant and non-equivariant part of a model on the operator level:

$$\hat{f} = \operatorname*{argmin}_{f \in \mathcal{F}} \hat{\mathcal{R}}(f) + \lambda_{\mathrm{G}} \|P_{\mathfrak{G}} f\|^2 + \lambda_{\perp} \|f - P_{\mathfrak{G}} f\|^2$$

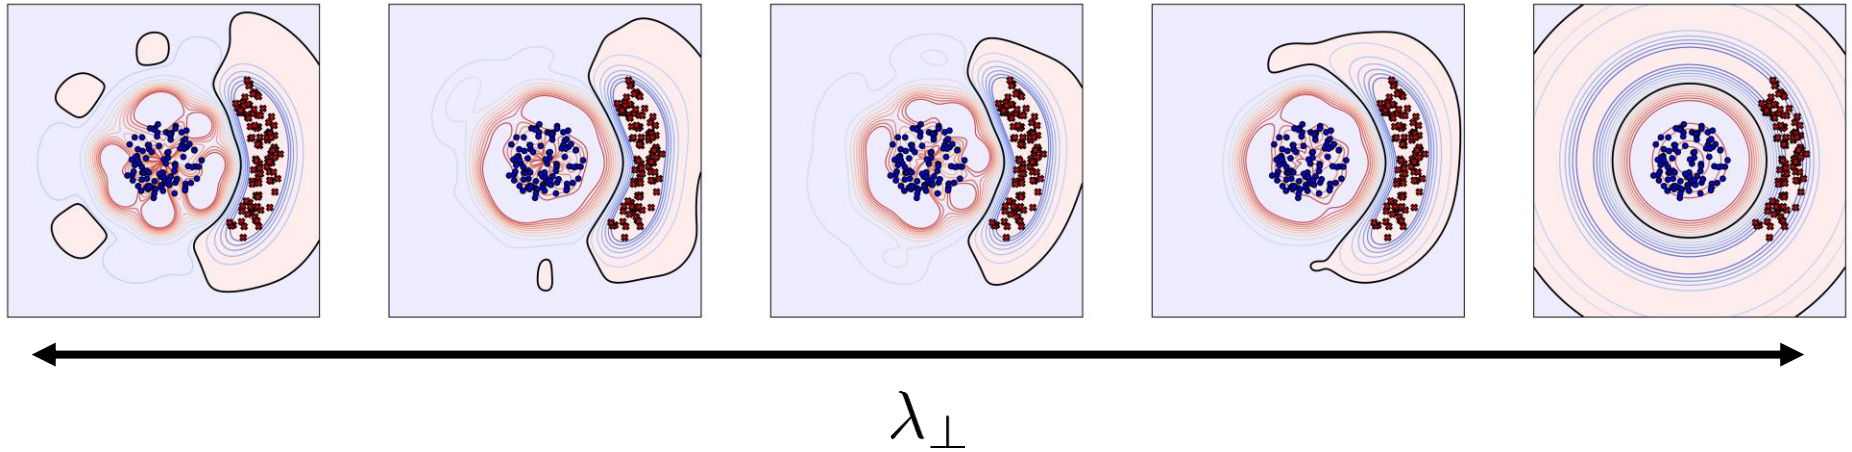

Achieves substantial runtime gains over sample-based regularisers

T. Berndt, J. Stühmer. "Approximate equivariance via projection-based regularization", ICML 2026

Equivariance Regularization

- Perform an orthogonal decomposition to regularize the equi equivariant part of a model on the operator level:

$$\hat{f} = \operatorname*{argmin}_{f \in \mathcal{F}} \hat{\mathcal{R}}(f) + \lambda_{\mathrm{G}} \|P_{\mathfrak{G}} f\|^2 + \lambda_{\perp} \|f - P_{\mathfrak{G}} f\|^2$$

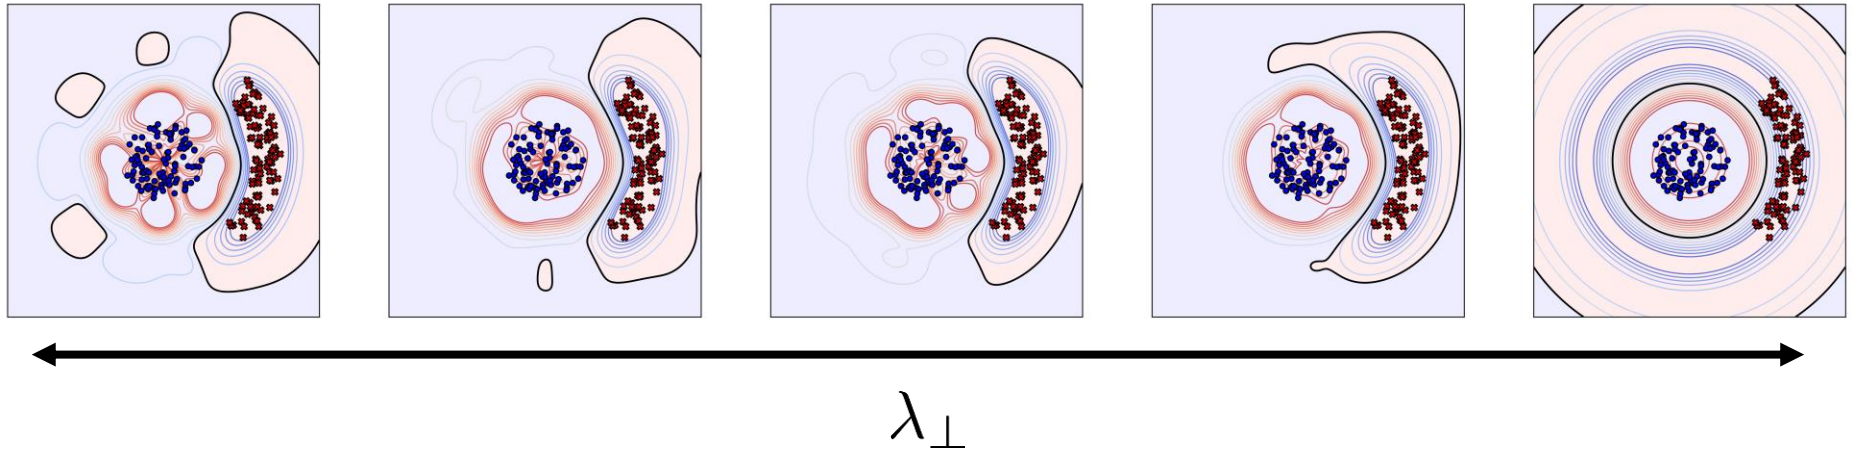

Achieves substantial runtime gains over sample-based regularisers

T. Berndt, J. Stühmer. "Approximate equivariance via projection-based regularization", ICML 2026

Experiments: Steerable 3D CNNs

Compared with

1. Non-equivariant 3D-CNN
   - no geometric inductive bias

2. $(S)O(3)$-Steerable 3D CNN
   - perfect equivariance undesirable in some cases

Network Group	Partial Equivariance	Synapse	Organ
CNN	N/A	0.716 ±0.008	0.920 ±0.003
$SO(3)$	None	0.738 ±0.009	0.607 ±0.006
	RPP	0.695 ±0.037	0.936 ±0.002
	P-SCNN	0.770 ±0.030	0.902 ±0.006
	ours	**0.774** ±0.025	0.933 ±0.000
$O(3)$	None	0.743 ±0.004	0.592 ±0.008
	RPP	0.722 ±0.023	0.940 ±0.006
	P-SCNN	0.769 ±0.005	0.905 ±0.004
	ours	0.773 ±0.025	**0.943** ±0.005

kidney left/right

femur left/right

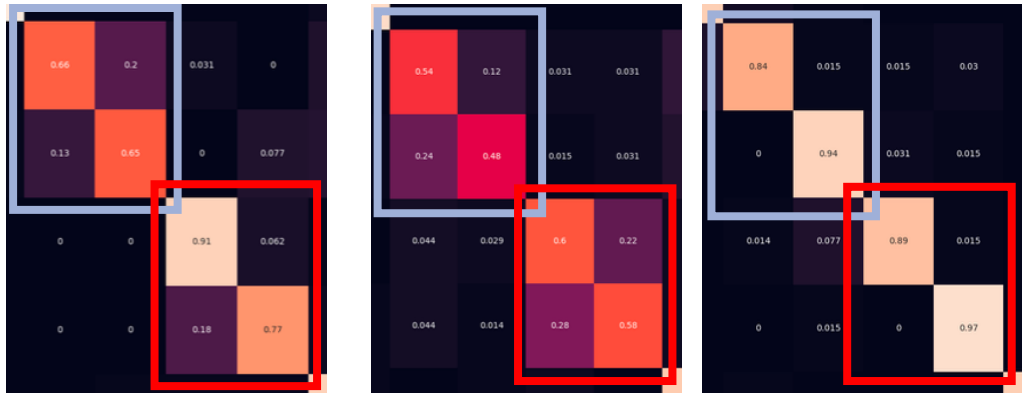

3D CNN S-CNN ours

Experiments on MedMNIST v2 organ classification tasks.

Summary - Approximate Equivariance

1. A theoretically-grounded approach to regularize general machine learning models towards exact equivariance
2. Use the orthogonal decomposition of functions into equivariant and non-equivariant components to penalize non-equivariance on an operator level over the whole group orbit
3. The closed-form projection can be computed efficiently in the frequency domain
4. Empirically demonstrate improvement over existing approaches for approximate equivariance

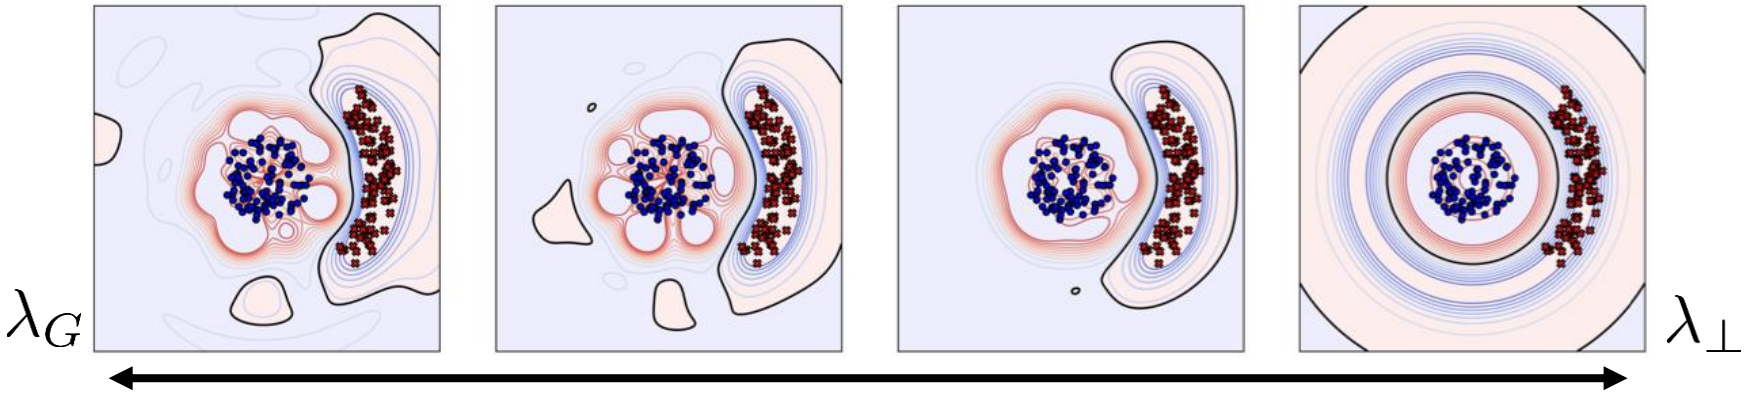

Protein Design

FrameDiff and FrameFlow [Yim et al. ‘23, ‘24]

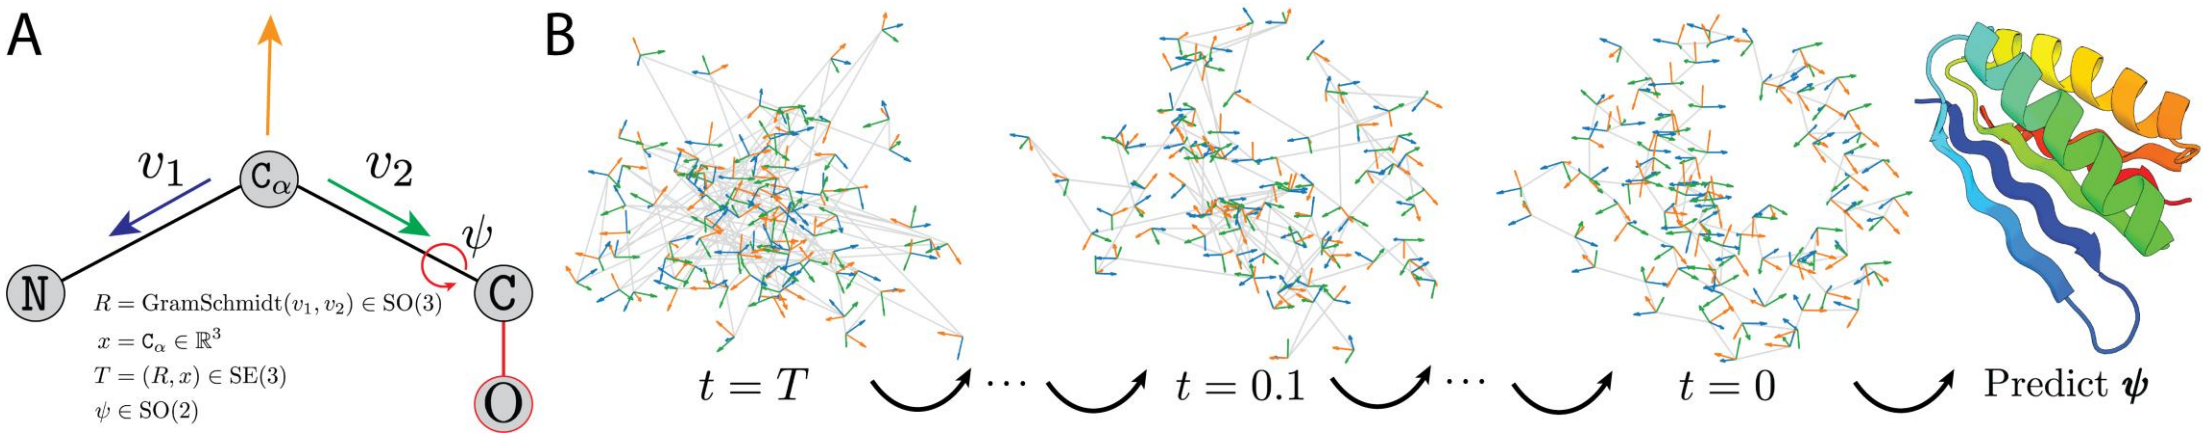

Yim et al., SE(3) diffusion model with application to protein backbone generation, ICML 2023

[Yim et al. 2023]

FrameDiff and FrameFlow [Yim et al. ‘23, ‘24]

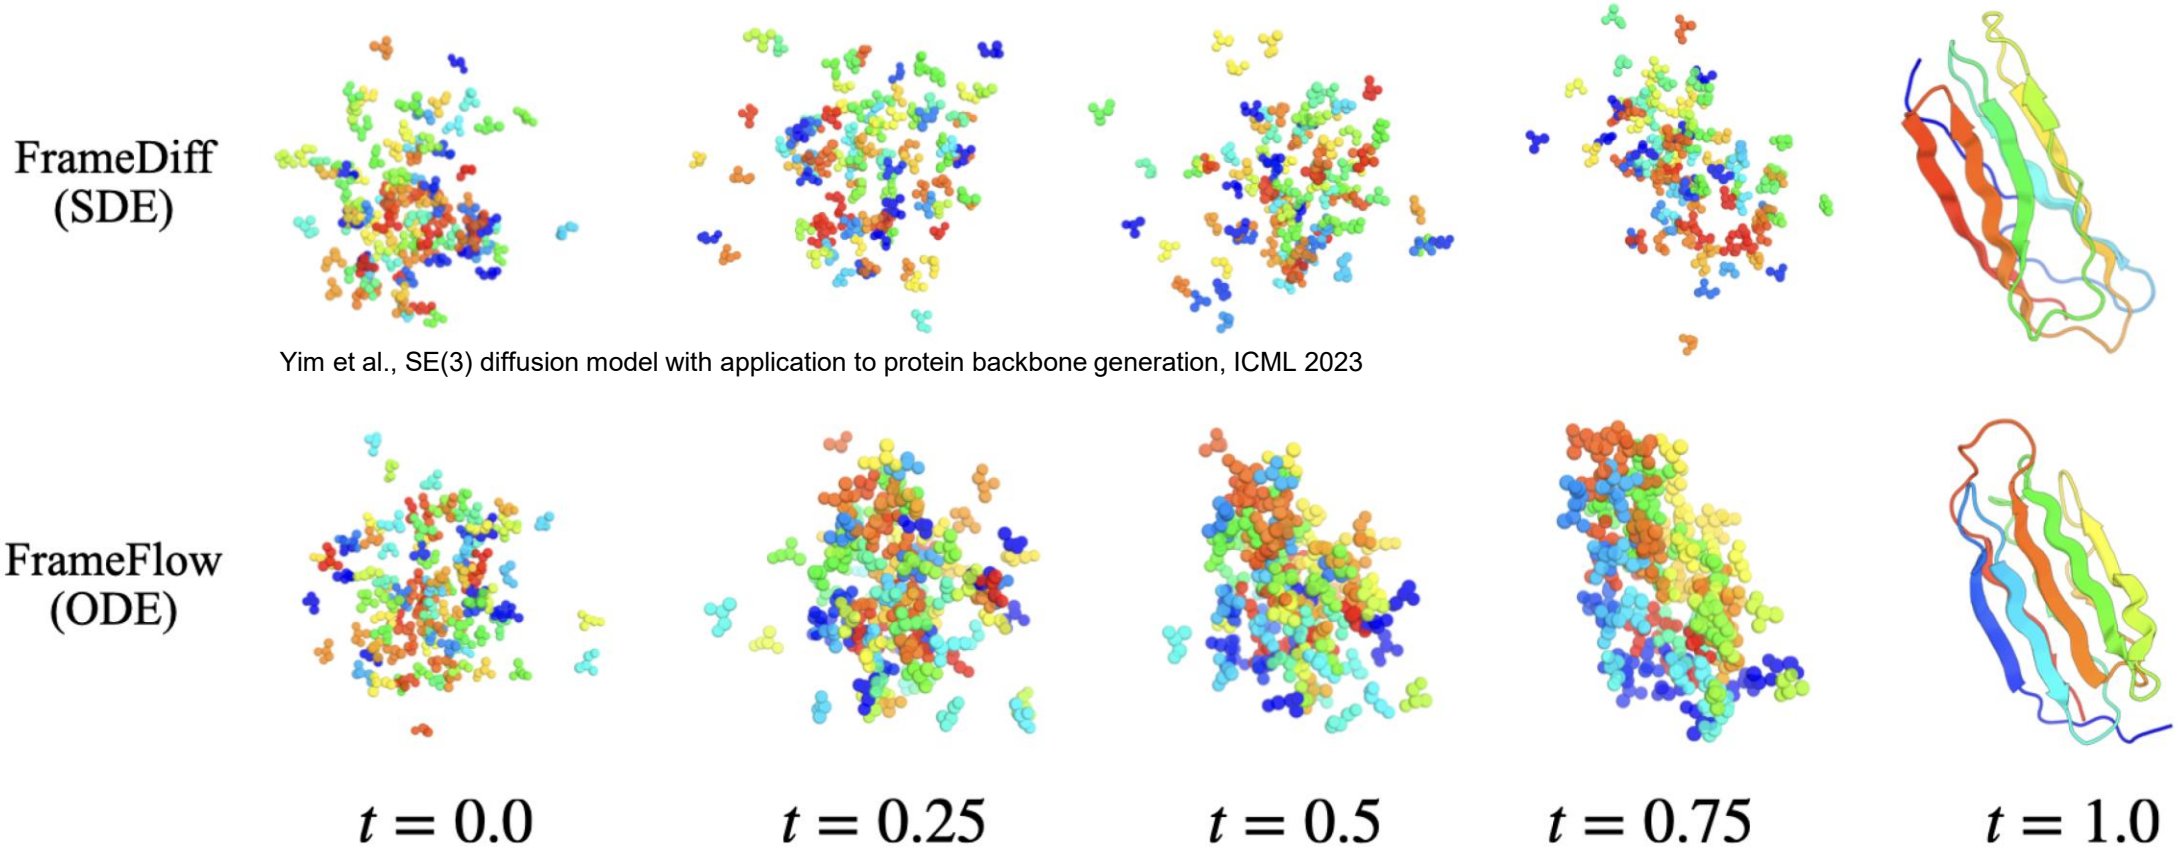

Yim et al., SE(3) diffusion model with application to protein backbone generation, ICML 2023

Yim et al., Fast protein backbone generation with SE (3) flow matching, arxiv 2023
Yim et al., Improved motif-scaffolding with {SE}(3) flow matching, TMLR 2024

[Yim et al. 2023]

Local Frames and Invariant Representations

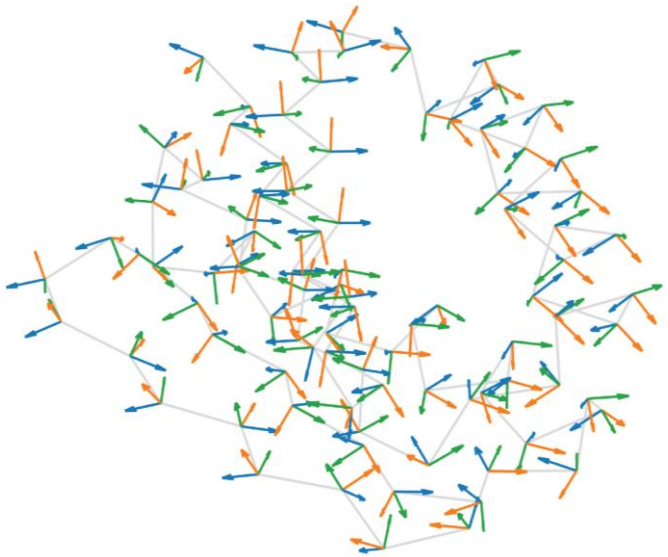

Protein backbone of oriented frames

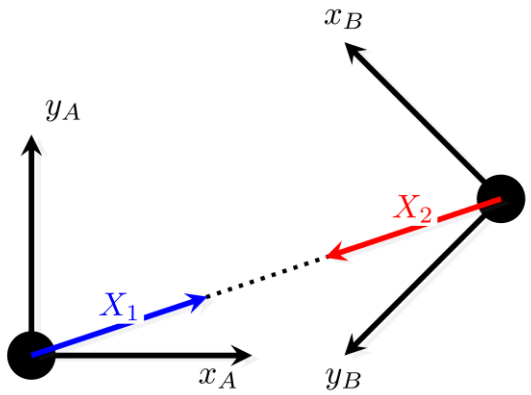

Nodes with attached local frame

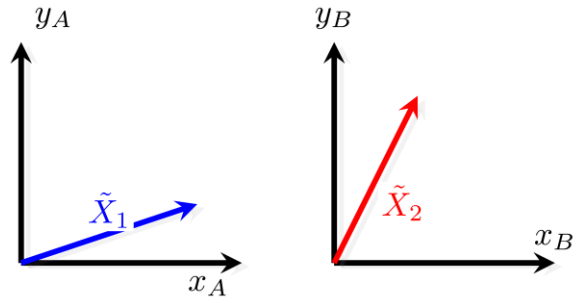

[Simon Wagner]

Invariant node features represented in their local frame

Flow Matching for Protein Backbone Generation

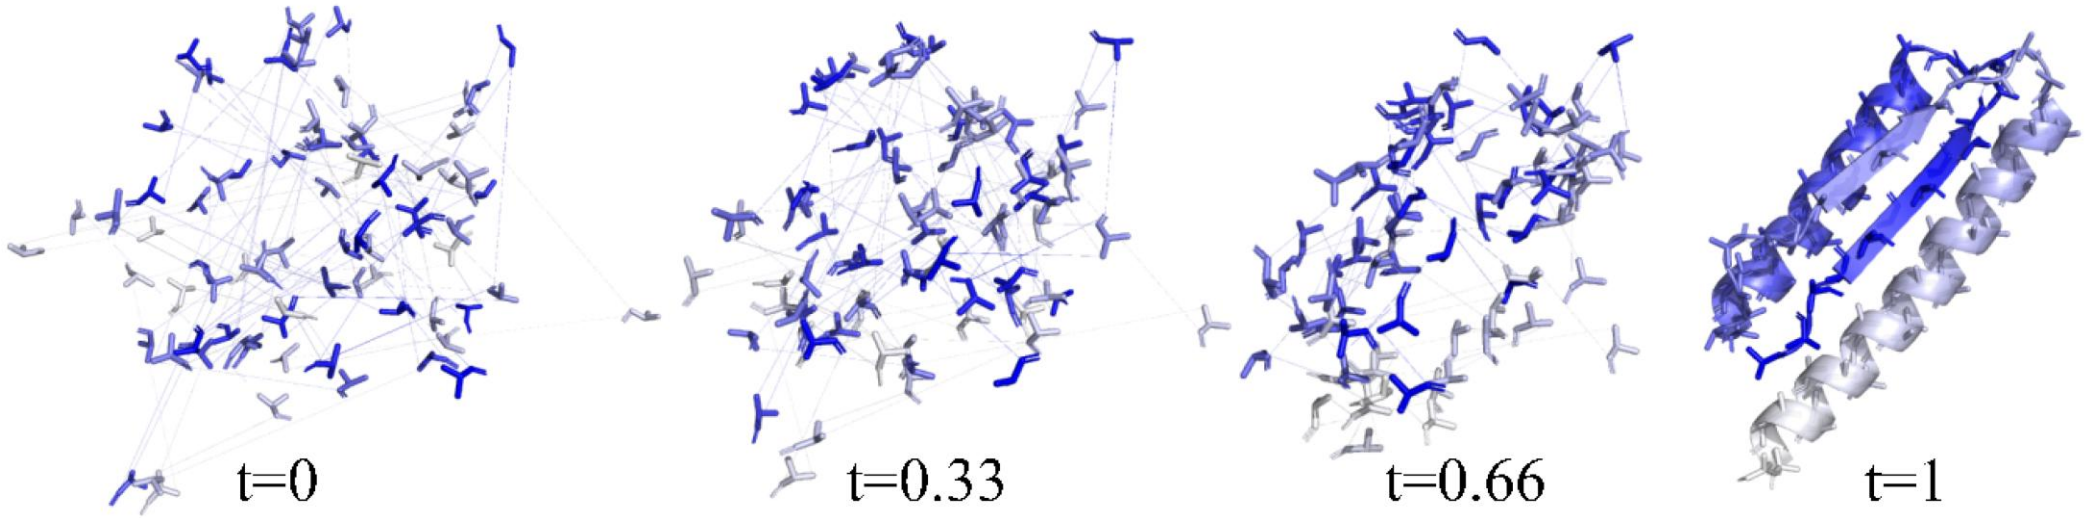

[Simon Wagner]

FrameFlow Architecture

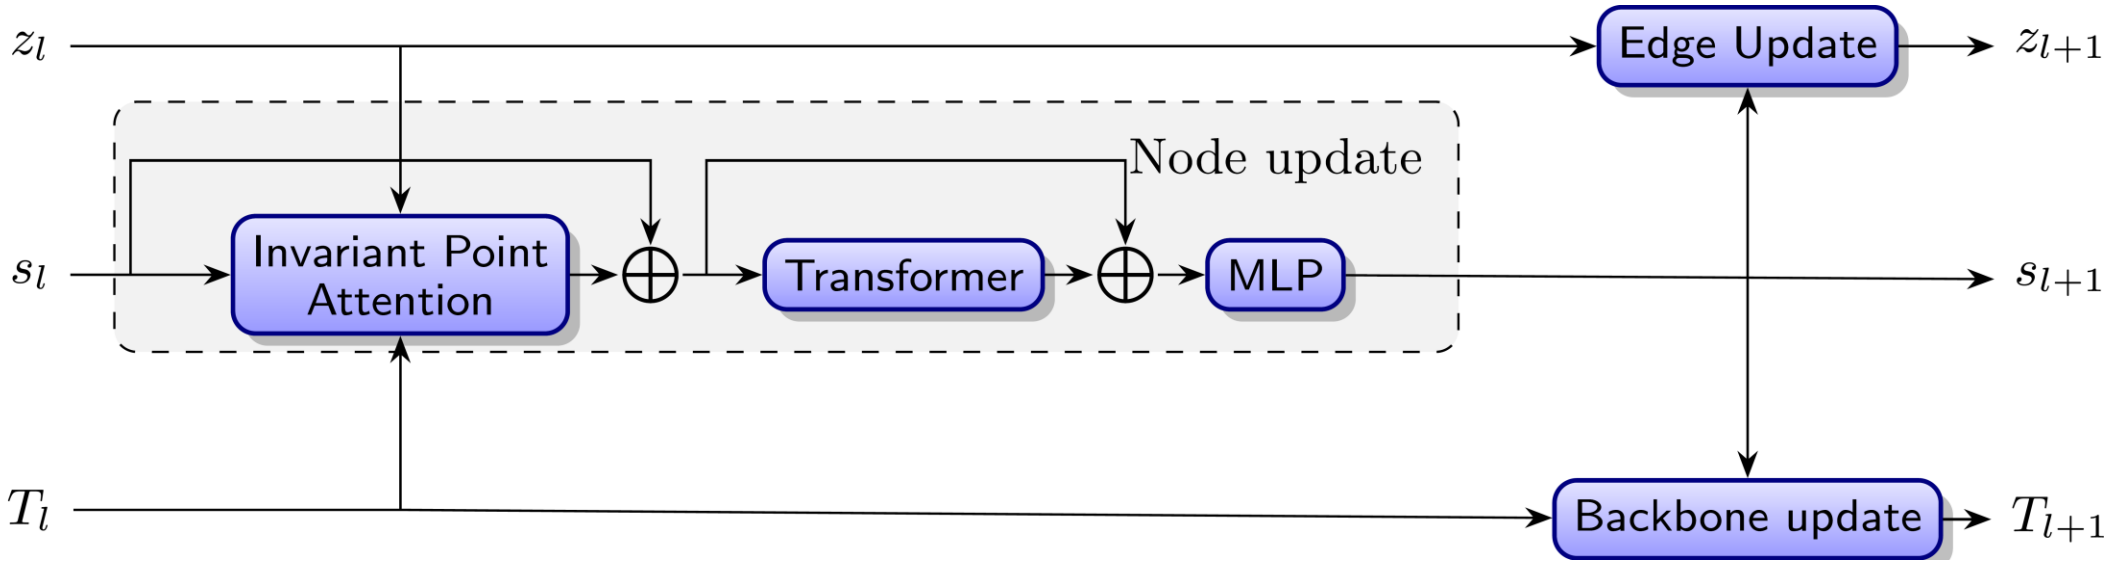

Geometric Algebra Flow Matching Architecture

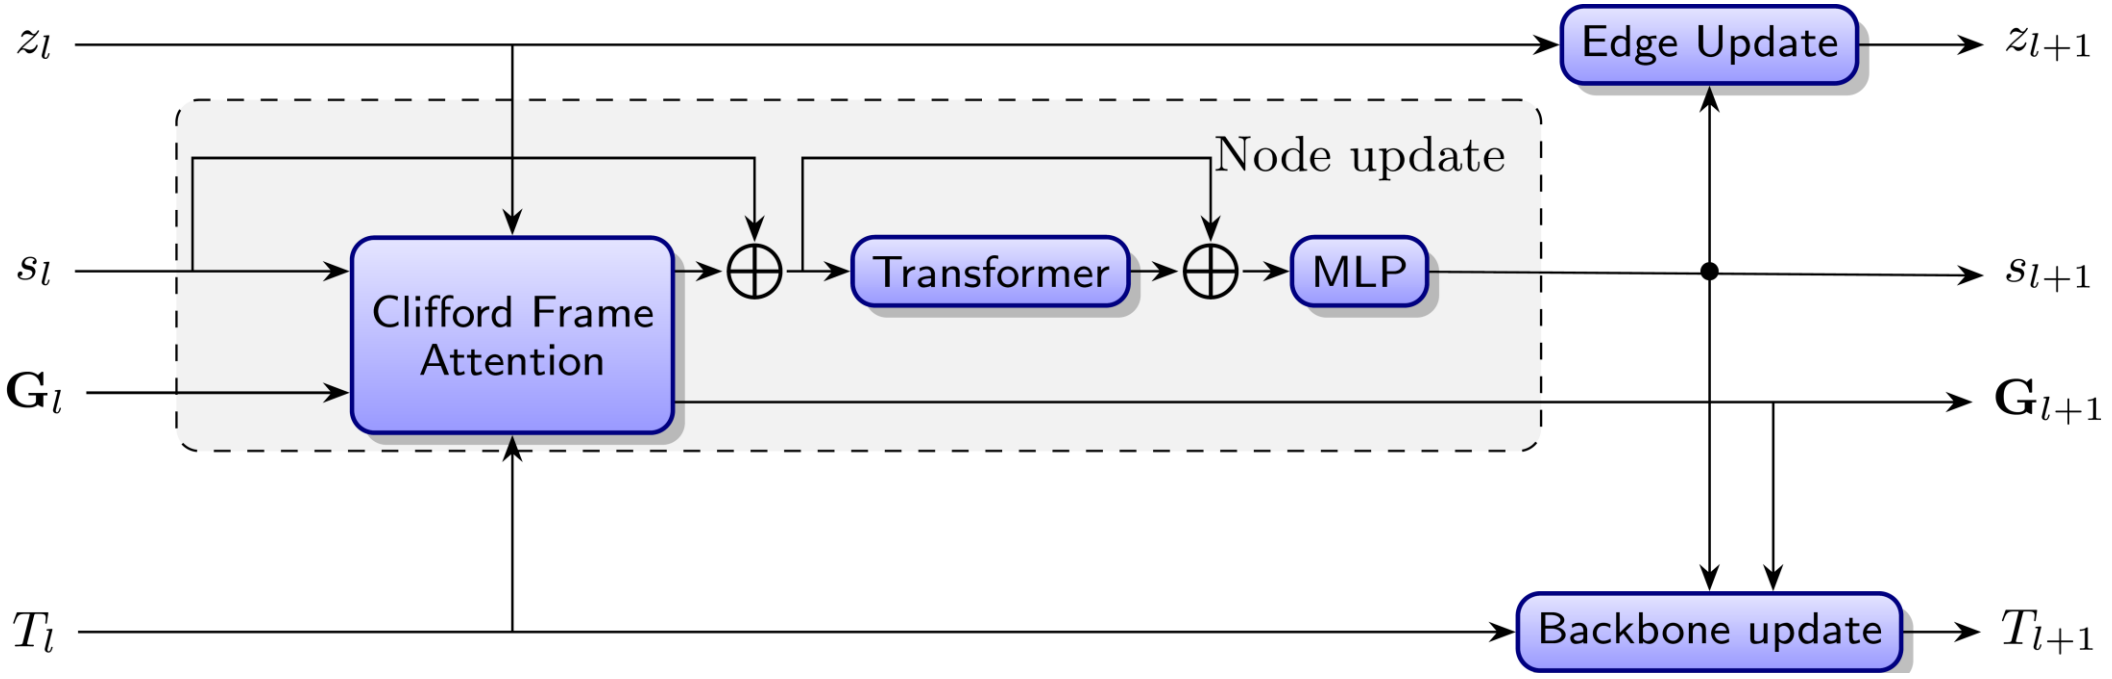

Geometric Algebra

- Recently introduced to ML/DL in the Geometric Algebra Transformer [Brehmer et al. 2023]
- Elements of the Euclidean geometric algebra $\mathbb{G}_3$ are *multivectors*

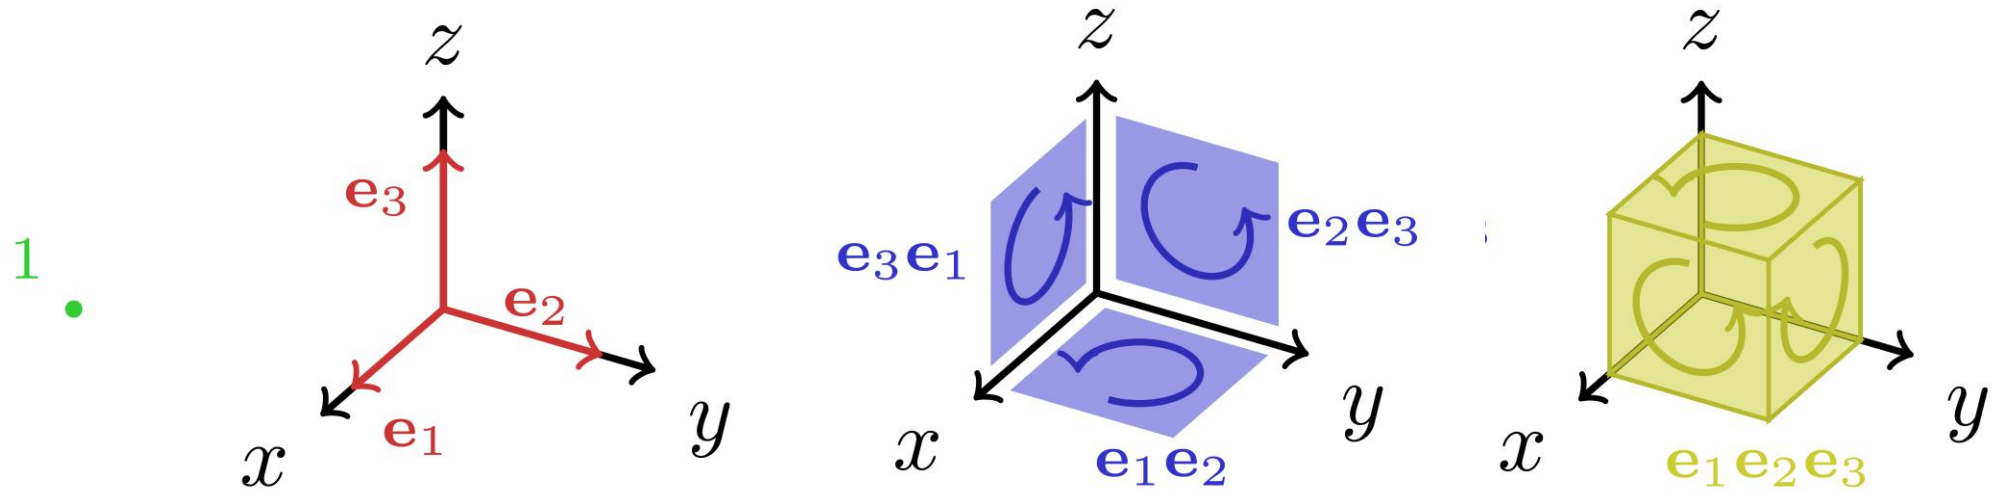

scalar vector bivector trivector

Geometric Product

$$\mathbf{ab} = \mathbf{a} \cdot \mathbf{b} + \mathbf{a} \wedge \mathbf{b}.$$

- A geometric algebra introduces a bilinear operation, the *geometric product*
- It is defined by its properties: Associativity, Distributivity, and that vectors square to scalars

$$aa = a^2 \in \mathbb{R}$$

- It can be separated into the inner product and the outer product

Projective Geometric Algebra – Bilinear Operations

JOIN

Join Operation	Illustration
Line containing points **p** and **q**. $\mathbf{p} \wedge \mathbf{q} = (p_w q_x - p_x q_w)\mathbf{e}_{41} + (p_w q_y - p_y q_w)\mathbf{e}_{42} + (p_w q_z - p_z q_w)\mathbf{e}_{43} + (p_y q_z - p_z q_y)\mathbf{e}_{23} + (p_z q_x - p_x q_z)\mathbf{e}_{31} + (p_x q_y - p_y q_x)\mathbf{e}_{12}$	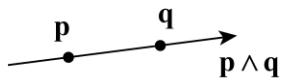
Plane containing line ***l*** and point **p**. $\boldsymbol{l} \wedge \mathbf{p} = (l_{vy} p_z - l_{vz} p_y + l_{mx} p_w)\mathbf{e}_{423} + (l_{vz} p_x - l_{vx} p_z + l_{my} p_w)\mathbf{e}_{431} + (l_{vx} p_y - l_{vy} p_x + l_{mz} p_w)\mathbf{e}_{412} - (l_{mx} p_x + l_{my} p_y + l_{mz} p_z)\mathbf{e}_{321}$	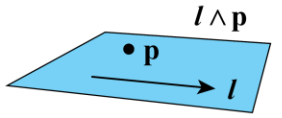

MEET

Meet Operation	Illustration
Line where planes **g** and **h** intersect. $\mathbf{g} \vee \mathbf{h} = (g_z h_y - g_y h_z)\mathbf{e}_{41} + (g_x h_z - g_z h_x)\mathbf{e}_{42} + (g_y h_x - g_x h_y)\mathbf{e}_{43} + (g_x h_w - g_w h_x)\mathbf{e}_{23} + (g_y h_w - g_w h_y)\mathbf{e}_{31} + (g_z h_w - g_w h_z)\mathbf{e}_{12}$	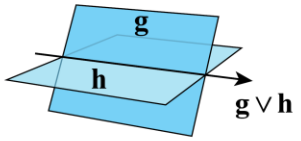
Point where plane **g** and line ***l*** intersect. $\mathbf{g} \vee \boldsymbol{l} = (g_z l_{my} - g_y l_{mz} + g_w l_{vx})\mathbf{e}_1 + (g_x l_{mz} - g_z l_{mx} + g_w l_{vy})\mathbf{e}_2 + (g_y l_{mx} - g_x l_{my} + g_w l_{vz})\mathbf{e}_3 - (g_x l_{vx} + g_y l_{vy} + g_z l_{vz})\mathbf{e}_4$	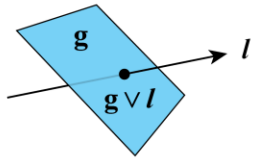

[Eric Lengyel, projectivegeometricalgebra.org]

Elements of the Projective Geometric Algebra (PGA)

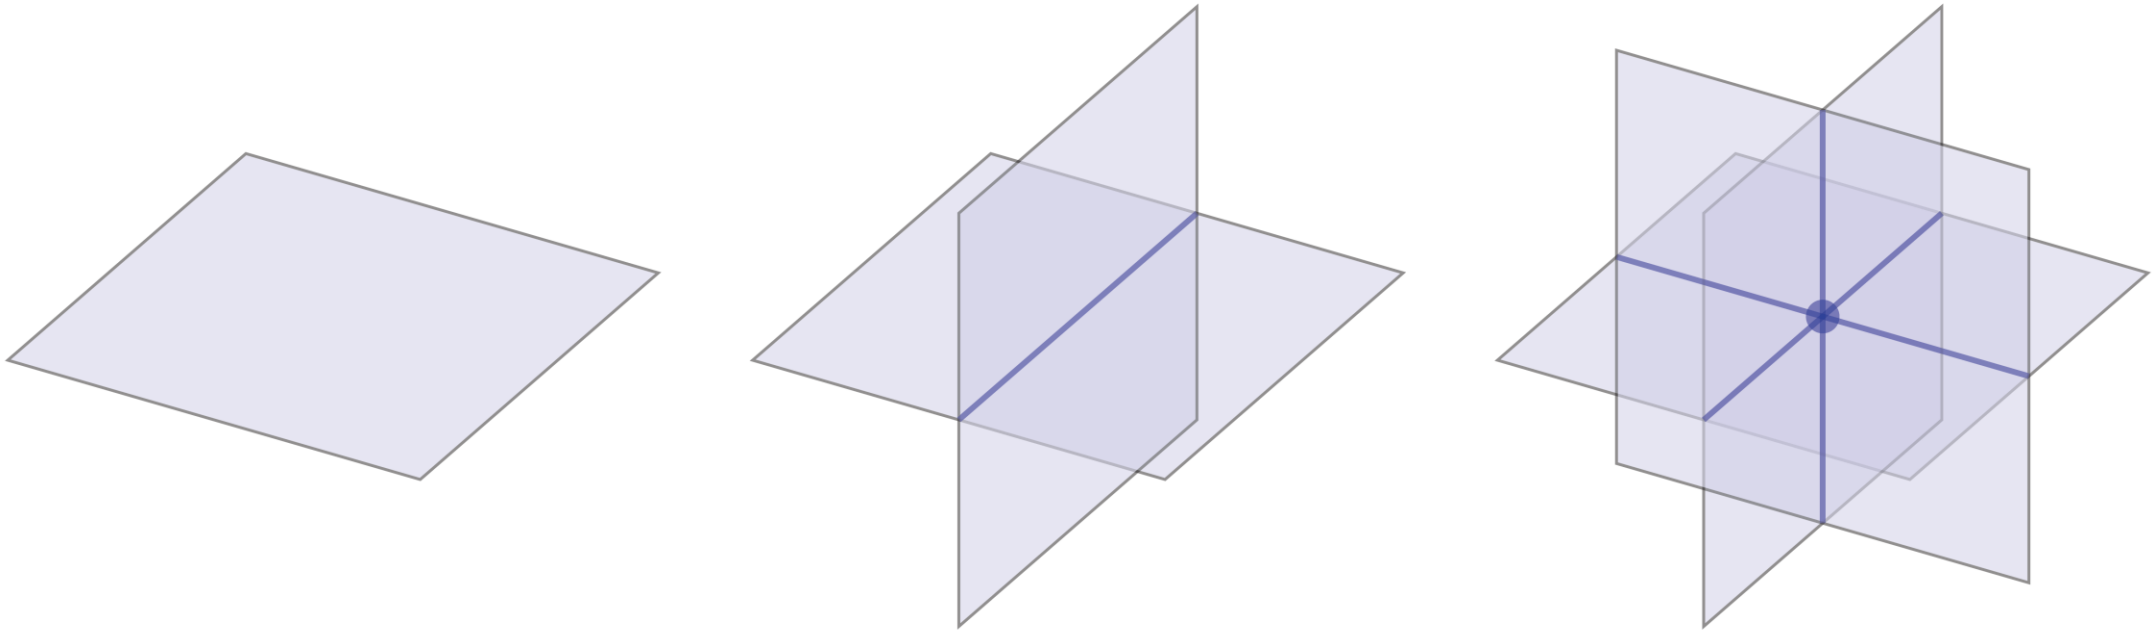

Vectors represent planes, 2-blades represent lines resulting from the intersecting of its generating planes and 3-blades represent points as the intersection of three planes.

Representing Frames in the Projective Geometric Algebra

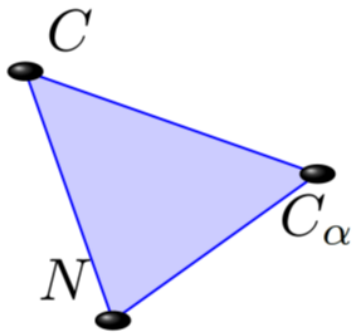

Protein Residue

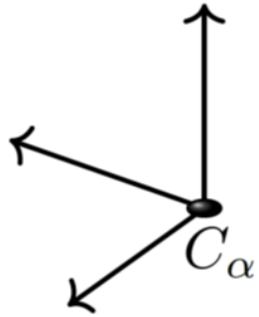

Residue Frame

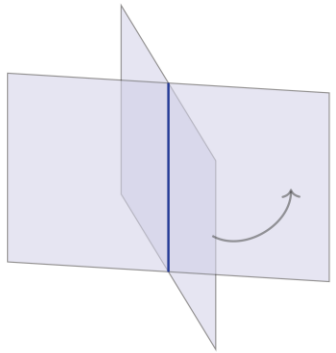

Rotation by reflection through two intersecting planes

and

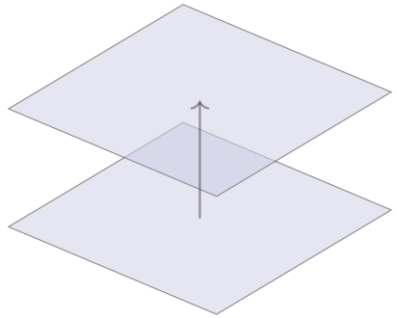

Translation by reflection through two parallel planes

Higher Order Messages

- The MACE-trick [1]: distributivity of the bilinear product can be used to construct higher order messages:

$$\mathbf{ab} = \mathbf{a} \cdot \mathbf{b} + \mathbf{a} \wedge \mathbf{b}.$$

$$\left(\sum_j m_{ij}\right)\left(\sum_k m'_{ik}\right) = \sum_{jk} m_{ij} m'_{ik} \equiv \sum_{jk} m^{(3)}_{ijk}$$

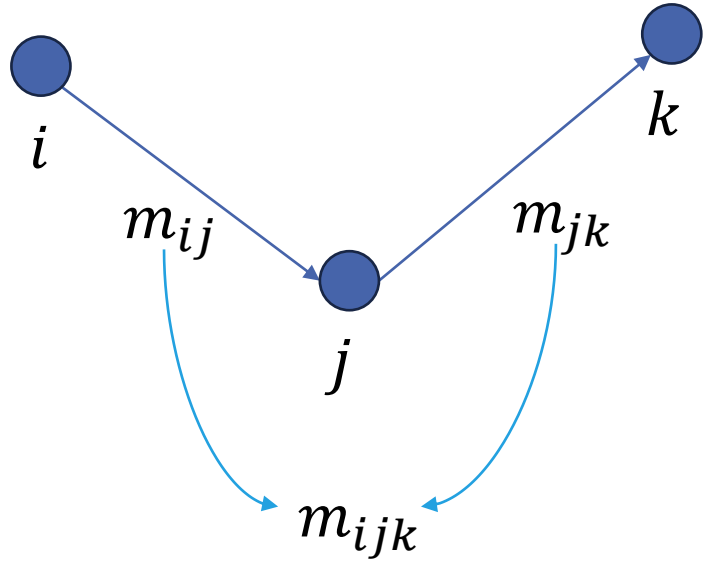

[1] Batatia, Kovacs, Simm, Ortner, Csányi. Mace: Higher order equivariant message passing neural networks for fast and accurate force fields. NeurIPS 2022

Designability

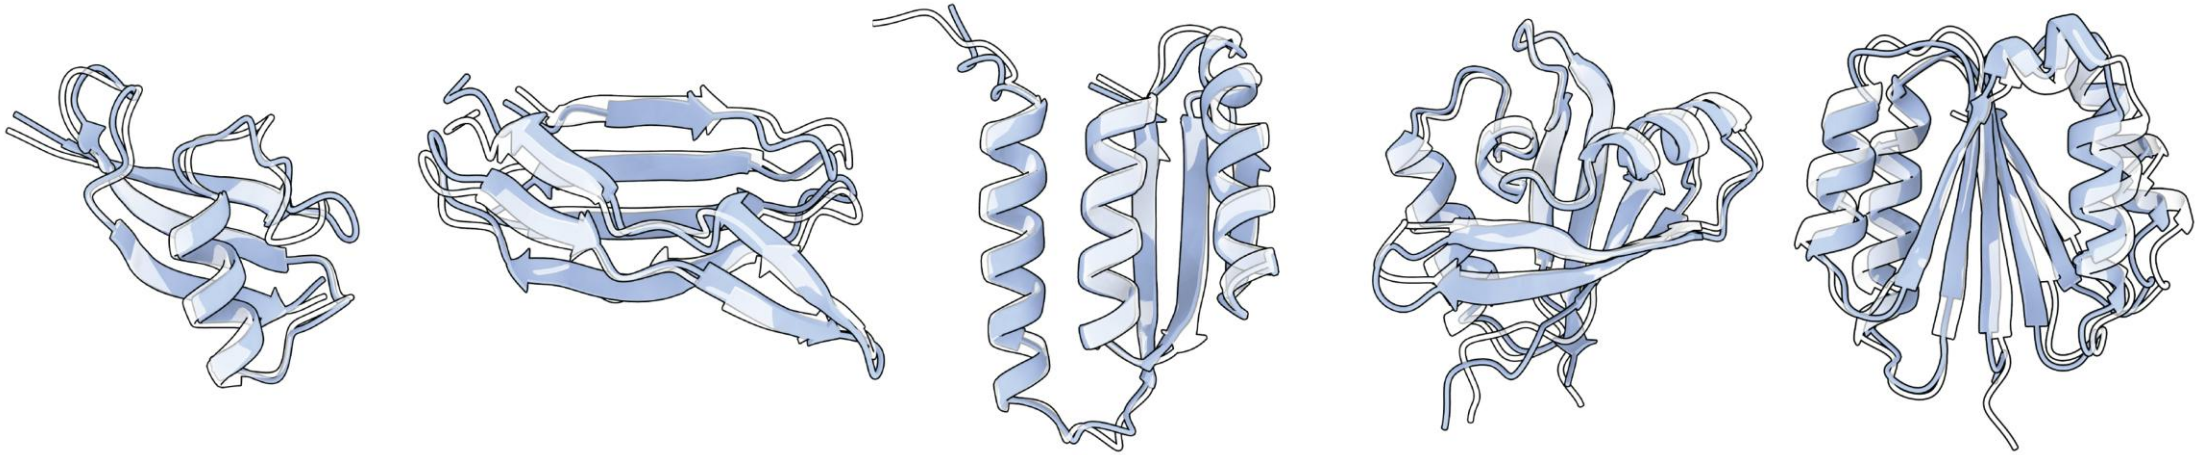

Blue color encodes the structure of a ProteinMPNN-generated sequence refolded with ESMfold and superposed onto the target sample colored in white

Designability and Secondary Structure (PDB)

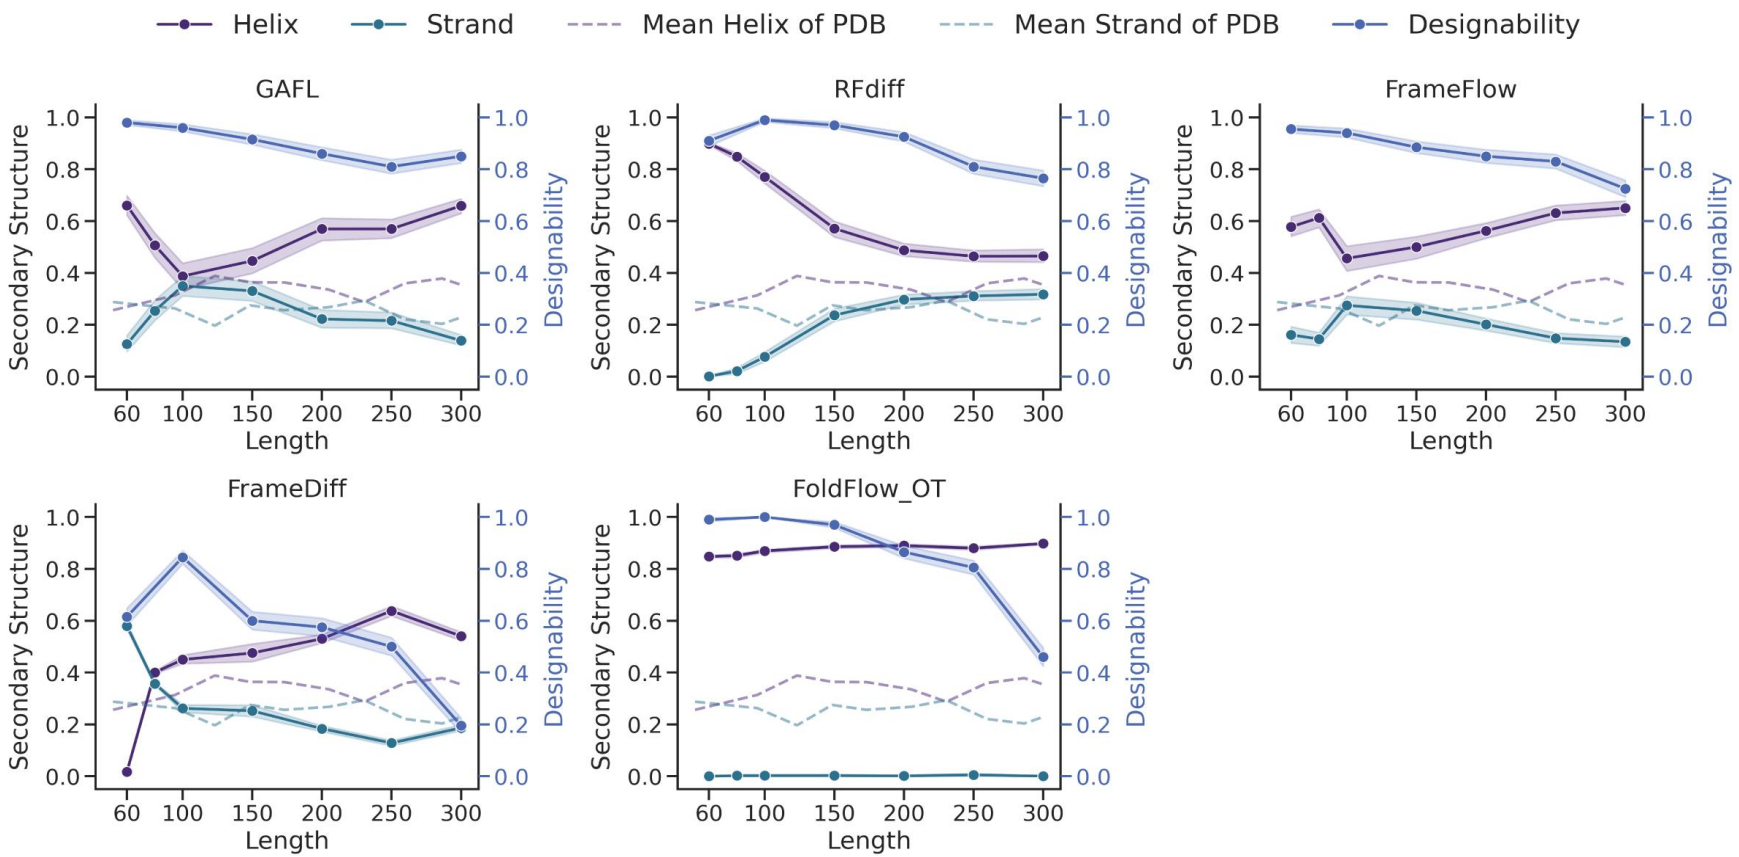

RFDiffusion over-represents α-helices for short protein lengths.

Distribution of Secondary Structure (PDB 300)

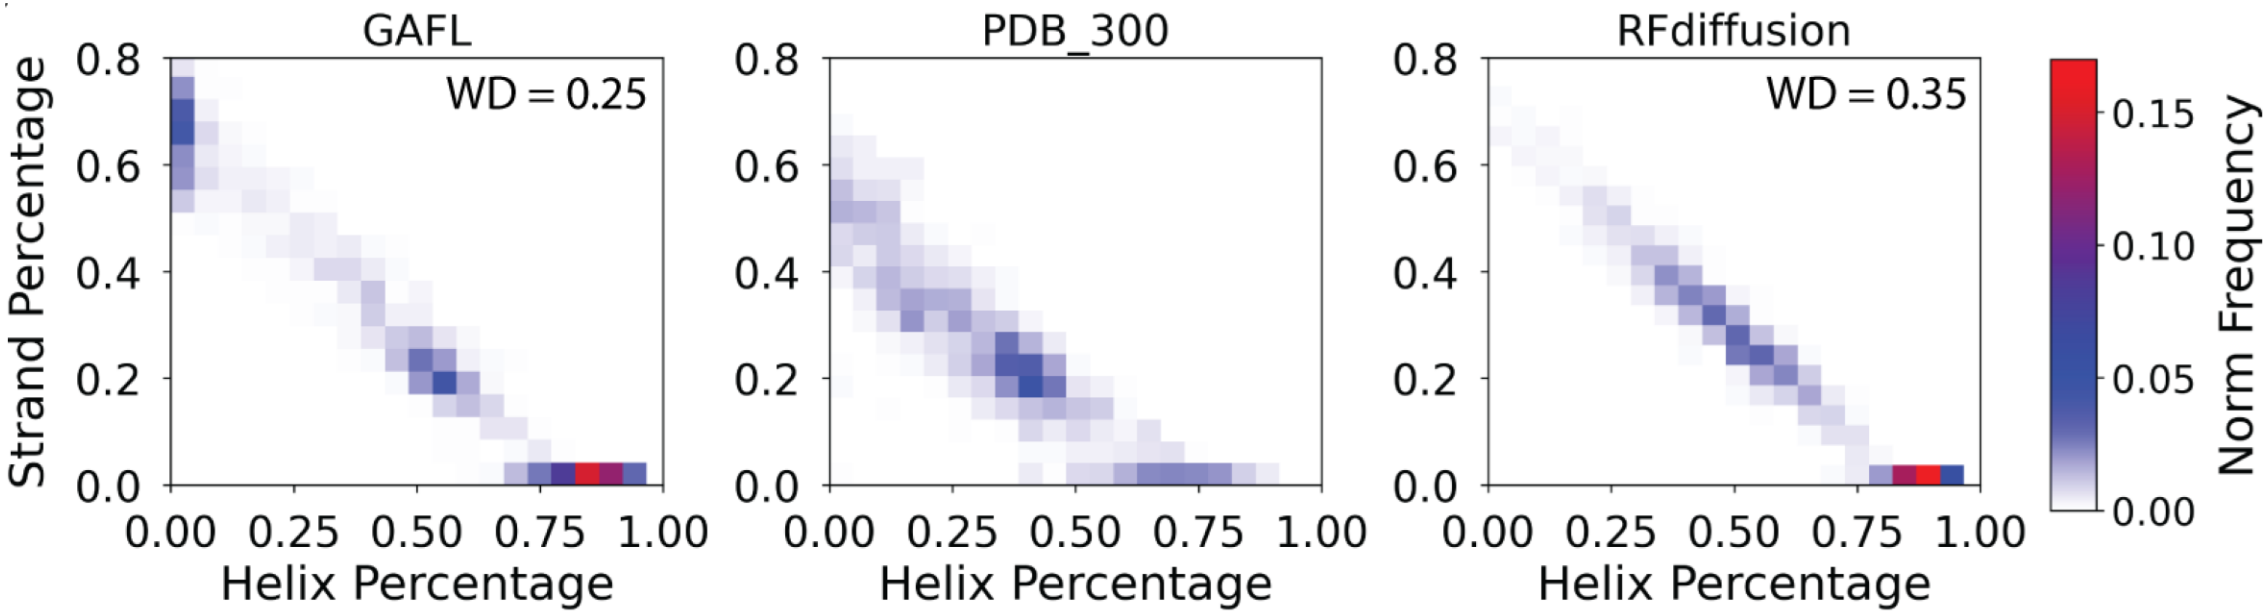

Distribution of β-strands and α-helices of sampled backbones generated by GAFL and.
Trained on PDB up to length of 300.

Evaluation on PDB (up to 300 residues)

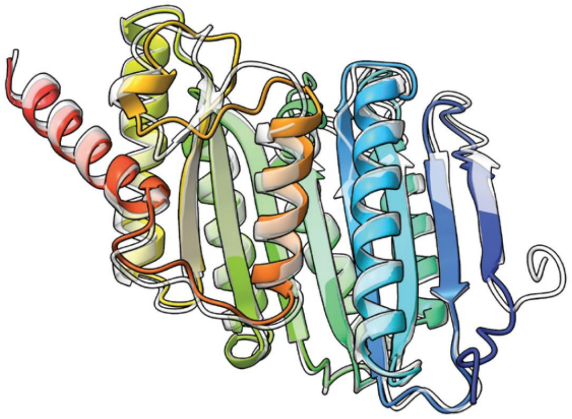

Method	Designability (↑)	Diversity (↓)	Novelty (↓)	Helix Content	Strand Content	Time [s]
PDB Dataset (300)	-	-	-	0.39 (0.00)	0.23 (0.00)	
FrameDiff	0.54 (0.02)	0.45 (0.00)	0.71 (0.00)	**0.53** (0.01)	0.20 (0.00)	24.3
FoldFlow-SFM	0.69 (0.01)	0.44 (0.00)	0.77 (0.00)	0.91 (0.00)	0.01 (0.00)	24.3
FoldFlow-OT	0.82 (0.01)	0.44 (0.00)	0.79 (0.00)	0.88 (0.00)	0.00 (0.00)	24.3
FrameFlow	0.85 (0.01)	**0.35** (0.00)	**0.70** (0.00)	0.56 (0.01)	0.20 (0.00)	6.6
RFdiffusion*	**0.89** (0.01)	0.37 (0.00)	0.74 (0.00)	0.58 (0.02)	**0.24** (0.02)	21.0
GAFL (ours)	**0.88** (0.01)	0.36 (0.00)	0.71 (0.00)	**0.53** (0.01)	**0.25** (0.01)	8.8

*Pretrained weights from folding model trained on dataset larger than PDB.

Protein backbones at lengths 100, 150, 200, 250, 300 (Sampled 200 each, in total 1,000 backbones).
Refolded by ESMfold.

Protein Samples (up to 300 residues)

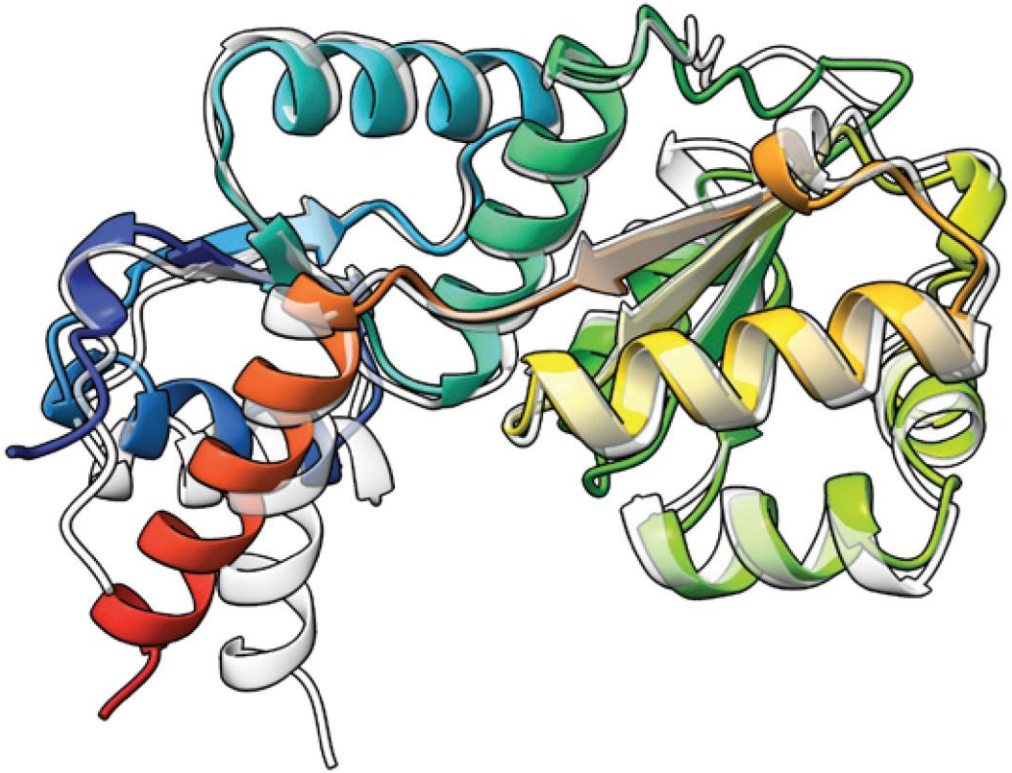

L = 200

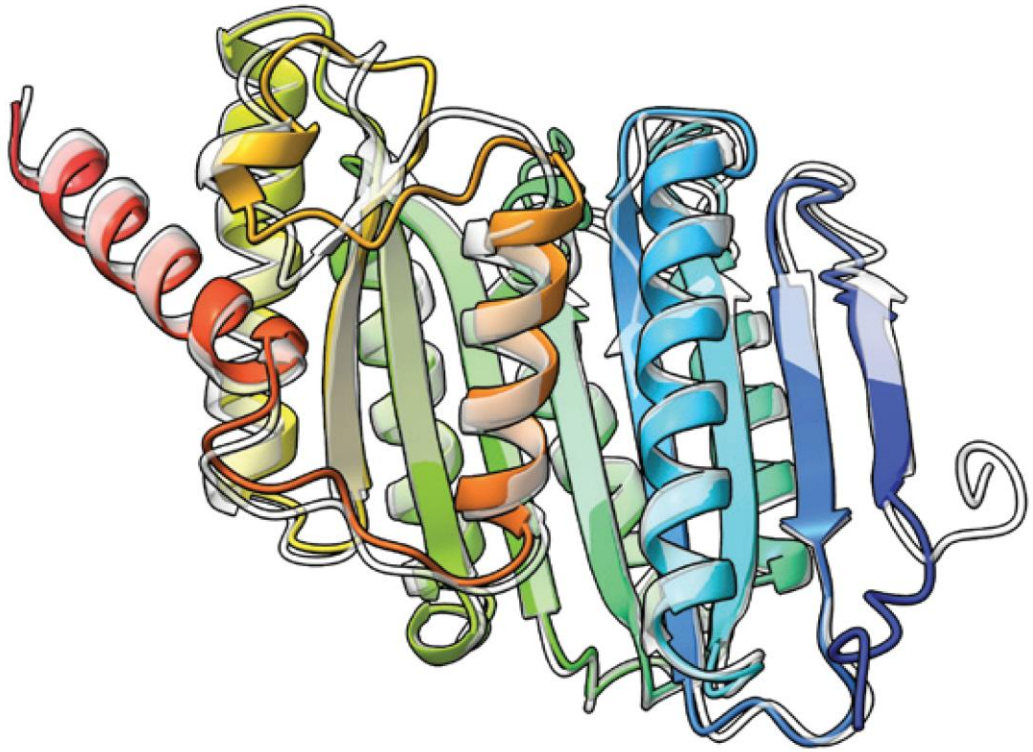

L = 300

Ablation Study (SCOPe-128)

PGA	Higher Ord. Msg.	Checkp. Selection	Designability [%] (↑)	Diversity (↓)
✓	✓	✓	$90.5\,^{90.6}_{89.6}$	$0.38\,^{0.39}_{0.36}$
✓		✓	$89.6\,^{90.3}_{89.0}$	$0.37\,^{0.40}_{0.37}$
		✓	$88.2\,^{88.6}_{86.7}$	$0.38\,^{0.39}_{0.38}$
FrameFlow-SCOPe			81.2	0.37

Ablation of GAFL models with different elements of the Clifford Frame Attention architecture
Trained on the SCOPe-128 dataset

Conditional Sampling of Conformational Ensembles

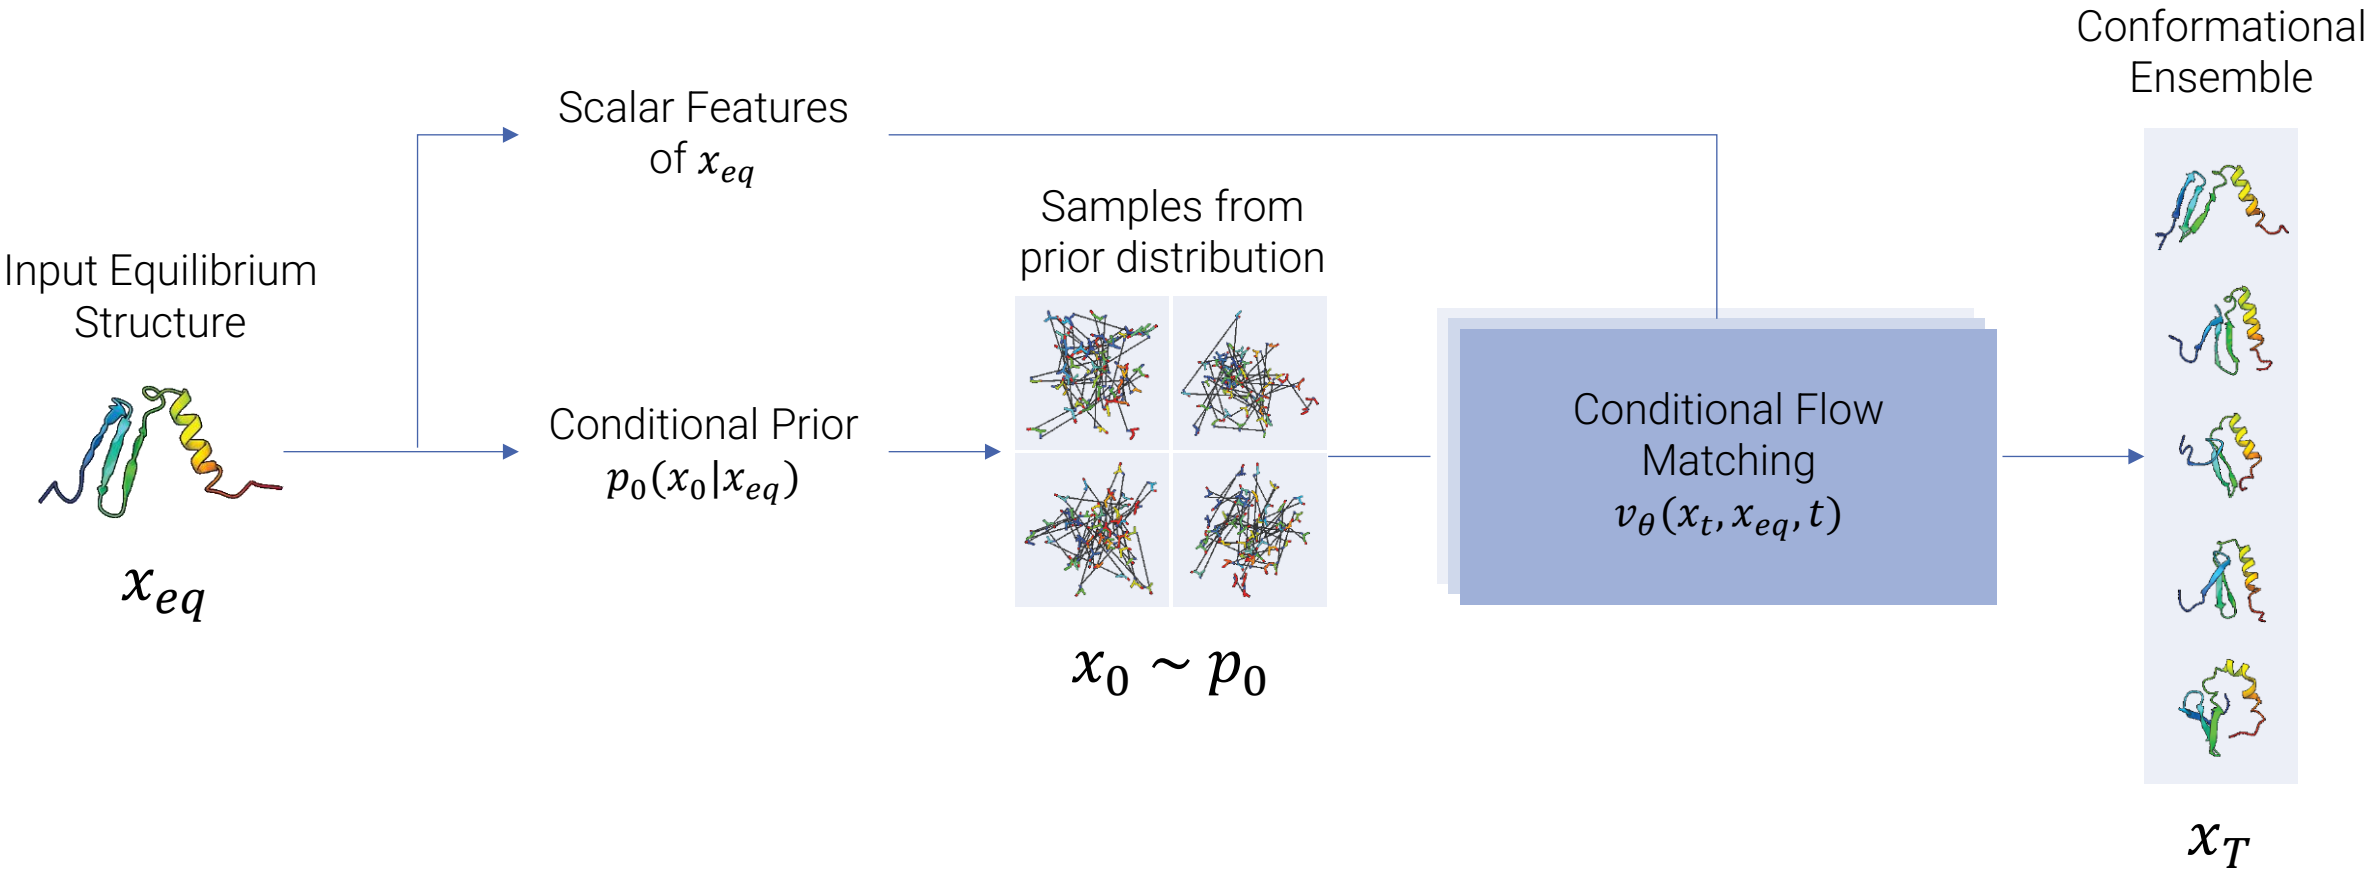

Examples of Sampled Ensembles

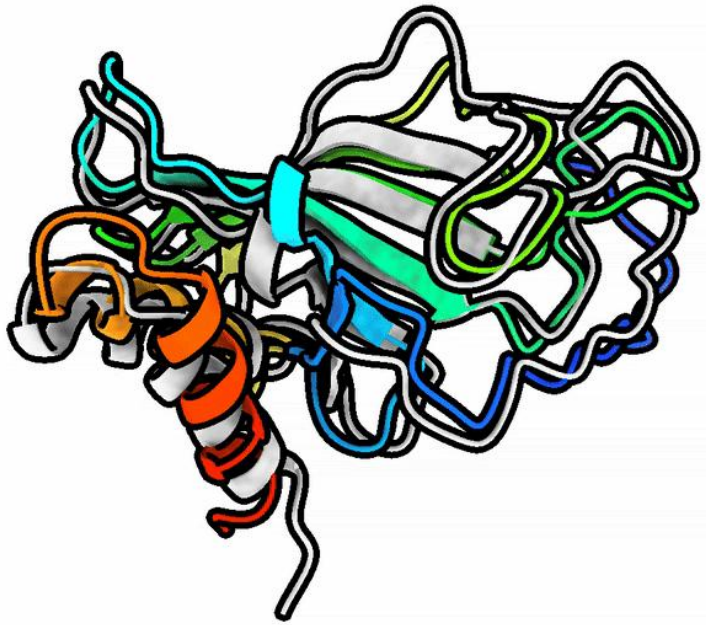

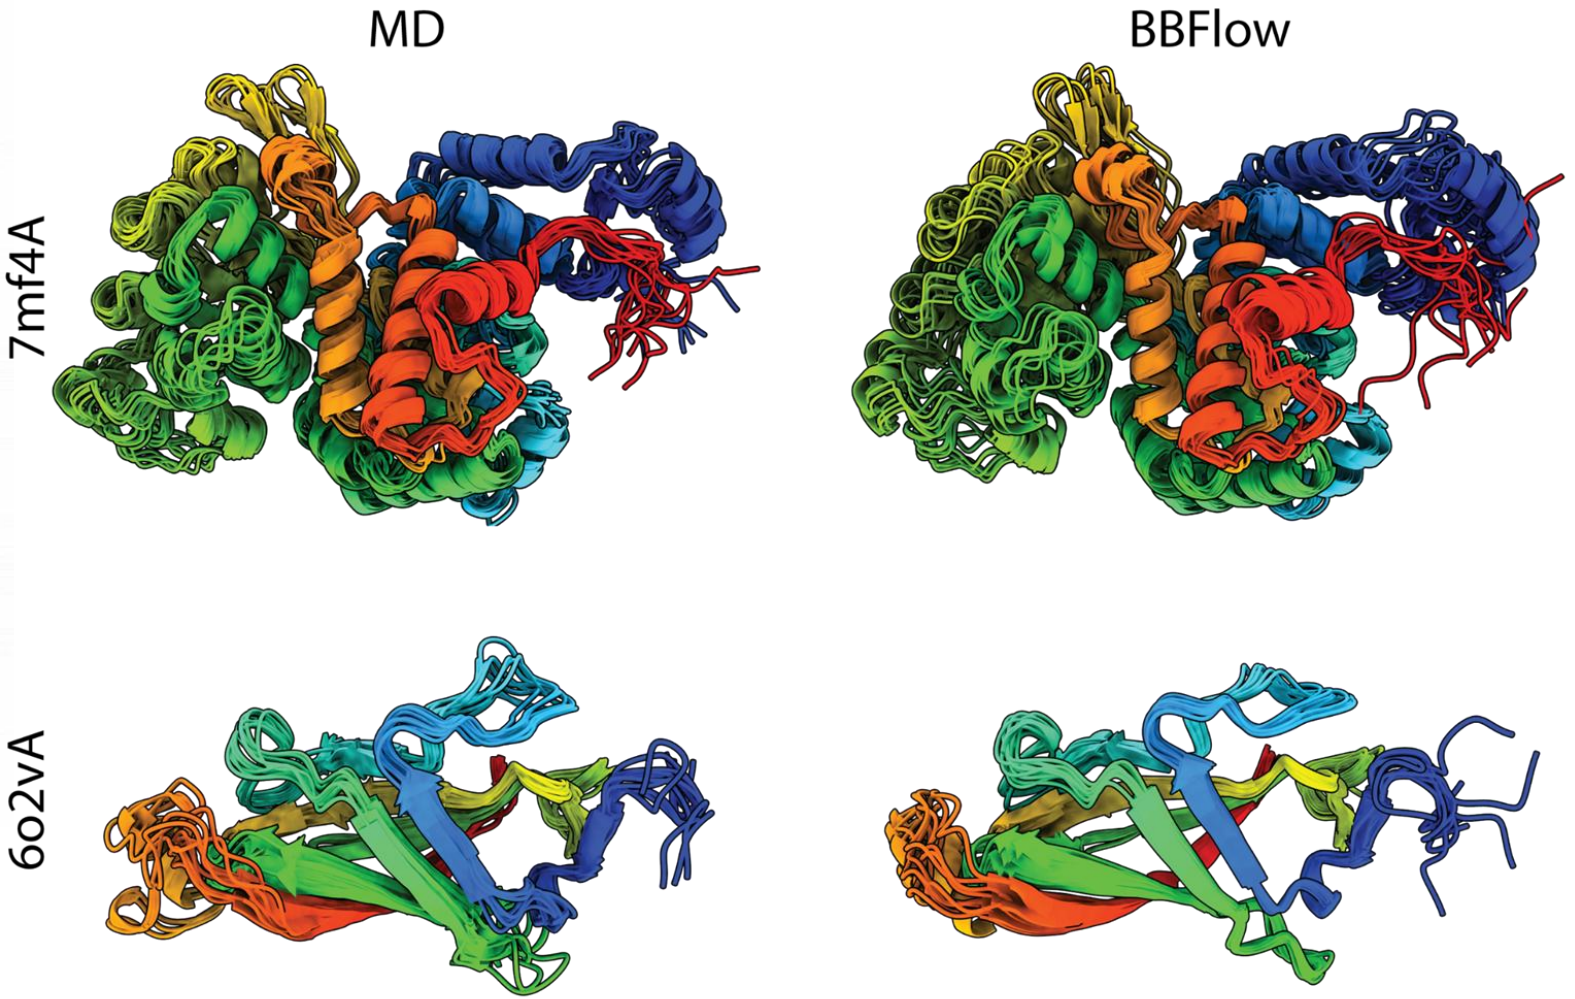

Flexibility conditioned protein design

- Current generative models for protein design sample overly rigid proteins
- Introduce a per-residue flexibility as condition for conditional sampling:

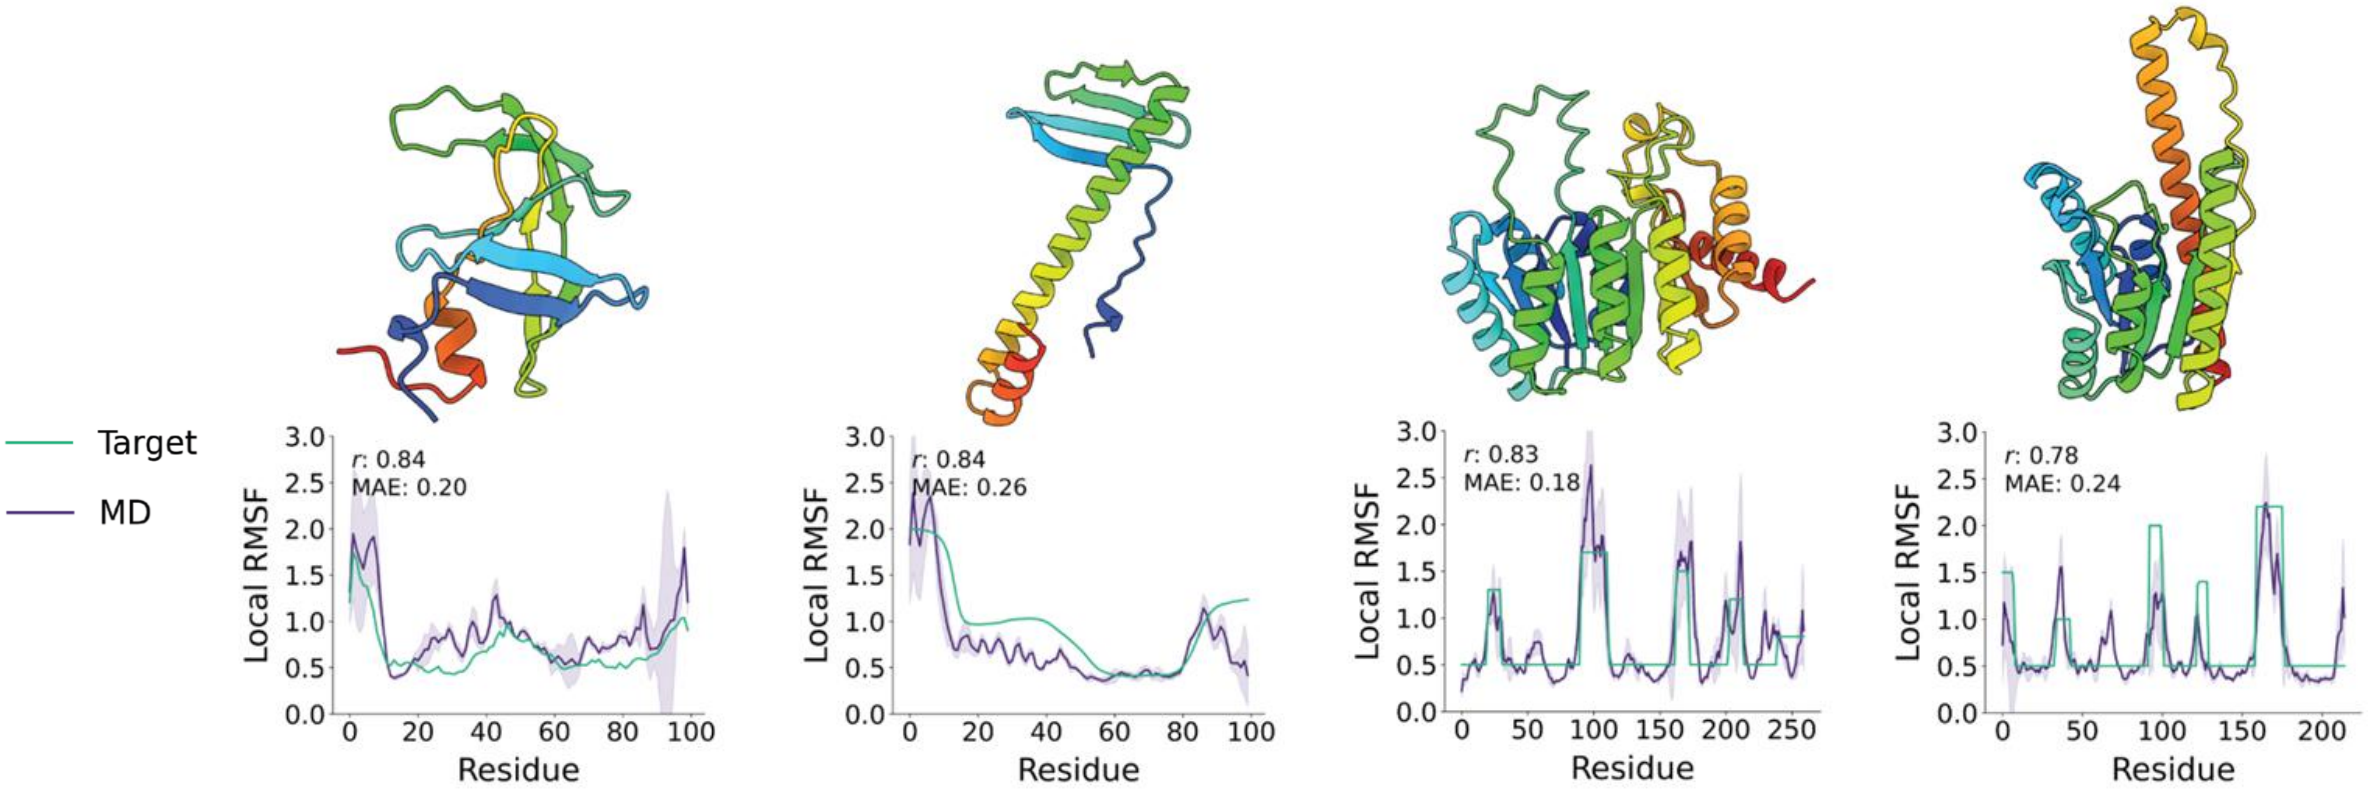

Conclusion

- Local frames provide a framework for building equivariant models by defining a local gauge
- Steerable representations can be used to define equivariant group convolution
- Multivector representations can be used as steerable basis:
  - spherical harmonics
  - geometric algebra
- This enables exciting applications in molecular Property prediction and protein design

References

- Dimenet
  - Gasteiger, J., Groß, J., & Günnemann, S. Directional Message Passing for Molecular Graphs. *International Conference on Learning Representations*. 2020.
- SEGNN
  - Brandstetter, J. et al. "Geometric and Physical Quantities improve E (3) Equivariant Message Passing." *International Conference on Learning Representations*. 2022.
- Approximate Equivariance
  - Berndt, T., Stühmer, J., Approximate Equivariance via Projection-Based Regularisation, *Proceedings of the 43rd International Conference on Machine Learning (ICML)*. 2026.
- FrameDiff
  - Yim, J et al. "SE (3) diffusion model with application to protein backbone generation." *Proceedings of the 40th International Conference on Machine Learning (ICML)*. 2023.
- FrameFlow
  - Yim, J., Campbell, A., Foong, A. Y., Gastegger, M., Jiménez-Luna, J., Lewis, S., ... & Noé, F. Fast protein backbone generation with se (3) flow matching. *arXiv preprint arXiv:2310.05297*. 2023
- GAFL
  - Wagner, S., Seute, L., Viliuga, V., Wolf, N., Gräter, F., & Stühmer, J. Generating highly designable proteins with geometric algebra flow matching. *Neurips 2024*