ConforFlux: Particle-Guided Trunk Repulsion for Diverse Protein Conformations

Shosuke Suzuki[1] Toshiyuki Amagasa[2]

[1]Graduate School of Science and Technology, University of Tsukuba, Japan
[2]Center for Computational Sciences, University of Tsukuba, Japan

suzuki-2001/conforflux

1. Default sampling returns one conformation

AlphaFold 3-type predictors keep returning one of a protein's conformations.

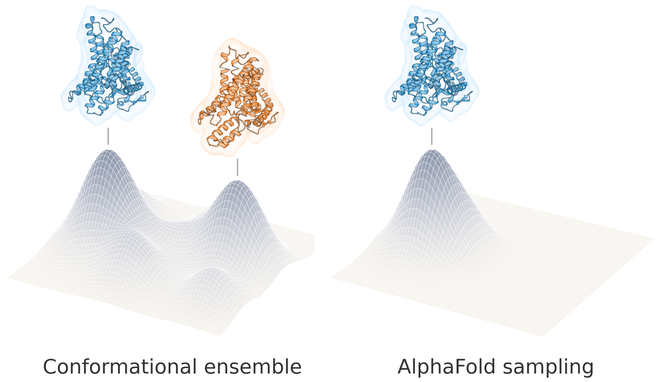
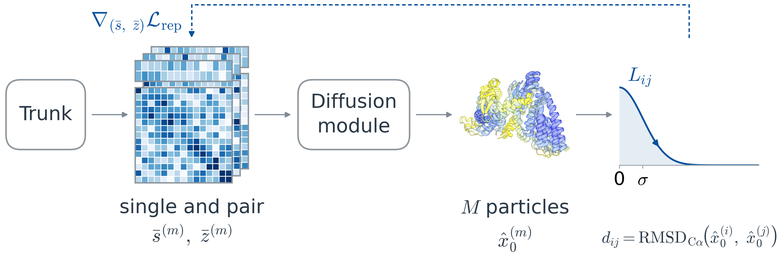


Figure 1. Schematic. States a protein occupies (left) and default samples (right).

2. Method: repulsion on the trunk embeddings

AlphaFold 3-type predictors run a heavy trunk once, producing single and pair representations $(\bar{s}, \bar{z})$, and a diffusion module F_θ that denoises coordinates conditioned on them at every block. The pair conditioning does not depend on the timestep, so a change to $(\bar{s}, \bar{z})$ acts on every remaining denoising step. We give each of M parallel trajectories its own copy $(s^{(m)}, z^{(m)}) \leftarrow (\bar{s}, \bar{z})$ and couple the particles only through a repulsion between their decoded structures.

$$\mathcal{L}_{\text{rep}} = \sum_{i<j} \exp\left(-d_{ij}^2 / 2\sigma^2\right)$$



At a guided step, each particle's denoised estimate $\hat{x}_0^{(m)} = F_\theta(x_t^{(m)}, \sigma_t; s^{(m)}, z^{(m)})$ is compared with the others by differentiable Kabsch Cα RMSD d_{ij} under an RBF kernel,

$$\mathcal{L}_{\text{rep}} = \sum_{i<j} L_{ij}, \qquad L_{ij} = \exp\left(-\frac{d_{ij}^2}{2\sigma^2}\right).$$

Its gradient updates each particle's copy by an RMS-normalised descent step,

$$z^{(m)} \leftarrow z^{(m)} - \eta\, \lambda_t\, \text{RMSnorm}\left(\nabla_{z^{(m)}} \mathcal{L}_{\text{rep}}\right),$$
$$\lambda_t = \max_{i \neq j} L_{ij} \cdot \left[1 + \log\left(1 + \max(0, \hat{t}_t - 1)\right)\right],$$

and likewise for $s^{(m)}$. RMS normalisation decouples step size from gradient magnitude, and the AdaLN and pair-bias conditioning is recomputed after each step. The first factor of λ_t falls to zero once the particles are σ apart; the second grows with the noise level $\hat{t}_t$. Guidance fires every 3 steps with $\eta = 0.02$, $M = 5$ and $\sigma \in \{0.5, \ldots, 2.5\}$ Å. The predictor's weights and inputs are unchanged.

3. Benchmark: more states reached on Boltz-2

Our method reaches more states than default Boltz-2 in every category.

	Domain motion			Transporter			Cryptic pocket		Fold switching		Geometry	
Method	FR	Succ.	RC50	FR	Succ.	RC50	Succ.	RC50	Succ.	RC50	Clash	Cα bond
Default (X-ray)	0.33	65.8 ± 5.5	0.58	0.24	57.1 ± 5.9	0.34	51.5 ± 6.2	0.36	56.7 ± 9.1	0.13	10.20	0.029
Default (MD)	0.36	64.5 ± 5.5	0.58	0.30	54.3 ± 5.9	0.46	54.5 ± 6.1	0.42	50.0 ± 9.1	0.13	10.97	0.033
CF-random	0.50	71.1 ± 5.2	0.71	0.45	70.0 ± 5.5	0.57	59.1 ± 6.0	0.48	60.0 ± 8.9	0.20	16.83	0.030
AFsample3	0.41	72.4 ± 5.1	0.74	0.38	74.3 ± 5.2	0.63	54.5 ± 6.1	0.55	60.0 ± 8.9	0.20	13.49	0.029
AF-cluster	0.45	65.8 ± 5.5	0.61	0.43	65.7 ± 5.7	0.51	56.1 ± 6.1	0.42	53.3 ± 9.1	0.13	24.12	0.030
Pair rep. scaling	0.40	68.4 ± 5.3	0.63	0.34	68.6 ± 5.5	0.49	54.5 ± 6.1	0.48	63.3 ± 8.8	0.20	11.99	0.039
ConforMix-B1	0.43	63.2 ± 5.5	0.63	0.36	54.3 ± 5.9	0.37	50.0 ± 6.1	0.36	50.0 ± 9.1	0.13	54.80	0.373
ConforMix-B2	0.36	65.8 ± 5.5	0.61	0.30	50.0 ± 6.0	0.29	56.1 ± 6.1	0.55	60.0 ± 8.9	0.13	62.34	11.656
BioEmu	0.44	67.1 ± 5.4	0.71	0.42	75.7 ± 5.1	0.57	48.5 ± 6.1	0.45	46.7 ± 9.1	0.13	10.27	0.164
Ours (λ_t)	0.49	75.0 ± 5.0	0.79	0.50	74.3 ± 5.2	0.60	62.1 ± 6.0	0.48	60.0 ± 8.9	0.27	16.75	0.039
Ours (λ)	0.52	76.3 ± 4.9	0.74	0.48	71.4 ± 5.4	0.60	62.1 ± 6.0	0.52	60.0 ± 8.9	0.20	13.50	0.034

Table 1. Succ., states reached (%). RC50, both states reached. FR, fill ratio. Clash, MolProbity clashscore. Cα bond, RMS deviation from 3.80 Å. Tinted, best. λ, gate without noise term.

4. One state by default, both with repulsion

Default sampling stays near one state; our method reaches both.

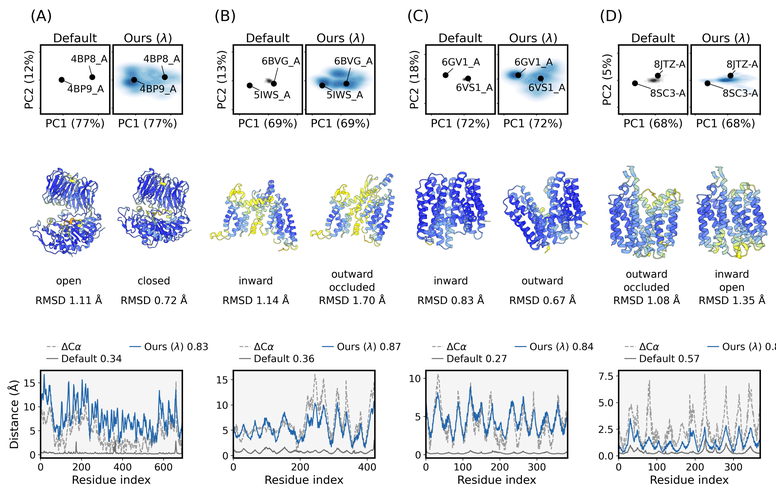


Figure 2. (D) OCT1, released after training. Top, conformational space. Middle, closest prediction per state. Bottom, ensemble motion (blue) and state difference (dashed).

5. The direction of the step matters

A random push does not help, and pushing the coordinates breaks the backbone.

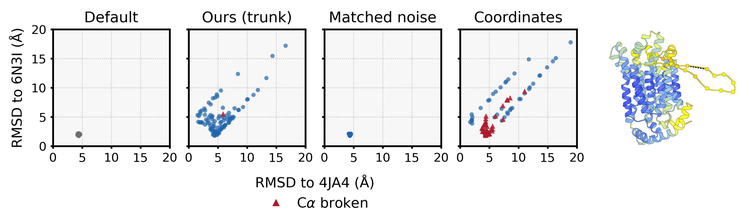


Figure 3. One target, four ways to push. Red, samples with a broken backbone.

6. Dopamine transporter: the inward state appears

Our method reaches the inward state, which default sampling never does.

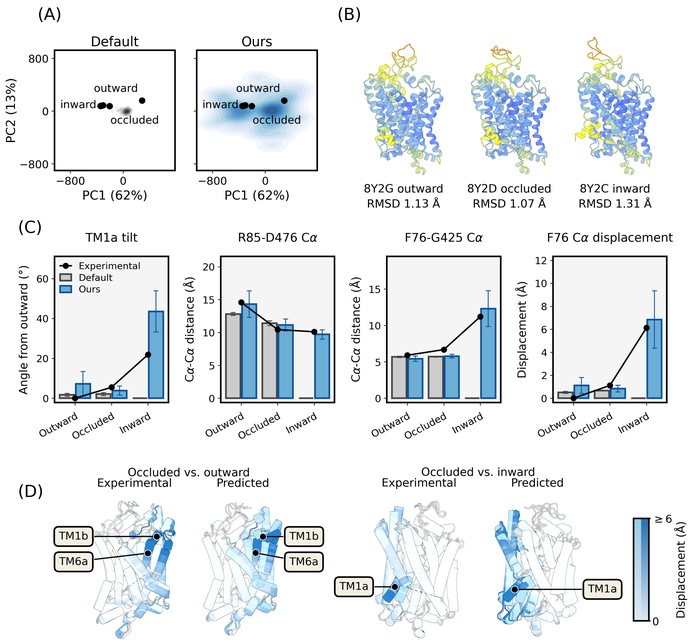


Figure 4. (A) Conformational space, experimental states in black. (B) Closest prediction per state. (C) Mechanism metrics. (D) Moving helices, experiment and prediction.

7. Other predictors, and structures newer than training

It also works on two more predictors and on structures newer than training.

		Domain motion		Transporter		Cryptic pocket		Fold switching	
Trunk	Method	Succ.	RC50	Succ.	RC50	Succ.	RC50	Succ.	RC50
OpenFold3-p2	Default	60.5 ± 5.6	0.45	67.1 ± 5.6	0.49	45.5 ± 6.1	0.21	53.3 ± 9.1	0.07
	Ours (λ_t)	71.1 ± 5.2	0.74	74.3 ± 5.2	0.63	45.5 ± 6.1	0.36	56.7 ± 9.1	0.13
	Ours (λ)	72.4 ± 5.1	0.74	75.7 ± 5.1	0.63	54.5 ± 6.1	0.36	50.0 ± 9.1	0.13
Protenix	Default	61.8 ± 5.6	0.42	67.1 ± 5.6	0.51	51.5 ± 6.1	0.27	40.0 ± 8.9	0.13
	Ours (λ_t)	71.1 ± 5.2	0.74	74.3 ± 5.2	0.71	53.0 ± 6.1	0.42	50.0 ± 9.1	0.20
	Ours (λ)	72.4 ± 5.1	0.71	74.3 ± 5.2	0.71	56.1 ± 6.1	0.48	43.3 ± 9.0	0.20

Table 2. The same guidance, unchanged, on OpenFold3-preview-2 and Protenix. 100 samples per target, against default at the same budget; metrics as in Table 1.

Trunk	Method	RC50	Alt-RMSD	Clashscore	Cα bond
Boltz-2	Default	4/12	2.63	21.29	0.022
	Ours (λ_t)	7/12	2.01	21.65	0.022
	Ours (λ)	6/12	2.12	21.29	0.022
OpenFold3-p2	Default	10/20	2.36	7.74	0.058
	Ours (λ_t)	12/20	1.85	8.73	0.064
	Ours (λ)	12/20	2.20	8.21	0.060
Protenix	Default	10/20	2.39	19.86	0.039
	Ours (λ_t)	12/20	2.02	21.73	0.044
	Ours (λ)	12/20	1.98	20.49	0.044

Table 3. Targets whose two references were both released after each predictor's training cutoff. RC50, targets with both states reached; Alt-RMSD (Å), mean distance to the harder state. Boltz-2 draws 500 samples, the others 100.