CryoBoltz

Multiscale guidance of protein structure prediction with heterogeneous cryo-EM data

Rishwanth Raghu, Axel Levy, Gordon Wetzstein, and Ellen D. Zhong

Presented by Ishan Saha and Karan Tandon

Background

Distinguishing Sequence from Structure



M = methionine

A = alanine; G = glycine

Sequence: an ordered chain of amino acids. Letters name the building blocks.

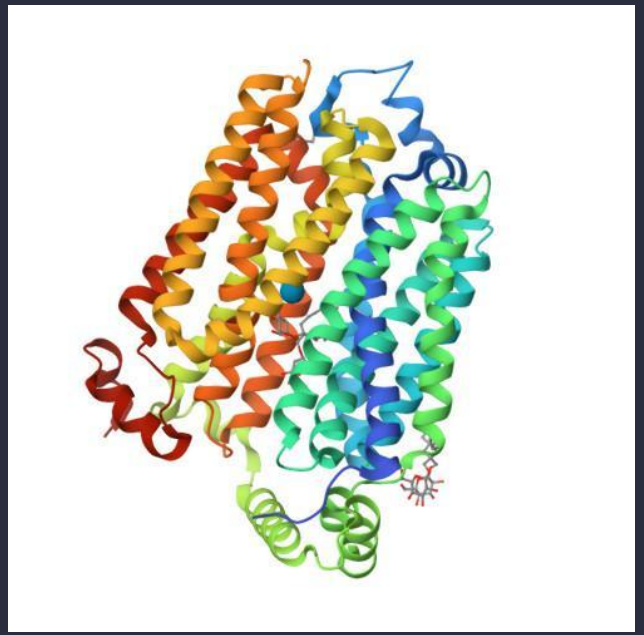

Structure: atom positions in 3D. STP10 is a plant sugar transporter.

STP10 sequence and structure: RCSB PDB 7AAR.

One protein can adopt several structural states

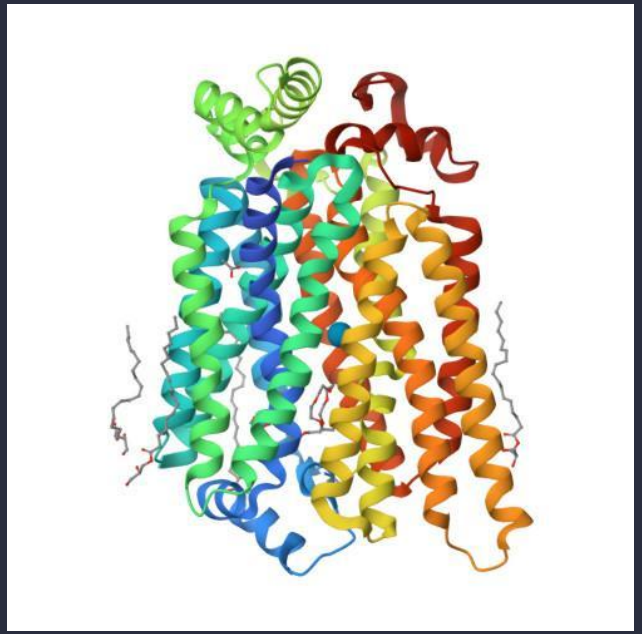

STP10: outward-facing, with the sugar-binding site closed.

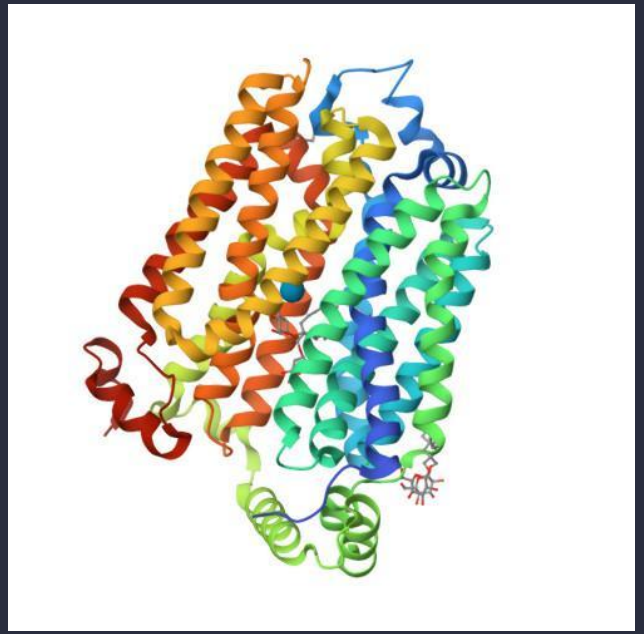

Same protein: sugar-binding site open toward the cell interior.

Structure images: RCSB PDB 7AAQ and 7AAR.

Cryo-EM reconstructs 3D density from 2D views

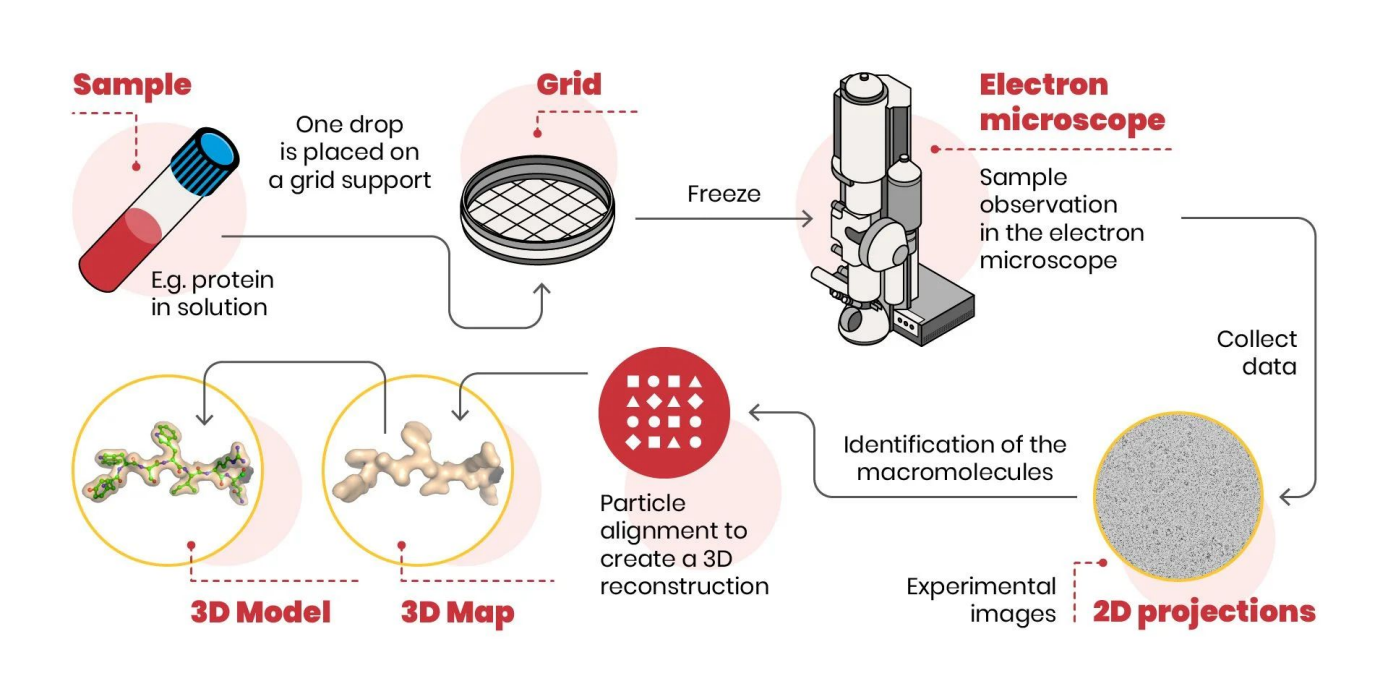

Many orientations provide the information needed for three-dimensional reconstruction.

Workflow illustration: UniSR.

CryoBoltz uses reconstructed maps as input

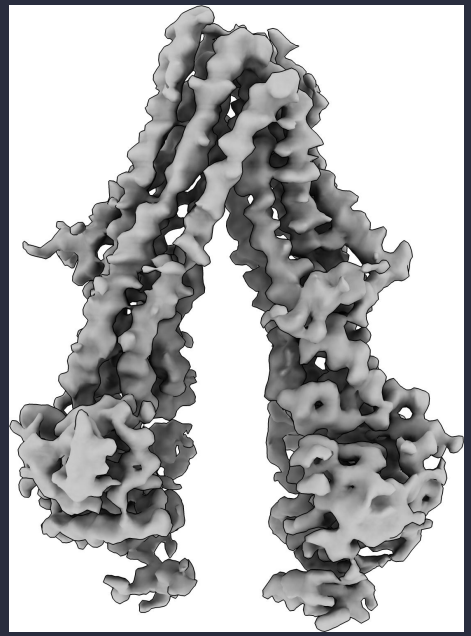

The surface displays a boundary within a three-dimensional density grid.

- A map stores a density value in each 3D cell, or **voxel**.
- Input: a reconstructed map, not raw microscope images.
- Different state-associated maps are fitted in separate runs.

Map illustration: authors' presentation.

Density does not uniquely determine atom positions

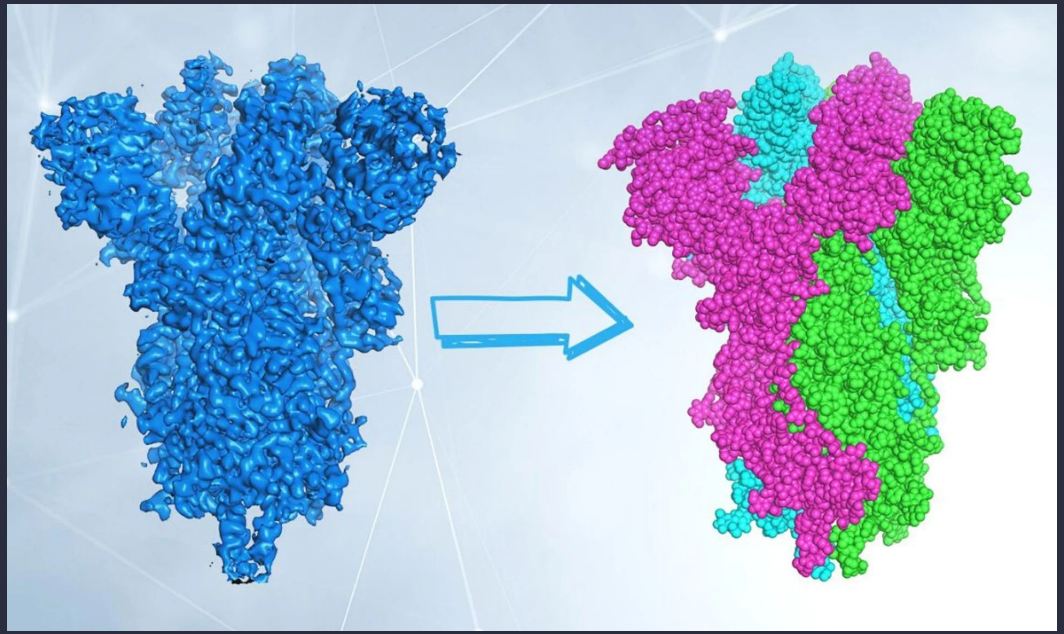

Map (left) and atomic model (right). Blurred density can support several atomic arrangements.

Background example: UniSR / Massimo Degano; SARS-CoV-2 spike.

Weak density can limit map-based model building

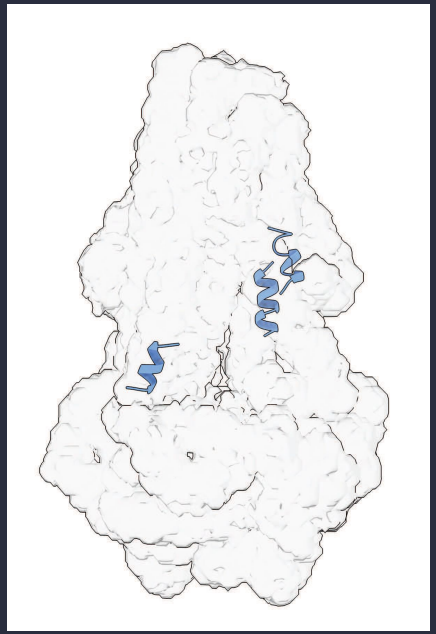

Example: only fragments could be built in this transporter map.

- **ModelAngelo** uses the map, sequence, and learned structural knowledge.
- Weak density can leave parts of the chain **unbuilt**.

Example from paper, Fig. 5; the result is examined later.

Boltz-1 generates structures by denoising coordinates

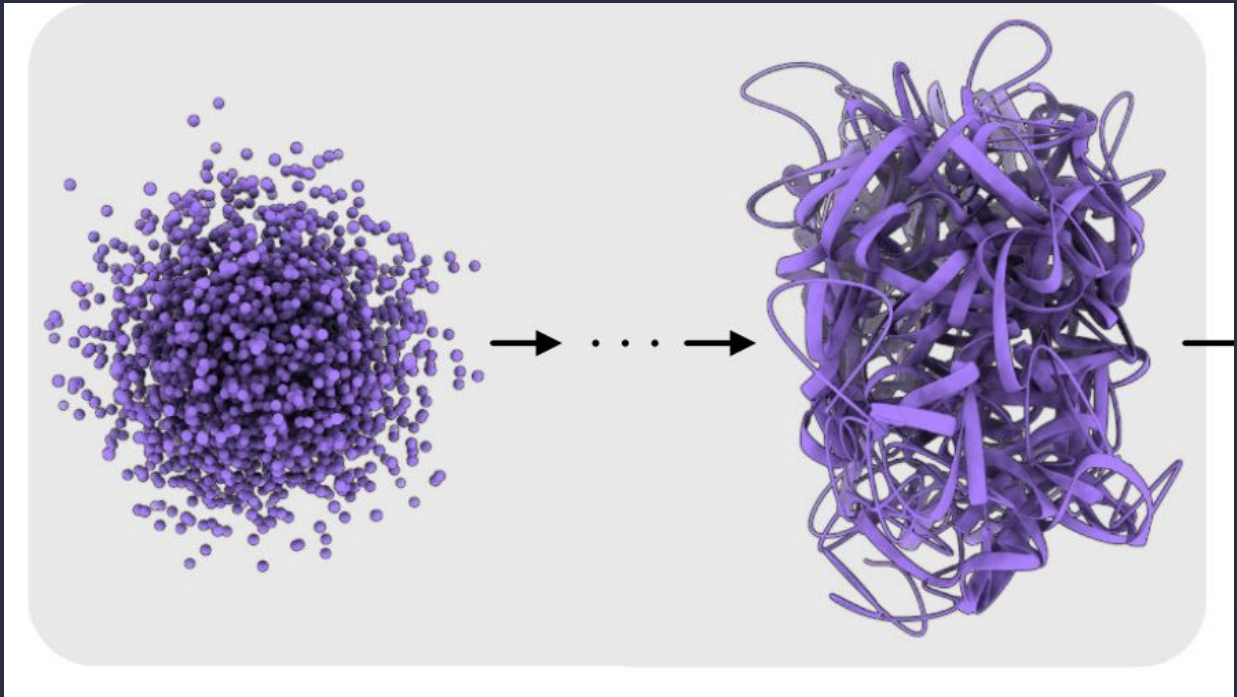

Boltz-1 uses sequence information to generate coordinates by diffusion.

Warm-up excerpt: paper, Fig. 2.

Unguided sampling may miss the measured state

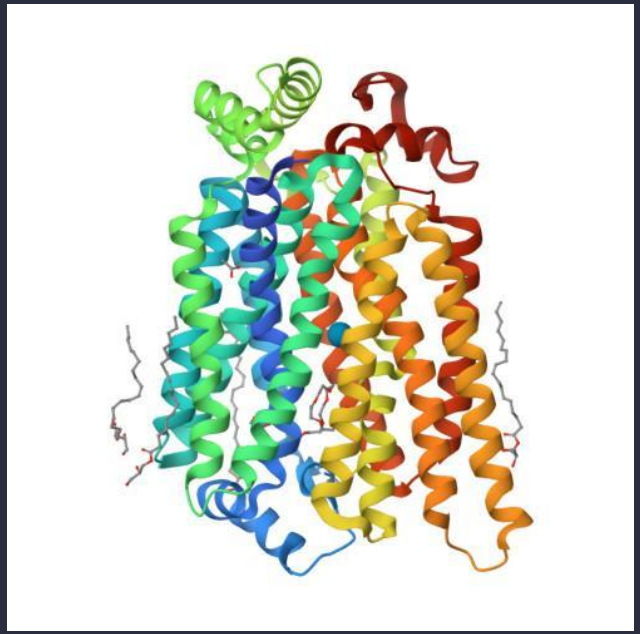

An outward-facing reference illustrates a possible competing state.

- Unguided Boltz-1 **does not see the map**.
- More samples help only if a suitable state is reached.

STP10 reference structure: RCSB PDB 7AAQ.

The map constrains the experimental state

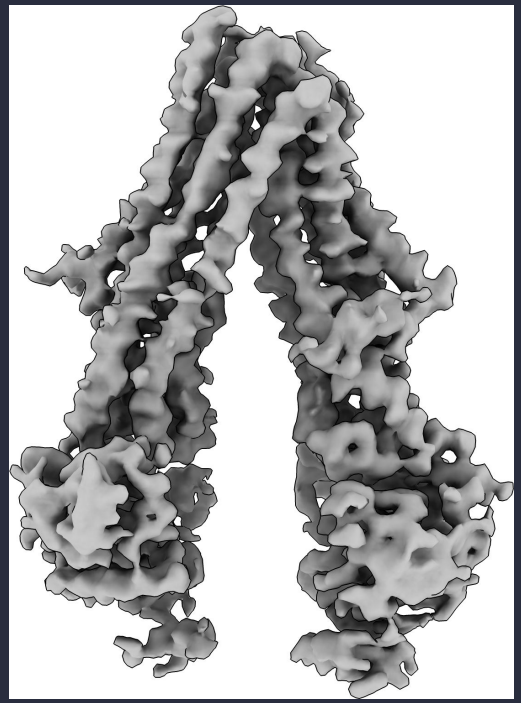

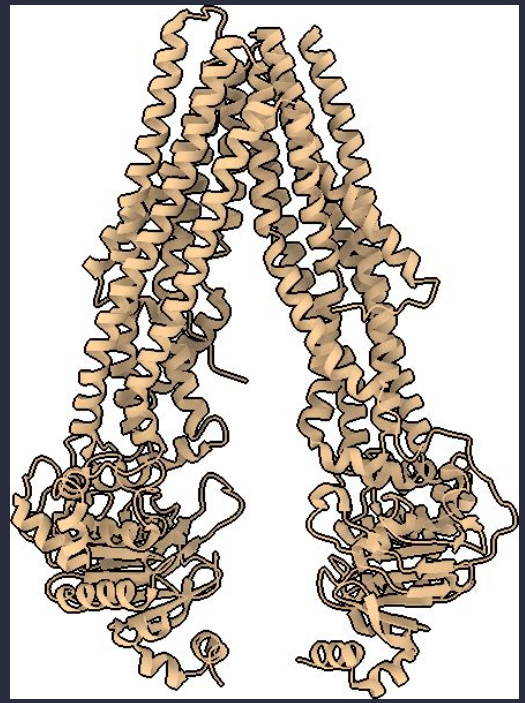

Cryo-EM constrains the arrangement.

Boltz-1 supplies plausible atomic detail.

Illustrations: authors' presentation.

Guidance changes generation, not just selection

Generate candidate structures → Then select best available fit

Selection cannot choose a state that was never generated.

- CryoBoltz lets the map influence **generation itself**.
- The map affects coordinate updates during sampling.

Guided inference

Each denoising step predicts a clean structure

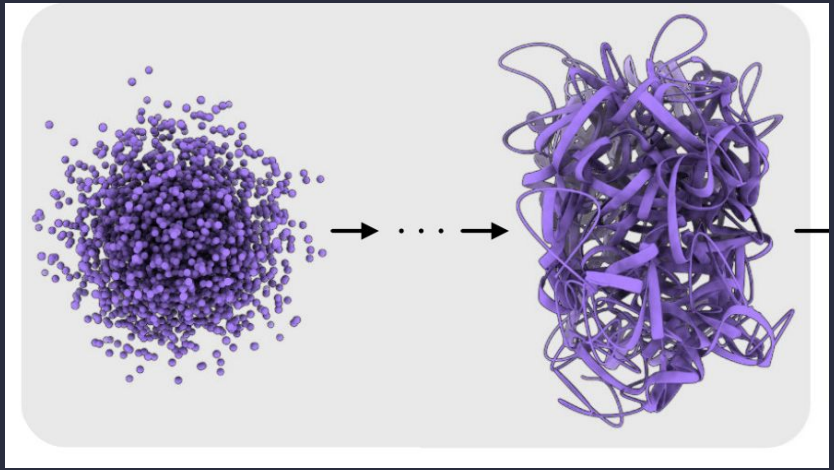

Start with noisy coordinates; repeatedly predict and move toward a clean structure.

- **Clean** means without added sampling noise, not necessarily correct.
- Sampling steps are **not physical protein motion**.

Paper, Fig. 2, warm-up excerpt.

The denoiser uses fixed network parameters

$$x_t \longmapsto \hat{x}_\theta(x_t, s, t)$$

- x_t: noisy coordinates at denoising stage t.
- s: sequence conditioning; θ: fixed network weights.
- The hat marks a **prediction**, not the known true structure.

The forward model converts coordinates to density

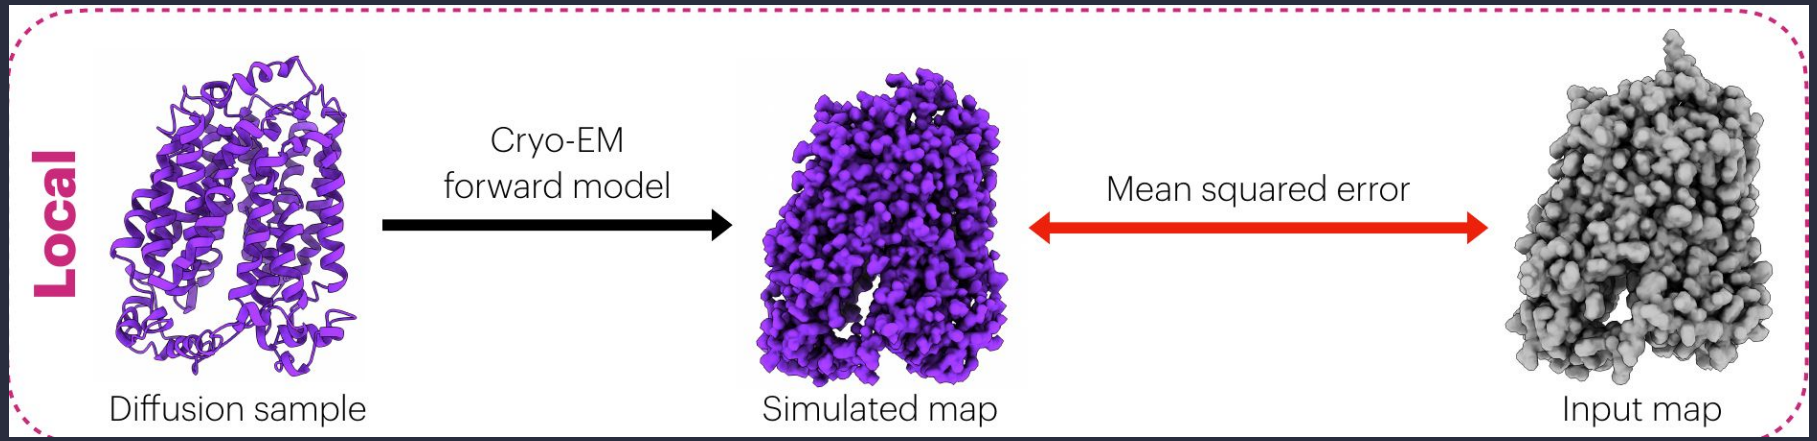

The forward model $F(x,s)$ predicts density from coordinates x and sequence s.

Local-guidance illustration: authors' presentation.

Atom positions determine the predicted density

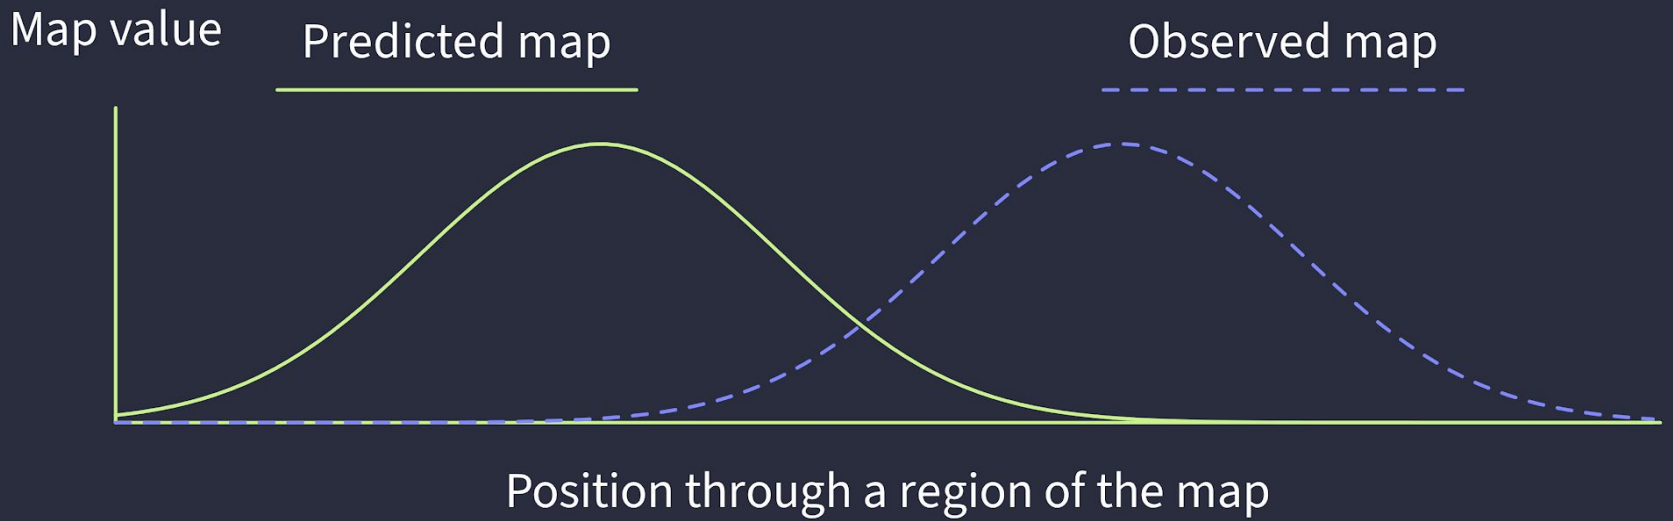

A displaced prediction leaves missing signal in one region and excess signal in another.

Illustrative 1D profiles, not measured data. The real maps are three-dimensional.

Inference combines a prior with map evidence

$$\text{Prior plausibility} \times \text{map agreement} \longrightarrow \text{updated support}$$

- A plausible structure can still disagree with the experiment.
- A map-fitting configuration can still have **unrealistic atomic geometry**.

The prior assigns support before observing the map

$$\text{Updated support} \propto \text{map agreement} \times p_0(x \mid s)$$

- x: a candidate clean structure; s: its sequence.
- $p_0(x \mid s)$: prior probability density over clean structures.
- $\mid$ means **conditioned on**; $\propto$ means **proportional to**.

The likelihood measures agreement with the map

$$\text{Updated support} \propto p(y \mid x, s) \times p_0(x \mid s)$$

- y: the measured map.
- $p(y \mid x,s)$: likelihood of this map if the candidate structure were correct.
- Both terms are evaluated for the **same candidate** x.

Bayes' rule defines the target posterior

$$p(x \mid y, s) \propto p(y \mid x, s) \times p_0(x \mid s)$$

- **Posterior:** support for structures after observing the map.
- $\propto$: the product must be normalized over possible structures.
- This is the **inference target**, not a guarantee about CryoBoltz samples.

Guidance evaluates the predicted clean structure

- The sampler holds noisy coordinates, not a finished molecular structure.
- Averaging over all possible clean structures is expensive.
- Approximation: evaluate map agreement at the current clean prediction.

The loss compares a candidate with the observed map

$$L(x, y)$$

- x: candidate clean coordinates; y: the fixed observed map.
- L: a comparison rule returning **one mismatch value**.
- Lower loss means a better fit according to that rule.

The clean prediction becomes the loss input

$$L(x, y) \quad \longrightarrow \quad L\big(\hat{x}_\theta(x_t, s, t), y\big)$$

- Replace candidate x with the clean prediction $\hat{x}_\theta(x_t, s, t)$.
- x_t: noisy coordinates; s: sequence; t: denoising stage.
- The observed map y and network weights θ stay fixed.

The negative gradient points toward lower loss

$$g_t = -\nabla_{x_t} L\big(\hat{x}_\theta(x_t, s, t), y\big)$$

- **Gradient:** sensitivity of the loss to each noisy coordinate.
- g_t: the correction direction; minus points toward **decreasing** loss.
- Backpropagation follows the path **through the denoiser**.

Guidance updates coordinates, not network weights

Sampling update = base-model update + **map correction**

- The network weights remain fixed.
- The guidance strength controls the correction's influence.
- Evaluating one clean prediction makes this **approximate inference**.

Method:

Using “Multiscale Guidance”

MULTISCALE GUIDANCE

Overview

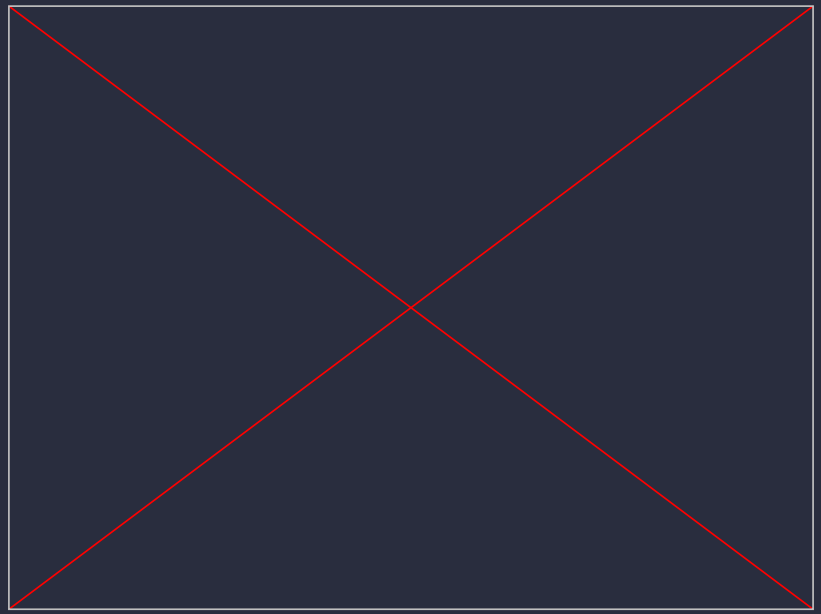

One sample through the four stages;
color marks the active stage.

Goal: Guide the Diffusion Model to sample a viable structure, that is also highly likely from cryo-EM map

Video: authors' project page. Figure: paper, Fig. 2 (CC BY 4.0).

MULTISCALE GUIDANCE

Overview

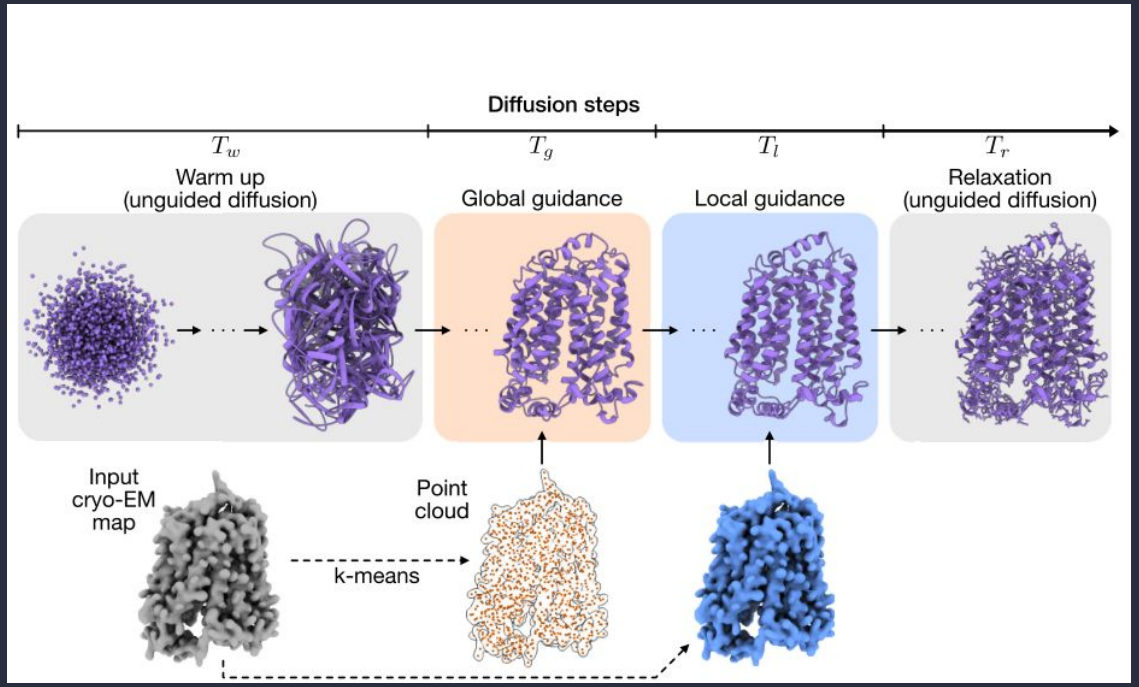

1. **Warm up:** Build an initial molecular arrangement
2. **Global Guidance:** Guide through the general structure
3. **Local Guidance:** Guide through the finer details
4. **Relaxation:** Denoise and address distortions

Video: authors' project page. Figure: paper, Fig. 2 (CC BY 4.0).

MULTISCALE GUIDANCE

Stage 1: warm up from Boltz-1's own guess

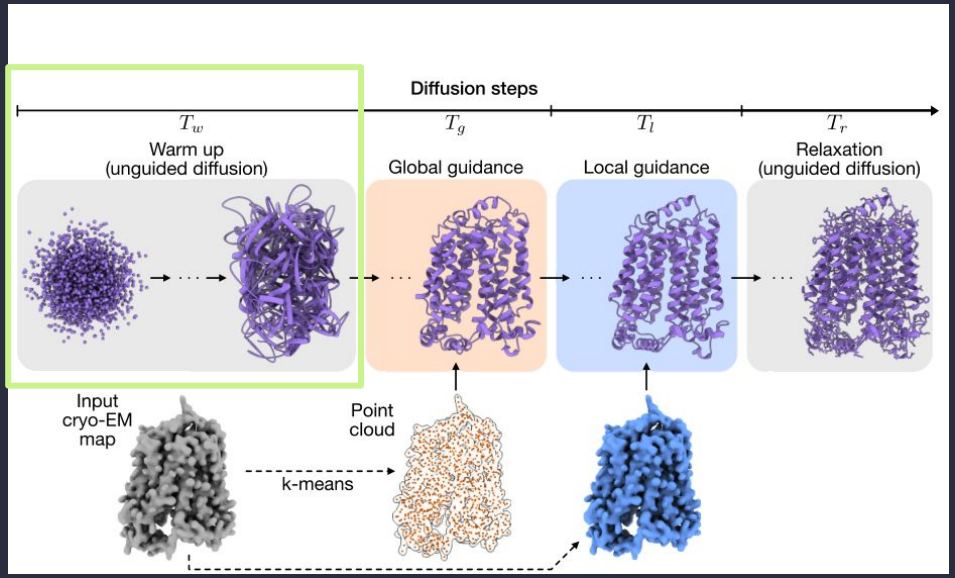

Eq. 2 $dx = [f - g^2 s\theta(x, s, t)]\, dt + g\, dw$

sθ: Boltz-1's learned score, from the sequence s only · no map term · 100 steps (real maps)

- Plain Boltz-1 denoising from random noise, using only the sequence.
- Early on the structure is a blur, so the map has nothing to grip; guidance waits.
- Ends near the state Boltz-1 prefers, which the next stages reshape.

- **dx:** how coordinates change in one small step.
- **f:** a drift term from the noise schedule; a bookkeeping term.
- **g:** how much noise exists at this stage.
- **sθ:** Boltz-1's score, the direction toward more plausible structures.
- **g dw:** fresh random noise injected each step. This is what makes samples differ from each other.

Figure: paper, Fig. 2 (CC BY 4.0). Equations: paper, Sections 2.1, 2.2, 3.2, 3.3; schedule from Appendix A.5.

MULTISCALE GUIDANCE

Why warm up first: Boltz-1 supplies the fold

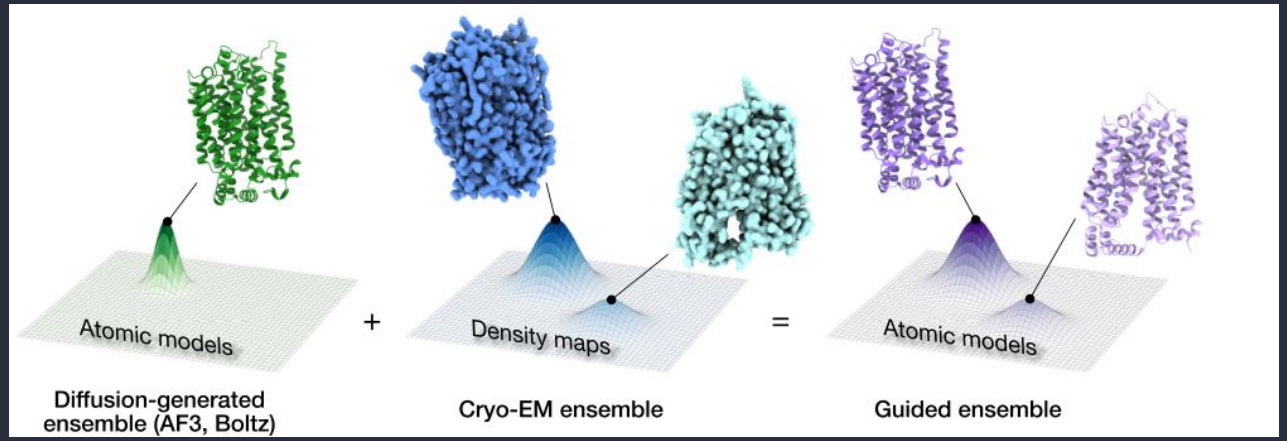

Boltz-1's samples cluster around one state (left). The cryo-EM data covers the states actually present (middle). Guidance moves samples toward the measured state (right).

- At high noise the predicted structure is a blur, so comparing it with the map gives no useful direction.
- Boltz-1 alone sets the fold, helices and chain path, which a ~4 Å map cannot pin down.
- Guidance only reshapes this start. If Boltz-1 starts far off, CryoBoltz fails (DSL1/SNARE complex, full IgG).

Figure: paper, Fig. 1 (CC BY 4.0). Warm-up rationale and failure cases: paper, Sections 3.1 and 5.

MULTISCALE GUIDANCE

Stage 2: global guidance matches the shape

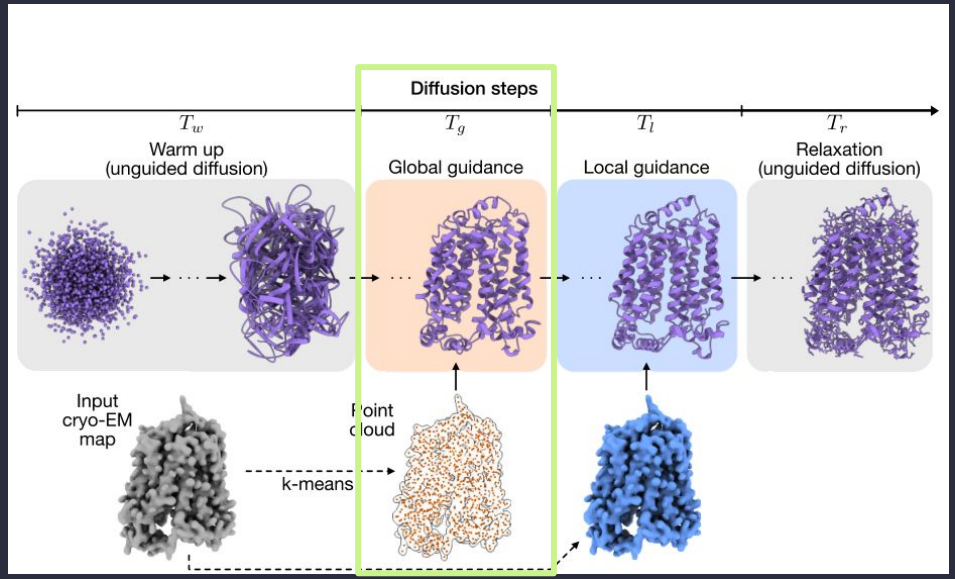

Eq. 4 $\tilde{s} = -\nabla_x D(\hat{x}\theta(x, s, t), Y)$

D: Sinkhorn divergence · Y: k-means point cloud of the map · $\hat{x}\theta$: predicted clean structure · 50 steps

- Turns the map into a coarse point cloud and pulls the predicted clean structure $\hat{x}\theta$ toward it.
- Optimal transport matches every atom to some point, so the pull works even without overlap: whole helices can move.
- Heuristic, not a likelihood. Its strength λ fades from 0.25 to 0.05.

Figure: paper, Fig. 2 (CC BY 4.0). Equations: paper, Sections 2.1, 2.2, 3.2, 3.3; schedule from Appendix A.5.

MULTISCALE GUIDANCE

Why global guidance works: every atom gets a pull

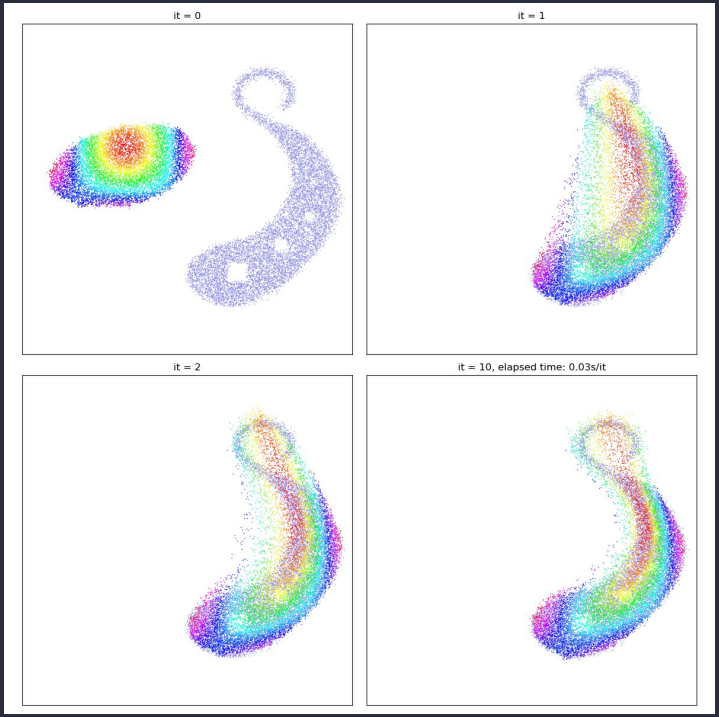

- Optimal transport pairs every atom with part of the map's point cloud; the cost is the total distance atoms must travel.
- Even an atom far from any density gets a gradient toward its partner, so misplaced helices move as blocks.

- Left: the same idea in 2D. One cloud flows into a different shape within a few steps.
- The point cloud keeps only coarse shape (about one point per 4 atoms), so this stage cannot chase fine noise.
- Preview: STP10 inward state goes from 3.82 Å unguided to 1.29 Å with global guidance alone.

Left: GeomLoss, gradient descent on Sinkhorn divergence (2D illustration; iterations 0, 1, 2, 10). STP10 numbers: paper, Table 1.

MULTISCALE GUIDANCE

Stage 3: local guidance fits the detail

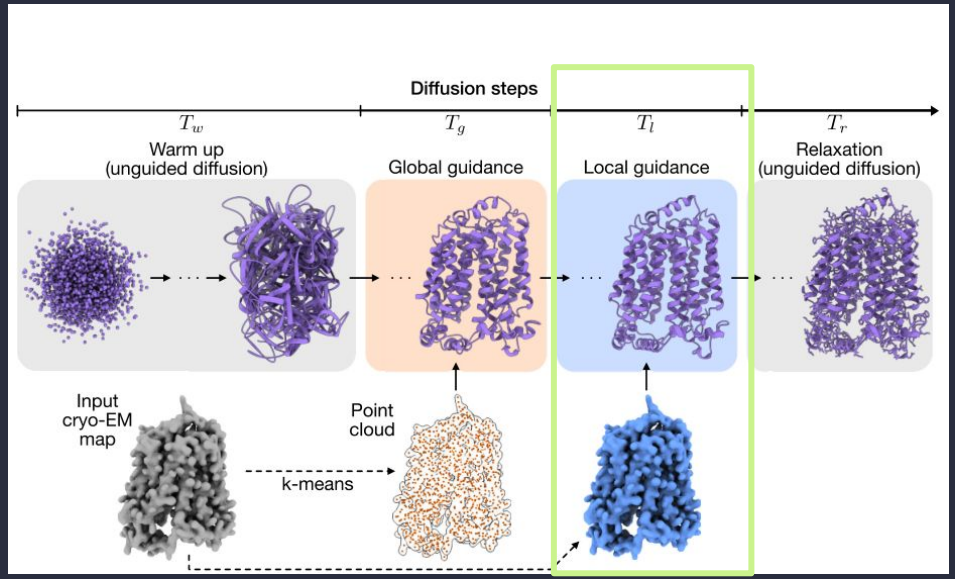

Eq. 5 $\tilde{s} = -\nabla_x \|y - B(\Gamma(\hat{x}\theta(x, s, t), s))\|^2$

Γ: one Gaussian per atom · B: blur to the map's resolution · y: the cryo-EM map · 25 steps

- Renders a map from $\hat{x}\theta$ and compares it with the real map voxel by voxel.
- Under Gaussian noise this squared error is the negative log-likelihood: the physics-based term.
- Sees fine detail but only pulls where densities already overlap, so it refines rather than moves. $\lambda = 0.5$.

Figure: paper, Fig. 2 (CC BY 4.0). Equations: paper, Sections 2.1, 2.2, 3.2, 3.3; schedule from Appendix A.5.

MULTISCALE GUIDANCE

Why local guidance refines: its gradient needs overlap

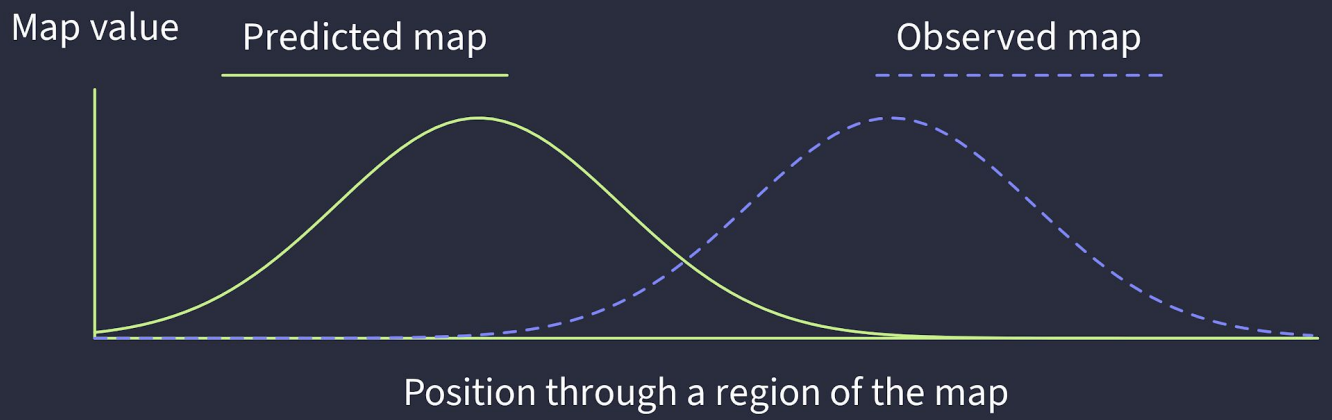

- Squared error only changes where predicted and observed density overlap; a helix displaced beyond its blur width gets almost no pull.
- After global guidance the densities already overlap, so each atom is pushed toward the denser side: loops and side chains settle.
- Preview: the CH67 binding loop goes from 1.90 Å with global guidance alone to 1.27 Å with global + local.

Illustrative 1D profiles, not measured data. CH67 numbers: paper, Table 1.

MULTISCALE GUIDANCE

Stage 4: relaxation cleans up geometry

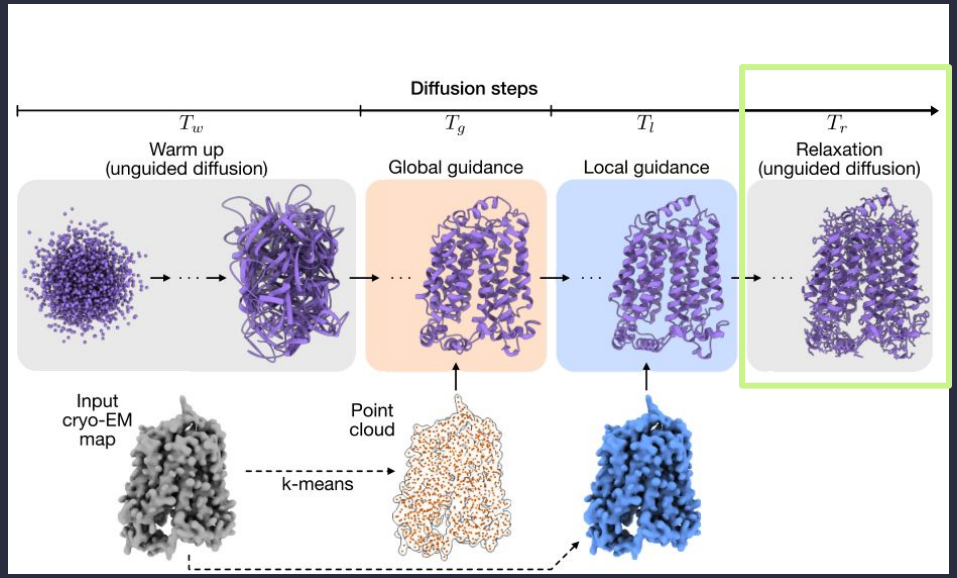

Eq. 2 again, with guidance off (λ = 0)

Boltz-1 alone for the last 25 steps

- Guidance turns off; Boltz-1 finishes the last denoising steps.
- Its prior repairs clashes and bond geometry that the map pull may have strained.
- Good geometry does not prove the right state; that still needs validation.

Figure: paper, Fig. 2 (CC BY 4.0). Equations: paper, Sections 2.1, 2.2, 3.2, 3.3; schedule from Appendix A.5.

MULTISCALE GUIDANCE

Why relax last: strong pulls can strain geometry

Guidance strength λ over the 200 sampling steps (real-map schedule)

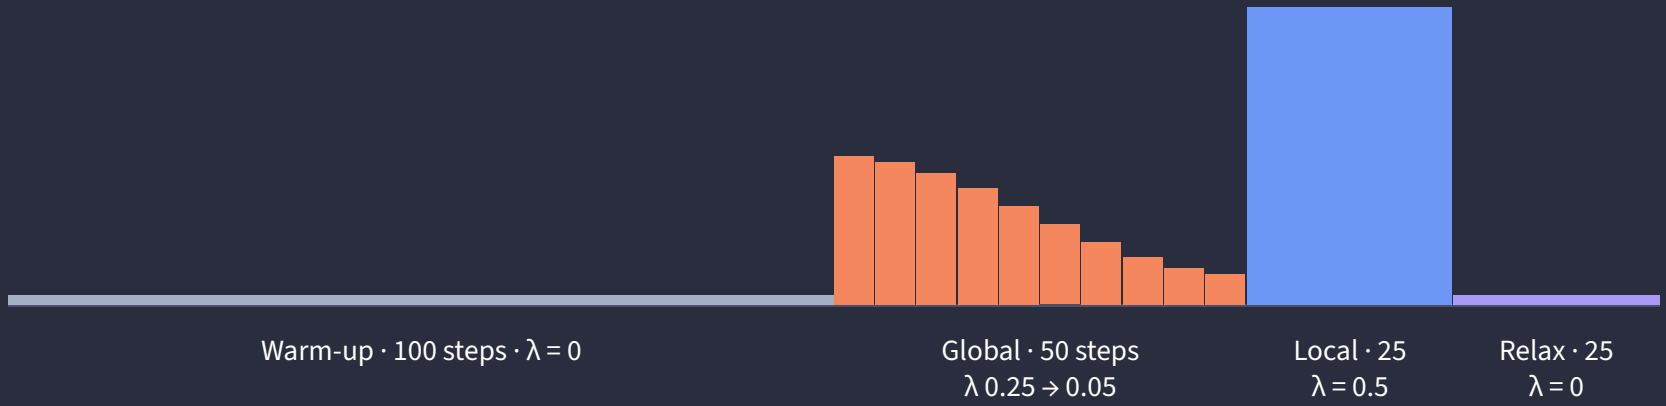

- Guidance adds a force Boltz-1 never saw in training; a strong pull can leave clashes or strained bonds.
- The last 25 steps are plain Boltz-1 at low noise: small moves that restore chemistry without undoing the fit.
- Relaxation fixes geometry, not state: a structure in the wrong state can still look chemically perfect.

Schedule: paper, Appendix A.5. The global and local losses have different scales, so their λ values are not directly comparable.

Results

RESULTS

Reading the tables: best of 25, averaged over 3 runs

Metric	What it measures	Better
All-atom RMSD (Å)	Average distance between matched atoms after superposition	Lower
Cα RMSD (Å)	Same, using one backbone atom per residue	Lower
TM score	Fold similarity from 0 to 1; less sensitive to a few bad regions	Higher
Local RMSD (Å)	Cα RMSD of the CH67 CDR H3 loop only	Lower

- Each method draws 25 samples in each of 3 replicate runs; the tables report the best sample per run, averaged over the 3 runs.
- "Best" means closest to the known structure, so absolute numbers are optimistic. The same rule applies to every method.
- Baselines: Boltz-1, Boltz-1 with MSA subsampling, AlphaFold3, and ModelAngelo (real maps only).

Paper, Section 4.1 and the captions of Tables 1 and 2.

RESULTS · SYNTHETIC MAPS

Table 1: CryoBoltz is best on both synthetic maps

Structure	Metric	CryoBoltz	Boltz-1	Boltz-1 + MSA sub.	AlphaFold3	Local only*	Global only*
STP10 (inward)	RMSD (Å) ↓	**1.057**	3.815	3.768	1.263	3.860	1.287
	Cα RMSD (Å) ↓	**0.371**	3.554	3.513	0.623	3.559	0.778
	TM score ↑	**0.998**	0.863	0.865	0.994	0.862	0.990
STP10 (outward)	RMSD	**0.888**	2.656	2.542	4.478	2.722	1.164
	Cα RMSD	**0.440**	2.419	2.295	4.228	2.458	0.779
	TM score	**0.997**	0.948	0.953	0.828	0.946	0.991
CH67 antibody	RMSD	**1.048**	1.961	1.954	1.887	1.443	1.281
	Cα RMSD	**0.637**	1.469	1.522	1.453	0.945	0.880
	Local RMSD	**1.269**	3.120	3.270	3.191	1.718	1.899
	TM score	**0.994**	0.972	0.969	0.971	0.990	0.988

Paper, Table 1. *Ablations: CryoBoltz with only one guidance stage. Best of 25 samples (lowest all-atom RMSD), mean of 3 replicates. MSA sub.: depths 64–1024.

RESULTS · SYNTHETIC MAPS

Table 1, STP10 rows: each predictor gets only one state

Structure	Metric	CryoBoltz	Boltz-1	Boltz-1 + MSA sub.	AlphaFold3	Local only*	Global only*
STP10 (inward)	RMSD (Å) ↓	**1.057**	3.815	3.768	1.263	3.860	1.287
	Cα RMSD (Å) ↓	**0.371**	3.554	3.513	0.623	3.559	0.778
	TM score ↑	**0.998**	0.863	0.865	0.994	0.862	0.990
STP10 (outward)	RMSD	**0.888**	2.656	2.542	4.478	2.722	1.164
	Cα RMSD	**0.440**	2.419	2.295	4.228	2.458	0.779
	TM score	**0.997**	0.948	0.953	0.828	0.946	0.991

- Inward map: Boltz-1 (3.82 Å) and MSA subsampling (3.77 Å) stay in the outward state; CryoBoltz reaches 1.06 Å.
- Outward map: AlphaFold3 is stuck inward (4.48 Å); CryoBoltz reaches 0.89 Å, with TM ≥ 0.997 in both states.

Paper, Table 1. *Ablations: CryoBoltz with only one guidance stage. Best of 25 samples (lowest all-atom RMSD), mean of 3 replicates. MSA sub.: depths 64–1024.

RESULTS · SYNTHETIC MAPS

STP10: one sequence, two maps, two correct states

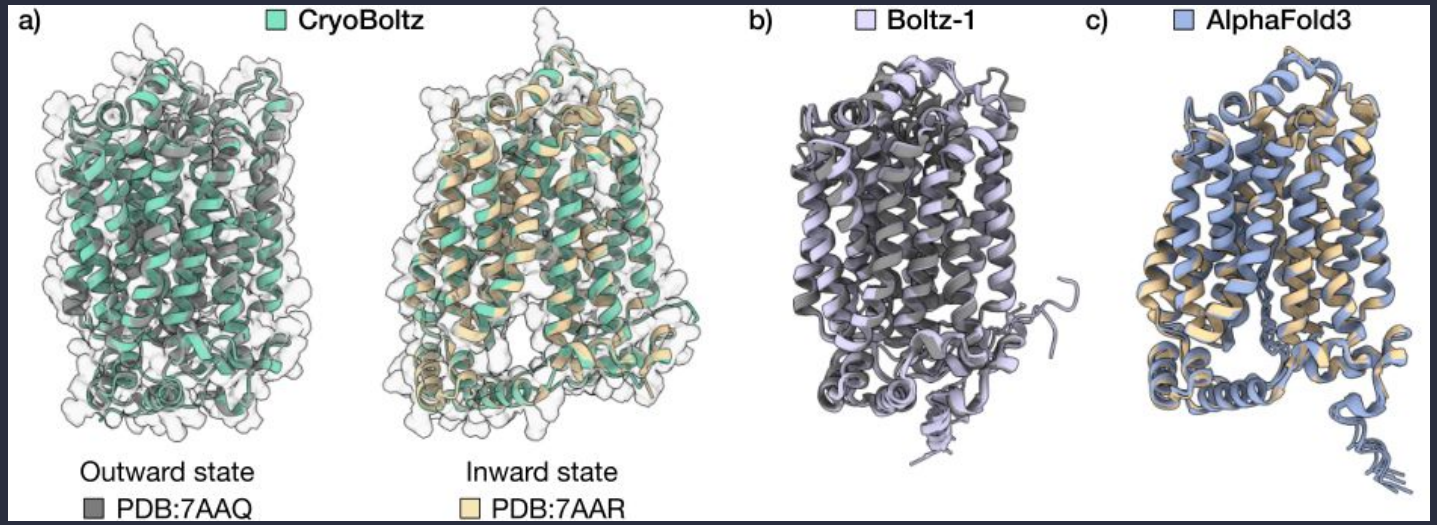

- (a) Guided by each map, CryoBoltz (green) overlays the outward (gray) and inward (tan) references.
- (b) Boltz-1's best samples all sit outward; (c) AlphaFold3's all sit inward.
- The switch is a coordinated helix rearrangement: the move global guidance is built for.

Paper, Fig. 3. Synthetic 2 Å maps from PDB 7AAQ (outward) and 7AAR (inward). Green: CryoBoltz; gray and tan: reference structures.

RESULTS · SYNTHETIC MAPS

Table 1, CH67 rows: predictors miss the binding loop

Structure	Metric	CryoBoltz	Boltz-1	Boltz-1 + MSA sub.	AlphaFold3	Local only*	Global only*
CH67 antibody	RMSD	**1.048**	1.961	1.954	1.887	1.443	1.281
	Cα RMSD	**0.637**	1.469	1.522	1.453	0.945	0.880
	Local RMSD	**1.269**	3.120	3.270	3.191	1.718	1.899
	TM score	**0.994**	0.972	0.969	0.971	0.990	0.988

- Whole Fab: 1.05 Å for CryoBoltz vs 1.96 Å (Boltz-1) and 1.89 Å (AlphaFold3).
- CDR H3 loop: 1.27 Å vs 3.12 Å (Boltz-1) and 3.19 Å (AlphaFold3); MSA subsampling does not help (3.27 Å).

Paper, Table 1. *Ablations: CryoBoltz with only one guidance stage. Best of 25 samples (lowest all-atom RMSD), mean of 3 replicates. MSA sub.: depths 64–1024.

RESULTS · SYNTHETIC MAPS

CH67: the 4 Å map places the binding loop

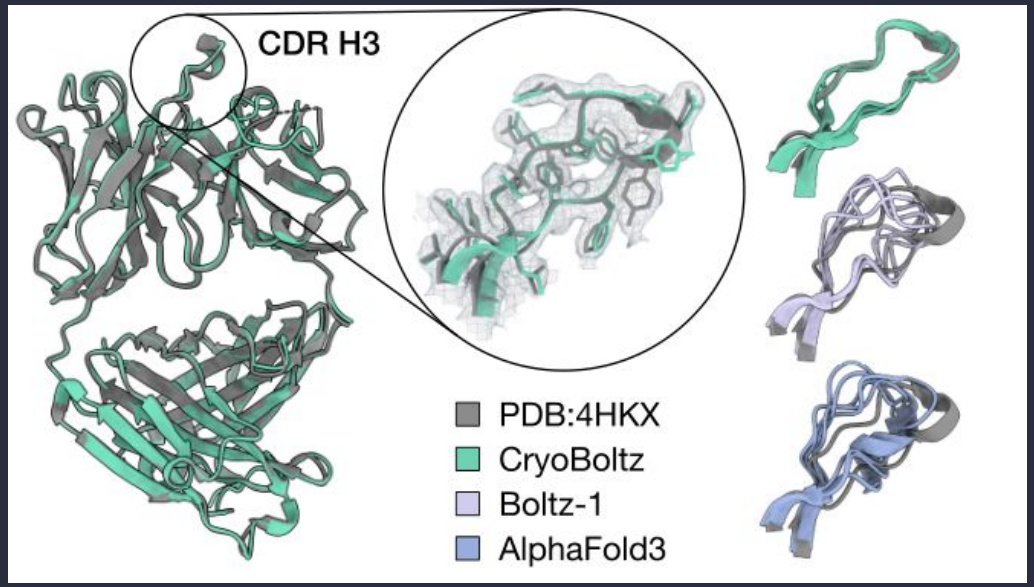

- CDR H3 is the loop that contacts the influenza target. It varies most between antibodies, so training data says little about it.
- Zoom: CryoBoltz (green) follows the reference (gray) inside the density, backbone and most side chains.
- Right: Boltz-1 and AlphaFold3 loops splay away from the reference.
- 4 Å is typical for antibody cryo-EM, so this is the realistic case.

Paper, Fig. 4. Synthetic 4 Å map from PDB 4HKX. Right: three best samples per method.

RESULTS · SYNTHETIC MAPS

Table 1, ablations: global moves, local refines

Structure	Metric	CryoBoltz	Boltz-1	Local only*	Global only*
STP10 (inward)	RMSD (Å) ↓	**1.057**	3.815	3.860	1.287
	Cα RMSD (Å) ↓	**0.371**	3.554	3.559	0.778
	TM score ↑	**0.998**	0.863	0.862	0.990
STP10 (outward)	RMSD	**0.888**	2.656	2.722	1.164
	Cα RMSD	**0.440**	2.419	2.458	0.779
	TM score	**0.997**	0.948	0.946	0.991
CH67 antibody	RMSD	**1.048**	1.961	1.443	1.281
	Cα RMSD	**0.637**	1.469	0.945	0.880
	Local RMSD	**1.269**	3.120	1.718	1.899
	TM score	**0.994**	0.972	0.990	0.988

Paper, Table 1. *Ablations: CryoBoltz with only one guidance stage. Best of 25 samples (lowest all-atom RMSD), mean of 3 replicates. MSA sub.: depths 64–1024.

- STP10 inward: local only stays at 3.86 Å, no better than unguided Boltz-1 (3.82 Å); global only reaches 1.29 Å; both give 1.06 Å.
- CH67 loop: local only (1.72 Å) beats global only (1.90 Å), because the loop needs refining, not moving; both give 1.27 Å.
- Each stage fixes what the other cannot.

RESULTS · EXPERIMENTAL MAPS

Table 2: CryoBoltz is most accurate on all 10 real maps

Structure	Res. (Å)	CryoBoltz			Boltz-1			AlphaFold3			ModelAngelo	
		RMSD all	RMSD Cα	TM	RMSD all	RMSD Cα	TM	RMSD all	RMSD Cα	TM	Comp. (%)	TM
Pgp (apo)	4.3	**1.382**	**1.208**	**0.989**	6.994	7.194	0.767	3.827	3.865	0.904	40.3	0.361
Pgp (inward)	4.4	**1.348**	**1.187**	**0.989**	5.630	5.692	0.828	2.692	2.663	0.947	18.3	0.134
Pgp (occluded)	4.1	**1.727**	**1.677**	**0.979**	2.929	2.904	0.942	3.440	3.420	0.921	2.3	0.010
Pgp (collapsed)	4.4	**1.309**	**1.261**	**0.988**	3.425	3.412	0.917	4.568	4.554	0.864	2.5	0.010
Pma1 (active)	3.25	**2.046**	**1.776**	**0.973**	2.987	2.752	0.935	6.628	6.389	0.769	91.5	0.889
Pma1 (inhibited)	3.52	**1.999**	**1.590**	**0.979**	6.140	5.829	0.794	8.017	7.776	0.723	72.8	0.721
CYP (open)	6.5	**4.167**	**3.946**	**0.957**	8.532	8.439	0.788	6.490	6.361	0.890	0.0	0.000
CYP (closed)	4.4	**2.004**	**1.552**	**0.990**	8.784	8.667	0.743	3.585	3.391	0.946	18.9	0.102
YbbAP (bound)	3.66	**1.320**	**0.678**	**0.997**	3.623	3.339	0.928	3.749	3.480	0.922	81.7	0.801
YbbAP (unbound)	4.05	**2.454**	**2.039**	**0.974**	7.842	7.654	0.776	4.022	3.744	0.913	55.4	0.548

RESULTS · EXPERIMENTAL MAPS

Table 2, P-glycoprotein: 1.3–1.7 Å in all four states

Structure	Res. (Å)	CryoBoltz			Boltz-1			AlphaFold3			ModelAngelo	
		RMSD all	RMSD Cα	TM	RMSD all	RMSD Cα	TM	RMSD all	RMSD Cα	TM	Comp. (%)	TM
Pgp (apo)	4.3	**1.382**	**1.208**	**0.989**	6.994	7.194	0.767	3.827	3.865	0.904	40.3	0.361
Pgp (inward)	4.4	**1.348**	**1.187**	**0.989**	5.630	5.692	0.828	2.692	2.663	0.947	18.3	0.134
Pgp (occluded)	4.1	**1.727**	**1.677**	**0.979**	2.929	2.904	0.942	3.440	3.420	0.921	2.3	0.010
Pgp (collapsed)	4.4	**1.309**	**1.261**	**0.988**	3.425	3.412	0.917	4.568	4.554	0.864	2.5	0.010

- Boltz-1 ranges 2.9–7.0 Å and AlphaFold3 2.7–4.6 Å across the same four states.
- ModelAngelo builds only 2.3–40.3% of the chain; TM scores near 0 mean it recovers almost none of the fold.

Paper, Table 2. Pgp: P-glycoprotein; CYP: CYP102A1. Comp.: share of deposited residues ModelAngelo built. Best of 25 samples, mean of 3 replicates.

RESULTS · EXPERIMENTAL MAPS

P-glycoprotein: four transport states from four maps

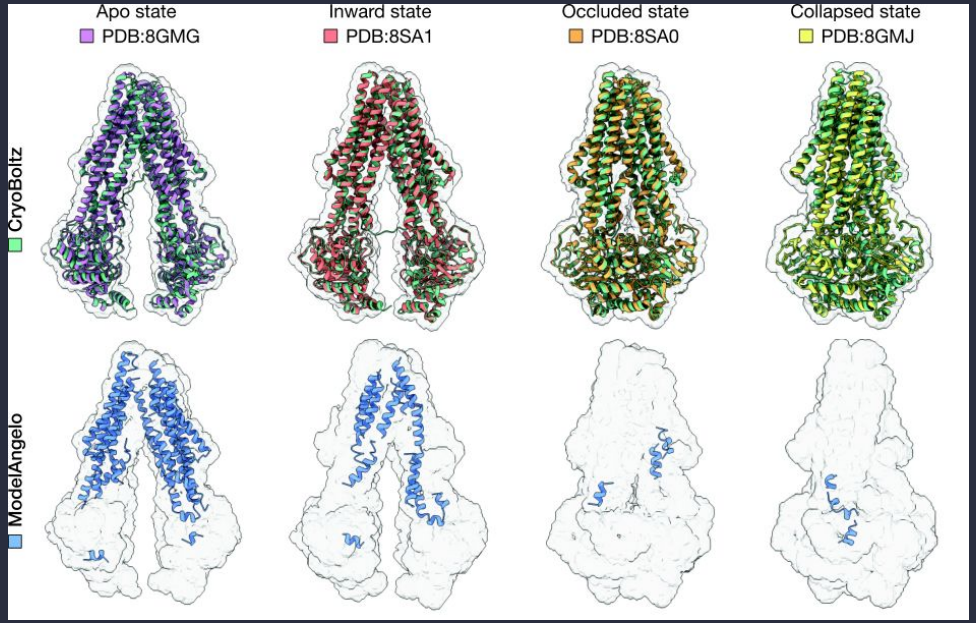

- P-glycoprotein pumps drugs, including chemotherapy drugs, out of cells. The four maps are snapshots of its transport cycle.
- Top: CryoBoltz (green) sits inside each map. Bottom: ModelAngelo builds only fragments from the same maps.
- Caveat: CryoBoltz's maps were masked around the deposited models to remove detergent density; ModelAngelo got the original maps.

Paper, Fig. 5. Maps EMD-40226, 40259, 40258, 40227 (4.1–4.4 Å).

RESULTS · EXPERIMENTAL MAPS

Table 2, other systems: guidance finds the second state

Structure	Res. (Å)	CryoBoltz			Boltz-1			AlphaFold3			ModelAngelo	
		RMSD all	RMSD Cα	TM	RMSD all	RMSD Cα	TM	RMSD all	RMSD Cα	TM	Comp. (%)	TM
Pma1 (active)	3.25	**2.046**	**1.776**	**0.973**	2.987	2.752	0.935	6.628	6.389	0.769	91.5	0.889
Pma1 (inhibited)	3.52	**1.999**	**1.590**	**0.979**	6.140	5.829	0.794	8.017	7.776	0.723	72.8	0.721
CYP (open)	6.5	**4.167**	**3.946**	**0.957**	8.532	8.439	0.788	6.490	6.361	0.890	0.0	0.000
CYP (closed)	4.4	**2.004**	**1.552**	**0.990**	8.784	8.667	0.743	3.585	3.391	0.946	18.9	0.102
YbbAP (bound)	3.66	**1.320**	**0.678**	**0.997**	3.623	3.339	0.928	3.749	3.480	0.922	81.7	0.801
YbbAP (unbound)	4.05	**2.454**	**2.039**	**0.974**	7.842	7.654	0.776	4.022	3.744	0.913	55.4	0.548

- The state Boltz-1 misses: Pma1 inhibited 2.00 Å (Boltz-1 6.14), YbbAP unbound 2.45 Å (7.84), CYP closed 2.00 Å (8.78).
- Resolution sets the ceiling: CYP open (6.5 Å) only reaches 4.17 Å and ModelAngelo builds nothing there; it does well only on the sharper Pma1 and YbbAP maps.

Paper, Table 2. Pgp: P-glycoprotein; CYP: CYP102A1. Comp.: share of deposited residues ModelAngelo built. Best of 25 samples, mean of 3 replicates.

RESULTS · EXPERIMENTAL MAPS

Pma1, CYP102A1, YbbAP: two states from two maps

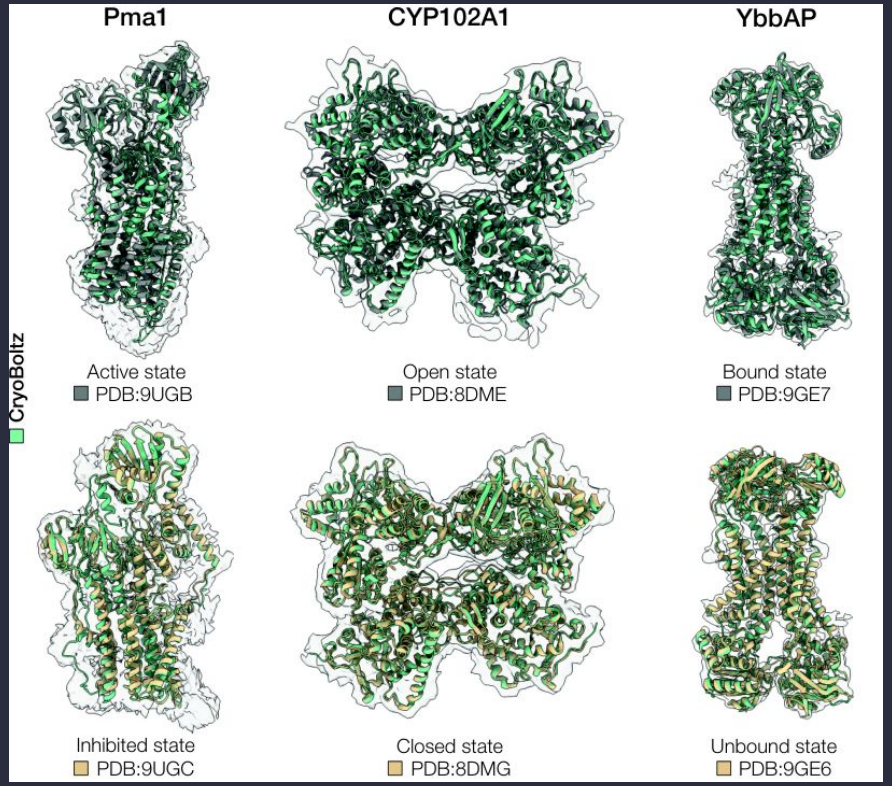

- Top row: one state of each complex; bottom row: the other, each guided by its own map.
- Pma1 is an ATPase (active vs inhibited), CYP102A1 an enzyme (open vs closed), YbbAP a transporter (with vs without ATP).
- Where the map is blurriest (CYP102A1 open, 6.5 Å), the prior fills in more, so atom positions there are a hypothesis.

Paper, Fig. 6. Green: CryoBoltz; gray and tan: deposited models.

DISCUSSION

Limitations: stability, initialization, and selection

- The map likelihood has many local optima, so guidance is unstable; CryoBoltz relies on many samples and keeps the best.
- It needs a reasonable warm-up structure. It failed where unguided Boltz-1 started far off (DSL1/SNARE complex, full IgG antibody).
- Reported samples are picked by distance to the true structure. Picking by map fit instead (RSCC, Table B2) keeps most of the gain.
- On blurry maps the prior fills the gaps: CYP102A1 at 6.5 Å still has 4.17 Å error.
- Stage lengths and guidance strengths are set by hand.

Paper, Section 5 and Appendix Table B2.

DISCUSSION

Takeaways

- A cryo-EM map can steer a pretrained structure predictor at sampling time, with no retraining.
- Global guidance changes the conformational state; local guidance fits the detail.
- CryoBoltz had the lowest error on all 13 maps tested, including states Boltz-1 and AlphaFold3 do not produce.
- It builds complete models where ModelAngelo builds fragments, but details the map cannot resolve still come from the prior and need validation.

DISCUSSION

Questions

1. Does CryoBoltz replace map-based builders like ModelAngelo, or complement them?
2. At low resolution, how do we tell which atoms the map supports and which Boltz-1 made up?
3. The stage schedule is hand-set. Could it be adaptive instead?
4. CryoBoltz fails when the warm-up structure is far off. Fundamental limit or fixable?
5. Would fine-tuning Boltz-1 on cryo-EM data beat guiding it at sampling time?