Experimental conditioning

How measurements inform protein models

A field guide

Mannat Vikramaditya Jain

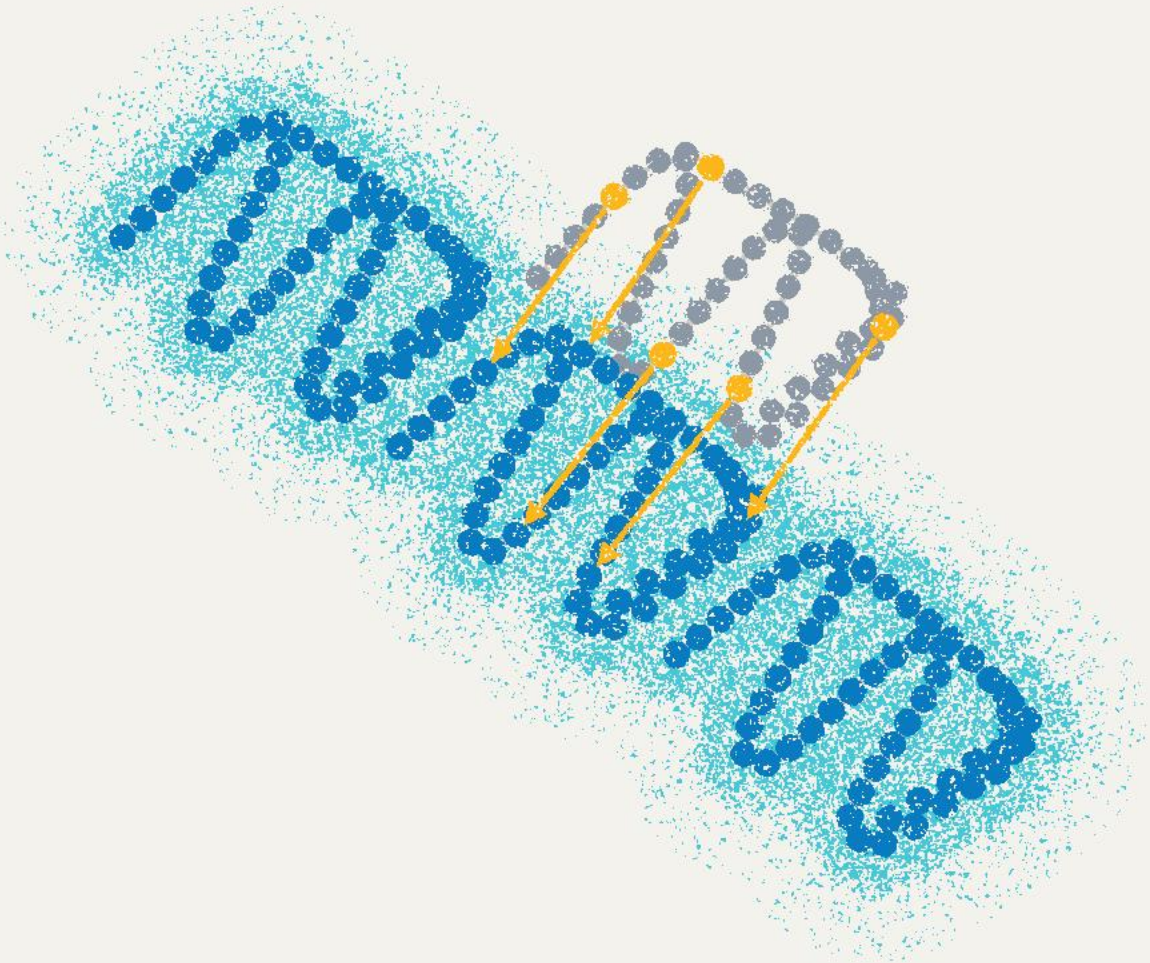

CONDITIONING

Experimental data can improve a predicted structure.

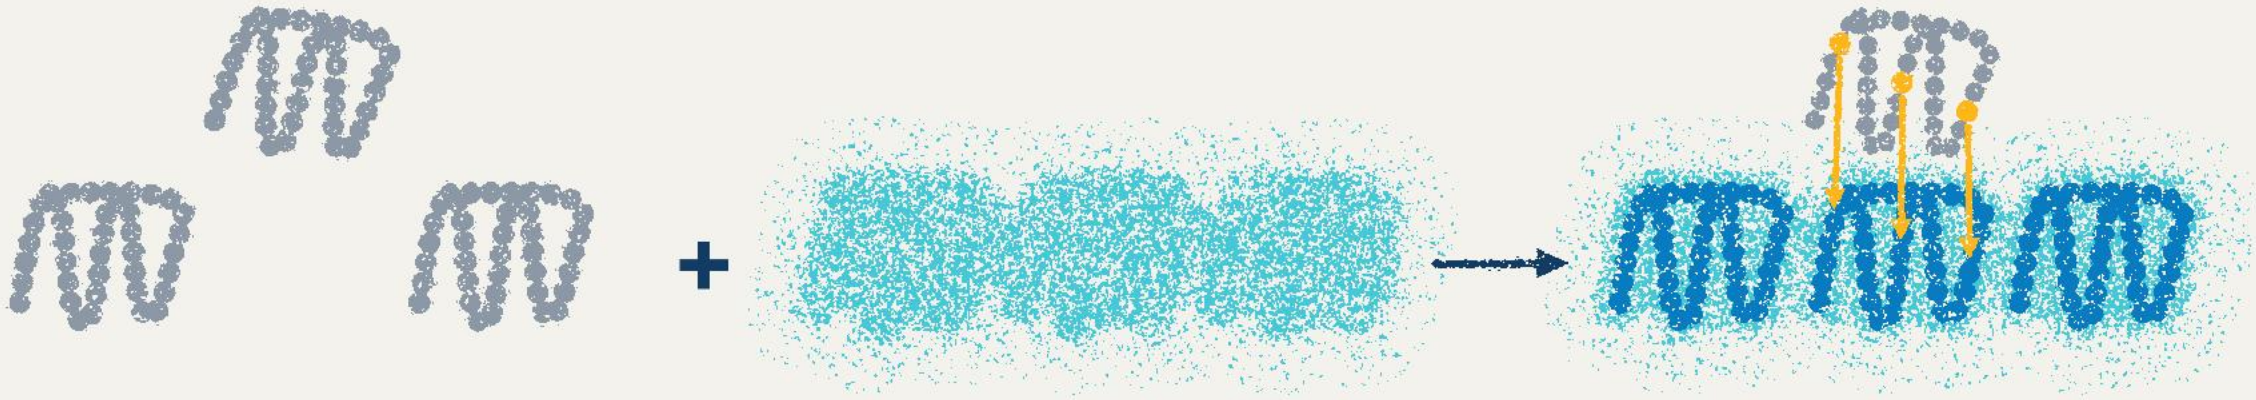

predicted ZPD trimer
AlphaFold2 Multimer

8.6 Å cryo-EM map
measured on ZPD itself

revised with ROCKET
moved to agree with the map

The prediction is inferred from sequences; the map is a measurement of the protein itself.

Fadini et al., Nature Methods 2026 (ROCKET): Results, ZPD; Fig. 4 · drawings schematic, not the ZPD map or models

The fitted model agreed better with a later map it never saw.

Map correlation (CC) with the later 4.6 Å map

higher is better

AlphaFold2 Multimer prediction, unrefined

0.29

ROCKET pipeline, including polishing

0.67



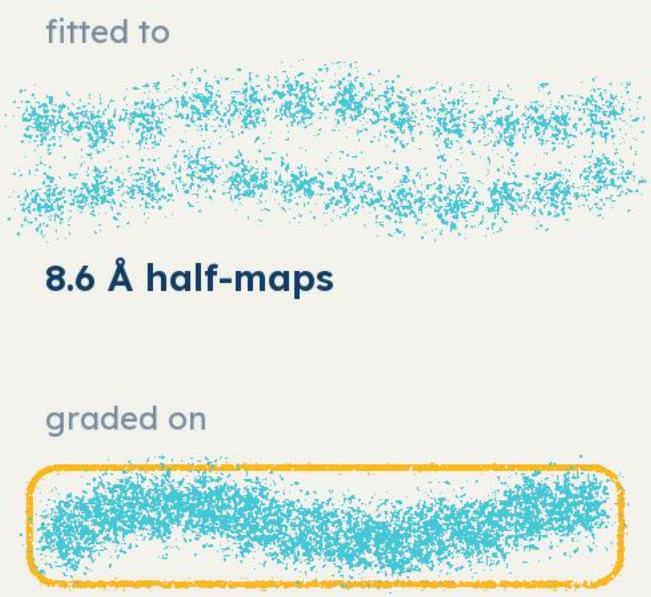

One case: the whole pipeline vs the unrefined prediction; no interval reported.

Fadini et al., Nature Methods 2026 (ROCKET): Fig. 4b; Results, ZPD · map glyphs schematic

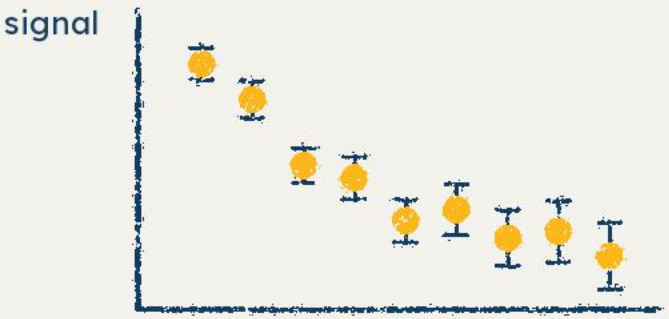

what an instrument records, with its errors

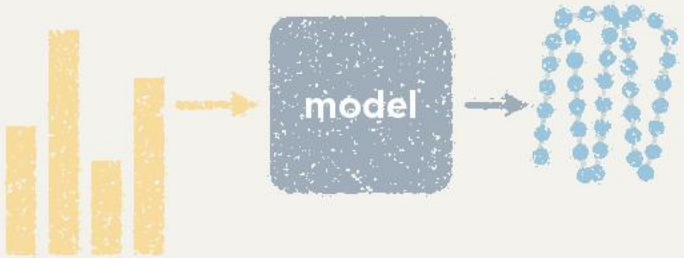

where the data enter a protein model

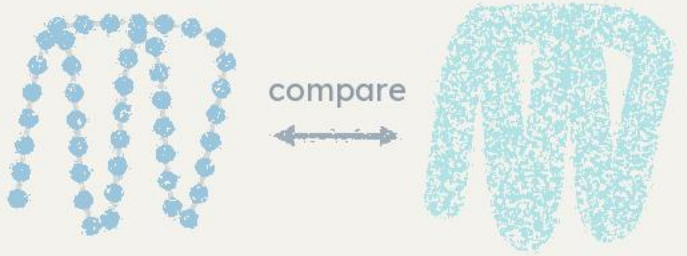

scored against data the fit never saw

I

Experimental methods

pages 6–12

II

Integrating

pages 15–23

III

Validation and transfer

pages 26–30

Coming up…

- What does each experiment actually record? pages 6–11
- Which parts of a structure can that signal pin down? pages 7–11
- When do two assay numbers measure the same thing? page 12

Experiments measure signals; atomic coordinates are inferred.

	Measured what the instrument records	Reconstructed a map, after processing	Inferred an atomic model that fits
X-ray crystallography	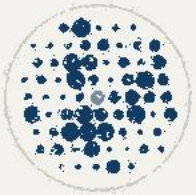 **spot intensities, each ± σ** phases are not measured	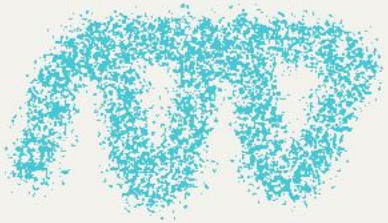 **density, with phases** often taken from a model	 **coordinates refined** against data and chemistry
Cryo-EM electron microscopy	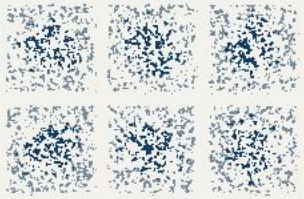 **noisy 2D images** orientations unknown	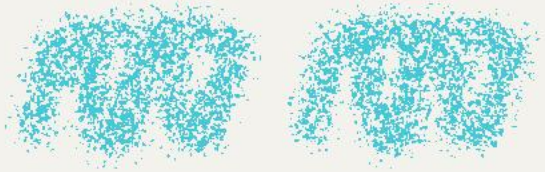 **two half-maps** agreement = signal	 **a model built in** into the map

Drawings schematic.

Fadini et al. 2026, Nature Methods (ROCKET): Methods (X-ray and cryo-EM inputs, half-map noise model)

NMR shifts constrain local angles; NOEs link residues far apart in sequence.

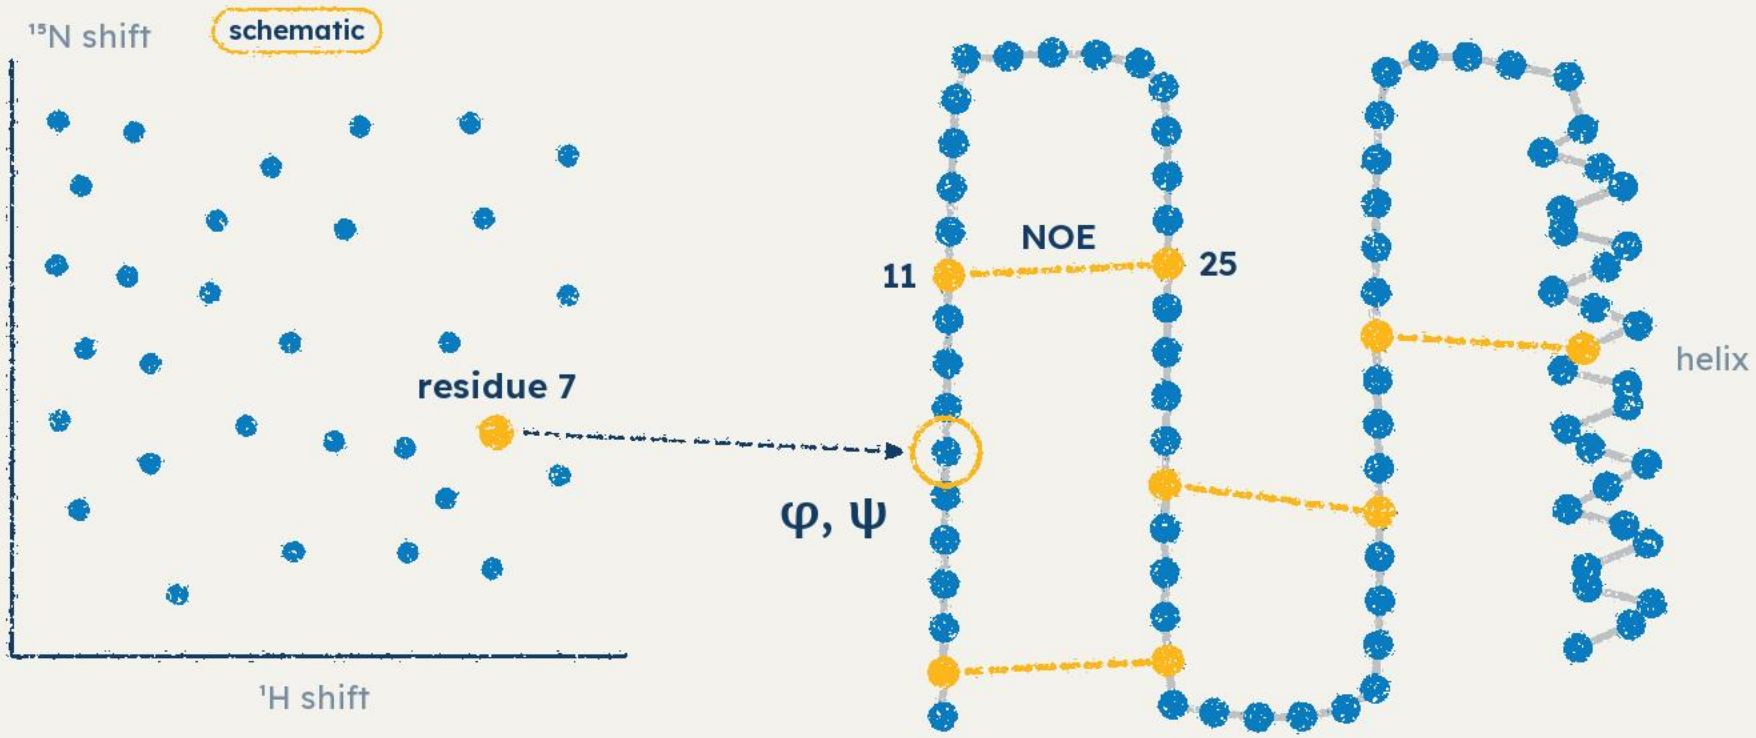

peak = one backbone amide; its shifts give φ, ψ

NOE: two protons a few Å apart, far apart in sequence

Shen & Bax 2013, J Biomol NMR (TALOS-N): backbone shifts → φ/ψ

A crosslink says two residues were in contact. Real ones helped most where AlphaFold2 was wrong.

What one crosslink records

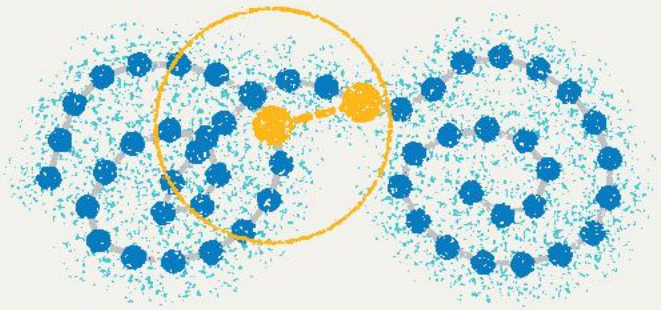

one residue pair

read as a Cα–Cα contact, ≈ 10 Å

mass spectrometry names the two linked residues

AlphaLink on real in-cell crosslinks, 31 proteins

relative TM-score gain over AlphaFold2 · mean and 95% CI · MSAs thinned to Neff 10

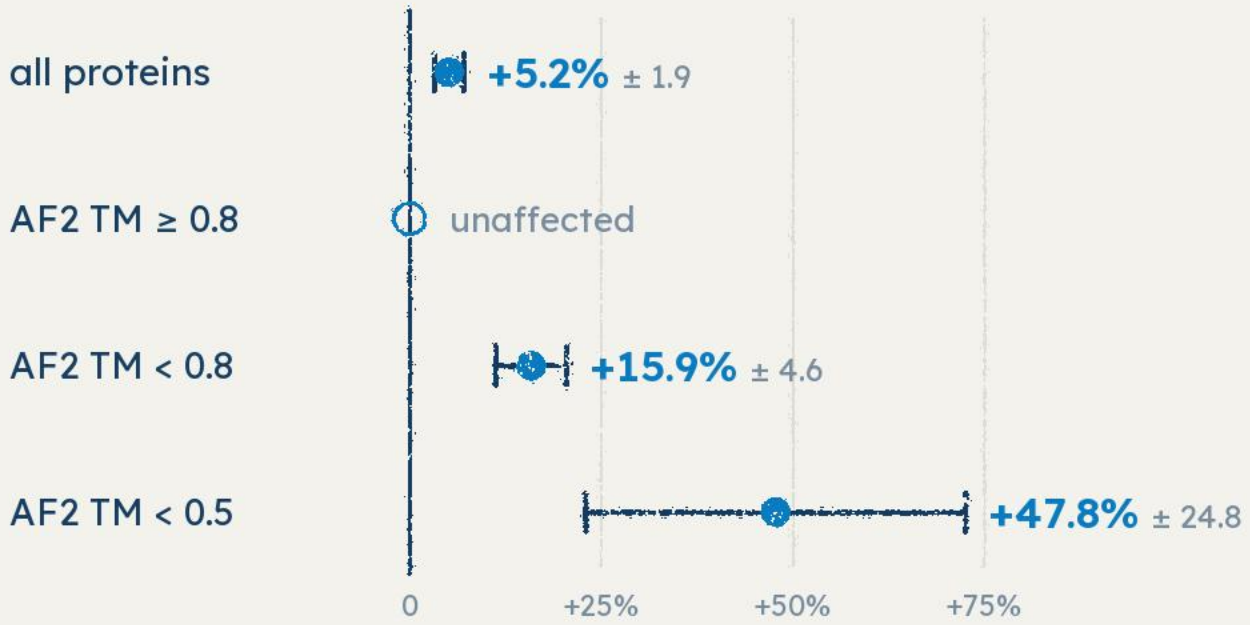

Stahl et al. 2023, Nature Biotechnology (AlphaLink): Results (real-data panel); ± = 95% CI as stated

SAXS sees an ensemble average, so a good fit is not unique.

1 Two states, half each



closed · 0.5 open · 0.5

2 One structure

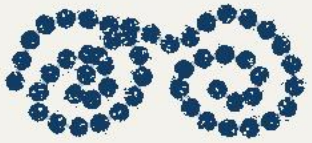

partly open · 1.0

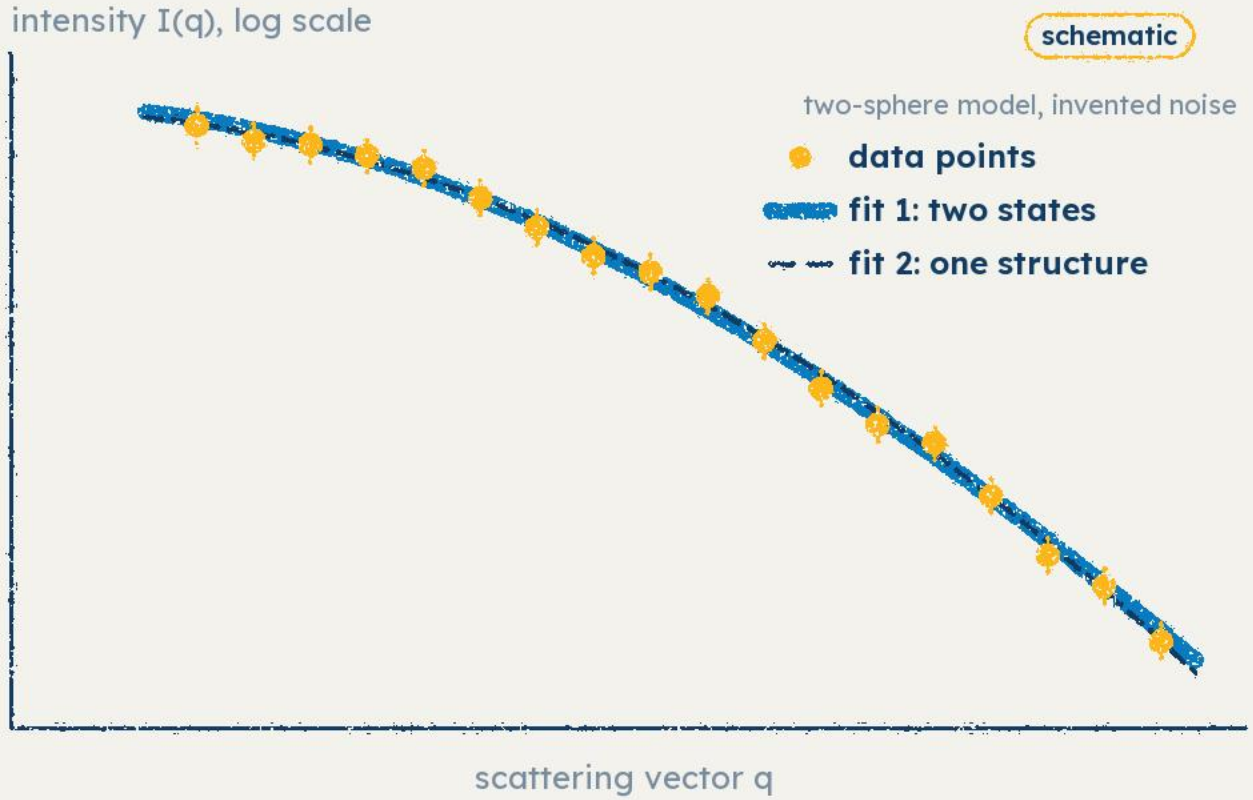

The sample holds no intermediate, yet this curve cannot tell 1 from 2.

Real-data cases: Potrzebowski et al. 2018 (calmodulin); Shevchuk & Hub 2017 (synthetic LBP test); both PLoS Comput Biol

HDX measures uptake per peptide; only a forward model connects it to structure.

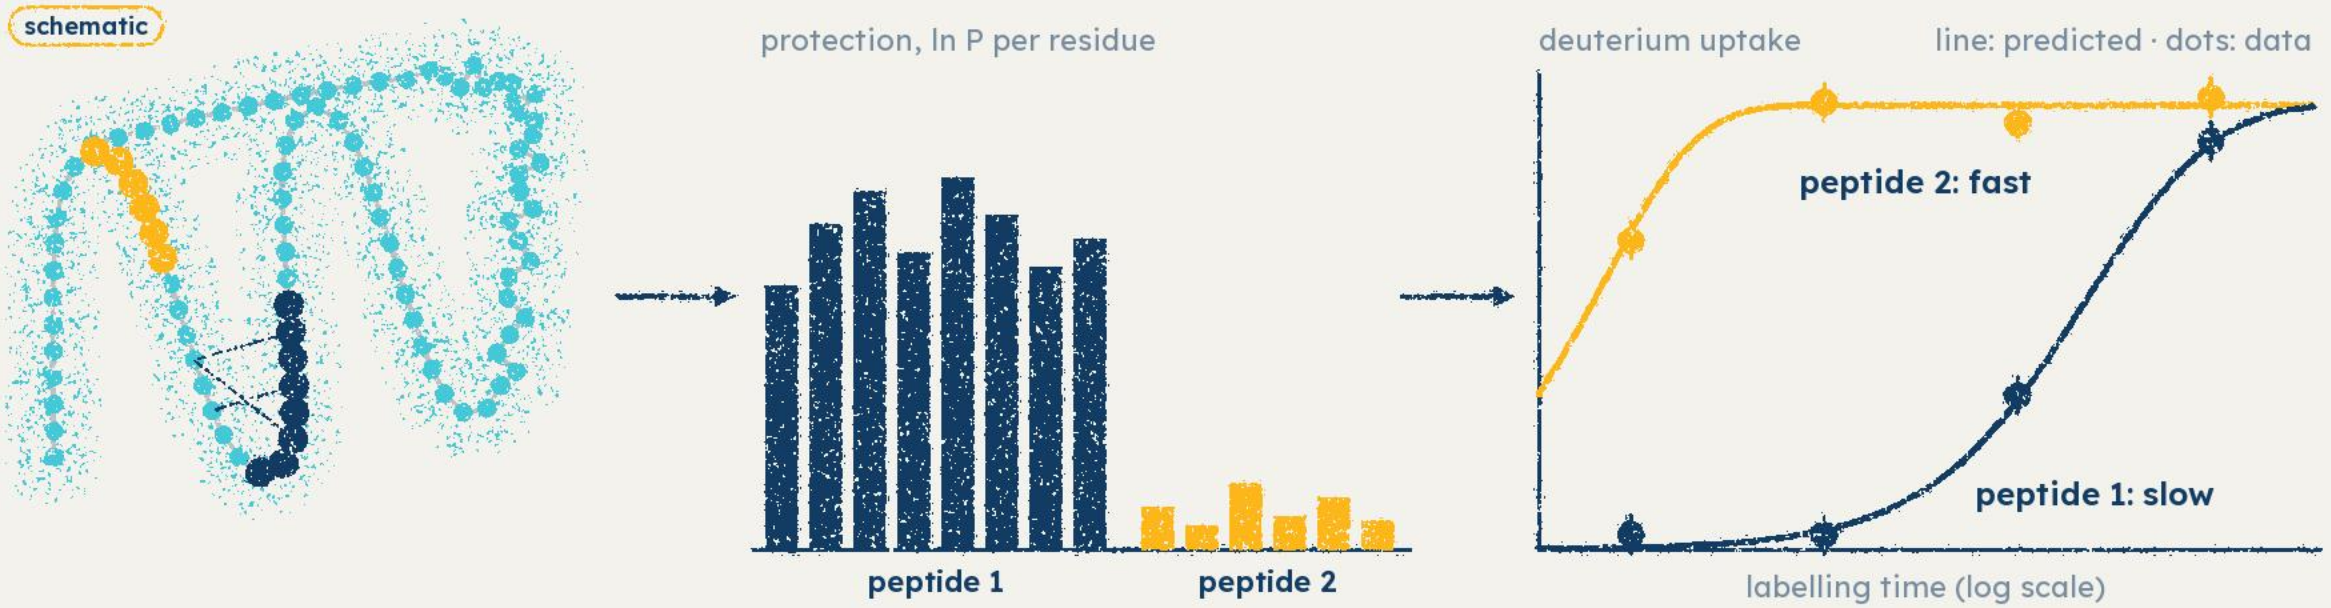

1 Structure
contacts and H-bonds per amide

2 Protection
ln P = 0.35 · contacts + 2.0 · H-bonds

3 Uptake per peptide
predicted, then compared with data

A structure read from HDX inherits this model's fitted weights and its peptide averaging.

Bradshaw et al. 2020, Biophys J (HDXer): Methods Eq. 1 (β = 0.35, 2.0)

Averaged FRET is a mean efficiency, not a mean distance.

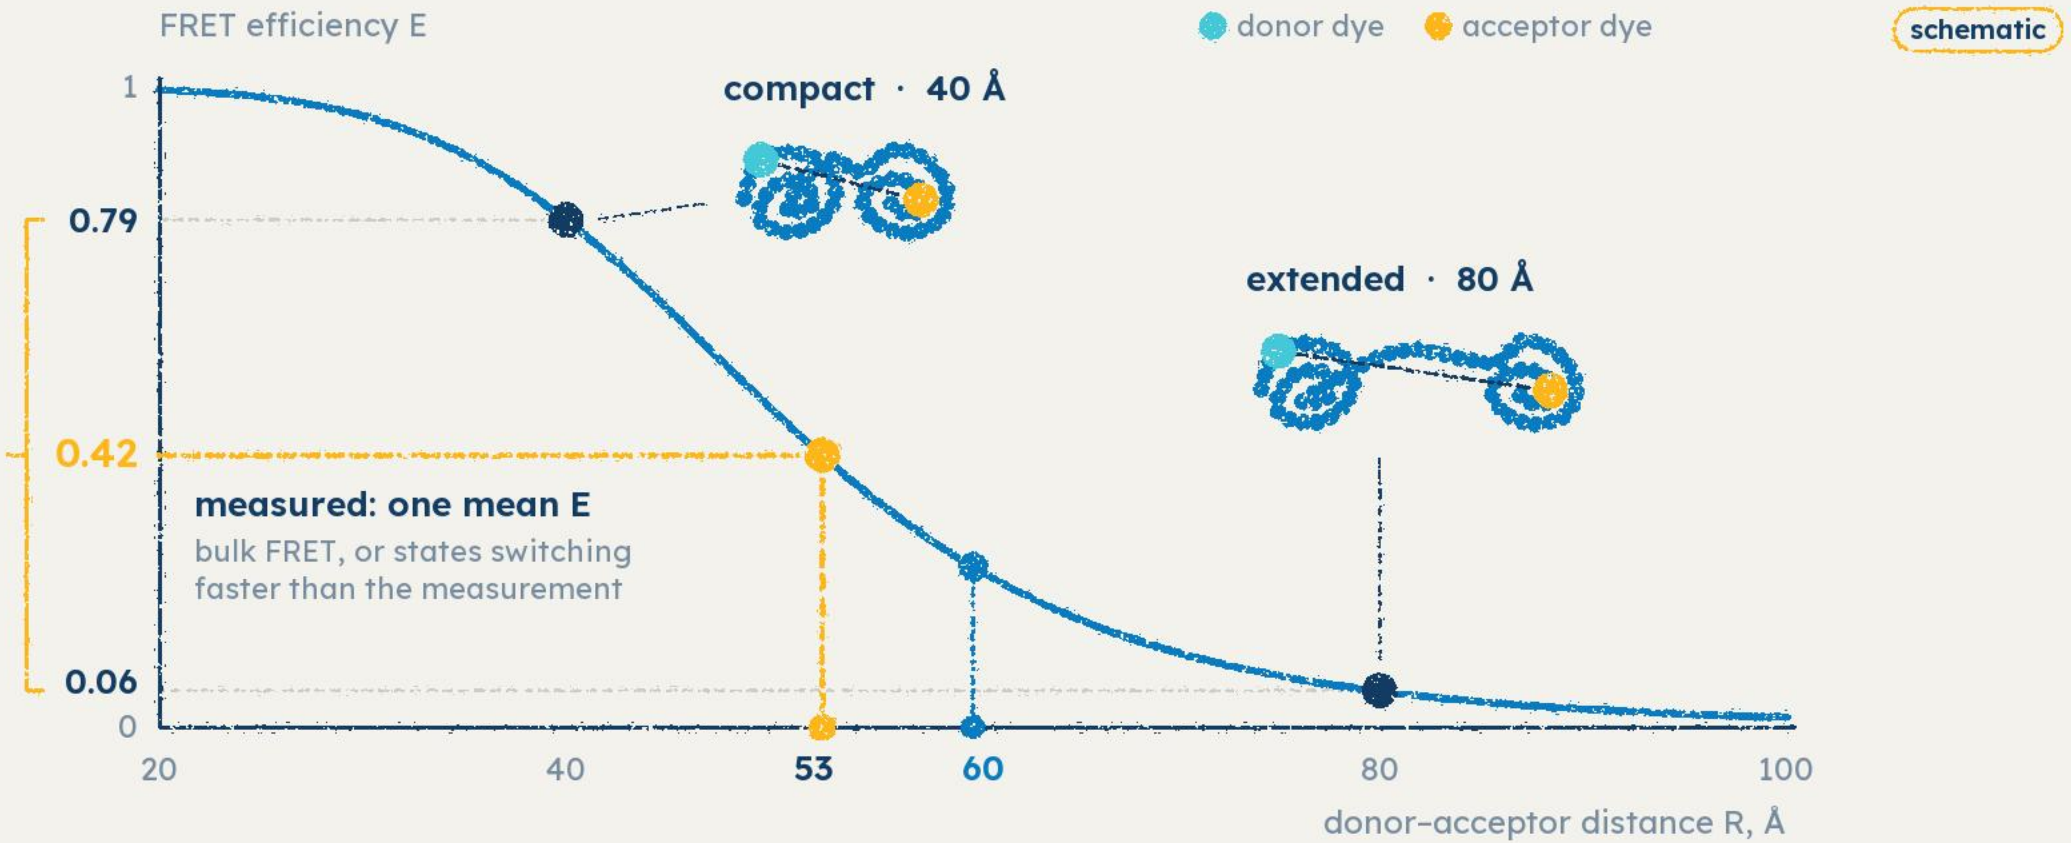

Read as one structure, E = 0.42 gives 53 Å; no molecule is there, and the mean distance is 60 Å.

Schematic: Förster relation $E = 1/(1 + (R/R_0)^6)$, illustrative $R_0 = 50$ Å, two states 50:50

An assay label needs a definition of what was measured.

Tite-Seq: one anti-fluorescein scFv, a full titration curve per variant

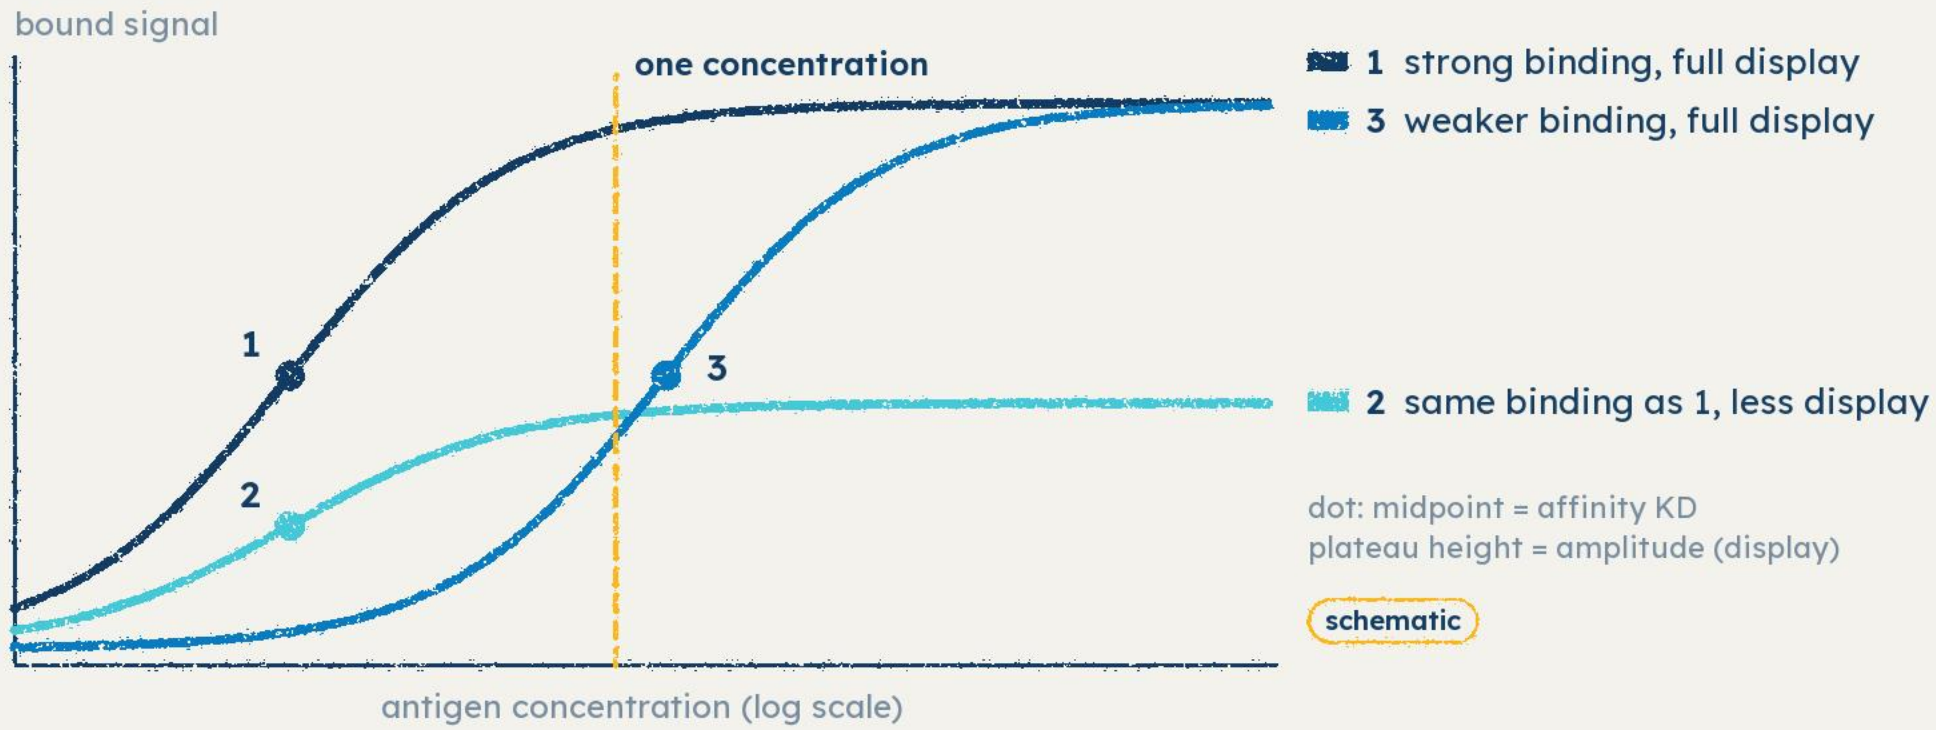

At one concentration variants 2 and 3 look alike; the full curve separates affinity (KD) from display.

Adams et al. 2016, eLife (Tite-Seq): Eq. 1 form f = A·c/(c + KD) + B; curves use illustrative parameters

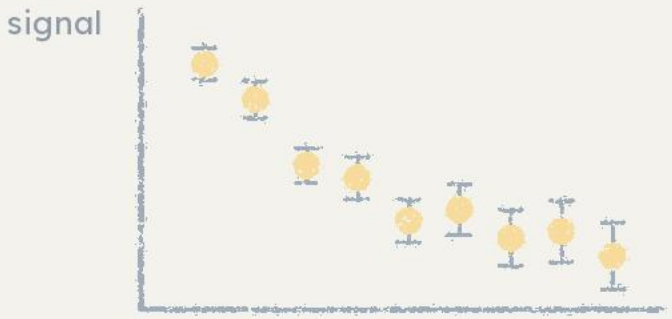

what an instrument records, with its errors

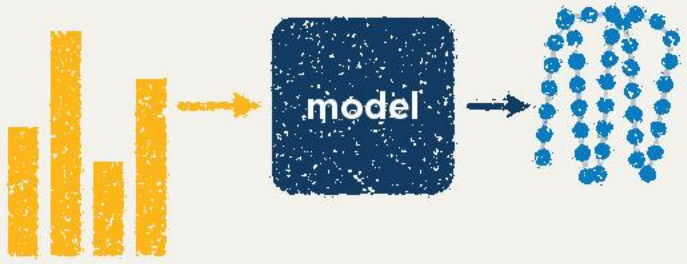

where the data enter a protein model

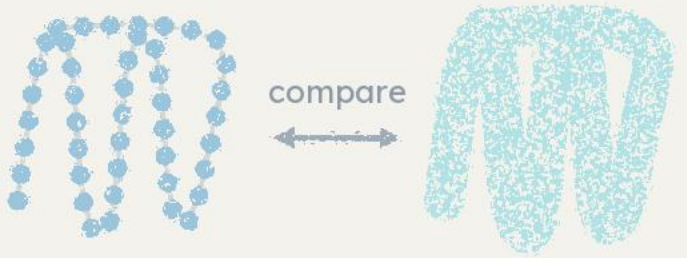

scored against data the fit never saw

I
Experimental methods
pages 6–12

II
Integrating
pages 15–23

III
Validation and transfer
pages 26–30

Coming up…

- Where does a measurement enter the model? pages 15–21
- Does a closer fit mean a better structure? pages 20, 22
- How hard should a signal pull, given its errors? page 23



A measurement can enter the model at five different points.

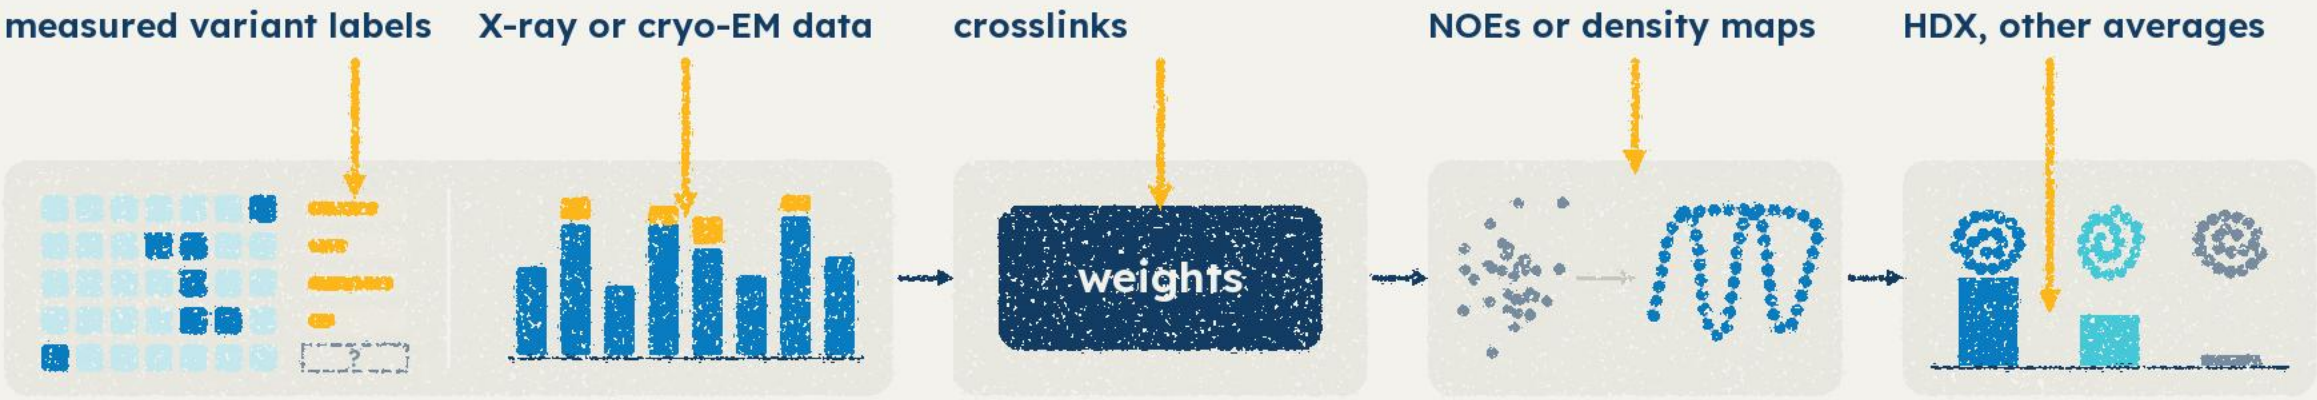

17 Labeled context
ProteinNPT

18 Optimize input
ROCKET

16 Train conditioner
AlphaLink

19 Guide sampler
Experiment-guided AF3

21 Reweight pool
HDXer, BioEn

Composite schematic: each study uses its own model.

Notin et al. 2023 (ProteinNPT) · Fadini et al. 2026 (ROCKET) · Stahl et al. 2023 (AlphaLink)
Maddipatla et al., accepted (EG-AF3) · Bradshaw et al. 2020 (HDXer) · Köfinger et al. 2019 (BioEn)

Train once on simulated crosslinks, then feed in real ones.

1 Fine-tune once, on simulated links



2 Predict, with the target's measured links

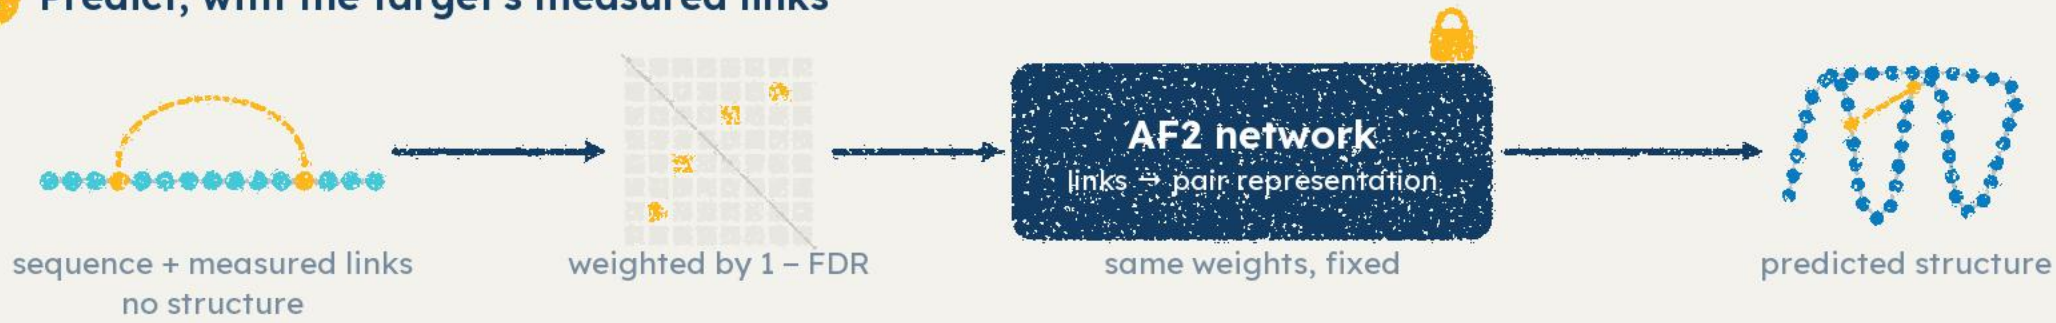

CASP14 control, no links given: fine-tuned network 0.701 mean TM, AF2 0.702.

Stahl et al. 2023, Nat. Biotechnol. (AlphaLink): Fig. 1; Methods; Extended Data Fig. 10

Labeled context can change a prediction without another weight update.

What the model reads

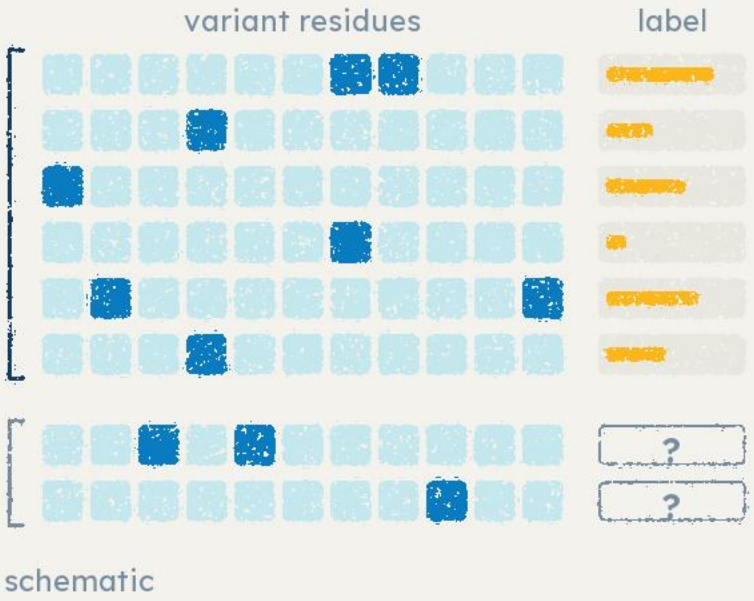

support rows
measured variants of the target

query rows
label masked

schematic

Average Spearman ρ, same trained model

support rows given at inference

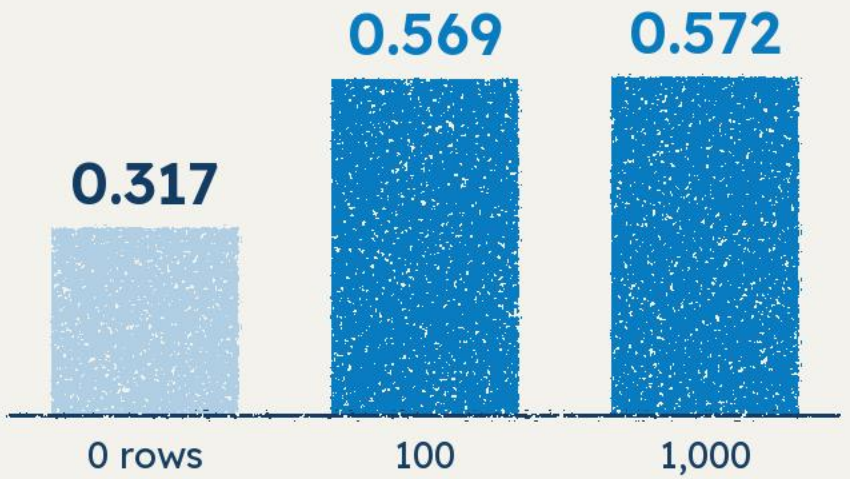

mean over 8 DMS assays and 3 CV schemes; no CI reported

Only the batch changed, whole labeled rows at a time; the weights, already trained on this assay, did not.

Notin et al. 2023, NeurIPS (ProteinNPT): §3; Table 2, PDF p. 10

Fit the input, not the weights: ROCKET steers a frozen predictor toward the measured data.

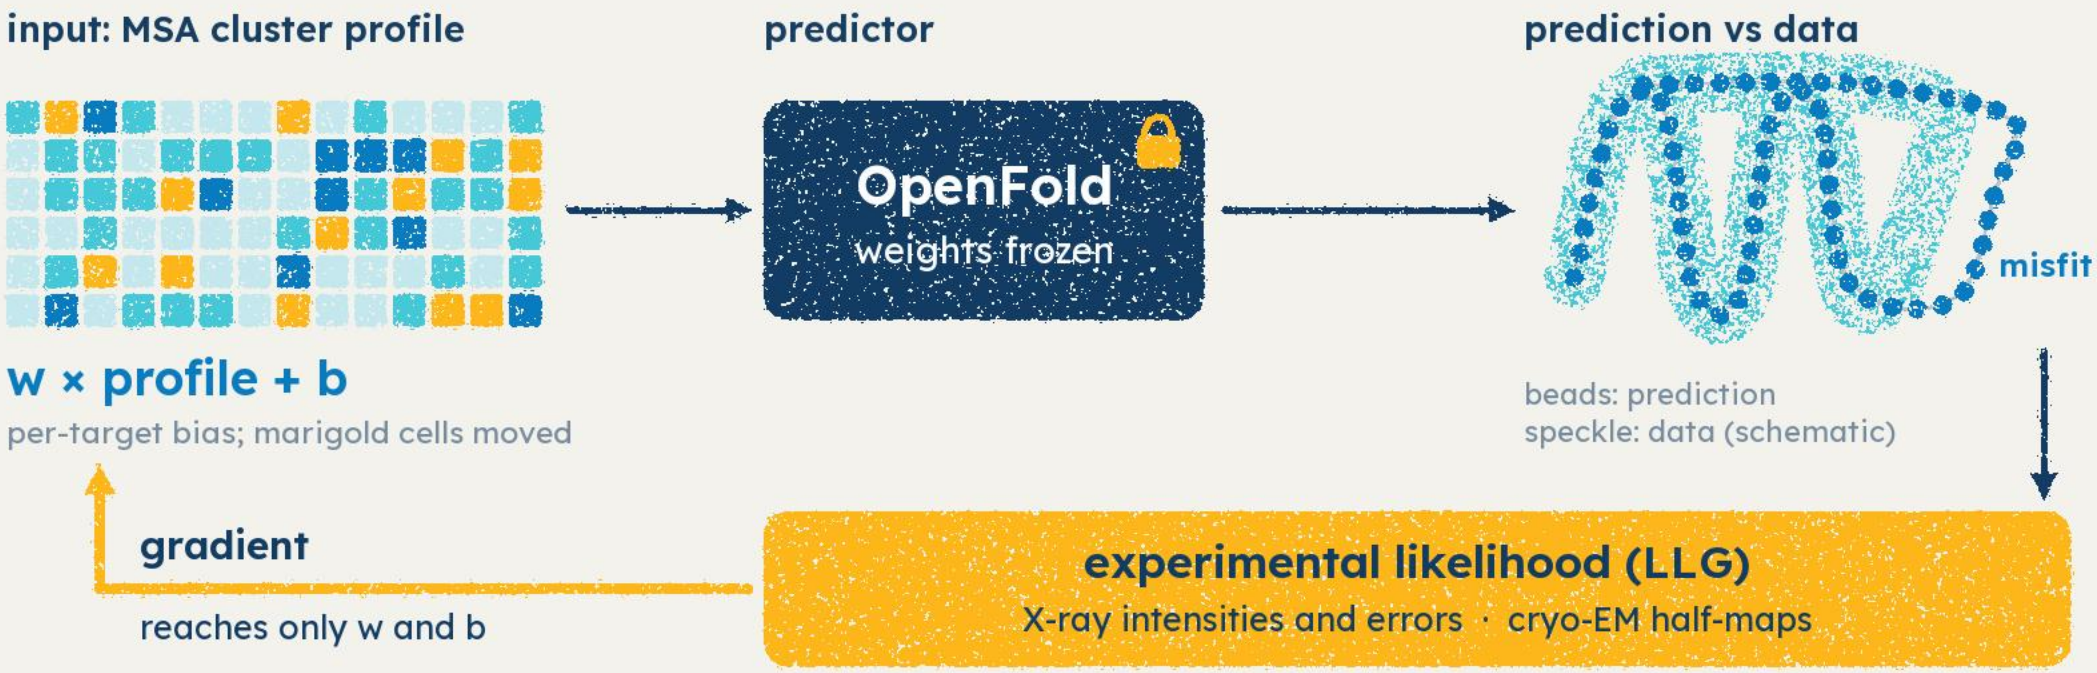

Worked cryo-EM case (ZPD, fitted at 8.6 Å, checked at 4.6 Å): pages 02–03.

Fadini et al. 2026, Nat. Methods (ROCKET): Fig. 1b; Methods

Guidance adds a data pull to every denoising step.

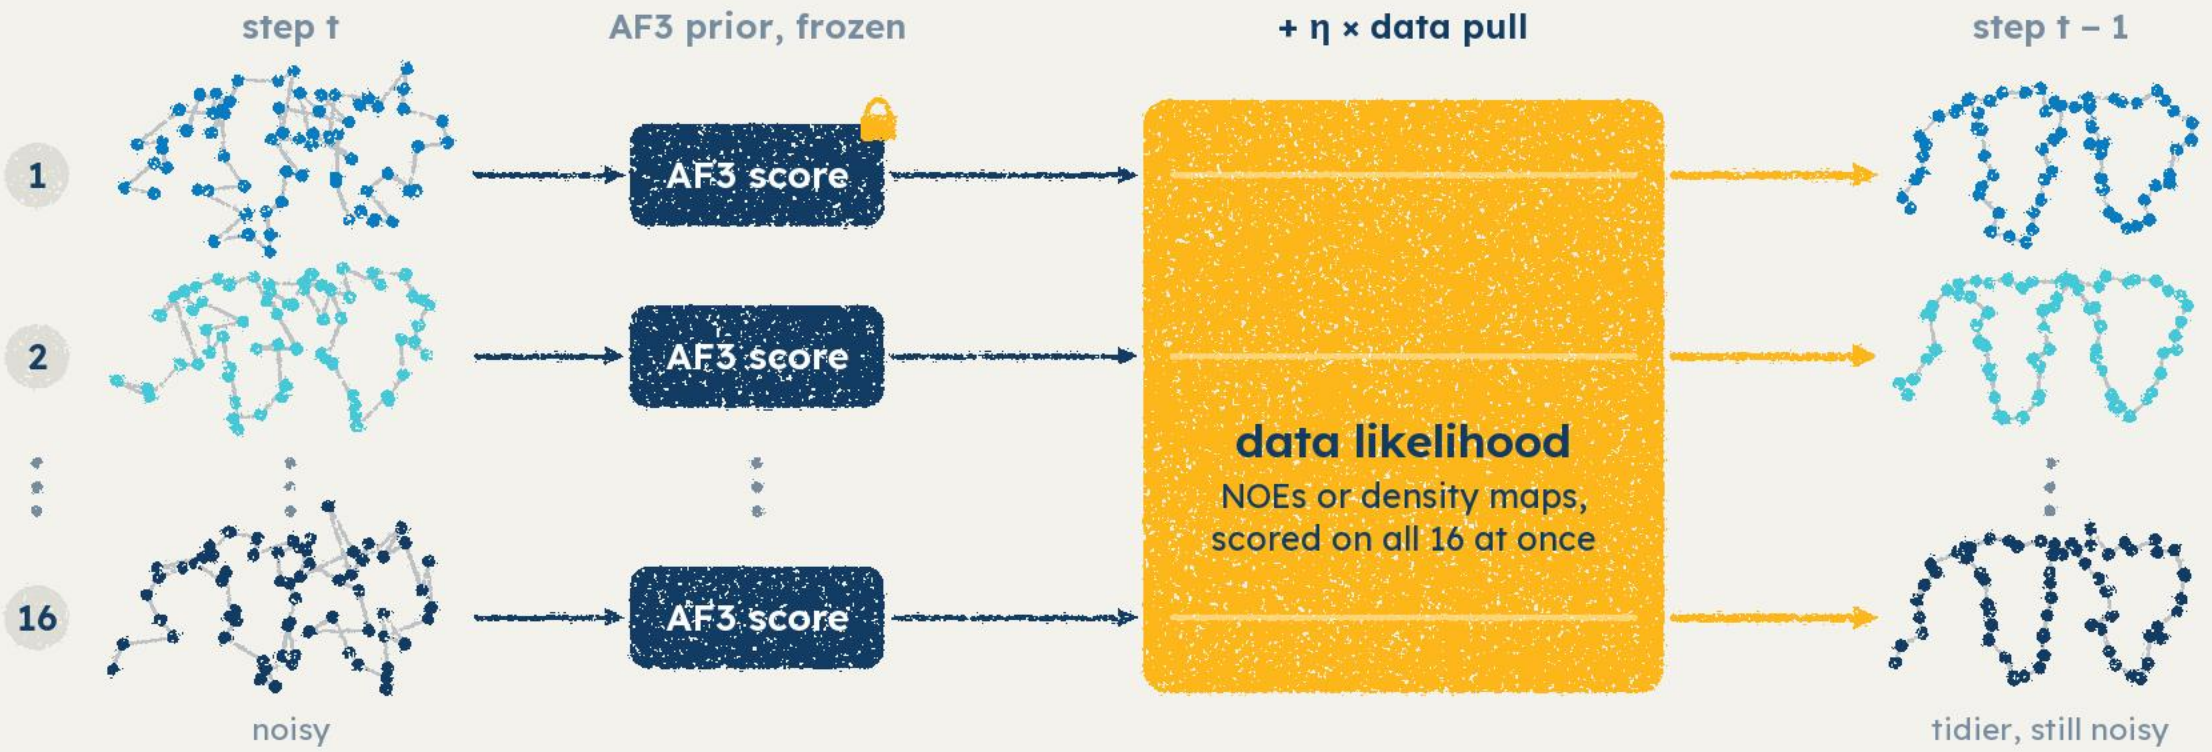

Each conformer's step = its own AF3 score + η × gradient of log p(data | all 16).

Maddipatla et al., Nat. Biotechnol., accepted manuscript (EG-AF3): Methods, Eq. 4

Fewer NOE violations, scored on the guiding NOEs.

Ubiquitin (PDB 1D3Z), one protein

NOE restraint groups violated, % of 1,212; lower is better

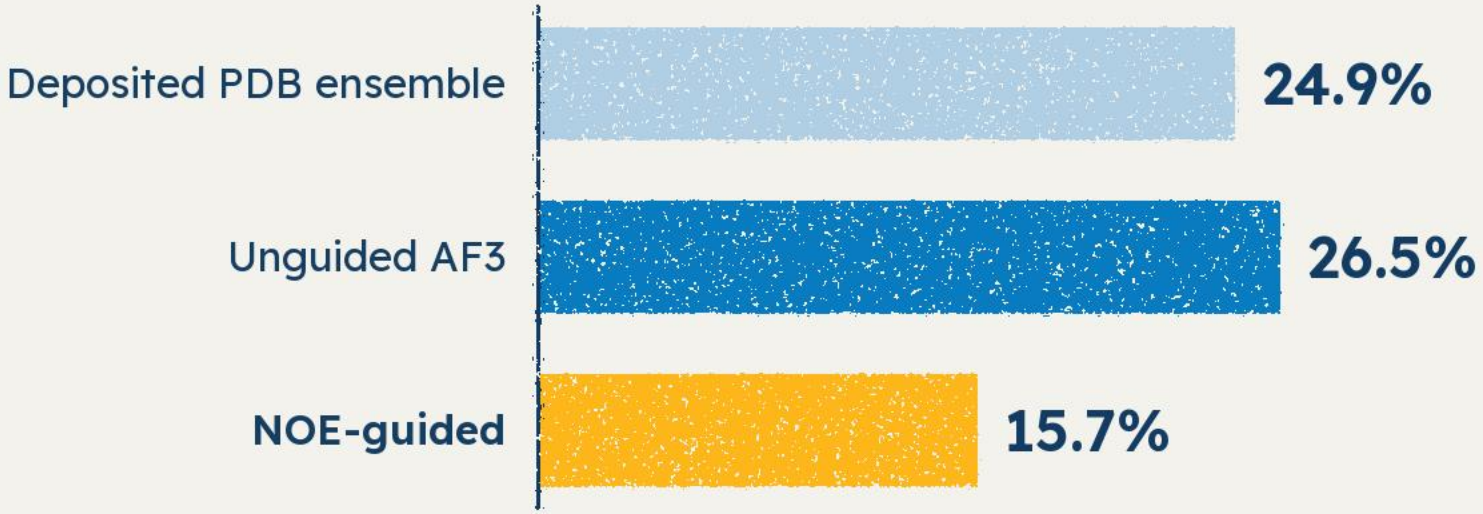

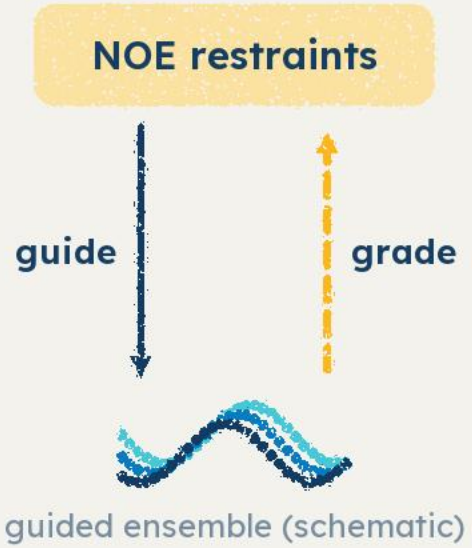

The grader is the guide: this shows the guidance took hold, not that the ensemble is right.

Maddipatla et al., Nat. Biotechnol., accepted manuscript (EG-AF3): Fig. 2c

Reweighting shifts weight among sampled structures; a missing state stalls the fit.

HDXer on synthetic HDX of TeaA (60% closed, 40% open), made by the same forward model it fits.

Simulated pool
decoys match neither state

schematic frames and weights



misfit to target HDX
mean squared deviation

Full pool
weight after fitting

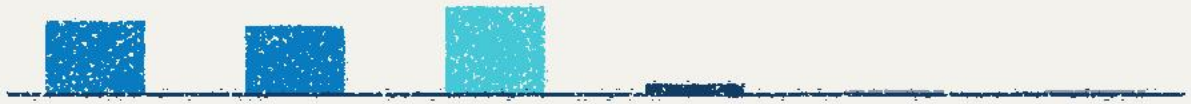

$\leq 10^{-7}$
from 2.2×10^{-3} before reweighting

Closed frames
removed before fitting

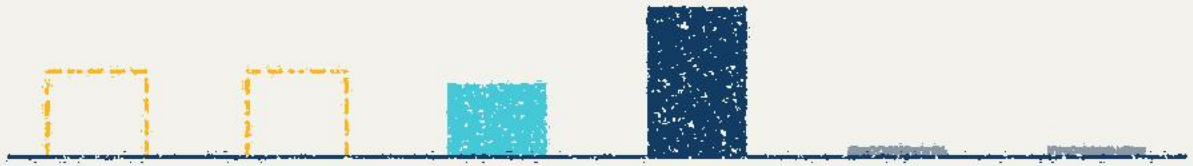

$\approx 10^{-4}$
stalls here

weight piles onto semi-open frames

Bradshaw et al. 2020, Biophys. J. (HDXer): Results, Figs 3, 7

A better map fit can come from a broken structure.

CryoSampler on TRPV3, trained with and without its geometry penalty.

schematic
cyan: map density
marigold: fails a check

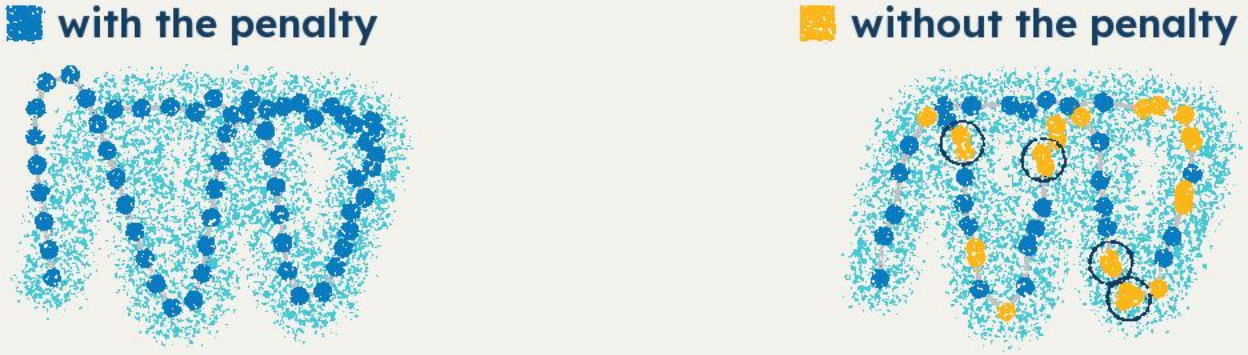

Map fit, CC_volume
higher is better

0.84 ± 0.02

0.93 ± 0.01

Ramachandran outliers
lower is better

0.00% ± 0.00

55.94% ± 1.94

TRPV3 only · mean ± SD over 3 runs, not CIs · fit scored on the four training maps · preprint

Shenoy et al. 2026, arXiv 2605.09832v1 (CryoSampler, preprint): Table 3, PDF p. 14, TRPV3

Weight each measurement by the error it is expected to carry.

schematic

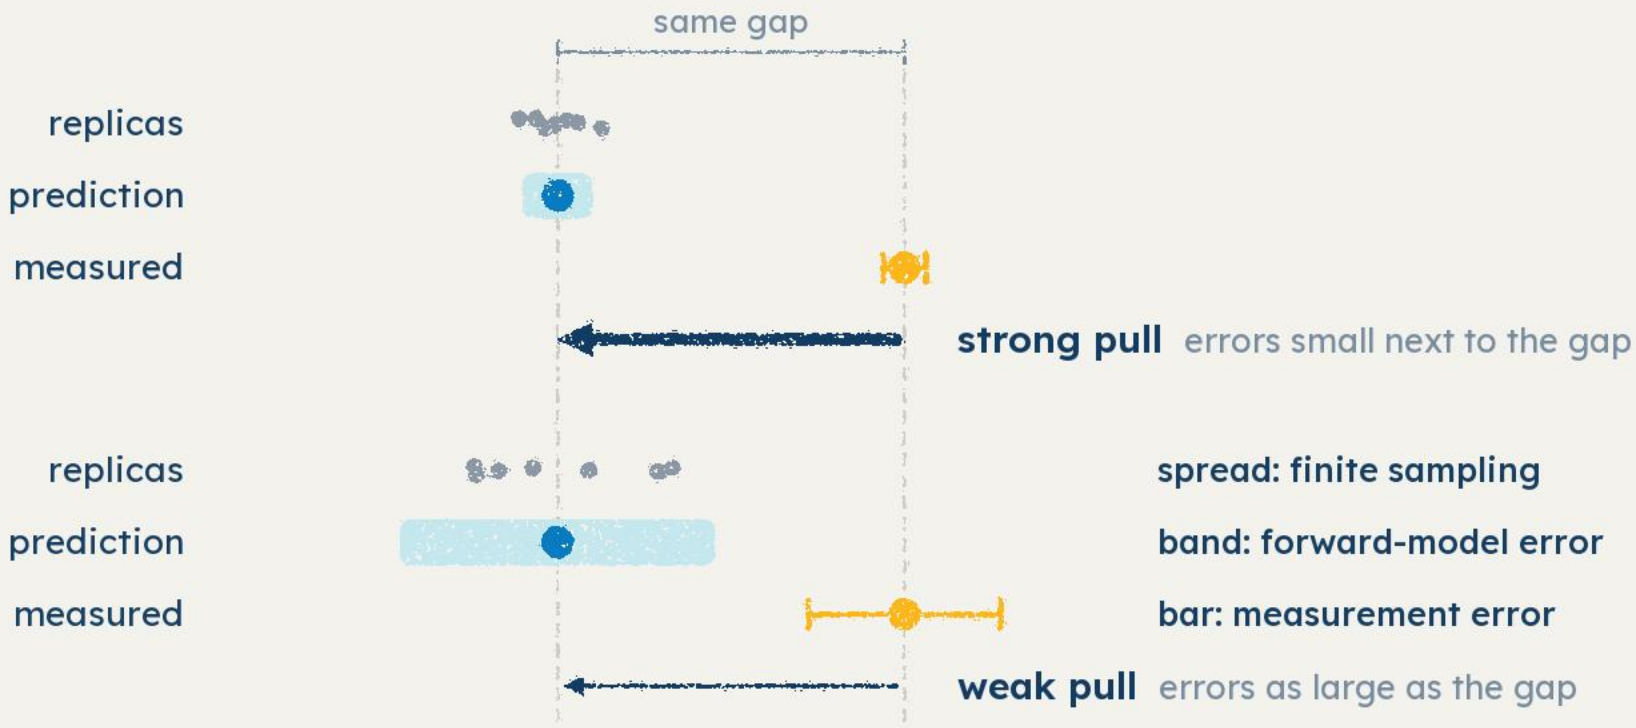

Metainference infers the data errors along with the ensemble; a gap inside them is not evidence against the model.

Bonomi et al. 2016, Sci. Adv. (metainference): Results

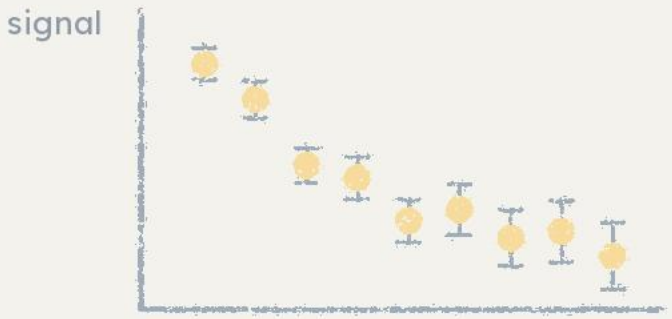

what an instrument records, with its errors

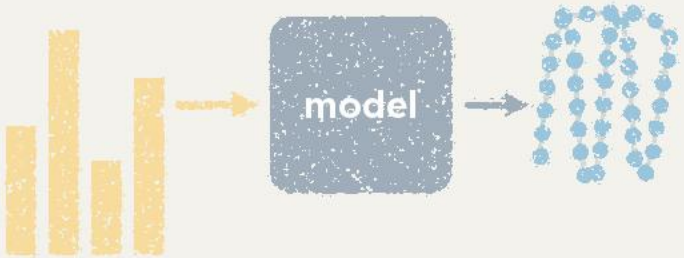

where the data enter a protein model

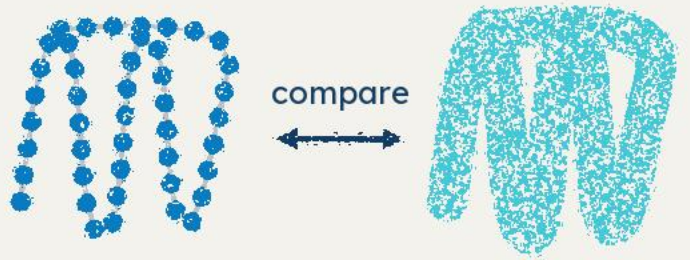

scored against data the fit never saw

I

Experimental methods

pages 6–12

II

Integrating

pages 15–23

III

Validation and transfer

pages 26–30

Coming up…

- Does a learned motion transfer to a related protein? pages 26–27
- Can a benchmark split hide the same protein? page 28
- Which checks are independent of the fit? page 29



TRPV3 maps train the sampler; TRPV5 maps only grade it.

CryoSampler (2026 preprint): a motion learned on one TRP channel, sampled in a related one.

TRPV3 maps
train the model

Input for TRPV5
one static Boltz-2 structure

TRPV5 samples
drawn by the model
16 sampled, 4 shown

TRPV5 maps
held out: grade only
same motion, same spacing

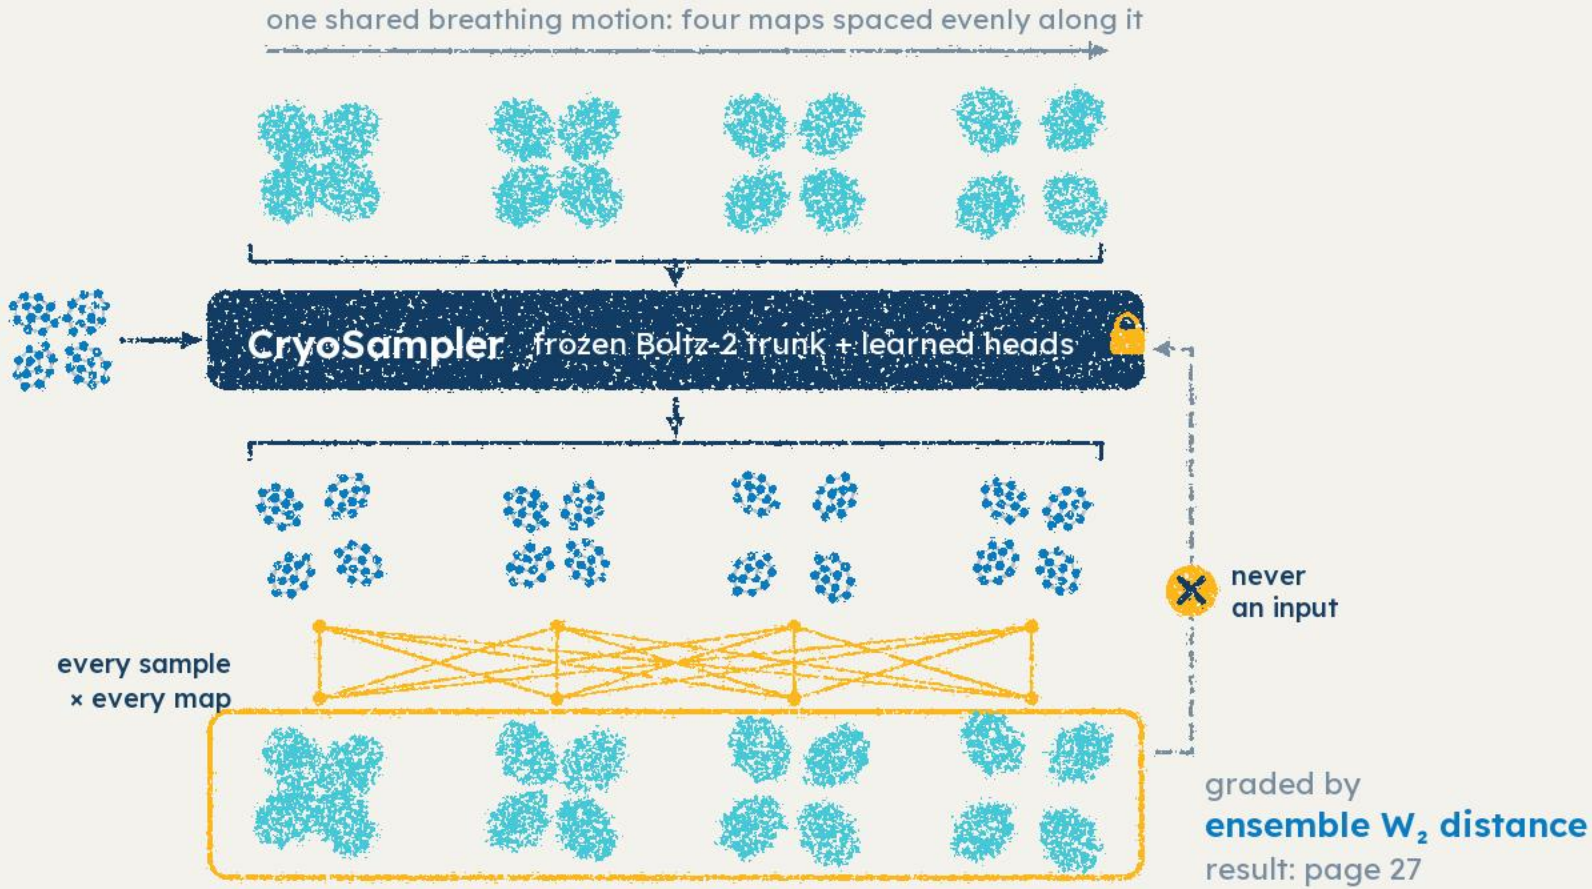

Schematic glyphs: channel top views; the pore widens along the motion.

CryoSampler: Shenoy, Astore, Levy, Poitevin, Hanson et al., arXiv 2605.09832v1, 2026 preprint · §4.2, App. A.4.2

Trained on TRPV3 maps, CryoSampler came closest to held-out TRPV5 maps.

Each method made 16 TRPV5 structures with no TRPV5 map as input.

Ensemble W_2 distance to the held-out TRPV5 maps

lower = closer · unitless, built from √(1 – CC_volume) per structure–map pair

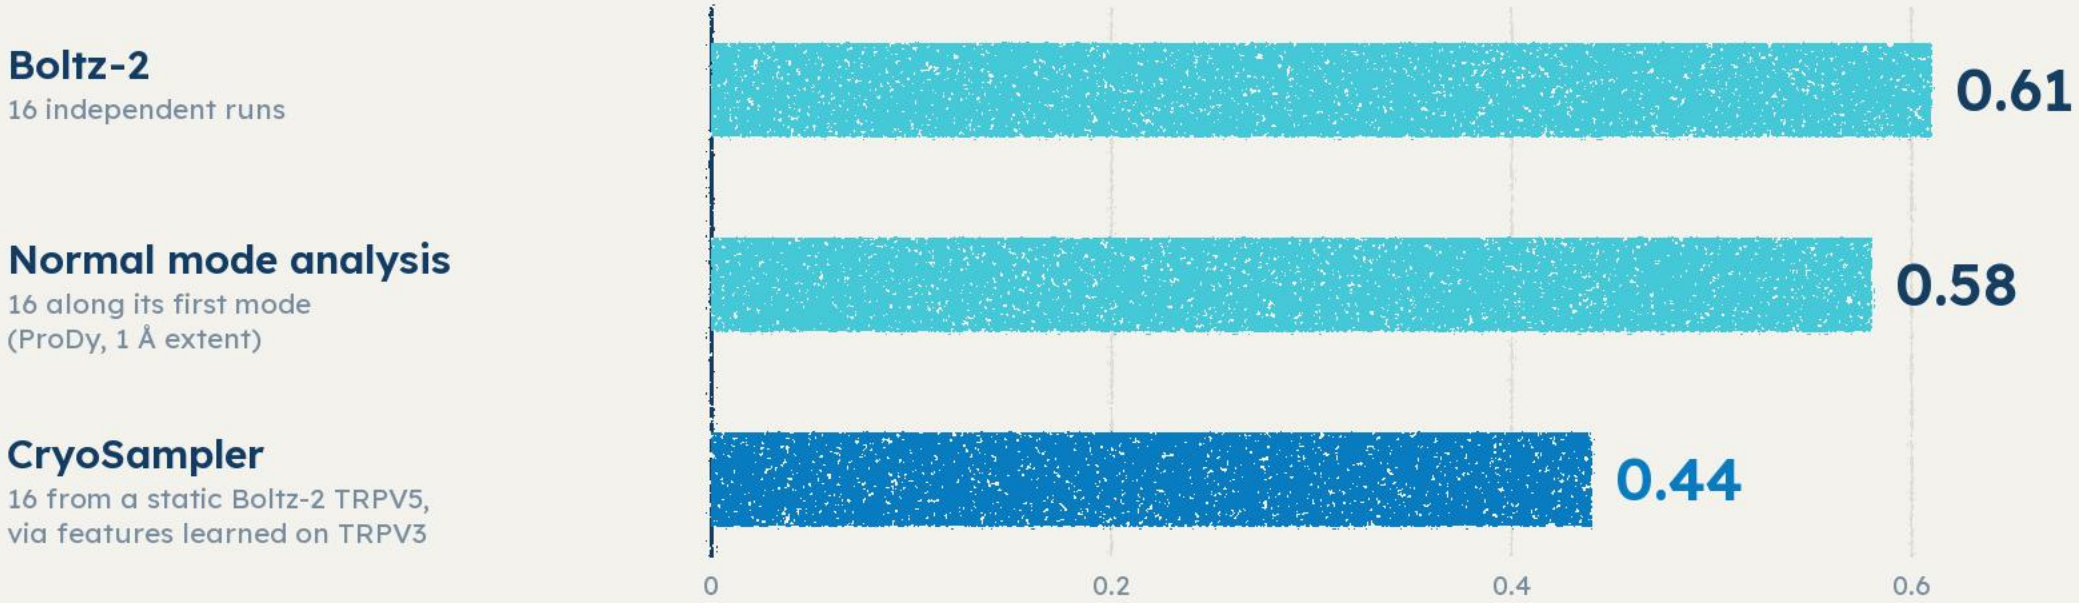

Transfer test: one related pair (TRPV3 → TRPV5), one shared motion. Preprint Table 1; no uncertainty reported.

CryoSampler: Shenoy, Astore, Levy, Poitevin, Hanson et al., arXiv 2605.09832v1, 2026 preprint · Table 1, PDF p. 8; §4.2; App. A.4.2

A held-out assay can still expose the same protein.

PRIMO, retrained under two training partitions and scored on one ubiquitin assay.

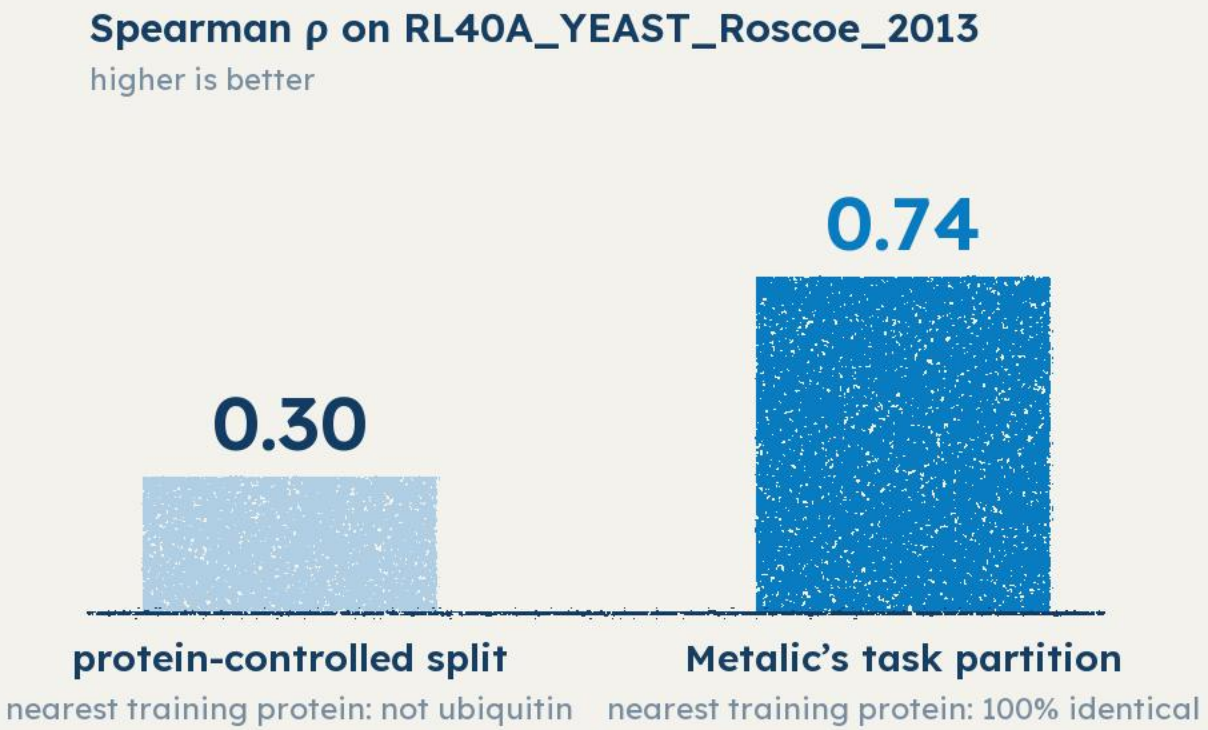

Both bars are PRIMO (0.74 is not a Metalic score); training composition also differs, so this is not a single-factor test. No CI reported.

Teufel et al. 2025, arXiv 2512.02315 (PRIMO, preprint): Table 1

Guiding-data scores show fit; held-out data test the structure.

score also steered the model score only grades

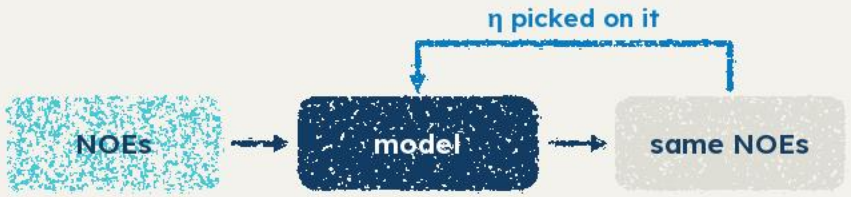

EG-AF3 · ubiquitin, one protein

26.5% → 15.7%

of 1,212 NOE restraint groups violated

Graded on the guiding data: the guidance took hold.

Page 20.

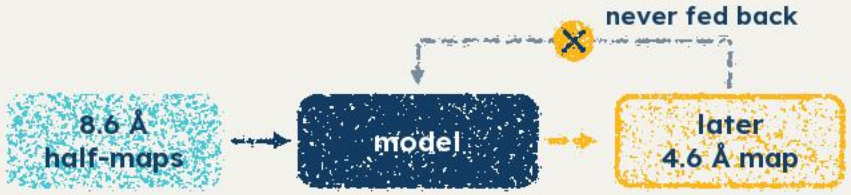

ROCKET · ZPD, one case

0.29 → 0.67

map correlation with a later map the fit never used

Graded on held-out data: the structure improved.

Whole pipeline; no CI. Page 03.

EG-AF3: Maddipatla et al., Nat. Biotechnol., accepted (Fig. 2c; Methods) · ROCKET: Fadini et al., Nat. Methods 2026 (Fig. 4b; Methods)

Measurements can move a protein model; independent checks show how far to trust the move.

1 Measurement fit

A closer fit can come from scoring on the guiding data or from broken geometry.

pages 20, 22

2 Integration choices

Data entered inputs, weights, sampling or pool weights; each route changes a different part.

pages 15–23

3 Validation and transfer

Transfer conclusions remain specific to the tested proteins, conditions and splits.

pages 26–28

Studies as cited on the pages named above; preprints labelled there.