Supplementary Information

Supplementary Table of Contents

1 Supplementary Figures and Tables

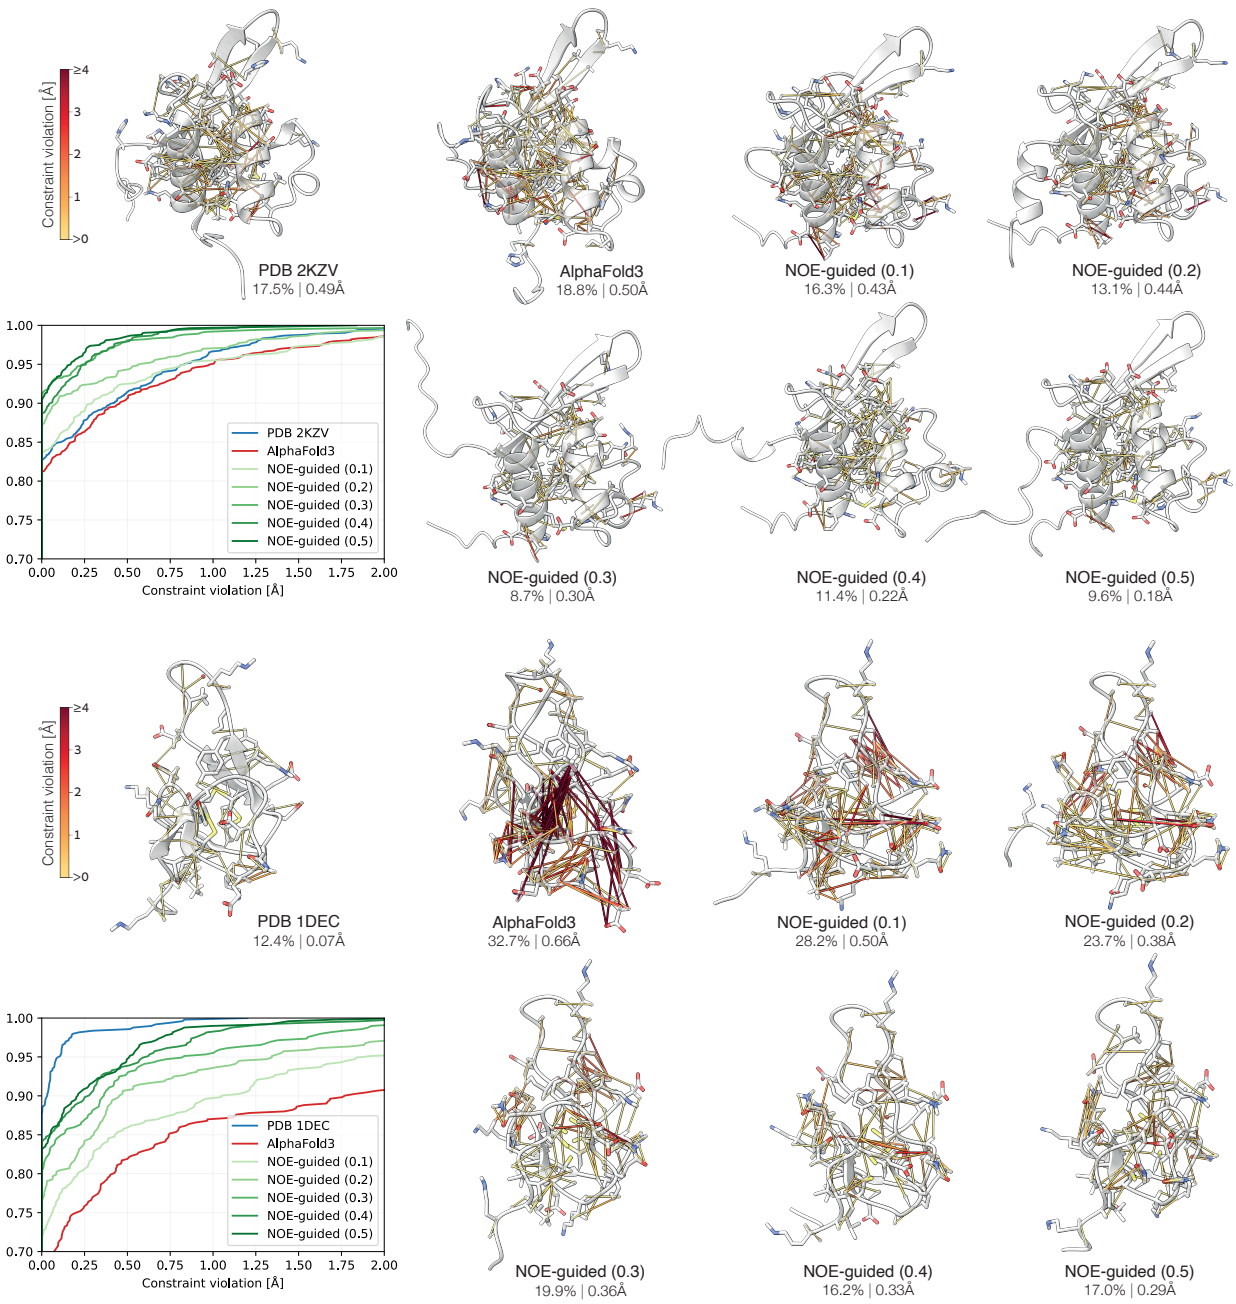

SI Figure 1: Guidance of `2KZV` (top) and `1DEC` (bottom) structure by NOE pairwise distance bounds with varying data term strength. Visualized are the pairwise distance constraint violations for the PDB ensemble, unguided AlphaFold3, and NOE-guided predictions with varying data term strength indicated in parenthesis. The fraction of the violated constraints and their median violation are reported below each structure. A single best-fitting structure from the ensemble is shown for clarity. Plots show cumulative distribution of constraint violations.

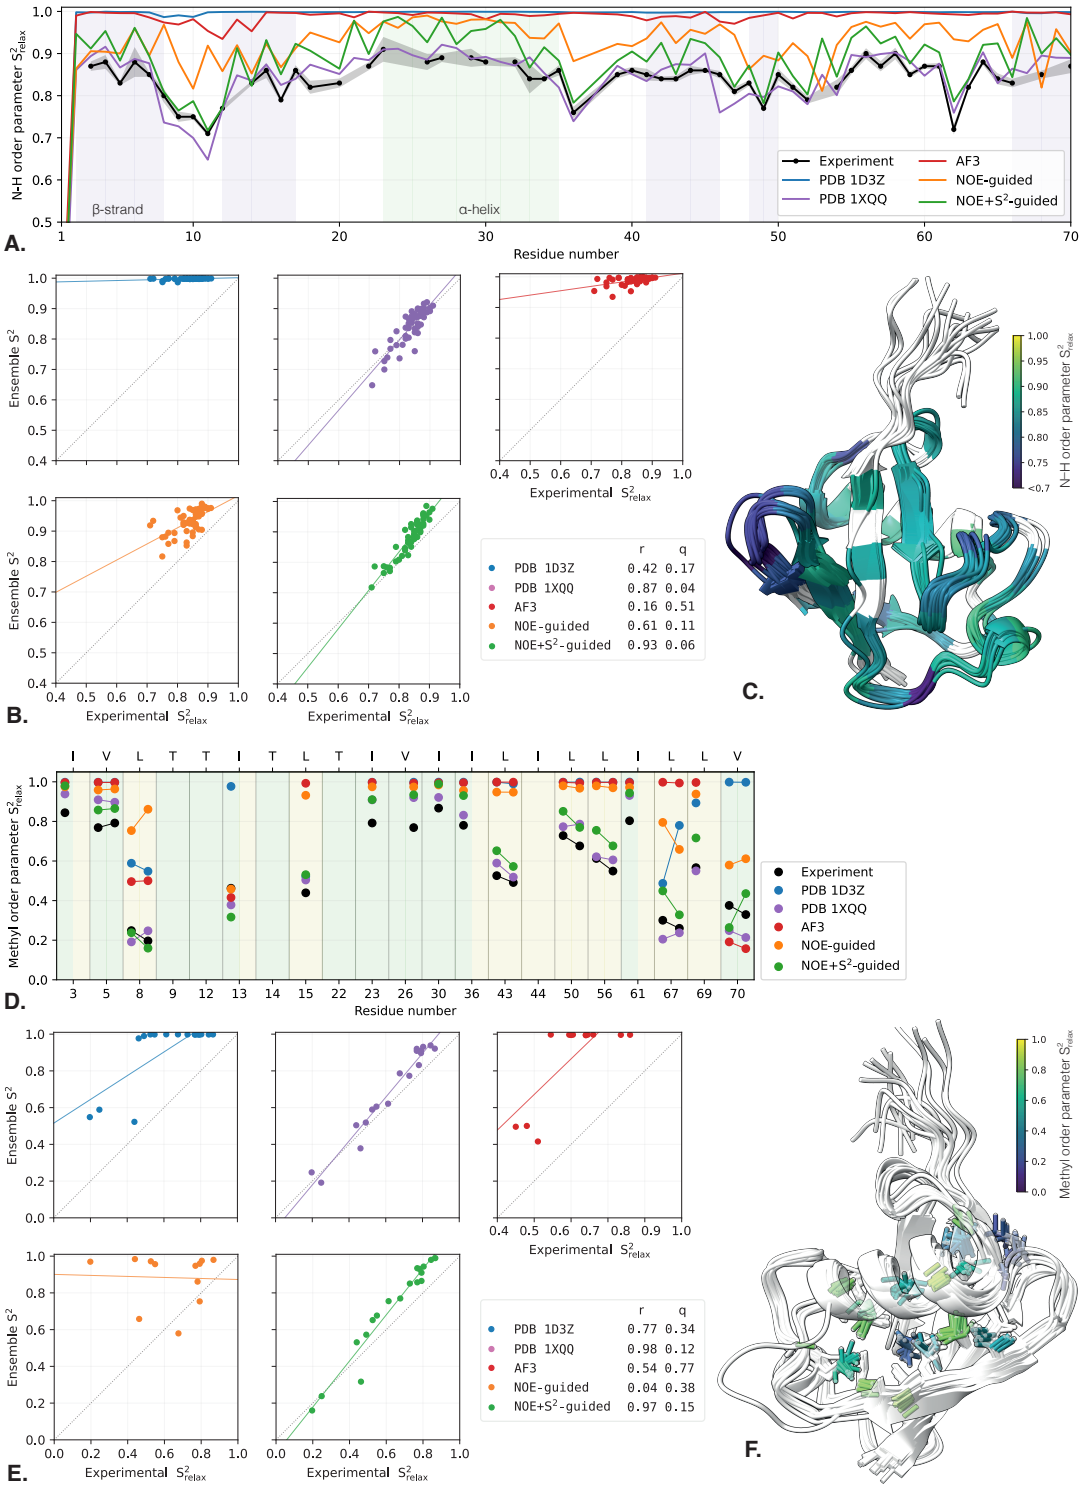

SI Figure 2: Guidance of ubiquitin structure with the combination of NOE pairwise distance bounds with amide (top) and methyl (bottom) order parameters obtained from relaxation data (ps-ns motion). (A) Experimentally measured amide order parameter S^2 from [22] (black) vs. order parameter calculated from the ensembles of different methods. (B) The same data visualized as a scatter plot with correlation (r, higher is better) and normalized fitting error (q, lower is better) factors. (C) Visualization of the disorder parameter on the NOE+S^2-guided ensemble. Lower values mean bigger motion. (D) Experimentally measured methyl order parameter S^2 from [21] (black) vs. order parameter calculated from the ensembles of different methods. C_γ and C_δ groups are indicated in green and yellow, respectively. For instances with two data points, these refer to the two methyl groups of Val (γ_1, γ_2), Leu (δ_1, δ_2) and Ile (γ_2, δ_1). (E) Same data visualized as scatter plots. (F) Visualization of the disorder parameter on the NOE+S^2-guided ensemble with relevant bonds colored by their S^2 values.

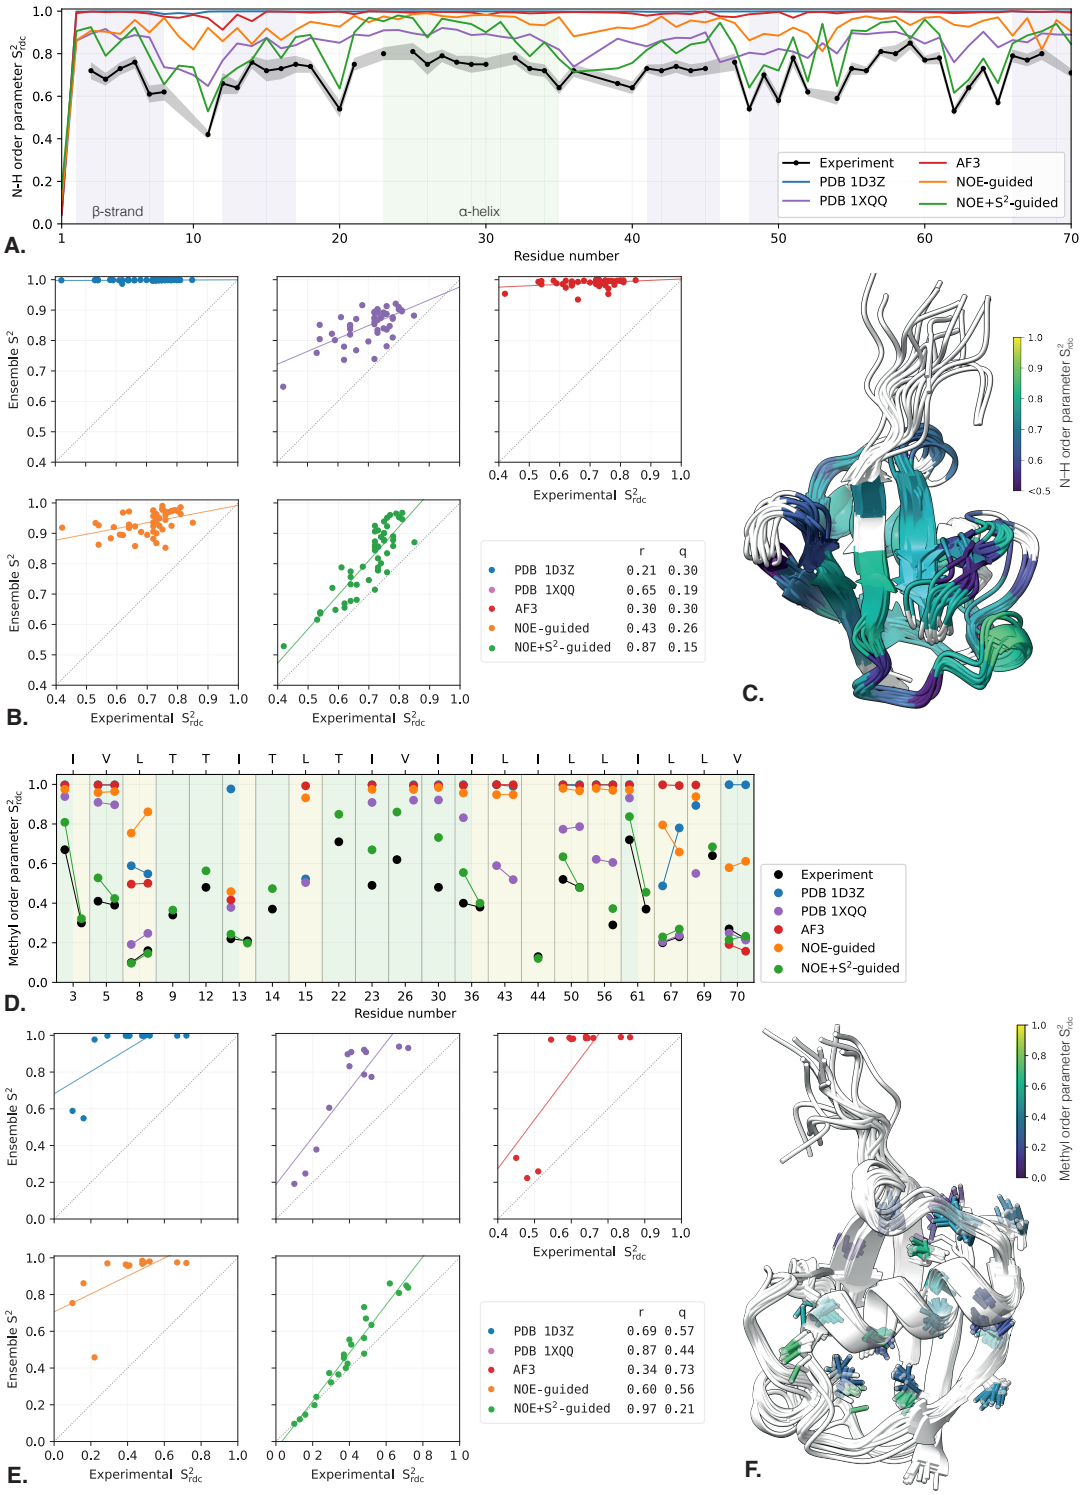

SI Figure 3: Guidance of ubiquitin structure with the combination of NOE pairwise distance bounds with amide (top) and methyl (bottom) order parameters obtained from residual dipolar coupling data (μs-ms motion). (A) Experimentally measured amide order parameter S^2 from [29] (black) vs. order parameter calculated from the ensembles of different methods. (B) The same data visualized as a scatter plot with correlation (r, higher is better) and normalized fitting error (q, lower is better) factors. (C) Visualization of the disorder parameter on the NOE+S^2-guided ensemble. Lower values mean bigger motion. (D) Experimentally measured methyl order parameter S^2 from [9] (black) vs. order parameter calculated from the ensembles of different methods. C_γ and C_δ groups are indicated in green and yellow, respectively. For instances with two data points, these refer to the two methyl groups of Val (γ_1, γ_2), Leu (δ_1, δ_2) and Ile (γ_2, δ_1). (E) Same data visualized as scatter plots. (F) Visualization of the disorder parameter on the NOE+S^2-guided ensemble with relevant bonds colored by their S^2 values.

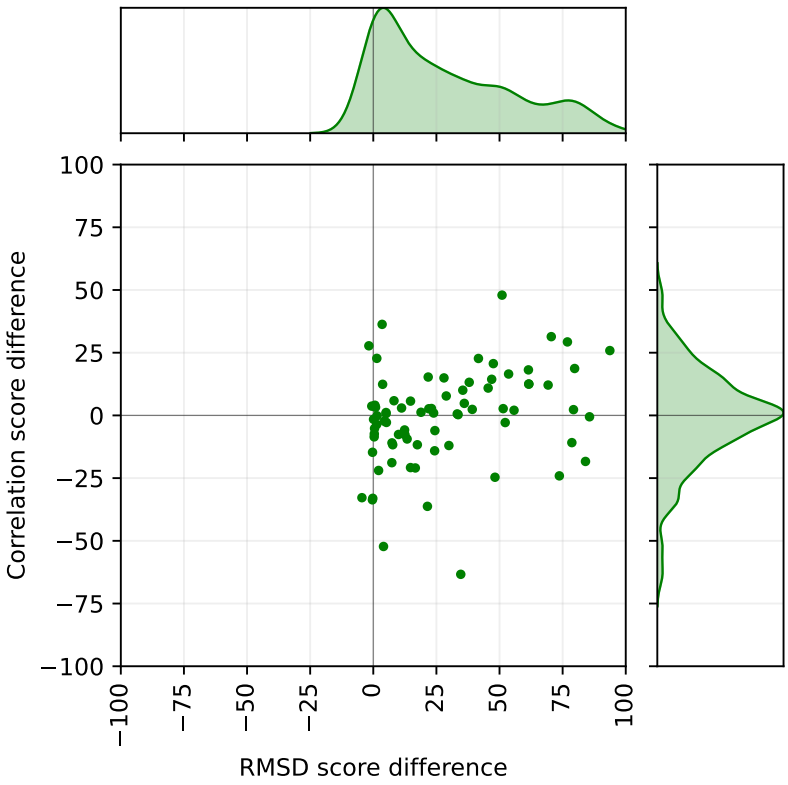

SI Figure 4: ANSURR [11] measures how well a structure's residue-by-residue rigidity profile agrees with that inferred from NMR chemical shifts: the correlation score captures similarity of the pattern (where flexible/rigid regions occur), while the RMSD score captures agreement in the magnitude of rigidity/flexibility across residues. Scores are normalized from 0 to 100 with higher numbers indicating better agreement. Depicted is the paired difference in the two ANSURR scores between the NOE-guided ensembles and the corresponding PDB structures. Each point corresponds to one of the evaluated 77 structures for which the median shift was above 80%. Marginal plots show kernel density estimates of each score individually. Our guidance method consistently improves the RMSD score while maintaining comparable correlation scores.

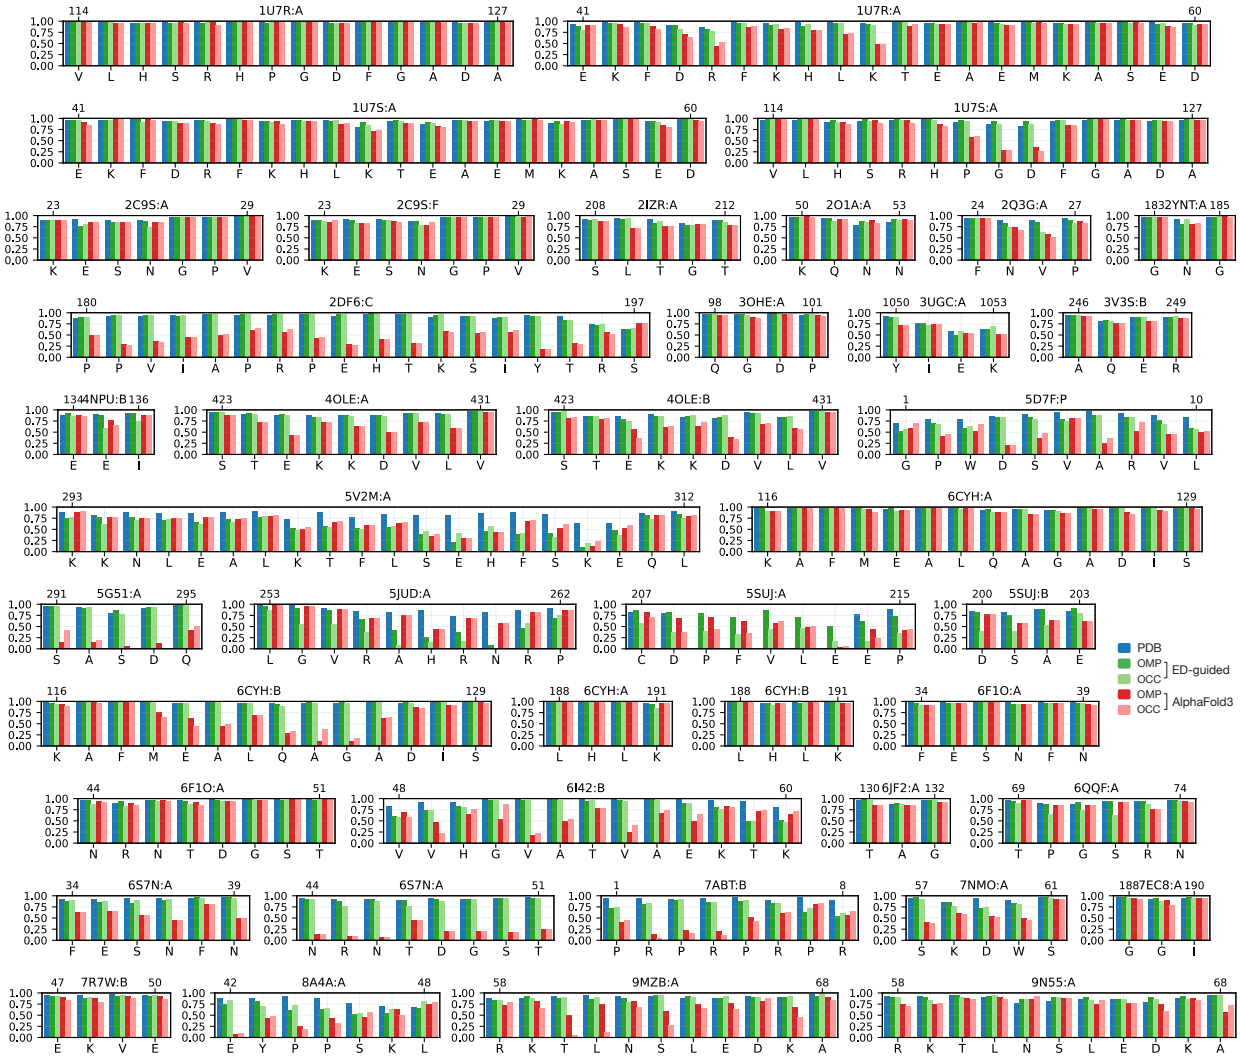

SI Figure 5: A summary of per-residue real-space correlation coefficients (RSCC) for the evaluated X-ray crystallographic structures reported in Table 2. Color-coded bars represent the RSCC of the deposited PDB structure (blue), electron density-guided AlphaFold3 with OMP (dark green) and occupancy optimization ((light green) ensemble pruning, and unguided AlphaFold3 with OMP (dark red) and occupancy optimization (light red) ensemble pruning.

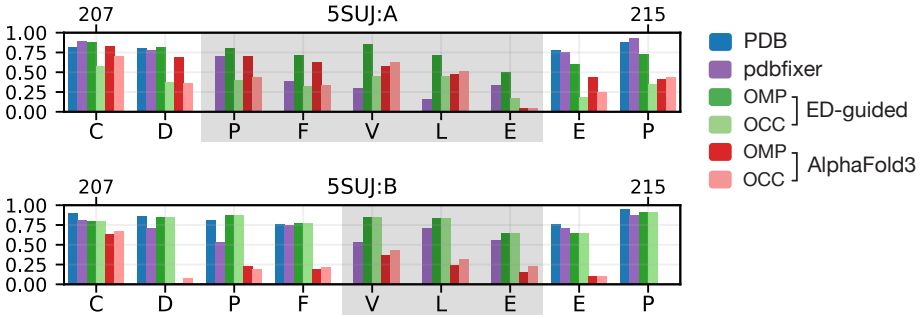

SI Figure 6: Per-residue real-space correlation coefficients (RSCC) for the `5SUJ` X-ray crystallographic structure with unmodelled density (highlighted in grey). Color-coded bars represent the RSCC of the deposited PDB structure (blue), completed structure with pdbfixer (purple), electron density-guided AlphaFold3 with OMP (dark green) and occupancy optimization ((light green) ensemble pruning, and unguided AlphaFold3 with OMP (dark red) and occupancy optimization (light red) ensemble pruning.

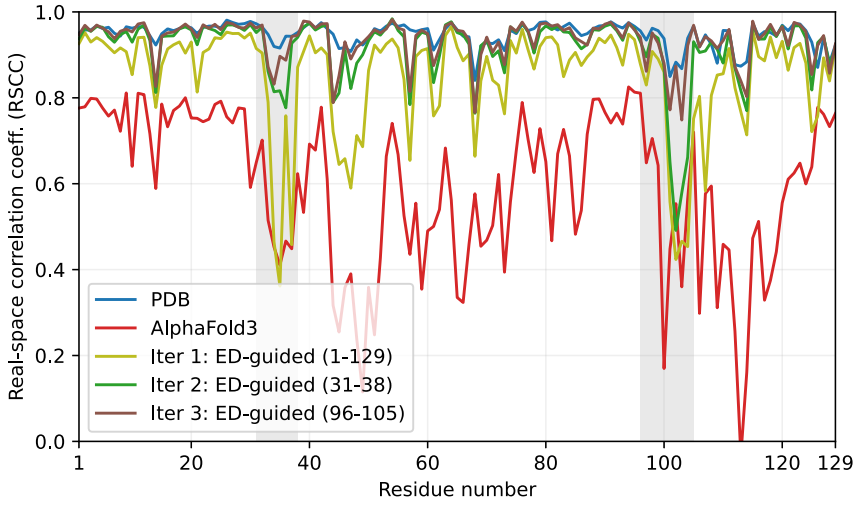

SI Figure 7: Per-residue real-space correlation coefficients (RSCC) for the `6S7N` X-ray crystallographic structure modeled with electron-density (ED) guided AlphaFold3 iterated with real-space refinement. Color-coded lines represent the RSCC of the deposited PDB structure (blue), unguided AlphaFold3 prediction (red), first iteration in which the entire structure was ED-guided (olive), followed by phenix real-space refinement and a second iteration of partial structure guidance for residues $1-38$ (green), followed by phenix refinement and a third iteration of partial structure guidance for residues $96-105$ (brown). Further iterations with partial guidance in worst-performing regions show further albeit diminishing improvement in RSCC.

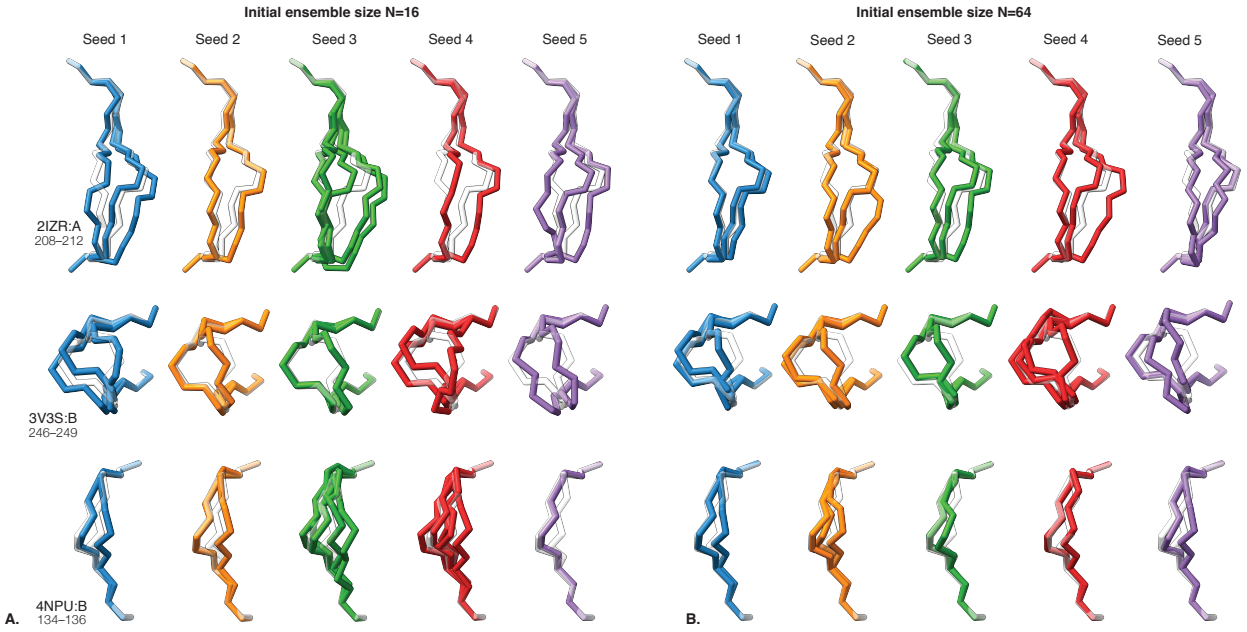

SI Figure 8: Repeatability of generated ensembles with five random seeds using initial ensemble sizes of 16 (A) and 64 (B) members. The generated backbones are overlaid on top of the corresponding PDB models depicted in white.

			AlphaFold3		NOE-guided		PDB	
pdb id	length	# constr.	viol. %	viol. Å	viol. %	viol. Å	viol. %	viol. Å
1D3Z	76	2727	26.5	0.34	**15.7**	**0.20**	24.9	0.39
1DEC	39	493	32.7	0.66	16.2	0.33	**12.4**	**0.07**
1PQX	91	1425	5.1	0.71	3.6	0.57	**1.2**	**0.05**
1RKL	36	218	62.4	0.30	50.9	**0.24**	**40.4**	0.26
1U0P	27	198	62.1	0.35	**54.5**	0.31	57.6	**0.13**
1YEZ	68	1177	13.8	0.41	**8.2**	**0.27**	11.9	0.46
2B3W	168	1214	39.8	1.29	**28.1**	0.65	35.7	**0.16**
2B5B	36	96	**24.0**	4.81	**24.0**	4.78	25.0	**0.35**
2HEQ	84	485	17.1	0.41	**11.8**	**0.28**	16.3	0.32
2HN8	38	719	54.0	0.34	43.7	0.29	**40.8**	**0.08**
2JPK	35	570	38.4	0.49	**25.4**	**0.30**	**25.4**	0.33
2JRM	60	1007	17.9	0.62	12.1	0.44	**10.9**	**0.43**
2JT1	77	1214	17.1	0.41	**11.4**	**0.20**	14.9	0.27
2JVD	48	1036	15.6	0.46	**7.0**	**0.18**	9.9	0.29
2JXP	155	2279	27.3	0.72	**16.6**	**0.41**	18.0	0.44
2K0M	104	1834	19.1	0.47	**11.9**	0.34	12.8	**0.29**
2K1G	135	2471	20.8	0.45	**13.9**	**0.31**	15.9	0.41
2K1S	149	1078	23.4	0.57	**15.7**	**0.29**	20.9	0.44
2K3A	155	1872	19.8	0.41	**11.2**	**0.22**	18.6	0.39
2K3D	87	993	24.6	0.41	**18.1**	**0.33**	23.2	0.50
2K50	115	2030	32.2	0.59	**18.3**	**0.30**	23.2	0.38
2K52	74	1107	31.6	0.83	**16.1**	**0.23**	23.9	0.48
2K53	70	932	45.1	0.48	36.8	**0.24**	**29.6**	0.35
2K57	55	982	17.3	0.46	**8.5**	**0.24**	11.4	0.47
2K5D	116	1658	19.4	0.55	**14.0**	**0.25**	16.4	0.35
2K5V	104	1694	15.8	0.40	**9.2**	**0.16**	14.9	0.31
2K75	106	1239	12.7	0.37	**8.6**	**0.16**	15.7	0.38
2KBN	109	1409	18.5	0.40	**11.4**	**0.22**	15.8	0.44
2KCD	120	1075	21.2	0.46	**16.8**	**0.25**	23.2	0.52
2KCT	94	1345	16.3	0.43	**10.0**	**0.23**	14.6	0.37
2KD0	85	1371	12.3	0.53	8.9	0.37	**8.4**	**0.31**
2KD1	118	2196	22.4	0.49	**14.0**	**0.31**	16.9	0.37
2KFP	125	1764	30.1	0.72	**19.1**	**0.25**	23.6	0.54
2KHD	108	784	16.7	0.47	**12.9**	**0.25**	17.9	0.46
2KIF	102	2359	21.3	0.42	**14.1**	**0.25**	16.7	0.37
2KIW	111	950	19.5	0.38	**12.0**	**0.21**	20.2	0.45
2KJR	95	1300	24.8	0.41	**14.3**	**0.17**	22.5	0.40
2KK8	84	1461	12.0	0.66	**5.1**	**0.30**	6.3	0.36
2KKL	140	1459	26.5	0.54	**19.4**	**0.30**	24.9	0.48
2KKZ	140	2307	23.2	0.53	**13.9**	**0.22**	16.1	0.35
2KL5	110	2125	31.3	0.72	**19.8**	**0.37**	22.4	0.50
2KL6	114	2643	27.8	0.52	**17.6**	**0.28**	21.8	0.39
2KNS	33	233	66.1	0.46	**54.5**	0.39	57.9	**0.33**
2KOB	108	1603	29.9	0.49	**20.4**	**0.26**	27.9	0.44
2KPB	26	267	58.8	1.01	43.4	0.92	**21.3**	**0.39**
2KPN	103	1806	22.6	0.41	**14.3**	**0.27**	20.0	0.31

			AlphaFold3		NOE-guided		PDB	
pdb id	length	# constr.	viol. %	viol. Å	viol. %	viol. Å	viol. %	viol. Å
2KRS	74	861	16.3	0.41	**10.0**	**0.19**	19.4	0.44
2KRT	121	2310	20.7	0.89	13.3	**0.40**	**11.9**	0.45
2KVO	120	1388	24.4	0.50	**16.0**	**0.24**	21.3	0.43
2KZV	92	835	18.8	0.50	**8.7**	**0.30**	17.5	0.49
2L06	155	1693	24.8	0.52	**20.3**	**0.35**	25.9	0.42
2L1P	83	1139	2.5	0.64	**1.9**	**0.34**	2.6	0.47
2L33	91	2365	24.9	0.49	**16.7**	**0.32**	17.0	0.39
2L3B	130	1242	23.8	0.56	**15.9**	**0.29**	23.5	0.47
2L3G	126	3359	26.4	0.52	**17.7**	**0.35**	20.5	0.42
2L82	162	4718	36.4	1.16	**15.6**	**0.29**	18.4	0.42
2L8V	143	1215	21.5	0.40	**15.0**	**0.22**	22.5	0.41
2L9R	69	1461	16.7	0.36	**9.9**	**0.28**	11.2	0.33
2LAH	160	4417	30.2	0.76	21.3	0.47	**20.2**	**0.46**
2LAK	160	1494	13.1	0.39	**8.9**	**0.19**	17.0	0.50
2LEA	135	2623	15.3	0.60	8.7	0.36	**0.5**	**0.26**
2LF2	175	1853	17.8	0.39	**12.2**	**0.18**	19.1	0.37
2LFI	122	1910	31.9	0.78	22.7	**0.37**	**20.4**	0.38
2LGH	144	1521	17.8	0.42	**10.1**	**0.18**	15.9	0.31
2LI3	30	295	28.5	4.01	8.1	**0.28**	**1.0**	0.28
2LK2	89	849	11.7	0.49	**8.7**	**0.15**	13.4	0.35
2LL8	101	1838	17.8	0.49	**13.1**	**0.38**	14.3	0.42
2LN3	83	2708	22.5	0.51	15.4	**0.30**	**15.1**	0.36
2LNA	99	1266	14.5	0.43	**8.5**	**0.22**	11.5	0.33
2LND	112	3137	40.0	1.95	**16.5**	**0.34**	18.0	0.40
2LRH	134	4193	23.7	0.64	16.0	**0.37**	**14.7**	0.42
2LTA	110	3115	23.0	0.79	13.3	**0.34**	**12.6**	0.37
2LTL	119	2428	15.0	0.44	**9.6**	**0.22**	11.9	0.35
2LTM	107	2694	18.9	0.45	**13.3**	**0.28**	15.1	0.36
2LVB	112	2853	21.4	0.55	**15.8**	**0.36**	16.6	0.45
2LX7	60	415	7.7	0.36	**5.8**	**0.24**	7.7	0.32
2LXU	108	1210	14.2	0.42	**9.2**	**0.22**	13.1	0.34
2M47	163	1549	16.8	0.51	**10.7**	**0.21**	13.9	0.42
2M4F	151	3310	42.6	0.63	**30.9**	**0.36**	33.3	0.53
2M5O	97	2928	24.4	0.47	**16.5**	**0.29**	17.8	0.41
2MA6	61	733	12.6	0.39	**8.6**	**0.19**	11.7	0.44
2MK2	109	1311	11.2	0.46	**7.6**	**0.20**	11.8	0.35
2MQL	105	2261	29.6	0.61	**17.9**	**0.40**	20.8	0.45
2MRW	40	159	37.1	0.14	25.8	0.34	**23.3**	**0.10**
2RN7	108	353	9.6	0.29	**7.4**	**0.23**	10.5	0.42
3BBG	40	396	39.6	1.35	23.2	0.25	**5.8**	**0.21**
5L82	37	485	57.9	0.88	51.1	0.66	**24.1**	**0.59**
6FIP	99	1615	37.1	0.66	**24.0**	**0.26**	25.4	0.42
6GT7	115	1997	47.6	0.72	**35.3**	**0.43**	36.1	0.52
6SOW	58	1120	44.6	0.64	32.8	**0.40**	**31.3**	0.48
6SVC	35	1147	79.7	0.63	**69.0**	**0.56**	83.7	0.66

SI Table 1: Summary of quantitative performance indicators of the evaluated NMR structures. Reported are the number of NOE constraints, percentage of constraint violation and median magnitude of the violated constraints for unguided and NOE-guided AlphaFold3 and the corresponding deposited PDB structures.

PDB id	Res [Å]	Chain	Location start	Location end	Sub-sequence	Ensemble size AF3 OMP	Ensemble size AF3 OCC	Ensemble size ED-guided OMP	Ensemble size ED-guided OCC	Cosine similarity (↑) AlphaFold3 OMP	Cosine similarity (↑) AlphaFold3 OCC	Cosine similarity (↑) ED-guided OMP	Cosine similarity (↑) ED-guided OCC	Cosine similarity (↑) PDB	R_{work} (↓) AlphaFold3 OMP	R_{work} (↓) AlphaFold3 OCC	R_{work} (↓) ED-guided OMP	R_{work} (↓) ED-guided OCC	R_{work} (↓) PDB	R_{free} (↓) AlphaFold3 OMP	R_{free} (↓) AlphaFold3 OCC	R_{free} (↓) ED-guided OMP	R_{free} (↓) ED-guided OCC	R_{free} (↓) PDB
Unimodal backbone																								
1U7R	1.15	A	114	127	VLHSRHPGDFGADA	2	1	3	1	0.941	0.933	**0.948**	0.941	0.943	0.172	0.174	**0.170**	0.172	**0.170**	0.192	0.195	**0.191**	0.191	**0.191**
			41	60	EKFDRFKHLKTEAEMKASED	2	1	2	1	0.840	0.820	**0.896**	0.867	0.889	0.183	0.185	0.172	0.173	**0.170**	0.206	0.205	0.192	0.194	**0.191**
6S7N	1.20	A	34	39	FESNFN	1	1	5	1	0.581	0.581	0.888	0.887	**0.894**	0.229	0.229	0.193	0.195	**0.191**	0.250	0.250	0.215	0.218	**0.208**
			44	51	NRNTDGST	1	1	3	1	0.376	0.376	**0.911**	0.867	0.896	0.242	0.242	0.192	0.197	**0.191**	0.266	0.266	0.208	0.218	**0.208**
1U7S	1.40	A	114	127	VLHSRHPGDFGADA	2	1	5	1	0.828	0.813	**0.930**	0.891	0.901	0.199	0.203	**0.189**	0.191	0.190	0.215	0.217	**0.205**	0.210	0.208
			41	60	EKFDRFKHLKTEAEMKASED	2	1	3	1	0.874	0.842	**0.931**	0.894	0.906	0.205	0.212	**0.188**	0.194	0.190	0.223	0.230	0.208	0.213	**0.208**
6CYH	1.49	A	116	129	KAFMEALQAGADIS	2	1	4	1	0.894	0.866	**0.945**	0.906	0.924	0.157	0.159	**0.154**	0.156	0.155	0.171	0.173	**0.170**	0.172	0.171
		B				2	1	2	1	0.782	0.761	**0.933**	0.918	0.908	0.164	0.165	**0.154**	0.155	0.155	0.182	0.181	**0.170**	0.171	0.171
9N55	1.65	A	58	68	RKTLNSLEDKA	4	2	3	4	0.267	0.277	0.856	**0.870**	0.837	0.259	0.258	0.210	0.212	**0.207**	0.289	0.286	**0.239**	0.245	0.240
9MZB	2.80	A				2	1	2	1	0.460	0.416	0.902	0.880	**0.911**	0.328	0.327	0.272	0.268	**0.265**	0.396	0.401	0.355	**0.351**	0.352
Multimodal backbone (altlocs)																								
2Q3G	1.11	A	24	27	FNVP	2	1	3	1	0.840	0.783	0.889	0.797	**0.894**	0.165	0.167	**0.163**	0.166	0.164	0.182	0.183	**0.180**	0.183	0.182
7R7W	1.17	B	47	50	EKVE	2	1	3	3	0.885	0.794	0.908	**0.908**	0.899	0.182	0.183	0.182	0.182	**0.182**	0.196	0.197	0.197	0.197	**0.196**
3OHE	1.20	A	98	101	QGDP	3	1	4	2	0.885	0.856	**0.901**	0.893	0.894	0.151	0.152	**0.150**	0.150	0.150	0.169	0.170	**0.168**	0.168	0.168
2IZR	1.30	A	208	212	SLTGT	1	1	4	2	0.795	0.795	**0.908**	0.891	0.902	0.171	0.171	0.170	0.170	**0.170**	0.191	0.191	0.191	**0.190**	0.190
3UGC	1.34	A	1050	1053	YIEK	1	1	3	2	0.658	0.658	**0.808**	0.801	0.790	0.166	0.166	0.164	**0.163**	0.163	0.183	0.183	0.180	**0.179**	0.179
7EC8	1.35	A	188	190	GGI	2	1	2	1	0.937	0.906	**0.951**	0.916	0.941	0.165	0.166	**0.165**	0.166	**0.165**	0.193	0.193	**0.193**	0.193	**0.193**
5G51	1.45	A	291	295	SASDQ	4	1	4	2	0.598	0.530	**0.928**	0.910	0.914	0.205	0.204	0.187	0.187	**0.187**	0.223	0.223	**0.206**	0.208	0.208
4NPU	1.50	B	134	136	EEI	2	1	3	3	0.837	0.792	**0.881**	0.804	0.872	0.190	0.190	0.189	0.191	**0.189**	**0.212**	0.212	0.212	0.214	0.212
2O1A	1.60	A	50	53	KQNN	2	1	3	1	0.879	0.836	**0.908**	0.845	0.844	**0.173**	0.174	0.174	0.174	0.174	0.205	0.210	**0.205**	0.211	0.207
2YNT	1.60	A	183	185	GNG	2	1	1	2	0.822	0.807	0.820	0.874	**0.889**	0.151	0.151	0.151	0.151	**0.151**	0.191	0.191	0.191	0.191	**0.191**
3AZY	1.65	A	157	163	GGASIGV	1	1	2	1	0.717	0.717	0.839	0.683	**0.848**	0.174	0.174	0.173	0.174	**0.173**	0.195	0.195	0.194	0.195	**0.194**
3V3S	1.90	B	246	249	AQER	1	1	5	3	0.645	0.645	**0.779**	0.761	0.759	0.163	0.163	0.164	0.163	**0.163**	0.202	0.202	0.203	**0.202**	0.202
6QQF	1.95	A	69	74	TPGSRN	1	1	2	3	0.807	0.809	0.861	0.840	**0.863**	0.163	0.162	0.162	0.169	**0.159**	0.211	0.210	0.208	0.215	**0.208**
5V2M	1.99	A	293	312	KKNLEALKTFLSEHFSKEQL	2	1	4	3	0.792	0.786	**0.824**	0.795	0.779	0.196	0.196	0.202	0.202	**0.186**	0.249	0.249	0.251	0.249	**0.240**
6JF2	2.00	A	130	132	TAG	1	1	5	2	0.862	0.865	**0.940**	0.939	0.940	0.189	0.189	0.187	0.187	**0.187**	0.218	0.218	**0.217**	0.217	**0.217**
5SUJ	2.36	B	200	203	DSAE	1	1	2	2	0.694	0.694	**0.832**	0.764	0.813	0.202	0.202	0.202	0.203	**0.202**	0.243	0.243	0.242	0.245	**0.241**
		B				2	1	2	2	0.718	0.704	0.890	**0.901**	0.884	0.178	0.176	0.176	0.175	**0.175**	0.205	**0.204**	0.205	0.205	0.205
4OLE	2.52	C	423	431	STEKKDVLV	2	1	3	1	0.769	0.722	0.861	0.766	**0.873**	0.183	0.181	0.184	0.180	**0.175**	0.216	0.216	0.214	0.216	**0.205**
		D				2	1	2	2	0.809	0.754	0.851	0.855	**0.888**	0.182	0.181	0.181	0.181	**0.175**	0.215	0.213	0.212	0.214	**0.205**
8A4A	2.52	A	42	48	EYPPSKL	3	1	1	3	0.745	0.670	0.764	**0.885**	0.878	0.209	0.205	0.201	0.205	**0.201**	0.260	0.257	**0.251**	0.263	0.260
5JUD	2.59	A	253	262	LGVRAHRNRP	1	1	2	2	0.746	0.746	0.815	0.778	**0.847**	0.189	0.189	0.195	0.199	**0.188**	**0.214**	**0.214**	0.220	0.224	0.216
Peptides																								
2DF6	1.30	C	180	197	PPVIAPRPEHTKSIYTRS	1	2	2	5	0.388	0.485	0.838	0.826	**0.854**	0.211	0.208	0.177	0.186	**0.174**	0.231	0.234	0.194	0.204	**0.192**
2DF6	1.30	D	180	191	PPVIAPRPEHTK	3	3	2	3	0.419	0.419	0.861	**0.877**	0.853	0.254	0.253	0.177	0.177	**0.174**	0.270	0.270	0.195	0.195	**0.192**
5D7F	1.30	P	1	10	GPWDSVARVL	3	1	4	5	0.435	0.408	0.724	0.734	**0.757**	0.197	0.199	0.188	0.189	**0.181**	0.198	**0.181**	0.204	0.206	0.194
7ABT	1.31	B	1	8	PRPRPRPR	2	1	3	3	0.515	0.487	0.805	0.803	**0.884**	0.208	0.210	0.184	0.183	**0.169**	0.228	0.230	0.207	0.208	**0.193**
6I42	1.38	B	48	60	VVHGVATVAEKTK	2	1	4	4	0.487	0.495	0.849	0.850	**0.908**	0.213	0.216	0.178	0.179	**0.162**	0.231	0.237	0.195	0.199	**0.178**
3AZY	1.70	A	157	163	GGASIGV	2	5	1	2	0.852	0.845	0.841	0.849	**0.853**	**0.173**	**0.173**	**0.173**	0.174	**0.173**	0.194	0.194	**0.194**	0.194	0.194
1DDV	1.90	B	1001	1006	TPPSPF	1	1	2	2	0.561	0.652	0.867	0.873	**0.945**	0.263	0.258	0.233	0.231	**0.210**	0.289	0.289	0.274	0.273	**0.257**
1CKB	1.90	B	1	8	PPPVPPRR	1	1	3	4	0.809	0.817	**0.882**	0.870	0.868	0.196	0.193	0.186	0.192	**0.177**	0.258	0.249	0.243	0.246	**0.235**
1CJR	2.30	A	1	15	KETAAAKFERQHMDS	1	1	4	2	0.672	0.690	0.893	0.892	**0.929**	0.279	0.271	0.227	0.218	**0.201**	0.346	0.353	0.290	0.287	**0.263**
Unmodeled altlocs																								
5NVJ	1.18	A	35	39	MIDDG	1	1	1	2	0.817	0.817	0.848	**0.867**	0.851	0.213	0.213	0.210	**0.209**	0.210	0.236	0.236	0.234	0.234	**0.233**
4OLE	2.52	A	423	431	STEKKDVLV	1	1	1	2	0.724	0.724	**0.916**	0.911	0.914	0.178	0.178	**0.175**	0.177	**0.175**	0.207	0.207	**0.205**	0.207	0.205
Unmodeled region																								
5SUJ	2.36	A	207	215	CDPFVLEEP	1	2	1	4	0.620	0.672	0.673	**0.804**	0.625	0.202	0.205	0.201	0.207	**0.200**	0.242	0.244	0.241	0.248	**0.240**
		B				2	2	1	1	0.448	0.456	**0.691**	**0.691**	0.634	0.204	0.204	**0.201**	**0.201**	0.202	0.243	0.242	**0.238**	**0.238**	0.241

SI Table 2: Summary of quantitative performance indicators of the evaluated X-ray crystallographic structures. Reported are the structure resolution, the location of the generated region, cosine similarity as the real-space metric and the R-factors (R_{work} and R_{free}) as the Fourier-space metrics for unguided and electron density-guided AlphaFold3 and the corresponding deposited PDB structures. Ensemble sizes produced by OMP and occupancy optimization (OCC) pruning are reported in the latter cases.

PDB id	Res [Å]	Chain	Location start	Location end	Sub-sequence	Ensemble size AF3 OMP	Ensemble size AF3 OCC	Ensemble size ED-guided OMP	Ensemble size ED-guided OCC	Cosine similarity (↑) AlphaFold3 OMP	Cosine similarity (↑) AlphaFold3 OCC	Cosine similarity (↑) ED-guided OMP	Cosine similarity (↑) ED-guided OCC	Cosine similarity (↑) PDB	R_{work} (↓) AlphaFold3 OMP	R_{work} (↓) AlphaFold3 OCC	R_{work} (↓) ED-guided OMP	R_{work} (↓) ED-guided OCC	R_{work} (↓) PDB	R_{free} (↓) AlphaFold3 OMP	R_{free} (↓) AlphaFold3 OCC	R_{free} (↓) ED-guided OMP	R_{free} (↓) ED-guided OCC	R_{free} (↓) PDB
1U7R	1.15	A	114	127	VLHSRHPGDFGADA	2	1	3	1	0.941	0.933	**0.948**	0.941	0.943	0.185	0.188	0.172	0.175	**0.170**	0.205	0.210	0.192	0.194	**0.191**
			41	60	EKFDRFKHLKTEAEMKASED	2	1	2	1	0.840	0.820	**0.896**	0.867	0.889										
6S7N	1.20	A	34	39	FESNFN	1	1	5	1	0.581	0.581	0.888	0.887	**0.894**	0.273	0.273	0.193	0.201	**0.191**	0.287	0.287	0.215	0.224	**0.208**
			44	51	NRNTDGST	1	1	3	1	0.376	0.376	**0.911**	0.867	0.896										
2DF6	1.30	C	180	197	PPVIAPRPEHTKSIYTR	1	2	2	5	0.388	0.485	0.838	0.826	**0.854**	0.281	0.277	0.180	0.189	**0.174**	0.301	0.301	0.195	0.204	**0.192**
		D	189	191	PPVIAPRPEHTK	3	3	2	3	0.419	0.419	0.861	**0.877**	0.853										
1U7S	1.40	A	114	127	VLHSRHPGDFGADA	2	1	5	1	0.828	0.813	**0.930**	0.891	0.901	0.214	0.221	**0.187**	0.195	0.190	0.228	0.237	**0.203**	0.214	0.208
			41	60	EKFDRFKHLKTEAEMKASED	2	1	3	1	0.874	0.842	**0.931**	0.894	0.906										
6CYH	1.49	A	116	129	KAFMEALQAGADIS	2	1	4	1	0.894	0.866	**0.945**	0.906	0.924	0.167	0.169	0.155	0.158	**0.155**	0.183	0.184	0.171	0.174	**0.171**
		B				2	1	2	1	0.782	0.761	**0.933**	0.918	0.908										
5SUJ	2.36	A	207	215	CDPFVLEEP	1	2	1	4	0.620	0.672	0.673	**0.804**	0.625	0.205	0.208	**0.200**	0.207	0.202	0.246	0.248	**0.239**	0.244	0.241
		B	200	203	DSAE	1	1	2	2	0.694	0.694	**0.832**	0.764	0.813										
		B	207	215	CDPFVLEEP	2	2	1	1	0.448	0.456	**0.691**	**0.691**	0.635										
4OLE	2.52	A	423	431	STEKKDVLV	1	1	1	2	0.724	0.724	**0.916**	0.911	0.914	0.182	0.178	0.179	0.179	**0.175**	0.210	0.207	**0.204**	0.206	0.205
		B				2	1	2	2	0.718	0.704	0.890	**0.901**	0.884										
		C				1	1	3	2	0.772	0.724	0.864	0.853	**0.876**										
		D				1	1	2	1	0.760	0.760	0.875	0.841	**0.888**										

SI Table 3: A subset of structures from Table 2 in which several locations in the same structure were merged and jointly refined.

PDB id	Res [Å]	Chain	Location start	Location end	Sub-sequence	Map	Cosine similarity (↑) OMP	Cosine similarity (↑) OCC	R_{work} (↓) OMP	R_{work} (↓) OCC	R_{free} (↓) OMP	R_{free} (↓) OCC
2IZR	1.30	A	208	212	SLTGT	2Fo–Fc	0.892 ± 0.049	0.805 ± 0.022	0.1725 ± 0.0004	0.1733 ± 0.0004	0.1893 ± 0.0005	0.1894 ± 0.0006
						END	**0.893 ± 0.008**	0.806 ± 0.043	0.1727 ± 0.0002	0.1732 ± 0.0005	0.1900 ± 0.0005	0.1899 ± 0.0006
						Omit	0.837 ± 0.004	0.831 ± 0.002	**0.1720 ± 0.0001**	0.1720 ± 0.0001	**0.1891 ± 0.0003**	0.1888 ± 0.0002
4NPU	1.50	B	134	136	EEI	2Fo–Fc	0.887 ± 0.015	**0.897 ± 0.005**	0.1892 ± 0.0003	**0.1888 ± 0.0003**	0.2122 ± 0.0007	0.2109 ± 0.0008
						END	0.860 ± 0.036	0.866 ± 0.007	0.1894 ± 0.0004	0.1893 ± 0.0006	0.2115 ± 0.0006	0.2125 ± 0.0010
						Omit	0.834 ± 0.003	0.821 ± 0.013	0.1889 ± 0.0003	0.1888 ± 0.0003	**0.2114 ± 0.0002**	0.2122 ± 0.0006
3V3S	1.90	B	246	249	AQER	2Fo–Fc	0.806 ± 0.025	0.816 ± 0.022	0.1628 ± 0.0002	0.1626 ± 0.0002	0.2038 ± 0.0006	0.2039 ± 0.0003
						END	0.817 ± 0.015	**0.827 ± 0.019**	0.1626 ± 0.0002	0.1627 ± 0.0004	0.2041 ± 0.0004	0.2046 ± 0.0004
						Omit	0.782 ± 0.027	0.789 ± 0.009	0.1623 ± 0.0003	**0.1623 ± 0.0001**	0.2035 ± 0.0006	**0.2034 ± 0.0004**
5SUJ	2.36	B	200	203	DSAE	2Fo–Fc	0.727 ± 0.041	0.698 ± 0.041	0.2005 ± 0.0001	0.2003 ± 0.0003	0.2403 ± 0.0006	0.2401 ± 0.0004
						END	0.735 ± 0.037	0.737 ± 0.055	0.2004 ± 0.0002	0.2005 ± 0.0002	**0.2400 ± 0.0007**	0.2400 ± 0.0008
						Omit	0.780 ± 0.009	**0.790 ± 0.010**	**0.2003 ± 0.0002**	0.2005 ± 0.0002	0.2406 ± 0.0009	0.2402 ± 0.0009
4OLE	2.52	B	423	431	STEKKDVLV	2Fo–Fc	0.832 ± 0.017	0.824 ± 0.017	0.1787 ± 0.0011	0.1769 ± 0.0007	0.2082 ± 0.0010	0.2056 ± 0.0008
						END	**0.893 ± 0.011**	0.871 ± 0.033	0.1774 ± 0.0008	0.1756 ± 0.0008	0.2059 ± 0.0016	0.2039 ± 0.0015
						Omit	0.799 ± 0.004	0.801 ± 0.010	0.1756 ± 0.0005	**0.1756 ± 0.0003**	**0.2039 ± 0.0006**	0.2042 ± 0.0010

SI Table 4: ED-guided AlphaFold3 reconstruction on a subset of structures from Table 2. Shown are average metrics ± standard deviation over 5 independent runs with different random seeds using three types of inputs: $2F_o - F_c$, END, and omit maps. Note that cosine similarities, being real-space metrics, are not directly comparable across different map types.

	ESP cross-correlation coeffs				% violated NOE constraints			Median NOE constraint violation [Å]			Dihedral MAE [°]	
Method	**Mask**	**Volume**	**Backbone**	**Sidechains**	**Intra-chain**	**Inter-chain**	**Total**	**Intra-chain**	**Inter-chain**	**Total**	**ϕ**	**ψ**
7DA4 (EM)	0.667	0.576	0.508	0.591	30	21.25	26.5	0.87	0.67	0.67	51.3	84.3
7DAC (NMR)					30.8	21.3	27.0	0.86	0.09	0.54		
AlphaFold3	0.575	0.537	0.424	0.511	35.0	21.3	29.5	1.06	0.14	0.53	24.5	25.8
Guided												
ESP	0.710	0.656	0.532	0.645	32.5	22.5	28.5	0.91	4.97	0.92	27.0	24.4
NOE	0.488	0.449	0.354	0.432	17.5	11.3	15.0	0.47	0.08	0.29	32.2	30.5
ESP+NOE (0.3)	0.630	0.571	0.470	0.558	20.8	12.5	17.5	0.45	0.06	0.23	31.3	29.3
ESP+NOE (0.5)	0.567	0.528	0.413	0.492	13.3	10.0	12.0	0.46	0.08	0.24	30.4	33.9
ESP+dihedrals	0.681	0.619	0.514	0.616	35.0	25.0	31.0	1.26	0.14	0.70	11.6	17.5
ESP+NOE+dihedrals	0.598	0.554	0.433	0.523	20.8	20.0	20.5	0.74	0.07	0.38	15.1	14.4

SI Table 5: Guidance of 7DA4 with combinations of ESP, NOE and dihedral angles. Reported are cross-correlations with the ESP image, NOE constraint violations (within each chain and between the chains), and mean absolute disagreement with the chemical shift-derived backbone dihedral angles.

2 Supplementary Methods

2.1 Models, resources, and packages

For all experiments, we used Protenix [32] – an open-sourced PyTorch re-implementation of AlphaFold3. For the AlphaFold3 baseline mentioned across experiments, we used the official AlphaFold3 weights and source code [2].

To implement the corresponding forward models we used PyTorch (2.3.1+cu121) [25]. However, to efficiently calculate the density for larger structures, we used PyKeOps [6] and PyTorch.% All computations were performed on NVIDIA H100 and L40 GPUs running on Debian GNU/Linux 12 distribution.

Across all experiments, we retrieve the multiple sequence alignments (MSAs) using a wrapper around the ColabFold MMseqs2 API [24]. The wrapper submits the query sequences to the remote MMseqs2 server via HTTP POST, polls for job completion, and downloads and extracts the results. The output is provided in A3M format, which is compatible with AlphaFold3.

2.2 Input data

2.2.1 X-ray crystallography

Atomic models In the X-ray crystallography workflow, the input consists of a PDB ID, a chain identifier, and an amino acid subsequence in the chain. For this modality, we restrict our modelling to single protein chains. The corresponding PDB structure and MTZ file is retrieved from PDB-Redo [15] and the specific chain is identified using Gemmi (v0.6.5) [37]. From this chain, we extract only the standard amino acid residues explicitly modeled in the PDB file, while excluding structural waters, hydrogen atoms, and all other non-standard amino acids, and save the resulting coordinates as a separate PDB file. In addition, we mutate Selenium Methionine (`MSE`) to Methionine (`MET`) by replacing Selenium (`SE`) atoms with Sulfur (`SD`). We also mutate S-hydroxycysteine (`CSO`) to Cysteine (`CYS`) by removing the Oxygen from the hydroxyl group (`OG`). If alternate conformations (altlocs) are present [28], the PDB is split into the corresponding altloc files. Since AlphaFold3 uses one-based residue indexing, we renumber the residues in each split PDB to start from 1. In many cases, the retrieved PDB files can contain incomplete or unmodelled residues due to poor quality density. To address this, we use PDBFixer [7] (v1.9.0), which employs standard residue templates from OpenMM [8] force fields, to model the missing atoms in all split PDBs. To avoid steric clashes after imputing the atoms with PDBFixer, we relax the structure using AMBER-14 force field [35] relaxation. The isotropic B-factor for these imputed atoms is set to 100.00. For AlphaFold3, the input sequence is limited to the residues explicitly modelled in the reference chain. To construct the atom mask for applying the substructure conditioner, we use the provided amino acid subsequence to identify residues that will primarily be optimized with the density-based loss versus those optimized using the substructure conditioner.

Electron density maps Our proposed method generalizes across different types of real-space electron density maps. Here, we considered three map types.

- $2F_o - F_c$ **maps:** These maps are rendered using the Phenix [3] `fft` command:

```
fft hklin mtz_path mapout map_output_path
LABIN F1=FWT PHI=PHWT SIG1=SIGFP VF000 volume f000
```

The unit cell volume (`volume`) and the total scattering factor at zero scattering angle $F(0,0,0)$ (`f000`) are derived from the input atomic model. For a triclinic cell with edge lengths a, b, c (in Å) and angles α, β, γ (in radians), the unit cell volume is:

$$\texttt{volume} = a \cdot b \cdot c \cdot \sqrt{(1 - \cos^2\alpha - \cos^2\beta - \cos^2\gamma + 2\cos\alpha\cos\beta\cos\gamma)}$$

Unit cell parameters $(a, b, c, \alpha, \beta, \gamma)$ are parsed from the input atomic model. $F(0,0,0)$ is obtained using `phenix.f000`, which estimates the zero-angle scattering factor by summing electron counts from both the atomic model and the bulk solvent. The bulk solvent contribution is calculated using the default average solvent density of 0.35 $e^-/Å^3$. The atomic and solvent electron contributions are then combined to yield $F(0,0,0)$.

- **Absolute-scale END maps:** These density maps are generated using the END map workflow of [20], which combines Phenix and CCP4 programs to produce absolute-scale electron density maps ($e^-/Å^3$ units). We run

```
phenix.refine atomic_model_file mtz_file strategy=rigid_body
./END_RAPID.com pdb_refine_001.eff -norapid
```

  The atomic model first undergoes rigid-body refinement to prevent artifacts from misaligned models. The END script then rescales the observed structure factors against the calculated ones and offsets the density by adding $F(0,0,0)$ to every voxel. This places the map on an absolute scale, where zero corresponds to vacuum.

- **Composite omit maps:** These density maps are computed in Phenix using,

```
phenix.composite_omit_map atomic_model.pdb reflection_data.mtz
```

  These maps reduce model bias while minimizing phase degradation.

To align the real-space density maps with the corresponding atomic models from Section 2.2.1, we used `phenix.map_box` with `selection_radius=10` to carve the map around the corresponding atomic model with sufficient padding.

Since our X-ray density guidance operates locally, we carved the input density map to include voxels within a radius of 5Å of all atoms in the residue range (and chain) of interest from the reference PDBs. This pre-processing reduces VRAM memory consumption. During F_c computation, the ensemble density is calculated at voxels within this 5Å local neighborhood.

2.2.2 NMR

We used three sources of NMR observables for guidance: NOE-implied distance restraints, chemical-shift–implied backbone dihedral angles, and per-residue order parameters.

NOE-implied distance restraints. NOE restraint lists were parsed from NMR-STAR files (via `pynmrstar`), selecting only NOE-derived distance restraints; other restraint types such as dihedral angles, H-bonds, and residual dipolar coupling (RDC), were excluded. *Ambiguous NOEs*—cross-peaks whose atom-pair assignment is uncertain—were retained and represented as disjunctive (`OR`) constraints: each ambiguity group defines multiple candidate atom pairs sharing the same bounds, and the restraint is satisfied if *any one* candidate meets those bounds (not necessarily all). Lower/upper bounds were read directly from file; missing lower bounds were set to 0.0 Å. We note

that, although NOE physics is intensity-based, we adopt the common semi-quantitative distance-bound representation here, as is standard practice in MD-resolved NMR ensemble workflows; a future extension of our pipeline may incorporate intensity averaging within ambiguity groups prior to distance conversion. For each system, the corresponding NMR-STAR restraint file was retrieved from the PDB entry on the RCSB website.

Chemical-shift-implied dihedral angles. Backbone $\Phi^{\text{tgt}}/\Psi^{\text{tgt}}$ angles used in combined cryo-EM and NMR guidance were inferred from raw chemical shifts using TALOS-N [30], a tool that infers potential ranges of backbone dihedral angles from chemical shifts. Per target, the chemical-shift lists were retrieved from the BMRB entry cross-referenced to the corresponding PDB ID.

Order parameters. Order-parameter (S^2) values were taken on a per-residue basis from the respective publications (or associated supplementary datasets) for each target. Across all observables, when required values or bounds were unavailable for a residue or atom pair, we masked that item and did not guide on it.

2.2.3 Cryo-EM

In the Cryo-EM workflow, the inputs include a reference atomic model, amino acid sequences, and the corresponding electrostatic potential (ESP) map from the EMDB [1].

Atomic models In cryo-EM, the objective is not to refine high-resolution reference structures, as in X-ray crystallography, but to fit amino acid sequences into experimentally resolved volumes. This distinction is because of the lower resolution of cryo-EM ESP maps, which makes them useful for studying larger protein structures. Here, reference PDB structures from RCSB [27] are primarily used to align chain labels, construct atom masks, and serialize outputs. Unlike crystallographic pre-processing pipelines, missing residues or atoms are not imputed with modeling software (like PDBFixer), and alternate conformations are not split into separate models. To ensure consistency with AlphaFold3, all residues are renumbered to be 1-indexed. Finally, because cryo-EM often targets large complexes, multimeric inputs along with multiple sequences can be provided. Relevant residue ranges and chain identifiers are specified alongside the sequences before guidance.

Electrostatic potential maps The input ESP volumes are obtained from the Electron Microscopy Data Bank (EMDB) [1] in '.map' or '.mrc' format. To reduce VRAM usage without losing resolvable information, ESP maps with large side length D are downsampled in Fourier space to a smaller side length $D' < D$. The reported resolution r (in Å) is used to verify that the Nyquist frequency of the downsampled grid remains sufficient. This is checked using

$$D' \geq \frac{2D}{rd}$$

where d is the original voxel size in Å. This relation follows from (i) conservation of the total physical side length $D \cdot d = D' \cdot d' \implies D' = \frac{Dd}{d'}$ and (ii) the Nyquist criterion $r'_{\text{nyquist}} = 2d' \leq r \implies \frac{1}{d'} \geq \frac{2}{r}$.

Lastly, atomic masks are generated using the reference protein to provide controlled guidance, but are constructed in a way that maintains the generality of the approach:

- **Global mask:** A global mask is defined to include all voxels within $r_{\max} = 2.1$ Å of the atoms in the reference protein. Empty regions of the masks are filled using convex hull implemented in `scipy.spatial.ConvexHull` [34].

- **Per-chain mask:** For multimeric structures (e.g., amyloids), masks are generated per chain using a similar $r_{\max}$ distance threshold. Since this is restricted to well-separable chains, user-defined masks could likewise be imported after manual generation. This is later used to obtain better-defined guidance directions.

2.3 Protein structure inverse problems

Notation: The amino acid sequence of a protein is denoted by $\mathbf{a}$ and the corresponding 3D Cartesian coordinates of all atoms is denoted by $\mathbf{X} = (\mathbf{x}_1, \ldots, \mathbf{x}_m)$, where $\mathbf{x}_i \in \mathbb{R}^3$ is the position of the i-th atom.

Problem Setting: Given a protein's amino acid sequence $\mathbf{a}$ and its corresponding experimental observation $\mathbf{y}$ (from crystallography, NMR, or Cryo-EM), our goal is to sample an experimentally faithful ensemble of n structures $\boldsymbol{\mathcal{X}} = (\mathbf{X}^1, \ldots \mathbf{X}^n)$ from the posterior distribution $p(\boldsymbol{\mathcal{X}} \mid \mathbf{a}, \mathbf{y})$.

The posterior distribution can be factorized as $p(\boldsymbol{\mathcal{X}} \mid \mathbf{a}, \mathbf{y}) \propto p(\mathbf{y} \mid \boldsymbol{\mathcal{X}}, \mathbf{a}) \cdot p(\boldsymbol{\mathcal{X}} \mid \mathbf{a})$ using Bayes Rule. The prior term is captured by $p(\boldsymbol{\mathcal{X}} \mid \mathbf{a})$ which is the likelihood of an ensemble conditioned on the amino acid sequence $\mathbf{a}$. Using probabilistic laws, individual terms can be factorized into a product of independent samples in the prior term,

$$p(\boldsymbol{\mathcal{X}} = (\mathbf{X}^1 \ldots \mathbf{X}^n) \mid \mathbf{a}) = \prod_{i=1}^{n} p(\mathbf{X}^i \mid \mathbf{a}) \tag{1}$$

In our case, the prior $p(\boldsymbol{\mathcal{X}} \mid \mathbf{a})$ is implicitly modelled by AlphaFold3 [2].

The data term $p(\mathbf{y} \mid \boldsymbol{\mathcal{X}}, \mathbf{a})$ represents the probability of the experimental data $\mathbf{y}$ given the protein structure ensemble $\boldsymbol{\mathcal{X}}$ and the amino acid sequence $\mathbf{a}$. Since the data term is conditioned on the entire ensemble, we cannot separate the likelihood into separate structures. In essence, the data term models the underlying physics of the experiment. In this paper, we model the data term for three experimental modalities:

1. *X-ray crystallography* where we use real-space electron density (ED) maps;
2. *NMR* where we use nuclear Overhauser effect (NOE) pairwise distance restraints, amide (N-H) and methyl (C-C) order parameter obtained using relaxation (sensitive to ps-ns motion timescale) and residual dipolar coupling (RDC, sensitive to μs-ms motion timescale), and backbone dihedral angles ϕ and ψ estimated from NMR chemical shifts using TALOS-N [30];
3. *cryoEM* where we use electrostatic potential (ESP) maps.

Additionally, in all modalities we use the following terms:

1. *Substructure conditioning* allowing to fix parts of the structure to follow an existing model;
2. *Validity* ensuring physically plausible bond lengths, valence angles and lack of steric intersections.

In the sequel, we describe in detail the construction of these likelihoods.

2.3.1 Guided diffusion and non-i.i.d. sampling

AlphaFold3 generates structures using a diffusion-based generative model [13] over all-atom coordinates. In its original formulation, a single structure $\mathbf{X}$ is produced by integrating a variance-preserving stochastic differential equation (SDE) [31].

$$d\mathbf{X} = -\left(\frac{1}{2}\mathbf{X} + \nabla_{\mathbf{X}} \log p_t(\mathbf{X} \mid \mathbf{a})\right) \beta_t \, dt + \sqrt{\beta_t}\, \mathbf{N} \tag{2}$$

where β_t is the noise schedule, $\mathbf{N} \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$ is sampled from an isotropic normal distribution, and $\nabla_{\mathbf{X}} \log p_t(\mathbf{X} \mid \mathbf{a})$ is the learned score function modelled by AlphaFold3 [16]. To predict a single structure, AlphaFold3 samples from the prior distribution $\mathbf{X}_T \sim \mathcal{N}(\mathbf{0}, \beta_0 \cdot \mathbf{I})$ and integrates the SDE in Equation 2 from $t = T$ to $t = 0$ to iteratively desnoise the diffusion variable $\mathbf{X}_T$ to a noiseless protein structure $\mathbf{X}_0$.

In our setting, we expand the current SDE to model ensembles instead of single structures. Formally, we sample an ensemble $\boldsymbol{\mathcal{X}} = (\mathbf{X}^1, \mathbf{X}^2, \ldots, \mathbf{X}^n)$ of n structures. Thus, we can generalize the SDE in Equation 2 to samples ensembles instead of a single structure.

$$d\begin{bmatrix}\mathbf{X}^1 \\ \vdots \\ \mathbf{X}^n\end{bmatrix} = -\left(\frac{1}{2}\begin{bmatrix}\mathbf{X}^1 \\ \vdots \\ \mathbf{X}^n\end{bmatrix} + \begin{bmatrix}\nabla_{\mathbf{X}^1} \log p_t(\mathbf{X}^1 \mid \mathbf{a}) \\ \vdots \\ \nabla_{\mathbf{X}^n} \log p_t(\mathbf{X}^n \mid \mathbf{a})\end{bmatrix}\right) \beta dt + \sqrt{\beta_t}\begin{bmatrix}\mathbf{N}^1 \\ \vdots \\ \mathbf{N}^n\end{bmatrix} \tag{3}$$

where $\mathbf{N}^k \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$ and the unconditional score term $\nabla_{\mathbf{X}^k} \log p_t(\mathbf{X}^k \mid \mathbf{a})$ can be separated into independent structures like in Equation 1. In order to sample a *non-I.I.D.* ensemble from the posterior distribution, we plug in the guidance score to Equation 3, obtaining,

$$d\begin{bmatrix}\mathbf{X}^1 \\ \vdots \\ \mathbf{X}^n\end{bmatrix} = -\left(\frac{1}{2}\begin{bmatrix}\mathbf{X}^1 \\ \vdots \\ \mathbf{X}^n\end{bmatrix} + \begin{bmatrix}\nabla_{\mathbf{X}^1} \log p_t(\mathbf{X}^1 \mid \mathbf{a}) \\ \vdots \\ \nabla_{\mathbf{X}^n} \log p_t(\mathbf{X}^n \mid \mathbf{a})\end{bmatrix}\right) \beta dt - \boldsymbol{\eta} \nabla_{\boldsymbol{\mathcal{X}}} \log p\left(\mathbf{y} | \boldsymbol{\mathcal{X}}, \mathbf{a}\right) \beta dt + \sqrt{\beta_t}\begin{bmatrix}\mathbf{N}^1 \\ \vdots \\ \mathbf{N}^n\end{bmatrix} \tag{4}$$

Unlike the unconditional score term, the guidance score term Equation 4 is not separable because it is conditioned on ensemble $\boldsymbol{\mathcal{X}}$. Also, the hyperparameter $\boldsymbol{\eta}$ can be used to scale the strength of the guidance score and direct the flow towards higher posterior likelihood regions. The Pseudocode for guided AlphaFold3 is provided in Algorithm 1.

Parameter settings For X-ray crystallography, we consistently set $\boldsymbol{\eta} = 0.1$. For NOE constraints under NMR, we experimented with values $\boldsymbol{\eta} \in \{0.1, 0.2, 0.3, 0.4, 0.5\}$ and selected the best-performing result. For order parameters under NMR, we fixed $\boldsymbol{\eta} = 0.3$. To ensure numerical stability during the diffusion process, we apply gradient clipping to clip the guidance score. For Cryo-EM, the used values were $\eta \in 0.1, 0.15, 0.2$ where the best values per structures were chosen.

2.4 X-ray crystallography models

Let $\boldsymbol{\mathcal{X}} = \{\mathbf{X}^k\}_{k=1}^n$ be an ensemble of n conformers and conformer $\mathbf{X}^k$ have atomic coordinates $\{\mathbf{x}_i^k\}$ and isotropic B-factors B_i^k. The observed electron density is denoted as $F_{\mathrm{o}} : \mathbb{R}^3 \to \mathbb{R}$. In order to restrict map comparisons around a region of interest, we use a spatial mask.

Forward model. To compute the crystallographic likelihood term, we first calculate the theoretical electron density map $F_{\mathrm{c}} : \mathbb{R}^3 \to \mathbb{R}$ for all structures in ensemble $\boldsymbol{\mathcal{X}}$. F_{c} can be calculated at any spatial coordinate location $\boldsymbol{\xi} \in \mathbb{R}^3$. For a single conformer $\mathbf{X}^k = (\mathbf{x}_1^k \dots \mathbf{x}_m^k)$, F_{c} is a sum over kernel-density estimates built from six spherical Gaussians centered at the symmetry operation [5] of every atom.

$$F_{\mathrm{c}}(\boldsymbol{\xi}; \mathbf{X}^k, \mathbf{a}) = \sum_{q=1}^{N_{\mathrm{s}}} \sum_{i=1}^{m} \sum_{j=1}^{6} a_{i,j} \cdot \left(\frac{4\pi}{b_{i,j} + B_i^k} \right)^{3/2} \exp\left(\frac{-4\pi^2}{b_{i,j} + B_i^k} \cdot \|(\mathbf{R}_q \mathbf{x}_i^k + \mathbf{t}_q) - \boldsymbol{\xi}\|_2^2 \right) \tag{5}$$

Here, N_{s} is the number of symmetry operations [5], m is the number of atoms in the asymmetric unit of conformer $\mathbf{X}^k$, $\mathbf{R}_q \in SO(3)$ is the rotation matrix of symmetry operation q, $\mathbf{t}_q \in \mathbb{R}^3$ is the translation vector of the symmetry operation q, $a_{i,j}$ and $b_{i,j}$ are tabulated atomic form-factor coefficients for each heavy atom $\mathbf{x}_i^k$ [26], and B_i^k is the B-factor which represents the isotropic displacement parameter [33]. Here, we consider the B-factor to be a bandwidth parameter.

We can extend the density calculation at spatial coordinates $\boldsymbol{\xi}$ to ensembles

$$F_{\mathrm{c}}(\boldsymbol{\xi}; \{\mathbf{X}^k\}_{k=1}^n, \mathbf{a}) = \frac{1}{n} \sum_{k=1}^{n} F_{\mathrm{c}}(\boldsymbol{\xi}; \mathbf{X}^k, \mathbf{a}) \tag{6}$$

$F_{\mathrm{c}}(\boldsymbol{\xi}; \mathbf{X}^k, \mathbf{a})$ is computed using Equation 5. For our experiments, we assume that each atom has a uniform B-factor that is inversely related to the size of the ensemble, $B_i^k = \frac{4}{n} \forall i, k$. Additionally, we use an ensemble size of $n = 16$ for all our experiments. Our guidance framework differentiably incorporates all symmetry mates of the non-hydrogen atoms belonging to the chain of interest. In subsequent post-processing steps, we further append structural waters, non-standard residues, ligands, ions, and atoms from other chains of the proteins together with their respective symmetry mates.

X-ray log likelihood. To quantify the agreement between observed and calculated electron density maps, we measure the L_1 norm of the difference between the calculated ensemble-averaged density F_{c} (from Equation 6) and the observed density F_{o}:

$$\log p(\mathbf{y} = F_{\mathrm{o}} \mid \boldsymbol{\mathcal{X}}, \mathbf{a}) = - \left\| F_{\mathrm{o}}(\cdot) - F_{\mathrm{c}}(\cdot; \{\mathbf{X}^k\}_{k=1}^n, \mathbf{a}) \right\|_1 \tag{7}$$

The log-likelihood computation has been summarized in Algorithm 2.

2.5 NMR models

Differentiable hydrogen placement. AlphaFold3 does not explicitly model hydrogen atoms, yet their positions are essential for accurately modeling NMR observables. To address this, at each sampling step we first reconstruct hydrogen atoms of each atom in an amino acid in a differentiable manner using a PyTorch port of Hydride's [19] hydrogen placement algorithm. Each hydrogen-bearing center is identified by its local bonding environment and assigned a matching reference fragment. The fragment's non-hydrogen atoms are then aligned to the target by a single rigid-body superposition, after which the same transformation is applied to its hydrogens. In this way, proton positions adjust smoothly with the heavy-atom coordinates, yielding accurate and fully differentiable NMR calculations.

NOE distance restraints. We used solution-state NOE restraints as ensemble-averaged distance constraints. Conceptually, NOEs arise from through-space dipolar interactions between protons with observable effects typically for inter-nuclear separations $<\sim 6$ Å; intensities in NOESY spectra reflect an (approximate) r^{-6} dependence and are commonly used semi-quantitatively to bound distances. For an ensemble $\boldsymbol{\mathcal{X}} = \{\mathbf{X}^k\}_{k=1}^n$, we modeled the ensemble-average distance:

$$d_{ij}(\boldsymbol{\mathcal{X}}) = \frac{1}{n}\sum_{k=1}^{n} \left\| \mathbf{x}_i^k - \mathbf{x}_j^k \right\|_2$$

and enforced lower/upper bounds $(\underline{d}_{ij}, \overline{d}_{ij})$ via the log-likelihood:

$$\log p(D \mid \boldsymbol{\mathcal{X}}; \mathbf{a}) = -\sum_{(i,j)\in D} \left(\left[\underline{d}_{ij} - d_{ij}(\boldsymbol{\mathcal{X}})\right]_+^2 + \left[d_{ij}(\boldsymbol{\mathcal{X}}) - \overline{d}_{ij}\right]_+^2 \right)$$

where $[x]_+ = \max(x, 0)$. Additionally, ambiguous NOE assignments were retained as OR groups; for evaluation, the group-level violation equals the minimum violation among its members (i.e., the group is satisfied if any member lies within bounds). For restraints involving methyl groups, we modeled rapid threefold internal rotation by averaging over the three methyl proton positions generated by differentiable H placement.

We guided the diffusion sampler with the NOE likelihood using non-i.i.d. ensemble sampling, which couples conformers through the ensemble-level likelihood. Substructure conditioning was not used for NOE.

Order-parameters. Order parameters S^2 for amide and methyl bonds represent ensemble flexibility (lower values imply higher disorder). At each sampling step, we rigidly align all conformers $\mathbf{X}^k \in \boldsymbol{\mathcal{X}}$ in the ensemble to a reference $\mathbf{X}^{\text{ref}}$ to remove global rotation and translation. For each bond p in ensemble member $\mathbf{X}^k$, we compute a unit bond vector $\mathbf{d}_p^k$ (amide uses the N→H vector while methyl uses the C→C vector pointing to a methyl carbon; some amino acids like leucine, isoleucine and valine have two methyl groups and, hence, two vectors per residue). Given these normalized vectors, the calculated bond-order $S_{c,p}^2$ of the bond p is given by:

$$S_{c,p}^2(\boldsymbol{\mathcal{X}}; \mathbf{a}) = \frac{1}{|\boldsymbol{\mathcal{X}}|(|\boldsymbol{\mathcal{X}}|-1)} \sum_{i\neq j} P_2\left(\mathbf{d}_p^i \cdot \mathbf{d}_p^j\right)$$

where $P_2(x) = \frac{1}{2}(3x^2 - 1)$ is the Legendre polynomial of order 2. The log-likelihood is computed according to

$$\log p(S^2 \mid \boldsymbol{\mathcal{X}}, \mathbf{a}) = -\sum_p \left(S_{c,p}^2(\boldsymbol{\mathcal{X}}; \mathbf{a}) - S_p^2\right)^2$$

where the summations is performed over the collection of bonds on which the experimental data $S^2 = \{S_p^2\}$ have been collected.

Dihedrals likelihood Backbone dihedral restraints assess the agreement between predicted Ramachandran angles $(\phi^{\text{pred}}, \psi^{\text{pred}})$ from ensemble $\boldsymbol{\mathcal{X}}$ and target angles ($\Phi^{\text{tgt}} = \{\phi_i^{\text{tgt}}, \delta_{\phi,i}\}_{i=1}^M$ and $\Psi^{\text{tgt}} = \{\psi_i^{\text{tgt}}, \delta_{\psi,i}\}_{i=1}^M$) defined over M residues. The target angles and their associated uncertainty half-widths are obtained from TALOS-N [30] predictions. The log-likelihood is computed using a 2π-periodic (wrapped) angular discrepancy function that respects the circular nature of angles.

This term encourages locally accurate backbone geometry and regular secondary structure, complementing global guidance terms derived using NOE distance restraints or Cryo-EM ESP maps. The dihedral log-likelihood is defined as:

$$\log p\big(\Phi^{\text{tgt}}, \Psi^{\text{tgt}} \mid \boldsymbol{\mathcal{X}}, \mathbf{a}\big) = -\frac{1}{2M}\sum_{r=1}^{M}\left[\frac{\Delta\big(\phi_r^{\text{pred}}(\boldsymbol{\mathcal{X}}), \phi_r^{\text{tgt}}\big)^2}{\delta_{\phi,r}^2} + \frac{\Delta\big(\psi_r^{\text{pred}}(\boldsymbol{\mathcal{X}}), \psi_r^{\text{tgt}}\big)^2}{\delta_{\psi,r}^2}\right]$$

Where, $\Delta(\alpha, \beta)$ denotes the wrapped angular difference (implemented as `angle_diff`). This log-likelihood is averaged over all chains of structures in the ensemble.

Overall NMR likelihood. When guiding with order-parameters, NOE observations, and dihedrals, the log-likelihoods are combined as

$$\begin{aligned}\log p(S^2, D, \Phi^{\text{tgt}}, \Psi^{\text{tgt}} \mid \boldsymbol{\mathcal{X}}, \mathbf{a}) = &-\lambda_S \sum_p \big(S^2_{c,p}(\boldsymbol{\mathcal{X}}; \mathbf{a}) - S^2_p\big)^2 \\ &- \lambda_D \sum_{(i,j)\in D} \big(\big[\underline{d}_{ij} - d_{ij}(\boldsymbol{\mathcal{X}})\big]_+^2 + \big[d_{ij}(\boldsymbol{\mathcal{X}}) - \overline{d}_{ij}\big]_+^2 \big) \\ &+ \lambda_T \log p(\Phi^{\text{tgt}}, \Psi^{\text{tgt}} \mid \boldsymbol{\mathcal{X}}, \mathbf{a})\end{aligned} \tag{8}$$

We guide the model either without order parameter information, or using one of two types: relaxation-derived order parameters (ps-ns motions) or residual dipolar couplings (RDC)-derived order parameters (μs-ms motions). Dihedral restraints are incorporated only in experiments that involve Cryo-EM ESP maps. The scaling parameters $(\lambda_S, \lambda_D, \lambda_T)$ are set as follows:

$$(\lambda_S, \lambda_D, \lambda_T) = \begin{cases} (0,\ \lambda_D \in \{0.1, \ldots, 0.5\}, 0), & \text{NOE}, \\ (0.30, 0.45, 0), & \text{Methyl relaxation \& NOE}, \\ (0.45, 0.45, 0), & \text{Amide relaxation \& NOE}, \\ (0.45, 0.40, 0), & \text{Methyl } S^2 \text{ from RDC \& NOE}, \\ (0.30, 0.40, 0), & \text{Amide } S^2 \text{ from RDC \& NOE}, \\ (0.0, 0.40, 0.30), & \text{NOE, Cryo ESP, \& Dihedrals.} \end{cases} \tag{9}$$

When guiding with NOE's alone, we selected the best-performing value of λ_D from the tested set.

2.6 CryoEM models

ESP forward model As discussed in [36], the forward model for cryoEM electrostatic potential (ESP) maps is approximated by atomic scattering potentials, following equations (5) and (6). Unlike X-ray crystallography, cryoEM maps typically represent a single particle, so the model reduces to

$$F_{\text{c}}(\boldsymbol{\xi}; \mathbf{X}^k, \mathbf{a}) = \sum_{i=1}^{m}\sum_{j=1}^{5} a_{i,j} \cdot \left(\frac{4\pi}{b_{i,j} + B_i^k}\right)^{3/2} \exp\left(\frac{-4\pi^2}{b_{i,j} + B_i^k} \cdot \|(\mathbf{R}^k \mathbf{x}_i^k + \mathbf{t}^k) - \boldsymbol{\xi}\|_2^2\right), \tag{10}$$

where $\phi = (\mathbf{R}^k, \mathbf{t}^k)$ represents the pose of conformation $\mathbf{X}^k$ instead of a crystallographic symmetry operator. In the presence of significant blurring, the ensemble contribution can be evaluated as an average using Equation 6.

Protein – ESP map alignment strategies AlphaFold3's noisy time marginals $\mathbf{X}_t$ (atom coordinates at timestep t of the diffusion process) and the denoised predictions $\hat{\mathbf{X}}_0(\mathbf{X}_t, t)$ evolve on the learned diffusion manifold and not in a fixed frame. As a result, accurately predicted structures will generally be misaligned with the ERP map's voxel grid. Additionally, the $\hat{\mathbf{X}}_0(\mathbf{X}_t, t)$ predictions have a random pose at every iteration. In X-ray crystallography, Kabsch Algorithm [17] is used to solve the alignment problem. On the other hand, in cryo-EM *ab initio* protein-density fitting, no reference protein structure is available, making this alignment method inapplicable. Instead, a pure protein-density alignment strategy has to be implemented. Depending on whether the diffusion is performed over the whole ESP map or over well-separable chains, two different alignment strategies have been implemented:

- **Protein – full ESP map rigid alignment**: Global pose determined by solving:

$$\phi^k = (\mathbf{R}^k, \mathbf{t}^k) = \arg\min_{\phi^k} d(\mathbf{R}^k \mathbf{X}^k + \mathbf{t}^k, F_\mathrm{o}) \tag{11}$$

  where $d(\mathbf{X}, F_\mathrm{o})$ is the error objective defined in the following sections. The resulting pose is applied as a rigid transformation, $\tilde{\mathbf{X}}^k = \mathbf{R}^k \mathbf{X}^k + \mathbf{t}^k$ and all subsequent guidance scores are evaluated with respect to the aligned structure.

- **Protein – per chain non-rigid ESP alignment**: In this case, the alignment is performed independently for each chain. A set of per-chain poses $\{\phi^{k,c}\}_{c=1}^{N_\mathrm{c}} = \{(\mathbf{R}^{k,c}, \mathbf{t}^{k,c})\}_{c=1}^{N_\mathrm{c}}$ is estimated and applied to obtain the aligned structure $\tilde{\mathbf{X}}^k = \cup_c \left(\mathbf{R}^{k,c} \mathbf{X}^{k,c} + \mathbf{t}^{k,c}\right)$, where N_c is the number of chains.

In AlphaFold3, a structure $\mathbf{X}_t$ is represented as a tensor of shape $(N_\mathrm{a}, 3)$ where $N_\mathrm{a} = \sum_{c=1}^{N_\mathrm{c}} N_{\mathrm{a}_c}$, denotes the total number of atoms and N_{a_c} denotes the number of atoms in a chain c. After concatenation of $\tilde{\mathbf{X}}^k$, the application of the per-chain poses is a non-rigid transformation with respect to the full structure $\mathbf{X}^k$.

The chosen objective d is the trilinear interpolation $-F_\mathrm{o}(\tilde{\mathbf{X}}^k)$, and the maximization of F_o was implemented as a gradient descent over the pose ϕ^k. For this, hyperparameters of number of iterations N_iter, N_init initial random rotation, learning rates α_R, α_t, the side length T of the box around the weighted centroid of F_o and the positive B-factor B to be applied to the volume to ensure non-zero gradients for initial poor alignments can be defined in the configuration file. For cases where predictions $\hat{\mathbf{X}}_0(\mathbf{X}_t, t)$ have a starkly different tertiary structure from the intended $\mathbf{X}^k$ and this creates numerous unstable local minima (e.g., like what we observed in the insulin receptor structures `8U4B` and `8U4E`, a momentum decay α_m can be specified. This adds decaying Kabsch-aligned anchoring to the pose from the previous diffusion iteration to prevent large fluctuations of poses between diffusion time steps.

After-alignment diffusion manifold correction of $\mathbf{X}_t$ and $\hat{\mathbf{X}}_0(\mathbf{X}_t, t)$ Since in the diffusion process, $\mathbf{X}_{t-1}$ is computed as a function of the noisy structure and the $t = 0$ prediction $\mathbf{X}_{t-1} = f(\mathbf{X}_t, \hat{\mathbf{X}}_0(\mathbf{X}_t, t))$, in case a transformation was applied to the aligned $\hat{\tilde{\mathbf{X}}}_0$, for the correct $\mathbf{X}_{t-1}$ evaluation, an equivalent change must be applied to $\mathbf{X}_t$ in its relative frame.

- **Rigid alignment correction:** here the correction can be avoided if the direction gradient is evaluated using autograd as $\frac{\partial}{\partial X_t} \hat{\tilde{\mathbf{X}}}_0(\mathbf{X}_t, t)$, and the aligned $\hat{\tilde{\mathbf{X}}}_0$ is aligned back through $\mathbf{T}^{-1}(\tilde{\mathbf{X}}) = \mathbf{R}^T(\tilde{\mathbf{X}} - \mathbf{t})$.

- **Non-rigid alignment correction:** For correction and preservation of the manifold, non-rigid alignment is applied in two steps: (1) a rigid alignment $\mathbf{T}_1(\mathbf{X}) = \mathbf{R}_1\mathbf{X} + \mathbf{t}_1$ that aligns the complete $\mathbf{X}$ with the complete density F_o, and (2) the non-rigid per chain c transformations $\mathbf{T}_2^c(\mathbf{R}_2^c\mathbf{X}^c + \mathbf{t}_2^c)$ computed using gradient ascent. In this way, movements equivalent to the additional non-rigid fine-detail alignments $\mathbf{T}_2$ of $\hat{\tilde{\mathbf{X}}}_0$ can be defined in the relative frame of $\mathbf{X}_t$ and $\hat{\mathbf{X}}_0$ as $\mathbf{T}_2^*$, where the corrected structures we will denote as $\mathbf{X}_t^*$ and $\hat{\mathbf{X}}_0^*$. In other words, $\mathbf{T}_2 \circ \mathbf{T}_1 = \mathbf{T}_1 \circ \mathbf{T}_2^*$. Thus, $\hat{\mathbf{X}}_{0,c}^* = \mathbf{T}_1^{-1} \circ \mathbf{T}_2 \circ \mathbf{T}_1(\hat{\mathbf{X}}_{0,c})$ and $\mathbf{X}_t^* = \mathbf{T}_1^{-1} \circ \mathbf{T}_2 \circ \mathbf{T}_1(\mathbf{X}_t)$

ESP L_1 likelihood The objective L_1 is defined similarly to equation (7) and is computed as $-\|F_\text{c}(\mathbf{X}) - F_\text{o}\|_1$. This likelihood is best suited for homogeneous proteins or for heterogeneous ones with modest flexibility.

Optimal transport likelihood Different cryoEM ESP maps for a single sequence or protein assembly typically represent distinct states of a heterogeneous protein. Depending on the range of movement between the conformations of such proteins and the original predictions of AlphaFold3, the L_1 density objective gradients are often unable to guide the AlphaFold3 prediction to the map. Instead, due to floating point limitations and the absence of gradients for corresponding atoms that are far apart, this guidance strategy tends to push the noisy structures $\mathbf{X}_t$ and the corresponding $\hat{\mathbf{X}}_0$ predictions off the diffusion manifold. Therefore, a better approach for evaluating guidance gradients is to use a likelihood that directly points towards the density mass — namely, optimal transport.

We model the protein and ESP map as discrete measures $\mu = \sum_{i=1}^N \alpha_i \, \delta_{\mathbf{x}_i}$ and $\nu = \sum_{j=1}^M \beta_j \, \delta_{\mathbf{y}_j}$, where $\mathbf{x}_i \in \mathbb{R}^3$ are all the atoms of a protein and $\mathbf{y}_j \in \mathbb{R}^3$ are all the voxel centers inside the global mask (or equivalently, per-chain masks and per-chain coordinates). α_i is proportional to the atomic number Z_i and β_j is proportional to the observed value in the center of the voxel of F_0. The total mass of $\boldsymbol{\alpha}$ and $\boldsymbol{\beta}$ is normalized as $\sum_{i=1}^N \alpha_i = \sum_{j=1}^M \beta_j = 1$. The ground cost matrix is defined as $\mathbf{C} \in \mathbb{R}^{N \times M}$ with $C_{ij} = c(\mathbf{x}_i, \mathbf{y}_j)$ proportional to the 1- or 2-Norm between these two vectors. The *entropy optimal transport* objective is then defined as the *primal problem*:

$$\text{OT}_\varepsilon(\boldsymbol{\alpha}, \boldsymbol{\beta}; \mathbf{C}) = \min_{\substack{\mathbf{P} \in \mathbb{R}^{N \times M}_{\geq 0} \\ \mathbf{P}\mathbf{1}_M = \boldsymbol{\alpha},\, \mathbf{P}^\top \mathbf{1}_N = \boldsymbol{\beta}}} \Big\{ \langle \mathbf{P}, \mathbf{C} \rangle \; - \; \varepsilon \, H(\mathbf{P}) \Big\}, \qquad H(\mathbf{P}) = -\sum_{i,j} P_{ij} \log P_{ij}, \tag{12}$$

with $\langle \mathbf{P}, \mathbf{C} \rangle = \sum_{i,j} P_{ij} C_{ij}$ and $\mathbf{1}_K$ the K-vector of ones. Here $\mathbf{P}$ is a transport plan, where P_{ij} specifies the proportion of mass from source atom $\mathbf{x}_i$ assigned to voxel $\mathbf{y}_j$. The entropy term $\varepsilon H(\mathbf{P})$ promotes smooth, stable couplings and well-behaved gradients via discretizing assignments that are too sparse or sharp (coctain a lot of zero entries) and makes the problem strictly convex.

In practice, to reduce computational cost, we use the *Sinkhorn entropy optimal transport loss* in its dual representation with atom/voxel potentials $\boldsymbol{f} \in \mathbb{R}^N$ and $\boldsymbol{g} \in \mathbb{R}^M$:

$$\begin{aligned} \mathcal{L}_{\text{Sinkhorn},\varepsilon}(\boldsymbol{\alpha}, \boldsymbol{\beta}; \mathbf{C}) \; &:= \; \text{OT}_\varepsilon(\boldsymbol{\alpha}, \boldsymbol{\beta}; \mathbf{C}), \\ &= \max_{\boldsymbol{f}, \boldsymbol{g}} \; \boldsymbol{\alpha}^\top \boldsymbol{f} + \boldsymbol{\beta}^\top \boldsymbol{g} - \varepsilon \sum_{i,j} \alpha_i \beta_j \Big(\exp\Big(\tfrac{f_i + g_j - C_{ij}}{\varepsilon} \Big) - 1 \Big). \end{aligned} \tag{13}$$

We chose a powerful optimized implementation of optimal transport from `geomloss` [10], particularly to avoid memory bottlenecks. The hyperparameters of choice were `p=2`, `blur=`ϵ`=0.005`, and the binary `reach` flag, depending on whether mass needed to be dropped for a structure.

In cases where ESP maps contain separable chains, the OT objective was also applied per chain. The rationale was to separate the guidance directions and to prevent unwanted misaligned stable

local minima. This was required for structures 6W0O, 7DA4, and 9FH1. A typical issue with those structures was the discrepancy in AlphaFold3's tensor representation and the spatial arrangement of the subunits, where those were not ordered correspondingly. To solve that, we tested both the Hungarian algorithm for the Linear Assignment Problem (LAP), defined by chainwise distances and brute-force approaches. However, the most optimal solution was the per-chain OT objective: such guidance naturally ordered the AlphaFold'3 tensor representation to the spatial chains and thereby optimally solved for the correspondence implicitly.

Recycling and extended optimization of proteins In cases where AlphaFold3 early-time or unguided predictions poorly match the protein of interest, the initial protein optimization required to reach sufficient quality for ESP map alignment may consume the majority of available time steps in the diffusion process. Such a trend was observed for structures 6W0O and 9FH1. This, in turn, can leave an insufficient number of steps in the diffusion process for fine-detail protein-ESP map fitting, or make it completely impossible due to the noise schedule being insensitive to any change after the noise level σ_t reaches its lowest values (in practice, no significant change is observed after 160 steps). Additionally, the aforementioned proteins suffer from strong, stable local minima in the OT likelihood, e.g., 180° rotations about any axis lying on the plane perpendicular to the first principal component of 6W0O. This significantly increases the average time it takes for the protein to align to the intended spot during the diffusion process.

To allow additional time for finding these optimal protein-ESP map arrangements, we implement recycling: once the noisy structure reaches a specified time step t, instead of proceeding with the backwards diffusion process to obtain $\mathbf{X}_{t-1} = f_\theta(\mathbf{X}_t, \hat{\mathbf{X}}_0)$, noise is added as in the forward diffusion process to advance from time step t to $t+k$:

$$\mathbf{X}_{t+k} = \mathbf{X}_t + \sqrt{\sigma_{t+k} - \sigma_t}\,\epsilon, \quad \epsilon \sim \mathcal{N}(0, I).$$

For example, if protein-ESP alignment of 6W0Oconverged to all the intended rotations for each chain instead of the 180° opposite-symmetry poses after 160 time steps, it would not be further optimized. Adding an additional small level of noise, typically equivalent to 25 time steps of the forward process, does not break the correctly found tertiary structure while giving extra time for fine-detail fitting.

Objective normalization To balance multiple objectives with the ESP guidance, we employ a specific normalization strategy: one objective, typically the ESP L_1 or OT likelihood, is chosen as the reference, and all other likelihood terms are rescaled at each iteration so that their contributions remain a fixed proportion $\lambda_i \in [0, 1]$ of this main objective. This prevents magnitude imbalances that are otherwise observed when combining losses from very different modalities, and ensures stable relative weighting throughout the diffusion process.

Overall Cryo-EM likelihood The overall (pure) Cryo-EM guidance objective is defined as a weighted combination of the density-based L_1 and the OT likelihood:

$$\log p(\mathbf{y} = F_\mathrm{o} \mid \boldsymbol{\mathcal{X}}, \mathbf{a}) = -\lambda_1(t) \left\| F_\mathrm{o}(\cdot) - F_\mathrm{c}(\cdot; \{\mathbf{X}^k\}_{k=1}^n, \mathbf{a}) \right\|_1 - \lambda_2(t)\mathrm{OT}(\boldsymbol{\mathcal{X}}, F_\mathrm{o}, \mathbf{C}), \tag{14}$$

where $\lambda_1(t)$ and $\lambda_2(t)$ are time-dependent weights normalized as described above: one of these weights is set to $\lambda_i = 1$ as the main objective while the other is some portion of it normalized as per previous section. For the pure Cryo-EM loss, we typically simply use $\lambda_1 = 0$ and $\lambda_2 = 1$ until the structure roughly aligns to the density mass and then switch to pure density loss to optimize fine-detail features and optimize B-factors, i.e., $\lambda_1 = 1$ and $\lambda_2 = 0$. Some more complicated mixtures

were explored; however, they showed better results for combined guidance with other modalities. For example, in the NMR setting, the likelihood is scaled to 0.15 of the active objective, first relative to the OT loss (when it is turned on), and later relative to the L_1 density loss (once it becomes active).

2.7 Additional data terms

Substructure conditioner. Specifically for crystallographic structures, we aim to optimize a certain region of the protein while the remainder of the structure is stabilized by bootstrapping the diffusion process to a set of reference atomic coordinates. This is similar to `SubstructureConditioner` in Chroma [14].

As input, consider a list of reference atom locations $Y = \{\mathbf{y}_i : i \in A\}$ for atom indices $A \subseteq \{1 \dots m\}$ where m is the number of atoms in conformer $\mathbf{X}^k$. The log-likelihood is as a quadratic penalty on the deviation from reference atom locations

$$\log p(Y \mid \boldsymbol{\mathcal{X}}, \mathbf{a}) = -\lambda_{\text{sub}} \frac{1}{n} \sum_{k=1}^{n} \sum_{i \in A} \|\mathbf{x}_i^k - \mathbf{y}_i\|_2^2 \tag{15}$$

Before computing the log-likelihood we align all structures in the ensemble to the reference atom locations (limited to indices in A) using Kabsch algorithm [17]. For crystallography, $\lambda_{\text{sub}} = 0.1$, while for NMR and Cryo-EM, $\lambda_{\text{sub}} = 0.0$.

Validity likelihood. In order to prevent ensembles with elongated bonds and steric clashes, we incorporate a validity log-likelihood as regularization akin to AlphaFold2's violation loss [16]. Consider an ensemble of n protein conformations $\boldsymbol{\mathcal{X}} = \{\mathbf{X}^k\}_{k=1}^n$, where each structure $\mathbf{X}^k \in \mathbb{R}^{m \times 3}$. We define a binary bond matrix $\mathbf{B} \in \{0,1\}^{m \times m}$ such that $B_{ij} = 1$ if atoms i and j are covalently bonded, and 0 otherwise. The bond length loss for conformation $\mathbf{X}^k$ over bonded atom pairs ($B_{ij} = 1$) is given as,

$$\mathcal{L}_{\text{bond}}(\mathbf{X}^k; \mathbf{a}) = \sum_{i=1}^{m} \sum_{j=i+1}^{m} B_{ij} \cdot (\max(0, |d_{ij}^k - d_{ij}^{\text{ideal}}| - \delta_{\text{bond}}))^2$$

where, d_{ij}^{ideal} is the ideal bond length approximated as the sum of covalent radii of atoms i and j, $d_{ij}^k = \|\mathbf{x}_i^k - \mathbf{x}_j^k\|_2$ is the between atoms i and j in conformation k, and $\delta_{\text{bond}} = 0.2$Å is a tolerance margin. For non-bonded pairs (between and within residues), steric clashes are penalized using a soft collision loss, defined as the maximum violation over neighbors of each atom,

$$\mathcal{L}_{\text{collision}}(\mathbf{X}^k; \mathbf{a}) = \sum_{i=1}^{m} \max_{j \neq i, B_{ij}=0} (\max(0, (d_{ij}^{\text{ideal}} + p^{\text{collision}}) - d_{ij}^k))$$

Here, we use a padding distance $p^{\text{collision}} = 0.4$ (in Å) to prevent over-penalization of near-contact atoms. Lastly, we define a bond-angle violation loss as

$$\mathcal{L}_{\text{angle}}(\mathbf{X}^k; \mathbf{a}) = \sum_{j=1}^{m} \sum_{i \neq j} \sum_{k > i} B_{ij} B_{jk} \cdot (\max(0, |\theta_{ijk} - \theta_{ijk}^{\text{ideal}}| - \delta_{\text{angle}}))$$

Where θ_{ijk} (in degrees) is calculated using the dot product of the bond vectors from central atom j to atoms i and k, $\theta_{ijk}^{\text{ideal}}$ (in degrees) is retrieved from the Valence Shell Electron Pair Repulsion

(VSEPR) theory [12], and $\delta_{\text{angle}} = 12°$ is a tolerance margin. The resulting validity log-likelihood across the ensemble is given as

$$\log(\mathbf{B} \mid \boldsymbol{\mathcal{X}}, \mathbf{a}) = -\sum_{k=1}^{n} \lambda_{\text{bond}} \mathcal{L}_{\text{bond}}(\mathbf{X}^k; \mathbf{a}) + \lambda_{\text{collision}} \mathcal{L}_{\text{collision}}(\mathbf{X}^k; \mathbf{a}) + \lambda_{angle} \mathcal{L}(\mathbf{X}^k; \mathbf{a})$$

Where, $\lambda_{\text{bond}}, \lambda_{\text{collision}}$, and λ_{angle} are scaling factors that control the contribution of bond, collision, and bond angle terms. For ensembles guided using crystallographic density maps, we use $\lambda_{\text{bond}} = \lambda_{\text{collision}} = \lambda_{\text{angle}} = 0.075$. For NMR-guided ensembles, the scaling factors $(\lambda_{\text{bond}}, \lambda_{\text{collision}}, \lambda_{\text{angle}})$ depend on the choice of $(\lambda_S, \lambda_D, \lambda_T)$: [1]

Guidance type	$(\lambda_S, \lambda_D, \lambda_T)$	$(\lambda_{\text{bond}}, \lambda_{\text{collision}}, \lambda_{\text{angle}})$
NOE	$(0, \lambda_D \in \{0.1, \ldots, 0.5\}, 0)$,	$\{(0, 1.0, 1.0),\ (0, 2.0, 2.0)\}$
Methyl Relaxation & NOE	$(0.30, 0.45, 0)$	$(0.10, 0.10, 0)$
Amide Relaxation & NOE	$(0.45, 0.45, 0)$	$(0.15, 0.15, 0)$
Methyl S^2 from RDC & NOE	$(0.45, 0.40, 0)$	$(0.15, 0.15, 0)$
Amide S^2 from RDC & NOE	$(0.30, 0.40, 0)$	$(0.20, 0.20, 0)$
Cryo ESP	$(0, 0, 0)$	$(0, 0, 0)$
Cryo ESP & Dihedrals	$(0, 0, 0.30)$	$(0, 0, 0)$
Cryo ESP & NOE	$(0, 0.40, 0.0)$	$(0, 0, 0)$
Cryo ESP & Dihedrals & NOE	$(0, 0.40, 0.30)$	$(0, 0, 0)$

SI Table 6: Scaling factors for NMR-guided ensembles.

2.8 Force-field relaxation

Despite applying validity likelihood, a select few samples have broken bonds or steric clashes. Due to the non-physical nature of the ensemble, we filter out these samples from the ensemble. We consider a sample to have broken bonds if the distance between any pair of bonded atoms exceeds τ_{bond}. Likewise, we consider a sample to have steric clashes if the distance between any two atoms is less than τ_{clash}. Across all experiments, $\tau_{\text{bond}} = 2.1$Å and $\tau_{\text{clash}} = 1.1$Å.

Despite this, certain geometries can go undetected. To resolve this while maximizing the log-likelihood of the experimental observations, we relax the remaining samples in the ensemble using an off-the-shelf harmonic force-field. Specifically, we use OpenMM's [8] implementation of the AMBER14 [35] force field parameters for energy minimization within ColabFold [24]. In ColabFold, the energy is minimized for a maximum of 2000 iterations with energy tolerance threshold of 2.39 kcal/mol and stiffness of 100.0 kcal / mol Å^2 for NMR and 10.0 kcal / mol Å^2 for the other modalities, thereby controlling the strength of positional restraints applied to the atoms.

2.9 Ensemble pruning

Note: Here, *ensemble* refers to the set of samples that remain after relaxation, from which we prune a non-redundant subset. After relaxation, we aim to report a non-redundant subset of samples $\boldsymbol{\mathcal{X}}_{\mathcal{I}} = \{\mathbf{X}^k : k \in \mathcal{I}\}$ that best explains experimental observation $\mathbf{y}$. $\mathcal{I}$ denotes the indices of selected samples from the full ensemble. We propose two complementary strategies for populating $\mathcal{I}$.

[1] For NOE-guided experiments, we report the best-performing configuration out of two settings.

Orthogonal matching pursuit (OMP). OMP [23] is a greedy selection algorithm that begins with $\mathcal{I} = \emptyset$. At each iteration, the candidate $k \notin \mathcal{I}$ that maximizes $\log p(\mathbf{y}|\boldsymbol{\mathcal{X}}_{\mathcal{I}\cup\{k\}}, \mathbf{a})$ is added to the support set $\mathcal{I}$ until the log likelihood no longer increases. After termination, each sample in $\boldsymbol{\mathcal{X}}_{\mathcal{I}}$ is assigned a uniform occupancy of $\frac{1}{|\mathcal{I}|}$.

To avoid overfitting to noise in $\mathbf{y}$ and reduce redundancy in the ensemble, we introduce two early-stopping criteria. (i) The ensemble size is capped at $n_{\max} = 5$. Once $|\mathcal{I}| = 5$, the algorithm terminates. (ii) Before including a candidate, we evaluate the ensemble's validation loss f_{val}, and the algorithm terminates if the loss increases. The OMP selection procedure is applied only to ensembles derived from X-ray crystallography.

The selection of the best-performing sample at each iteration is based on Equation 6. The log-likelihood is computed over all voxels $\boldsymbol{\xi}$ in F_{o} that are within 2.5 Å of atoms in the relevant residue range of the reference PDB files (Section 2.2.1). Before a candidate is added to $\boldsymbol{\mathcal{X}}_{\mathcal{I}}$, a scalar B-factor B (applied uniformly to all atoms in the structure) is optimized to maximize the log-likelihood. This effectively tunes the bandwidth of $F_{\text{c}}(\boldsymbol{\xi}; \boldsymbol{\mathcal{X}}_{\mathcal{I}}, \mathbf{a})$ to best match the observed density $F_{\text{o}}(\boldsymbol{\xi})$. Optimization of B is carried out using Adam [18] for 100 iterations with a learning rate of 1.0. As a validation loss, we use R_{free} (Section 2.10), computed with Phenix's `phenix.model_vs_data`.

Occupancy optimization (OCC). *In crystallography*, we propose an alternative pruning strategy since OMP does not always yield the minimal subset explaining F_{o}. Instead, we optimize conformer occupancies in the relaxed ensemble $\boldsymbol{\mathcal{X}}_{\mathcal{I}}$ to best fit F_{o}. Each candidate initially assigned a uniform weight $p_k = \frac{1}{n}$ ($\sum_k p_k = 1$) with n samples in the ensemble $\boldsymbol{\mathcal{X}} = (\mathbf{X}^1 \dots \mathbf{X}^n)$.

In this algorithm, given $\boldsymbol{\mathcal{X}}$ and amino acid sequence $\mathbf{a}$, the *weighted average* calculated electron density is defined as

$$F_{\text{c}}(\boldsymbol{\xi}; \boldsymbol{\mathcal{X}}, \mathbf{a}) = \sum_k p_k \cdot F_{\text{c}}(\boldsymbol{\xi}; \mathbf{X}^k, \mathbf{a}),$$

where $F_{\text{c}}(\boldsymbol{\xi}; \mathbf{X}^k, \mathbf{a})$ is the calculated density of conformer $\mathbf{X}^k$ weighted by occupancy p_k. The occupancy weights are optimized by minimizing the following objective:

$$\mathcal{L} = \|F_{\text{o}}(\boldsymbol{\xi}) - F_{\text{c}}(\boldsymbol{\xi}; \boldsymbol{\mathcal{X}}, \mathbf{a})\|_1 + \lambda \sum_k p_k \log p_k,$$

The first term maximizes the log likelihood $\log p(F_{\text{o}}|\boldsymbol{\mathcal{X}}, \mathbf{a})$, while the second term serves as a sparsity regularizer that promote a small subset of conformer retain significant occupancy. The objective is minimized using Adam optimizer for $T_{\max} = 300$ iterations with a learning rate of $\alpha = 10^{-3}$ and $\lambda = 0.7$. We observe that, most occupancies decay to zero thereby best explaining F_{o} with minimal redundancy. After convergence, the pruned ensemble is defined as

$$\mathcal{I} = \{k : p_k > 0.1\}, \boldsymbol{\mathcal{X}}_{\mathcal{I}} = \{\mathbf{X}^k : k \in \mathcal{I}\}$$

A detailed pseudocode is provided in Algorithm 6. In practice, we apply both algorithms and ultimately select the minimal ensemble that provides the best fit to F_{o}.

Refinement *In crystallography*, after the termination of the ensemble pruning algorithms, the selected ensemble $\boldsymbol{\mathcal{X}}_{\mathcal{I}}$ is merged into a single PDB file. In this merged structure, atoms included in the set A (outside residue range of interest) are retained from the corresponding chain of the PDB-Redo entry [7], while atoms not included in A (in residue range of interest) are replaced with those modeled in $\boldsymbol{\mathcal{X}}_{\mathcal{I}}$. Each conformer in $\boldsymbol{\mathcal{X}}_{\mathcal{I}}$ is assigned a unique alternate conformer (altloc) identifier (`A`, `B`, etc.). Altloc occupancies are then set according to the probabilities estimated by

the selection algorithms (see Algorithms 6, 5). Other structural attributes, such as isotropic and anisotropic B-factors, are retained from the relaxed conformers.

The merged structure is then refined against the experimental diffraction data using `REFMAC5`, part of the CCP4 library [4]. `REFMAC5` refines a structural model by minimizing the difference between the observed structure factors (provided in the MTZ file from PDB-Redo) and the calculated structure factors (computed from the input model). The refinement program adjusts atomic coordinates and atomic displacement parameters (B-factors) to better fit the observed structure factors. During optimization, `REFMAC5` applies stereochemical restraints to maintain chemically plausible geometry. Since the coordinate updates are typically small, the refinement step preserves the conformational diversity introduced by our guidance framework while ensuring statistical consistency with the diffraction data. In practice, we perform 5 cycles.

2.10 Metrics

2.10.1 X-ray crystallography

To evaluate how well our density-guided structures agree with experimental X-ray crystallographic data, we report three metrics: R values, RSCC, and cosine similarity.

R-factors The R-factor (or crystallographic R-value) is a standard metric for assessing how well a crystal structure's model explains experimental X-ray diffraction data. This data is typically stored in MTZ files. The metric measures the alignment between the observed structure factor amplitudes $|F_{\text{obs}}|$ and the calculated structure factor amplitudes $|F_{\text{calc}}|$. Formally,

$$R = \frac{\sum_{\mathbf{h}} \big||F_{\text{obs}}(\mathbf{h})| - |F_{\text{calc}}(\mathbf{h})|\big|}{\sum_{\mathbf{h}} |F_{\text{obs}}(\mathbf{h})|}$$

Here, $\mathbf{h}$ denotes a Miller index triple (h, k, l) which denotes a discrete reflection in the reciprocal (Fourier) space. To avoid overfitting, R_{free} value is calculated in a similar manner, but using a small subset T of reflections that are excluded from the refinement pipeline and consequently serve as a validation set. Hence,

$$R_{\text{free}} = \frac{\sum_{\mathbf{h} \in T} \big||F_{\text{obs}}(\mathbf{h})| - |F_{\text{calc}}(\mathbf{h})|\big|}{\sum_{\mathbf{h} \in T} |F_{\text{obs}}(\mathbf{h})|}$$

It is important to note that $|F_{\text{obs}}|$ and $|F_{\text{calc}}|$ are structure factor amplitudes in reciprocal space. On the other hand, F_{o} and F_{c} used in our density-guidance framework denote the real space $3D$ volumetric electron density grids. Electron density maps are obtained by estimating phase to the structure factor amplitudes followed by an inverse Fourier transform into real space. For each experiment, we report both the R and R_{free} values as calculated by `REFMAC5` after refinement.

Real Space Correlation Coefficient (RSCC) The real-space correlation coefficient is a local validation metric that measures the alignment between a structural model and the observed experimental electron density map. Unlike R-values, RSCC is calculated directly in real space by comparing observed electron density $F_{\text{o}}(\boldsymbol{\xi})$ with calculated electron density $F_{\text{c}}(\boldsymbol{\xi})$ over an amino acid. Formally, RSCC for a grid of voxels $\mathcal{G}$ around a residue $a \in \mathbf{a}$ is given by,

$$\text{RSCC}_a = \frac{\sum_{\boldsymbol{\xi} \in \mathcal{G}} (F_{\text{o}}(\boldsymbol{\xi}) - \overline{F_{\text{o}}})(F_{\text{c}}(\boldsymbol{\xi}) - \overline{F_{\text{c}}})}{\sqrt{\sum_{\boldsymbol{\xi} \in \mathcal{G}} (F_{\text{o}}(\boldsymbol{\xi}) - \overline{F_{\text{o}}})^2} \sqrt{\sum_{\boldsymbol{\xi} \in \mathcal{G}} (F_{\text{c}}(\boldsymbol{\xi}) - \overline{F_{\text{c}}})^2}}$$

Where, $\overline{F_{\mathrm{o}}} = \mathbb{E}_{\mathcal{G}}[F_{\mathrm{o}}]$, and $\overline{F_{\mathrm{c}}} = \mathbb{E}_{\mathcal{G}}[F_{\mathrm{c}}]$. RSCC ranges from -1 to 1, with values closer to 1 indicating a residue that is well supported by the experimental density. For our use case, we used `phenix.real_space_correlation` to measure the RSCC for each residue (and altloc) in our protein.

Cosine Similarity. Whereas RSCC provides a residue-level measure of model-to-map alignment in real space, we are also report in a single global score over the entire residue range of interest. For this purpose, we use cosine similarity to quantify the overall alignment between the observed density F_{o} and the calculated density F_{c}. To compute this metric, we first identify the set of spatial coordinates $\boldsymbol{\xi} \in \mathbb{R}^3$ by selecting all voxels from F_{o} that lie within 2.5Å of atoms in the relevant residue range of the reference PDB structures. Formally,

$$\text{Cosine Similarity} = \frac{\sum_{\boldsymbol{\xi}} F_{\mathrm{o}}(\boldsymbol{\xi}) F_{\mathrm{c}}(\boldsymbol{\xi})}{\sqrt{\sum_{\boldsymbol{\xi}} F_{\mathrm{o}}^2(\boldsymbol{\xi})} \cdot \sqrt{\sum_{\boldsymbol{\xi}} F_{\mathrm{c}}^2(\boldsymbol{\xi})}}$$

where $F_{\mathrm{o}}(\boldsymbol{\xi})$ and $F_{\mathrm{c}}(\boldsymbol{\xi})$ denote vectors of density values at the selected voxel coordinates. As with RSCC, values closer to 1 indicate strong agreement between calculated and observed densities, signaling a good overall structural fit to the experimental map.

2.10.2 NMR

Percentage of violated NOE constraints Given a set of NOE distance-bound constraints grouped into ambiguous OR-groups, $G_m = \{\mathbf{r} = (i, j, \underline{d}_{ij}, \overline{d}_{ij})\}$, define the ensemble-averaged inter-proton distance

$$d_{ij}(\boldsymbol{\mathcal{X}}; \mathbf{a}) = \frac{1}{|\boldsymbol{\mathcal{X}}|} \sum_{k=1}^{|\boldsymbol{\mathcal{X}}|} \|\mathbf{x}_i^k - \mathbf{x}_j^k\|_2,$$

and the per-restraint violation magnitude

$$v_{ij} = \max\{\underline{d}_{ij} - d_{ij}(\boldsymbol{\mathcal{X}}; \mathbf{a}), d_{ij}(\boldsymbol{\mathcal{X}}; \mathbf{a}) - \overline{d}_{ij}, 0\}.$$

The group-level violation is $v_{G_m} = \min_{\mathbf{r} \in G_m} v_{ij}$, so an OR-group is satisfied if any member falls within bounds. The percentage of violated NOE constraints is

$$100 \times \frac{|\{\, m : \; v_{G_m} > 0 \,\}|}{M},$$

where M is the number of OR-groups. Ambiguity and methyl restraints are treated as in the NOE likelihood (ensemble-averaged distances; methyls averaged over the three proton positions).

Median violation magnitude Given the group-level violations $\{v_{G_m}\}_{m=1}^{M}$ defined above, the median violation magnitude is

$$\text{median}\big(\{\, v_{G_m} : \; v_{G_m} > 0 \,\}\big) \;\; [\text{Å}],$$

i.e., the median taken over violated groups only.

Order parameter correlation (r-factor) Given a set of experimentally determined order parameters, $\{S_p^2\}$ for some collection of bonds $p \in P$, and a corresponding set of order parameters $\{S_{c,p}^2\}$ estimated from the generated ensemble, the correlation coefficient is defined as

$$r = \frac{\sum_p \left(S_{c,p}^2 - \overline{S_c^2}\right)\left(S_p^2 - \overline{S^2}\right)}{\sqrt{\sum_p \left(S_{c,p}^2 - \overline{S_c^2}\right)^2} \cdot \sqrt{\sum_p \left(S_p^2 - \overline{S^2}\right)^2}}$$

where $\overline{S^2} = \frac{1}{|P|}\sum_p S_p^2$ and $\overline{S_c^2} = \frac{1}{|P|}\sum_p \S_{c,p}^2$ are the averages of the experimental and estimated order parameters, respectively. The coefficient ranges from -1 to 1, with higher values indicating better agreement with the experiment.

Order parameter normalized fitting error (q-factor) Similary to the correlation coefficient, the normalized fitting error is defined as

$$q = \frac{\sqrt{\sum_p \left(S_{c,p}^2 - S_p^2\right)^2}}{\sqrt{\sum_p S_{c,p}^2}}.$$

Lower values of q indicate a better agreement with the experiment. While r is insensitive to a constant offset in the estimated values, q is sensitive both to the error standard deviation and bias.

ANSURR metrics ANSURR [11] (Accuracy of NMR Structures Using RCI and Rigidity) validates a protein model by comparing per-residue flexibility inferred from backbone chemical shifts with rigidity predicted from the 3D structure. Flexibility from NMR data is estimated via the Random Coil Index (RCI), while structural rigidity is computed from the network of covalent and hydrogen-bond constraints using rigidity theory (e.g., FIRST). ANSURR reports two centile-ranked scores (0-100): a *correlation* score that reflects agreement in the pattern of flexible vs. rigid regions, and an *RMSD* score that reflects agreement in the magnitude of rigidity/flexibility. For structural ensembles, ANSURR is computed separately for each ensemble member, and the median of RMSD and correlation scores across the ensemble are used as representative metrics.

2.10.3 Cryo-EM

For Cryo-EM, the same metrics as for Crystallography were taken, mainly the cosine similarity and the per-residue model-map cross-correlation (CC) scores from Phenix from section 2.10.1.

2.11 Minimal end-to-end code invocation

2.11.1 X-ray crystallography

To guide AlphaFold3 predictions using real-space crystallographic electron density maps, run the following command in your terminal:

```
python3 run_xray.py <pdb_id> <chain_id> <region_subseq> \
--ccp4_setup_sh <CCP4 SETUP PATH> --phenix_setup_sh <PHENIX SETUP PATH>
```

Here, `<pdb_id>` is the four-character PDB identifier of the protein, `<chain_id>` specifies the target chain, `<region_subseq>` is the subsequence of the amino acid sequence to be guided. The `--ccp4_setup_sh` and `--phenix_setup_sh` arguments point to the respective CCP4 and PHENIX installation scripts.

2.11.2 NMR

To guide AlphaFold3 predictions using NOE distance restraints, run the following command in your terminal:

```
python3 run_nmr.py <pdb_id>
```

Here, `<pdb_id>` is the four-character PDB identifier of the protein. Optionally, users may also supply methyl RDC and relaxation files (`--methyl_relax_file` and `--methyl_relax_file`), as well as amide RDC and relaxation files (`--amide_rdc_file` and `--amide_relax_file`), to provide additional experimental guidance during structure prediction.

2.11.3 CryoEM

To guide AlphaFold3 predictions using CryoEM electrostatic potential maps, run the following command in your terminal:

```
python3 run_em.py <pdb_id> <emdb_id> <renumbered_pdb_file> <assembly_identifier> \
--phenix_setup_sh <PHENIX SETUP PATH> --dihedrals_file <TALOS dihedrals> \
--noe_restraints_file <NOE restraints> --noe_pdb_file <renumbered NOE file path> \
--sequences <space separated sequence> --counts <#chains for each sequence>
```

Here, `<pdb_id>` is the four-character PDB identifier of the protein, `<emdb_id>` is the EMDB [1] identifier of the cryo-EM electrostatic potential map and `<renumbered_pdb_file>` is the path to a PDB file whose residues have been renumbered to match the 1-indexed amino acid sequences. In other words, if residue A in the input sequence is at index 5, the corresponding residue in the PDB must also be renumbered to index 5. Missing or unresolved atoms are permitted. Since the modeling is performing *ab-initio*, the PDB is not used during guidance. Instead, it is used exclusively for evaluation and for constructing the initial volume mask when the default density when the default density-based alignment is selected (as opposed to RMSD alignment to atomic coordinates). Multi-chain compositions are specified with `--sequences` (space-separated amino-acid sequences for each unique chain) and `--counts` (the corresponding copy numbers, in order). `<assembly_identifier>` specifies the biological assembly or asymmetric unit to model. The `-phenix_setup_sh` argument point to the PHENIX installation script. Our framework also supports joint guidance using ESP maps. These optional restraints can be provided via `--dihedrals_file` (e.g., TALOS-N dihedrals), `--noe_restraints_file` (NOE distance restraints), and `--noe_pdb_file` – in the similar fashion to the renumbered pdb file, all the residues in these files should be reordered to match the 1-indices in the sequences given to our guidance model.

Algorithm 1 AlphaFold3 Guidance

Input: $\{\mathbf{f}^*\}$, $\{\mathbf{s}_i^{\text{inputs}}\}$, $\{\mathbf{s}_i^{\text{trunk}}\}$, $\{\mathbf{z}_{ij}^{\text{trunk}}\}$, noise schedule $[\beta_0, \beta_1, \ldots, \beta_T]$, noise factor $\gamma_0 = 0.8$, minimum noise factor $\gamma_{\min} = 1.0$, noise scale $\lambda = 1.003$, step scale $\kappa = 1.5$, experimental observation $\mathbf{y}$, guidance scale $\boldsymbol{\eta}$, reference structure $\mathbf{r}$, substructure conditioner flag b, substructure indices I, ensemble size n, number of atoms m.

Output: Guided Ensemble $\boldsymbol{\mathcal{X}}_l$

$\boldsymbol{\mathcal{X}}_l \sim \beta_0 \cdot [\mathbf{N}^1, \ldots, \mathbf{N}^n]^T$ ▷ $\mathbf{N}^i \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$, $\boldsymbol{\mathcal{X}} \in \mathbb{R}^{n \times m \times 3}$

for $\beta_\tau \in [\beta_1, \ldots, \beta_T]$ **do**

  $\boldsymbol{\mathcal{X}}_l \leftarrow \text{CentreRandomAugmentation}(\boldsymbol{\mathcal{X}}_l)$

  $\gamma \leftarrow \begin{cases} \gamma_0, & \beta_\tau > \gamma_{\min} \\ 0, & \text{otherwise} \end{cases}$

  $\hat{t} \leftarrow \dfrac{\beta_{\tau-1}(\gamma + 1)}{\sqrt{t^2 - \beta_\tau^2 / \beta_{\tau-1}}}$

  $\boldsymbol{\xi}_l \leftarrow \lambda \sqrt{\hat{t}^2 - \beta_\tau^2} [\mathbf{N}^1, \ldots, \mathbf{N}^n]^T$

  $\boldsymbol{\mathcal{X}}_l^{\text{noisy}} \leftarrow \boldsymbol{\mathcal{X}}_l + \boldsymbol{\xi}_l$

  $\boldsymbol{\mathcal{X}}_l^{\text{denoised}} \leftarrow \text{AlphaFold3}(\{\boldsymbol{\mathcal{X}}_l^{\text{noisy}}\}, \hat{t}, \{\mathbf{f}^*\}, \{\mathbf{s}_i^{\text{inputs}}\}, \{\mathbf{s}_i^{\text{trunk}}\}, \{\mathbf{z}_{ij}^{\text{trunk}}\})$

  if b **then**

    for $i \in I$ **do**

      $\boldsymbol{\mathcal{X}}_i^{\text{denoised}} \leftarrow [\mathbf{r}_i, \ldots, \mathbf{r}_i]$ ▷ Repeated n times

    end for

  end if

  $\delta_l \leftarrow (\boldsymbol{\mathcal{X}}_l - \boldsymbol{\mathcal{X}}_l^{\text{denoised}}) / \hat{t}$

  $\mathcal{L} \leftarrow \log p(\mathbf{y} \mid \boldsymbol{\mathcal{X}}, \mathbf{a})$

  $\mathbf{g} \leftarrow \dfrac{\partial \mathcal{L}}{\partial \boldsymbol{\mathcal{X}}_l^{\text{noisy}}}$ ▷ Guidance score wrt ensemble

  $\mathbf{g} \leftarrow \mathbf{g} \cdot \dfrac{\|\delta_l\|_2}{\|g\|_2}$ ▷ Gradient Scaling

  $\delta_l \leftarrow \delta_l + \mathbf{g} \cdot \boldsymbol{\eta}$

  $dt \leftarrow \beta_\tau - \hat{t}$

  $\boldsymbol{\mathcal{X}}_l \leftarrow \boldsymbol{\mathcal{X}}_l^{\text{noisy}} + \kappa \cdot dt \cdot \delta_l$

end for

return $\boldsymbol{\mathcal{X}}_l$

Algorithm 2 Likelihood for X-ray crystallography

Input: Observed map F_{o} on grid $\mathcal{G}$; ensemble $\boldsymbol{\mathcal{X}} = \{\mathbf{X}_k\}_{k=1}^{n}$, Crystallographic form-factors $\{a_{i,j}, b_{i,j}\}_{j=1}^{6}$ for each atom i across all conformers, Per-atom isotropic B_i^k for atom i in conformer k; unit cell $\mathbf{U}$; symmetry ops $\{(\mathbf{R}_q, \mathbf{t}_q)\}_{q=1}^{N_{\mathrm{s}}}$, ensemble, Optional voxel mask $W(\boldsymbol{\xi}) \in \{0, 1\}$ for $\boldsymbol{\xi} \in \mathcal{G}$ (default $W = 1$), amino acid sequence $\mathbf{a}$

Output: Scalar loss $\mathcal{L}$

function LogLikelihood-Xray($F_{\mathrm{o}}, \boldsymbol{\mathcal{X}} = \{\mathbf{X}_1 \dots \mathbf{X}_n\}$, $\mathbf{a}$)
  for $k = 1$ to n **do**
    for $\boldsymbol{\xi} \in \mathcal{G}$ **do**
      $F_{\mathrm{c}}^{(k)}(\boldsymbol{\xi}) \leftarrow 0$
      for $q = 1$ to N_{s} **do**
        for $i = 1$ to m **do**
          for $j = 1$ to 6 **do**
            $F_{\mathrm{c}}(\boldsymbol{\xi}; \mathbf{X}^k, \mathbf{a}) += \left[a_{i,j} \left(\frac{4\pi}{b_{i,j} + B_i^k} \right)^{3/2} \exp\left(\frac{-4\pi^2}{b_{i,j} + B_i^k} \|(\mathbf{R}_q \mathbf{x}_i^k + \mathbf{t}_q) - \boldsymbol{\xi}\|_2^2 \right) \right]$
          end for
        end for
      end for
    end for
  end for
  $F_{\mathrm{c}}(\boldsymbol{\xi}; \boldsymbol{\mathcal{X}}, \mathbf{a}) \leftarrow \frac{1}{n} \sum_{k=1}^{n} F_{\mathrm{c}}(\boldsymbol{\xi}; \mathbf{X}^k, \mathbf{a})$
  $\mathcal{L} \leftarrow \sum_{\boldsymbol{\xi} \in \mathcal{G}} W(\boldsymbol{\xi}) \left| F_{\mathrm{o}}(\boldsymbol{\xi}) - F_{\mathrm{c}}(\boldsymbol{\xi}; \boldsymbol{\mathcal{X}}, \mathbf{a}) \right|$
  return $\mathcal{L}$
end function

Algorithm 3 NOE distance-bound likelihood (ensemble-averaged; ambiguous OR-groups)

Input: Ensemble $\boldsymbol{\mathcal{X}} = \{\mathbf{X}^k\}_{k=1}^{n}$ (non-hydrogen atoms), Amino acid sequence $\mathbf{a}$; NOE dataset $D = \{G_1, \dots, G_M\}$ (OR-groups) each G contains restraints $\mathbf{r} = (i, j, \underline{d}, \overline{d})$;

Output: Scalar loss $\mathcal{L}$

function LikelihoodLoss-NOE($\boldsymbol{\mathcal{X}} = \{\mathbf{X}^k\}_{k=1}^{n}, \mathbf{a}, D$)
  $\mathcal{L} \leftarrow 0$
  for $k = 1$ to n **do**
    $\mathbf{H}^k \leftarrow$ `PlaceHydrogensDifferentiable`$(\mathbf{X}^k, \mathbf{a})$ ▷ Differentiable H placement per conformer
  end for
  for each OR-group $G \in D$ **do**
    $v_{\min} \leftarrow +\infty$
    for $\mathbf{r} = (i, j, \underline{d}, \overline{d}) \in G$ **do**
      `sum_d` $\leftarrow 0$
      for $k = 1$ to n **do**
        $\mathbf{x}_i^k, \mathbf{x}_j^k \leftarrow$ `GetAtomCoordinates`$(\mathbf{r}, \mathbf{X}^k, \mathbf{H}^k)$ ▷ Endpoint coordinates for restraint $\mathbf{r}$
        $d_k \leftarrow \|\mathbf{x}_i^k - \mathbf{x}_j^k\|_2$
        `sum_d` $\leftarrow$ `sum_d` $+ d_k$
      end for
      $d_{\mathrm{ens}} \leftarrow \frac{1}{n}$ `sum_d`
      $v \leftarrow \left(\max(\underline{d} - d_{\mathrm{ens}}, 0)\right)^2 + \left(\max(d_{\mathrm{ens}} - \overline{d}, 0)\right)^2$
      $v_{\min} \leftarrow \min(v_{\min}, v)$ ▷ OR-group: satisfied if any member fits
    end for
    $\mathcal{L} \leftarrow \mathcal{L} + v_{\min}$
  end for
  return $\mathcal{L}$
end function

Algorithm 4 Order-parameter (S^2) likelihood

Input: Ensemble $\boldsymbol{\mathcal{X}} = \{\mathbf{X}^k\}_{k=1}^n$, Amino acid sequence $\mathbf{a}$, Dataset $\mathcal{P} = \{(p, S_p)\}$ of measurement sites p with target S_p

function LIKELIHOODLOSS-ORDERPARAM($\boldsymbol{\mathcal{X}} = \{X_k\}_{k=1}^n, \mathbf{a}, \mathcal{P}$)
  $\mathcal{L} \leftarrow 0$
  $\mathbf{X}^{\text{ref}} \leftarrow \mathbf{X}^1$ ▷ First conformer as reference
  for $k = 1$ to n **do**
    $\mathbf{X}^k \leftarrow \texttt{AlignToRef}(\mathbf{X}^k, \mathbf{X}^{\text{ref}})$ ▷ Rigid alignment to reference pose
    $\mathbf{H}^k \leftarrow \texttt{PlaceHydrogensDifferentiable}(\mathbf{X}^k, \mathbf{a})$ ▷ Differentiable H placement per conformer
  end for
  for each $(p, S_p) \in \mathcal{P}$ **do**
    $S^2_{c,p} \leftarrow 0$
    for $k = 1$ to n **do**
      $\mathbf{d}_p^k \leftarrow \texttt{SiteVector}(p, \mathbf{X}^k, \mathbf{H}^k)$ ▷ unit bond vector $\mathbf{d}_p$ (N→H for amide; C→C for methyl)
    end for
    for $k = 1$ to n **do**
      for $l = 1$ to n **do**
        if $k \neq \ell$ **then**
          $x \leftarrow \mathbf{d}_p^k \cdot \mathbf{d}_p^l$
          $P_2 \leftarrow \frac{1}{2}(3x^2 - 1)$
          $S_c^2 \leftarrow S_c^2 + P_2$
        end if
      end for
    end for
    $S_c^2 \leftarrow \dfrac{1}{n(n-1)} S_c^2$
    $\mathcal{L} \leftarrow \mathcal{L} + \left(S_c^2 - S_p\right)^2$
  end for
  return $\mathcal{L}$
end function

Algorithm 5 Selecting samples using matching pursuit [23]

Input: Samples in ensemble $\boldsymbol{\mathcal{X}} = \{\mathbf{X}^1, \ldots, \mathbf{X}^n\}$; experimental observation $\mathbf{y}$; amino acid sequence $\mathbf{a}$; likelihood function to be maximized $\log p(\mathbf{y} \mid \boldsymbol{\mathcal{X}}, \mathbf{a})$; maximum samples to select $n_{\max}$; validation loss function $f_{\text{val}}(\boldsymbol{\mathcal{X}})$

Output: Pruned ensemble set $\boldsymbol{\mathcal{X}}_{\mathcal{I}} = \{\mathbf{X}^k : \ k \in \mathcal{I}\}$ and occupancies $\{p_k\}_{k \in \mathcal{I}}$

$\mathcal{I} = \emptyset$
$s_{\text{current}} = 0$
$\mathcal{L}_{\text{val}} = +\infty$ ▷ Validation loss
while $|\mathcal{I}| < n_{\max}$ **do**
  $L = \{\ell_k = \log p(\mathbf{y} \mid \boldsymbol{\mathcal{X}}_{\mathcal{I}\cup\{k\}}, \mathbf{a}) : k \in \mathcal{I}^{\text{c}}\}$
  $k^* \leftarrow \arg\max_k L \quad s^* \leftarrow \max_k L$ *Maximize* $\log p(\mathbf{y}|\boldsymbol{\mathcal{X}}_{\mathcal{I}\cup\{k^*\}}, \mathbf{a})$
  $\mathcal{L}^*_{\text{val}} \leftarrow f_{\text{val}}(\boldsymbol{\mathcal{X}}_{\mathcal{I}\cup\{k^*\}})$ ▷ Cross validation
  if $(s^* < s_{\text{current}})$ or $(\mathcal{L}^*_{\text{val}} > \mathcal{L}_{\text{val}})$ **then**
    break
  end if
  $\mathcal{I} \leftarrow \mathcal{I} \cup \{k^*\}$ *Add best sample*
  $s_{\text{current}} \leftarrow s^*$
  $\mathcal{L}_{val} \leftarrow \mathcal{L}^*_{\text{val}}$
end while
$\boldsymbol{\mathcal{X}}_{\mathcal{I}} = \{\mathbf{X}^k : k \in \mathcal{I}\}$
return $\boldsymbol{\mathcal{X}}_{\mathcal{I}}, \{p_k = \frac{1}{|\mathcal{I}|} : i \in \mathcal{I}\}$

Algorithm 6 Occupancy optimization (logit parameterization with entropy sparsity)

Input: Observed map $F_{\text{o}}(\boldsymbol{\xi})$ on grid $\mathcal{G}$, amino acid sequence $\mathbf{a}$, ensemble $\boldsymbol{\mathcal{X}} = (\mathbf{X}^1 \ldots \mathbf{X}^k)$, per-conformer maps $\{F_{\text{c}}(\boldsymbol{\xi}; \mathbf{X}^k, \mathbf{a})\}_{k=1}^n$, Mask $W(\boldsymbol{\xi}) \in \{0, 1\}$, sparsity weight $\lambda > 0$; step size $\alpha > 0$, max iterations $T_{\max}$, stability $\varepsilon_p = 10^{-8}$

Output: Pruned ensemble set $\boldsymbol{\mathcal{X}}_{\mathcal{I}} = \{\mathbf{X}^k : \ k \in \mathcal{I}\}$ and occupancies $\{p_k\}_{k \in \mathcal{I}}$

Initialize logits $z_k \leftarrow \frac{1}{n}, \forall k \in [1, \ldots, n]$
for $t = 1$ to $T_{\max}$ **do**
  $p \leftarrow \text{softmax}(z)$
  $F_{\text{c}}(\boldsymbol{\xi}; \boldsymbol{\mathcal{X}}, \mathbf{a}) \leftarrow \sum_{k=1}^n p_k \cdot F_{\text{c}}(\boldsymbol{\xi}; \mathbf{X}^k, \mathbf{a})$
  $\mathcal{L}_{\text{data}} \leftarrow \sum_{\boldsymbol{\xi} \in \mathcal{G}} W(\boldsymbol{\xi}) \cdot \left|F_{\text{o}}(\boldsymbol{\xi}) - F_{\text{c}}(\boldsymbol{\xi}; \boldsymbol{\mathcal{X}}, \mathbf{a})\right|$
  $\mathcal{L}_{\text{entropy}} \leftarrow -\sum_{k=1}^n p_k \log\big(\max(p_k, \varepsilon_p)\big)$
  $\mathcal{L} \leftarrow \mathcal{L}_{\text{data}} + \lambda \mathcal{L}_{\text{entropy}}$
  $g_z \leftarrow \nabla_z \mathcal{L}$
  $z \leftarrow z - \alpha \cdot g_z$
  if Converged **then break** ▷ Gradient Thresholding: $\|\nabla_z \mathcal{L}\|_2 < 10^{-3}$
end for
$p \leftarrow \text{softmax}(z)$
$\mathcal{I} \leftarrow \{k \in \{1, \ldots, n\} : p_k > 0.1\}$
$\boldsymbol{\mathcal{X}}_{\mathcal{I}} \leftarrow \{\mathbf{X}^k : k \in \mathcal{I}\}$
return $\boldsymbol{\mathcal{X}}_{\mathcal{I}}, \{p_k : k \in \mathcal{I}\}$

References

[1] Emdb—the electron microscopy data bank. *Nucleic acids research*, 52(D1):D456–D465, 2024.

[2] Josh Abramson, Jonas Adler, Jack Dunger, Richard Evans, Tim Green, Alexander Pritzel, Olaf Ronneberger, Lindsay Willmore, Andrew J Ballard, Joshua Bambrick, et al. Accurate structure prediction of biomolecular interactions with alphafold 3. *Nature*, 630(8016):493–500, 2024.

[3] Paul D Adams, Pavel V Afonine, Gábor Bunkóczi, Vincent B Chen, Ian W Davis, Nathaniel Echols, Jeffrey J Headd, L-W Hung, Gary J Kapral, Ralf W Grosse-Kunstleve, et al. Phenix: a comprehensive python-based system for macromolecular structure solution. *Biological crystallography*, 66(2):213–221, 2010.

[4] Jon Agirre, Mihaela Atanasova, Haroldas Bagdonas, Charles B Ballard, Arnaud Baslé, James Beilsten-Edmands, Rafael J Borges, David G Brown, J Javier Burgos-Mármol, John M Berrisford, et al. The ccp4 suite: integrative software for macromolecular crystallography. *Biological Crystallography*, 79(6):449–461, 2023.

[5] Carolyn Pratt Brock, T Hahn, H Wondratschek, U Müller, U Shmueli, E Prince, A Authier, V Kopskỳ, D Litvin, E Arnold, et al. *International tables for crystallography volume A: Space-group symmetry*. Wiley Online Library, 2016.

[6] Benjamin Charlier, Jean Feydy, Joan Alexis Glaunès, François-David Collin, and Ghislain Durif. Kernel operations on the gpu, with autodiff, without memory overflows. *Journal of Machine Learning Research*, 22(74):1–6, 2021.

[7] Peter Eastman. Pdbfixer. https://github.com/openmm/pdbfixer, 2013.

[8] Peter Eastman, Jason Swails, John D Chodera, Robert T McGibbon, Yutong Zhao, Kyle A Beauchamp, Lee-Ping Wang, Andrew C Simmonett, Matthew P Harrigan, Chaya D Stern, et al. Openmm 7: Rapid development of high performance algorithms for molecular dynamics. *PLoS computational biology*, 13(7):e1005659, 2017.

[9] Christophe Fares, Nils-Alexander Lakomek, Korvin FA Walter, Benedikt TC Frank, Jens Meiler, Stefan Becker, and Christian Griesinger. Accessing ns–μs side chain dynamics in ubiquitin with methyl RDCs. *J. Biomol. NMR*, 45(1):23–44, 2009.

[10] Jean Feydy, Thibault Séjourné, François-Xavier Vialard, Shun-ichi Amari, Alain Trouve, and Gabriel Peyré. Interpolating between optimal transport and mmd using sinkhorn divergences. In *The 22nd International Conference on Artificial Intelligence and Statistics*, pages 2681–2690, 2019.

[11] Nicholas J. Fowler, Adnan Sljoka, and Mike P. Williamson. A method for validating the accuracy of NMR protein structures. *Nature Communications*, 11:6321, 2020.

[12] Ronald J Gillespie. The vsepr model revisited. *Chemical Society Reviews*, 21(1):59–69, 1992.

[13] Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. *Advances in neural information processing systems*, 33:6840–6851, 2020.

[14] John B Ingraham, Max Baranov, Zak Costello, Karl W Barber, Wujie Wang, Ahmed Ismail, Vincent Frappier, Dana M Lord, Christopher Ng-Thow-Hing, Erik R Van Vlack, et al. Illuminating protein space with a programmable generative model. *Nature*, 623(7989):1070–1078, 2023.

[15] Robbie P Joosten, Fei Long, Garib N Murshudov, and Anastassis Perrakis. The pdb_redo server for macromolecular structure model optimization. *IUCrJ*, 1(4):213–220, 2014.

[16] John Jumper, Richard Evans, Alexander Pritzel, Tim Green, Michael Figurnov, Olaf Ronneberger, Kathryn Tunyasuvunakool, Russ Bates, Augustin Žídek, Anna Potapenko, et al. Highly accurate protein structure prediction with alphafold. *nature*, 596(7873):583–589, 2021.

[17] Wolfgang Kabsch. A solution for the best rotation to relate two sets of vectors. *Foundations of Crystallography*, 32(5):922–923, 1976.

[18] Diederik P. Kingma and Jimmy Ba. Adam: A method for stochastic optimization, 2015.

[19] Patrick Kunzmann, Jacob Marcel Anter, and Kay Hamacher. Adding hydrogen atoms to molecular models via fragment superimposition. *Algorithms for Molecular Biology*, 17(1):7, 2022.

[20] P Therese Lang, James M Holton, James S Fraser, and Tom Alber. Protein structural ensembles are revealed by redefining x-ray electron density noise. *Proceedings of the National Academy of Sciences*, 111(1):237–242, 2014.

[21] Xinli Liao, Dong Long, Da-Wei Li, Rafael Bruschweiler, and Vitali Tugarinov. Probing side-chain dynamics in proteins by the measurement of nine deuterium relaxation rates per methyl group. *J. Phys. Chem. B*, 116(1):606–620, 2012.

[22] SF Lienin, T Bremi, B Brutscher, R Brüschweiler, and RR Ernst. Anisotropic intramolecular backbone dynamics of ubiquitin characterized by NMR relaxation and MD computer simulation. *J. Am. Chem. Soc.*, 120(38):9870–9879, 1998.

[23] Stéphane G Mallat and Zhifeng Zhang. Matching pursuits with time-frequency dictionaries. *IEEE Transactions on signal processing*, 41(12):3397–3415, 1993.

[24] Milot Mirdita, Konstantin Schütze, Yoshitaka Moriwaki, Lim Heo, Sergey Ovchinnikov, and Martin Steinegger. Colabfold: making protein folding accessible to all. *Nature methods*, 19(6):679–682, 2022.

[25] Adam Paszke, Sam Gross, Soumith Chintala, Gregory Chanan, Edward Yang, Zachary DeVito, Zeming Lin, Alban Desmaison, Luca Antiga, and Adam Lerer. Automatic differentiation in pytorch. In *NIPS-W*, 2017.

[26] Edward Prince. *International Tables for Crystallography, Volume C: Mathematical, physical and chemical tables*. Springer Science & Business Media, 2004.

[27] Peter W Rose, Bojan Beran, Chunxiao Bi, Wolfgang F Bluhm, Dimitris Dimitropoulos, David S Goodsell, Andreas Prlić, Martha Quesada, Gregory B Quinn, John D Westbrook, et al. The rcsb protein data bank: redesigned web site and web services. *Nucleic acids research*, 39(suppl_1):D392–D401, 2010.

[28] Aviv A Rosenberg, Sanketh Vedula, Alex M Bronstein, and Ailie Marx. Seeing double: Molecular dynamics simulations reveal the stability of certain alternate protein conformations in crystal structures. *bioRxiv*, pages 2024–08, 2024.

[29] T Michael Sabo, Colin A Smith, David Ban, Adam Mazur, Donghan Lee, and Christian Griesinger. ORIUM: Optimized RDC-based Iterative and Unified Model-free analysis. *J. Biomol. NMR*, 58(4):287–301, 2014.

[30] Yang Shen and Ad Bax. Protein backbone and sidechain torsion angles predicted from nmr chemical shifts using artificial neural networks. *Journal of biomolecular NMR*, 56(3):227–241, 2013.

[31] Yang Song and Stefano Ermon. Generative modeling by estimating gradients of the data distribution. *Advances in neural information processing systems*, 32, 2019.

[32] ByteDance AML AI4Science Team, Xinshi Chen, Yuxuan Zhang, Chan Lu, Wenzhi Ma, Jiaqi Guan, Chengyue Gong, Jincai Yang, Hanyu Zhang, Ke Zhang, Shenghao Wu, Kuangqi Zhou, Yanping Yang, Zhenyu Liu, Lan Wang, Bo Shi, Shaochen Shi, and Wenzhi Xiao. Protenix - advancing structure prediction through a comprehensive alphafold3 reproduction. *bioRxiv*, 2025.

[33] KN Trueblood, H-B Bürgi, H Burzlaff, JD Dunitz, CM Gramaccioli, HH Schulz, U Shmueli, and SC Abrahams. Atomic dispacement parameter nomenclature. report of a subcommittee on atomic displacement parameter nomenclature. *Foundations of Crystallography*, 52(5):770–781, 1996.

[34] Pauli Virtanen, Ralf Gommers, Travis E Oliphant, Matt Haberland, Tyler Reddy, David Cournapeau, Evgeni Burovski, Pearu Peterson, Warren Weckesser, Jonathan Bright, et al. Scipy 1.0: fundamental algorithms for scientific computing in python. *Nature methods*, 17(3):261–272, 2020.

[35] Junmei Wang, Romain M Wolf, James W Caldwell, Peter A Kollman, and David A Case. Development and testing of a general amber force field. *Journal of computational chemistry*, 25(9):1157–1174, 2004.

[36] Stephanie A Wankowicz, Ashraya Ravikumar, Shivani Sharma, Blake Riley, Akshay Raju, Daniel W Hogan, Jessica Flowers, Henry van den Bedem, Daniel A Keedy, and James S Fraser. Automated multiconformer model building for x-ray crystallography and cryo-em. *Elife*, 12:RP90606, 2024.

[37] Marcin Wojdyr. Gemmi - library for structural biology [software]. https://github.com/project-gemmi/gemmi, 2021.