

Reshaping Biomolecular Structure Prediction through Strategic Conformational Exploration with HelixFold-S1

Lihang Liu^{1†}, Yang Liu^{1†}, Xianbin Ye^{1†}, Shanzhuo Zhang¹, Yuxin Li¹, Kunrui Zhu¹, Yang Xue¹, Jingbo Zhou¹, Xiaonan Zhang¹, Xiaomin Fang^{1*}

¹PaddleHelix Team, Baidu Inc., Shenzhen 518000, China.

*Corresponding author(s). E-mail(s): fangxiaomin666@gmail.com;

[†]These authors contributed equally to this work.

Abstract

Generating large ensembles of candidate conformations is standard for improving biomolecular structure prediction. Yet aimless sampling is inefficient and costly, producing many redundant conformations with limited diversity, particularly for complex multimeric assemblies. Here, we present HelixFold-S1, a guided planning approach specifically designed to enhance the structural prediction of biomolecular complexes by strategically targeting the most informative regions of conformational space to produce accurate conformations. For each complex, predicted inter-chain contact probabilities serve as a blueprint of the conformational space, guiding computational effort toward higher-probability, low-redundancy contacts that constrain structure generation. Across diverse biomolecular complex benchmarks, HelixFold-S1 achieves markedly higher structural accuracy than traditional unguided methods while reducing sampling requirements by an order of magnitude. Predicted contact probabilities also provide a rough indicator of prediction difficulty and sampling utility. These results demonstrate that guided planning reshapes conformational exploration and enables more efficient and accurate structural inference.

Keywords: Biomolecular structure prediction, Conformation space exploration, Planning-driven, Computational efficiency, Resource-aware modeling

Introduction

Biomolecular structure prediction lies at the heart of computational biology, enabling progress in drug discovery, protein engineering, and the study of molecular interactions. Deep learning has transformed this field, with recent breakthroughs [1–11] exemplified by the AlphaFold [1–3] and RoseTTAFold series [4, 5]. Despite these advances, accurate prediction of biomolecular complexes remains challenging.

To improve AlphaFold-like models of complexes, several specialized approaches have been proposed, such as denoising MSA profiles [12], integrating experimental constraints [13], or optimizing interaction interfaces through architectural advancements [14, 15]. Despite this progress, achieving consistent accuracy across diverse biomolecular systems remains a significant challenge. In parallel, improving the accuracy of a single prediction has proved difficult, and generating large ensembles of candidate conformations has therefore emerged as a simple yet powerful way to boost performance. The AlphaFold series improves precision through model ensembling, while AlphaFold3 [3] demonstrates that extensive sampling markedly enhances protein–antibody interface modeling. Recent studies [16–19] have generated thousands of candidates to expand conformational diversity. For instance, AFSample [16, 17] introduces stochasticity via dropout and random template masking to improve accuracy. Other strategies [20–23] diversify predictions through repeated MSA sub-sampling to probe alternative structural states.

Although generating large ensembles of candidate structures can improve predictive accuracy, it comes at substantial computational cost. For instance, performing a thousand samplings may require an entire GPU day. The main reason such massive sampling is needed is that current sampling methods are largely

42 unguided, leading to aimless exploration of conformational space. Specifically, predicting a biomolecular
43 complex is essentially a search problem within a vast inter-chain conformational space, where only a few
44 high-probability regions that correspond to potential binding modes are structurally plausible. The produced
45 conformations often cluster around a single local region of the conformational landscape (Fig. 1a), produc-
46 ing redundant conformations that offer limited additional information. Consequently, sampling efficiency
47 remains low. It is often unclear if further sampling is beneficial, as simpler targets gain little from extensive
48 ensembling, rendering much computation redundant. These challenges highlight the need for planning-driven
49 strategies that can first illustrate the most potentially informative regions, and then intelligently navigate
50 the conformational space to focus computational effort where it matters most, thereby achieving higher
51 accuracy at far lower computational cost.

52 Here, we introduce HelixFold-S1, a guided sampling strategy designed for biomolecular complexes that
53 explores conformational space with exceptional efficiency and precision. In contrast to traditional stochastic
54 search, HelixFold-S1 adopts a planning-based strategy to navigate the multi-modal landscape of complex
55 structures. This approach strategically targets the most informative regions, which are conceptualized as
56 high-probability modes within the conformational landscape, to facilitate the deliberate exploration of
57 diverse and plausible binding modes (Fig. 1b). For each target complex, we first predict inter-chain inter-
58 residue contact probabilities. Interpreting these distributions as a coarse-grained structural blueprint of the
59 conformational space provides a reduced representation that highlights where meaningful interaction pat-
60 terns are most likely to occur (Fig. 1c). This structural map then guides the sampling, with computational
61 effort concentrated in regions associated with higher-probability, low-redundancy contacts that act as spatial
62 constraints during structure generation. The resulting conformations are evaluated and ranked by their con-
63 fidence scores, enabling both accurate and resource-efficient modeling of complex biomolecular assemblies
64 (Fig. 1d).

65 We evaluated HelixFold-S1 across diverse biomolecular interfaces, including protein-protein, pro-
66 tein-antibody, protein-ligand, protein-RNA, and protein-DNA. Compared with conventional unguided
67 sampling, HelixFold-S1 substantially improves prediction accuracy at equivalent sampling effort, with the
68 largest gains observed for challenging protein-antibody complexes. Remarkably, it achieves comparable accu-
69 racy to traditional methods using an order of magnitude fewer sampling steps. The approach is generalizable
70 to other folding models to enhance structural precision. Predicted inter-chain contact probabilities correlate
71 strongly with structural difficulty and provide a practical guide for allocating sampling effort across tar-
72 gets. Guided sampling also improves exploration of the conformational space, producing more diverse and
73 higher-quality ensembles.

74 Results

75 Guided Sampling Strategy and Architecture of HelixFold-S1

76 The inference pipeline of HF-S1 comprises two stages. In stage 1, HF-S1 predicts a coarse blueprint of the
77 conformational space (Fig. 1c) to identify the most critical inter-chain interaction sites. An additional Con-
78 tact Prediction Module (CPM) is employed to estimate inter-chain inter-token contact probabilities based
79 on the pair representation, representing the likelihood that any two tokens are in spatial proximity, i.e., any
80 atom pair lies within 5\AA . This contact probability map provides a reduced, coarse-grained representation
81 of the high-dimensional conformational space, serving as a structural blueprint to guide subsequent con-
82 formation sampling. In stage 2, conformations are generated guided by the predicted blueprint. Contacts
83 are prioritized according to their predicted probabilities and sequentially selected from the blueprint. Each
84 selected contact is introduced as an additional constraint via the Contact Conditioning Module (CCM),
85 which processes these constraints and generates structures that attempt to satisfy them. This design aligns
86 with conditional structure prediction [24–26], where geometric priors bias the model toward specific states.
87 Our CCM uses the inter-chain contacts from Stage 1 as constraints to effectively steer subsequent struc-
88 ture prediction. By focusing on contacts with higher predicted probabilities, the model is more likely to
89 produce multiple high-quality conformations that capture the most informative interactions. To reduce sam-
90 pling redundancy, we employ a strategy called redundancy contact pruning (RCP). Once a conformation
91 is generated using a particular contact, any contacts that are already satisfied within the resulting struc-
92 ture are excluded from further selection. All generated conformations are subsequently ranked according
93 to model confidence scores, with the top-ranked prediction designated as the final output. Notably, RCP
94 is an optional optimization; it enhances efficiency by sequentially pruning redundant contacts but can be
95 bypassed for fully independent parallel generation. This allows HelixFold-S1 to balance high throughput via
96 parallelism with refined, low-redundancy sampling, depending on available resources.

97 The architecture of HelixFold-S1 (HF-S1) builds upon that of HelixFold3 (HF3) [9], which reproduces
98 the biomolecular structure prediction capabilities of AlphaFold3 (AF3) [3], and further extends its capacity
99 to more intelligently explore conformational space through the addition of two contact-centric modules: the
100 CPM and the CCM (Fig. 1e). A contact is considered to exist between two tokens if any pair of atoms
101 from the corresponding tokens lies within 5Å. The CPM computes a contact probability matrix across
102 all token pairs using the pair representations produced by the Pairformer, a specialized attention-based
103 module designed to capture long-range pairwise token dependencies. The CCM, incorporated into the Input
104 Embedder, enables the model to learn contact constraints and generate structures conditioned on these
105 constraints.

106 During inference, the two modules are employed sequentially: in the first stage (Fig. 1c), the CPM esti-
107 mates the overall contact landscape, and in the second stage (Fig. 1d), the CCM generates conformations
108 guided by the predicted contacts. In contrast, during the training process, the model is optimized within a
109 multi-task framework that alternates between two primary objectives: the contact prediction task (utiliz-
110 ing the CPM) and the contact-conditioned structure generation task (utilizing the CCM). This dual-task
111 optimization allows the model to jointly learn inter-chain interactions and effectively leverage them during
112 inference.

113 Improved Structural Accuracy across Complex Types

114 To systematically evaluate the performance of HF-S1, we constructed a test set comprising protein–antibody
115 (n=221), protein–protein (n=198), protein–ligand (n=238), protein–RNA (n=177), and protein–DNA
116 (n=254) complexes, collected between 2022.01.01 and 2024.12.31 from the RCSB PDB [27]. To minimize
117 overlap with the training data, all test samples were selected to have low sequence identity to the training
118 set. Sequences were clustered by similarity [28], and one representative per cluster was randomly chosen to
119 ensure diversity and reduce redundancy. For protein–ligand complexes, any ligands that appeared in the
120 training data were further excluded.

121 We compared AF3-like models, namely HF3 [9], Protenix [29], and Chai-1 [30], under different sampling
122 strategies. For HF3, we examined two sampling-augmented variants: one following the AFSample strategy
123 [16] (HF3 w/ AFSample) and another implementing our guided planning approach (HF-S1). To assess cross-
124 model compatibility, we integrated HF-S1-predicted contacts into Protenix and Chai-1 (Protenix/Chai-1
125 w/ S1). This evaluation prioritizes sampling generalizability across architectures over absolute rankings.
126 To manage the computational cost of benchmarking multiple external models, these variants were run in
127 fully parallel mode without RCP; crucially, this trade-off does not compromise the fundamental assessment
128 of our guided sampling strategy. Because HF-S1 may employ a distinct definition of contacts from those
129 folding models, the baseline models may not reach their best performance under this evaluation. Moreover,
130 Chai-1 currently supports contact constraints only for protein–protein complexes, and results are therefore
131 reported for protein–antibody and protein–protein systems only. For all methods, we generated 25 predicted
132 structures per target and ranked them using each method’s confidence score.

133 Top-5 precision (Fig. 2a) reveals that HF-S1 consistently outperforms both HF3 and HF3 w/ AFSam-
134 ple. Improvements are most significant for protein–antibody complexes, exceeding 55% over HF3 and 33%
135 over HF3 w/ AFSample. This is followed by protein–protein and protein–RNA/DNA interfaces, while
136 protein–ligand binding sites show the smallest gains. This likely reflects the structural diversity of antigen-
137 binding sites, which makes conventional sampling less effective at capturing native-like geometries, whereas
138 HF-S1’s guided sampling strategy increases the likelihood of generating correct interfacial contacts. For
139 protein–protein and protein–RNA/DNA interfaces, HF-S1 provides moderate but clear improvements. By
140 enforcing constraints on individual interfacial contacts, guided sampling helps capture correct local interac-
141 tions, although single-contact constraints alone are not sufficient to fully determine the global geometry of
142 these larger, more complex interfaces. Finally, protein–ligand binding sites are highly constrained, so while
143 HF-S1 can still improve predictions, the relative gains are smaller compared with other complex types.
144 Moreover, across different folding models, including HF3, Protenix, and Chai-1, incorporating the guided
145 sampling strategy of HF-S1 leads to consistent performance gains, underscoring its broad effectiveness. In
146 Fig. 2b, we further demonstrate that the performance advantage of HF-S1 over HF3 is consistently observed
147 across different Top-K precisions of the sampled conformations. The improvements are pronounced for Top-
148 1, Top-5, and Top-25, underscoring HF-S1’s robustness across multiple ranking thresholds, although the
149 Top-1 gains exhibit greater variability, potentially reflecting sensitivity to the selection of conformations
150 based on the confidence score.

151 Extensive sampling is particularly effective for protein–antibody interfaces [29], and HF-S1 shows the
152 largest gains on these targets. To systematically evaluate the benefit of extensive sampling on these interfaces,

153 we curated a dataset of 74 complexes released after 2024 and with residues less than 800 for efficiency,
154 ensuring no overlap with Boltz-2’s training data [31]. Sampling was conducted with top-performing methods
155 evaluated 1,000 times per complex, while others were sampled 100 times per complex. Top-1 precision on
156 protein–antibody interfaces (Fig. 2c) generally improves with increased sample size across all methods,
157 particularly for those incorporating advanced strategies such as HF-S1 or AFsample. Importantly, HF-S1
158 achieves comparable precision to HF3 with AFsample using only 10 samples. This indicates that 10 HF-
159 S1 samples are sufficient to reach the precision level that HF3 with AFsample requires 1,000 samples to
160 achieve, representing just 1% of the computational cost. To further assess performance across different
161 antibody types, Fig. 2d presents Top-1 precision for HF-S1 and HF3 with AFsample across canonical
162 antibodies, nanobodies, and scFv based on the full 1,000-sample evaluation. HF-S1 demonstrates consistent
163 improvements in Top-1 precision over HF3 w/ AFsample across all categories.

164 HF-S1 introduces two key components built on HF3: the Contact Prediction Module (CPM) and Con-
165 tact Conditioning Module (CCM). To evaluate the CPM, we assessed the accuracy of its predicted contact
166 probability matrices. For each target, token pairs corresponding to true contacts in experimental structures
167 were treated as positives, and all others as negatives. We then computed the area under the precision–recall
168 curve (AUPRC) between predicted and ground-truth contact maps, averaging across all targets. As a base-
169 line, a posterior contact probability matrix for HF3 was derived by aggregating sampled conformations,
170 with each entry defined as the inverse of the minimum inter-residue distance. Compared with HF3, HF-S1
171 consistently produced more accurate contact probabilities across diverse molecular types (Fig. 2e). Pro-
172 tein–antibody complexes showed the lowest AUPRC, reflecting the difficulty of this prediction scenario,
173 whereas protein–ligand complexes achieved the highest AUPRC, indicating a relatively simpler task. This
174 aligns with Fig. 2a: protein–antibody interfaces benefit most from guided sampling, while protein–ligand
175 interfaces show smaller relative gains. To examine the CCM, we evaluated whether HF-S1–predicted con-
176 formations adhered to specified contact constraints. We defined the contact satisfaction rate as the fraction
177 of predicted structures in which the given contacts were realized. Across most test cases, HF-S1 achieved
178 satisfaction rates above 70% (Fig. 2f), demonstrating its ability to effectively incorporate contact priors into
179 structure prediction.

180 Fig. 2g illustrates a representative case. Two inter-chain residue pairs (R43–N26 and V103–L28) were
181 selected from HF-S1’s predicted contact map and applied as spatial restraints during folding, yielding
182 two distinct binding conformations that both satisfied the defined contact distances (R43–N26 \approx 3.1 Å;
183 V103–L28 \approx 4.2 Å).

184 Contact Probability as an Indicator of Prediction Difficulty and Sampling 185 Utility

186 The contact probability matrix predicted by HF-S1 serves a dual role: it not only informs the structural
187 sampling strategy but also provides insight into the intrinsic difficulty of the structure prediction task, as
188 well as the potential benefits of multiple sampling.

189 We first show that predicted contact probabilities can serve as a proxy for estimating target difficulty.
190 Specifically, we defined the target-level contact probability as the maximum value in the predicted contact
191 probability matrix for each target and analyzed its relationship with the Top-5 precision of HF-S1 using 25
192 samples (Fig. 3a). Targets were stratified into high, medium, and low groups, revealing a strong correlation
193 across datasets: lower-probability targets consistently yielded poorer predictions, while higher-probability
194 targets were predicted with high precision. This indicates that low contact probability generally corre-
195 sponds to more difficult targets, which often require additional sampling to achieve accurate predictions.
196 When grouped by complex type (Fig. 3b), a similar trend is observed: Protein–protein targets typically
197 show higher target-level contact probabilities, while protein–antibody complexes cluster in the lower range,
198 indicating greater structural uncertainty. Protein–ligand complexes consistently exhibit high contact prob-
199 abilities, suggesting that HF-S1 can often localize ligand binding sites with high confidence. In contrast,
200 protein–RNA and protein–DNA complexes display a broader distribution, reflecting greater variability in
201 prediction difficulty across these categories.

202 We next examined whether target-level contact probabilities predict sampling utility. Targets were
203 grouped based on their target-level predicted contact probability, and we analyzed the precision improve-
204 ments of the best-performing multi-sample prediction relative to single-sample predictions (Fig. 3c). Targets
205 with intermediate contact probabilities achieved the greatest improvements, while those with low or high
206 probabilities saw more modest gains. Notably, although lower-probability targets did benefit from sampling,
207 their improvements were generally smaller than those in the intermediate group. This trend is intuitive: for
208 targets with high contact probabilities, accurate structures can often be recovered from a single sample,

209 leaving limited room for further enhancement. For lower-probability targets, the predicted contact maps are
210 weak across the board, suggesting that a much larger number of samples may be needed to identify accurate
211 structures. In contrast, intermediate cases offer partial yet informative contact signals, enabling the model
212 to better explore the structural landscape and refine its predictions through sampling. As protein–ligand
213 complexes consistently exhibit high target-level contact probabilities, they did not follow this trend.

214 We further examined performance across different contact probability groups by comparing HF-S1
215 (guided sampling) with HF3 (aimless sampling) (Fig. 3d). HF-S1 achieves the largest gains in the lower-
216 probability group, while improvements are smaller in the higher-probability group. This trend suggests
217 that targets with low contact probabilities are generally more challenging, and thus benefit more from
218 guided exploration. In contrast, high-probability targets are easier to fold accurately, leaving less room for
219 improvement.

220 Finally, we analyzed predicted contact probabilities across residue–residue pairs in the protein–protein
221 dataset. Left panel of Fig. 3e: the averaged contact probability matrix across all residue pairs, with the
222 pairs showing the highest probabilities indicated and those with the lowest probabilities highlighted. Right
223 panel of Fig. 3e: compositional differences between the top 10 and bottom 10 pairs ranked by mean proba-
224 bility. Hydrophobic pairs such as Leu–Leu and Met–Met tend to show higher predicted probabilities, which
225 could reflect their relatively frequent occurrence in tightly packed, energetically favorable environments.
226 Conversely, polar or charged residues, such as Asp, Glu, and Lys, often exhibit lower probabilities, perhaps
227 due to their general solvent exposure or context-specific interaction tendencies. These patterns indicate
228 that contact prediction models may capture fundamental physicochemical preferences underlying residue
229 interactions [32–35].

230 Guided Sampling Improves Exploration of Conformation Space

231 To test if guided sampling improves conformational exploration, we compared HF-S1 (guided planning) with
232 HF3 (aimless sampling). We calculated the standard deviation (std) of precision scores per target as a proxy
233 for structural diversity (Fig. 4a), where higher values indicate a broader exploration of the landscape. HF-
234 S1 generally shows higher std values, with data concentrated in the upper-left region of the plot, reflecting
235 more diverse ensembles. Targets were then grouped based on the difference in std between HF-S1 and
236 HF3 (Fig. 4b): for protein–ligand complexes, a difference ≥ 1 indicates HF-S1 std higher, ≤ -1 indicates
237 HF3 std higher, and intermediate values are comparable; for all other complex types, the threshold is 0.1.
238 Across complex types, most targets fall into the HF-S1 std higher category, particularly protein–antibody,
239 protein–RNA, and protein–DNA complexes, where the HF-S1 std higher category accounts for roughly twice
240 as many targets as the HF3 std higher category. This enhanced diversity suggests that guided planning
241 effectively directs sampling toward multiple plausible interaction geometries.

242 To examine how guided planning shapes the sampling trajectory, we tracked the average precision of
243 all conformations accumulated up to each sampling step K (Fig. 4c). For HF3, precision remains nearly
244 constant across steps, as the order of sampled conformations is effectively random. In contrast, HF-S1
245 exhibits a gradual decline in precision as sampling proceeds, reflecting its design to explore progressively
246 less probable contact configurations. The higher precision observed in the early sampling stages arises from
247 the use of higher-probability contacts as structural constraints, which guide the generation of more accu-
248 rate conformations. Notably, HF-S1’s initial conformations already outperform HF3 across most interface
249 types, excluding protein–DNA complexes. This demonstrates that leveraging contact probabilities as struc-
250 tural constraints not only improves efficiency but also prioritizes structurally meaningful regions of the
251 conformational landscape.

252 We compared the aimless sampling of HF3 with the guided strategy of HF-S1 using two representative
253 examples (Fig. 4d): the Gel4–Nb4 nanobody complex (PDB ID: 8pe1) [36] and the Cas12m2–crRNA protein-
254 RNA heterodimer (PDB ID: 8hhm) [37]. Analysis of predicted confidence and interface precision shows
255 that HF-S1 consistently generates a broader spectrum of structures. These samples span a wide range of
256 confidence and precision values, including a higher fraction of high-accuracy predictions. In contrast, HF3
257 samples cluster within narrower ranges, indicating more limited structural diversity and fewer high-quality
258 candidates.

259 Discussion

260 Extensive sampling is standard for structural accuracy, yet unguided methods are computationally inefficient
261 and lack insight into target difficulty. We propose that intelligent exploration of the conformational land-
262 scape enables efficient resource allocation. By directing sampling toward informative regions, our approach

263 enhances both efficiency and accuracy. Furthermore, predicted contact probabilities serve as a proxy for
264 target difficulty, guiding users on the required sampling scale for convergence.

265 Beyond efficiency, this Search-then-Filter paradigm avoids local minima and consistently achieves
266 superior Top-1 accuracy. For challenging targets, such as multi-epitope antibodies, cryptic pockets, or
267 transitioning GPCRs, this prioritized ensemble ensures rare but biologically active conformations are not
268 overlooked, providing value far beyond a single structure.

269 Despite these advances, assessing structural confidence remains challenging. Current metrics frequently
270 misidentify the most accurate structures. High-confidence predictions may lack accuracy, while low-
271 confidence sets often contain near-native conformations (Supplementary Figure 3). Refining these via
272 sophisticated scoring or ensemble-based calibration could optimize candidate selection and prevent over-
273 looking high-quality structures. Moreover, guided planning can exhibit diminishing returns. Greedy contact
274 selection from highest to lowest probability may progressively sample less informative regions (Fig. 4c).
275 Implementing early-stopping based on contact saturation or plateauing structural quality could sustain
276 sampling efficiency. Finally, maintaining diversity is critical. Current sequential contact exclusion reduces
277 redundancy but remains computationally serial. Although this strategy is effective (Supplementary Figure
278 4), more proactive methods are needed to preserve diversity more efficiently.

279 Intelligently navigating the conformational space remains a critical challenge. Sophisticated strate-
280 gies, such as Monte Carlo Tree Search or reinforcement learning, could dynamically balance exploration
281 and exploitation. Such methods promise to identify high-quality, diverse conformations more efficiently,
282 enhancing the robustness of structure prediction across complex biomolecular targets.

283 Methods

284 Model Architecture

285 HF-S1 builds upon the HF3 architecture and is designed to support two complementary tasks: inter-chain
286 contact prediction and contact-conditioned structure prediction. To this end, HF-S1 introduces two addi-
287 tional components: the Contact Prediction Module and the Contact Conditioning Module, which extend the
288 base architecture to enable contact-level reasoning and constraint-based structure generation, respectively.

289 The contact prediction task aims to estimate the inter-chain contact distribution of a given protein
290 complex. Various input features—including sequence, multiple sequence alignment (MSA), and template
291 information—are first encoded and then processed by a Pairformer module to generate single and pair
292 representations. These pairwise representations are subsequently passed to a dedicated Contact Prediction
293 Module (CPM), which includes a Pairformer Stack and performs a binary classification task on each element
294 of the pairwise representation to output contact probabilities for each inter-chain token pair. A contact
295 is defined as the presence of any atom pair between two tokens within 5\AA in 3D space. The resulting
296 pairwise contact probability matrix captures the inter-chain contact distribution of the complex and serves
297 as an informative intermediate representation. By operating in the simplified space of contacts—rather
298 than directly in the complex, high-dimensional structural space—the model achieves greater computational
299 efficiency and facilitates the contact-conditioned structure prediction task.

300 The contact-conditioned structure prediction task introduces a Contact Conditioning Module (CCM) to
301 incorporate external contact constraints. These constraints are represented as a binary matrix $\{c_{ij}\}$, where
302 $c_{ij} \in \{0, 1\}$ indicates whether token pair (i, j) is in contact (1 if any atom pair is within 5\AA , and 0 otherwise).
303 This matrix is projected through a linear layer and then fused into the pairwise activations within the Input
304 Embedder of the model. During training, for each complex, 0–10 inter-chain contacts are randomly sampled
305 from the contact set extracted from its ground-truth structure and provided as input. During inference,
306 contact constraints are sampled from the contact probability matrix produced by the contact prediction
307 task; the specific sampling strategy is described in a later section. The model learns to utilize the provided
308 contact information to enhance structure prediction accuracy.

309 Training Regime

310 Parameters of the newly introduced Contact Prediction Module and Contact Conditioning Module are
311 randomly initialized, while the remaining parts of the model inherit weights from the pre-trained HF3.
312 Fine-tuning is performed using the same training dataset as HelixFold3, which includes Protein Data
313 Bank (PDB) [27] structures released before September 30, 2021, supplemented with self-distillation data to
314 enhance generalization. The training follows a three-stage fine-tuning strategy: the first stage focuses on the

315 contact-conditioned structure prediction task to improve complex structure accuracy with inter-chain con-
 316 tact constraints; the second stage adds the contact prediction task, jointly optimizing the model for both
 317 tasks; the third stage follows the same setting of the second stage but extends to larger crop size.

318 In the first fine-tuning stage, the model is trained exclusively on the contact-conditioned structure
 319 prediction task. For each training sample, inter-chain contacts are extracted from experimentally determined
 320 complex structures to form a ground-truth contact set \mathcal{C} . This set contains all inter-chain token pairs where
 321 at least one atom from each token lies within 5Å in three-dimensional space. Contact conditioning is applied
 322 with 70% probability: 1–10 token pairs are uniformly sampled from \mathcal{C} and provided as binary contact
 323 constraints. In the remaining 30% of samples, no contact constraints are used, which helps maintain the
 324 model’s ability to predict structures without external guidance. While trained on individual experimental
 325 structures, the CPM learns underlying interaction probabilities by aggregating structural patterns from
 326 the vast diversity of homologous complexes and motifs in the training set. This allows the module to map
 327 a protein pair to a multi-modal probability landscape during inference, representing a global statistical
 328 distribution of plausible binding modes rather than a single deterministic state.

In the second and third fine-tuning stages, the model is optimized within a multi-task framework that
 simultaneously addresses both the contact-conditioned structure prediction task and the contact prediction
 task. The latter aims to estimate the probability that each inter-chain token pair is in atomic contact,
 serving as a basis for generating contact constraints during inference. During the specific optimization of
 the contact prediction task in these stages, the Contact Conditioning Module (CCM) is intentionally not
 activated; this is a deliberate design to ensure the CPM learns to infer intrinsic interaction patterns directly
 from raw features (MSA and sequences) without the "leakage" of ground-truth constraints. To supervise
 this task, a binary classification loss is applied over all inter-chain token pairs:

$$\mathcal{L}_{\text{contact}} = \frac{1}{|\mathcal{P}|} \sum_{(i,j) \in \mathcal{P}} \text{cross_entropy}(p_{ij}^{\text{contact}}, y_{ij}^{\text{contact}}).$$

329 Here, \mathcal{P} denotes the set of all token pairs (i, j) such that token i and token j belong to different chains.
 330 p_{ij}^{contact} is the predicted probability of contact between tokens i and j . $y_{ij}^{\text{contact}} = 1$ if $(i, j) \in \mathcal{C}$ (i.e., the
 331 token pair is in contact), and $y_{ij}^{\text{contact}} = 0$ otherwise. During training, half of the samples are used for the
 332 contact-conditioned structure prediction task, following the protocol established in the first stage, while
 333 the other half are dedicated to training the contact prediction task, which guides the model to estimate
 334 inter-chain contact probability distributions.

The loss function largely follows the original AF3/HF3 formulation, with an additional contact loss term
 introduced during fine-tuning:

$$\mathcal{L}_{\text{loss}} = \alpha_{\text{confidence}} \mathcal{L}_{\text{confidence}} + \alpha_{\text{diffusion}} \mathcal{L}_{\text{diffusion}} + \alpha_{\text{distogram}} \mathcal{L}_{\text{distogram}} + \alpha_{\text{contact}} \mathcal{L}_{\text{contact}},$$

335 with hyperparameters $\alpha_{\text{confidence}} = 0.01$, $\alpha_{\text{diffusion}} = 4$, and $\alpha_{\text{distogram}} = 0.3$. The contact loss coefficient
 336 α_{contact} is set to 1 during training samples used for the contact prediction task and 0 during samples used for
 337 the contact-conditioned structure prediction task. The definitions of all other loss terms remain consistent
 338 with those in AF3.

339 All stages use the Adam optimizer [38] with parameters $\beta_1 = 0.9$, $\beta_2 = 0.95$, and $\epsilon = 10^{-8}$, and a learning
 340 rate of 2×10^{-4} . The mini-batch size is fixed at 128 for all stages. The first fine-tuning stage consists of
 341 10,000 training steps with a crop size of 384. The second fine-tuning stage extends to 20,000 steps, also with
 342 a crop size of 384. The third stage continues training for an additional 3,000 steps with an increased crop
 343 size of 640. We conducted model training on 128 Nvidia A100 GPUs, with a total duration of approximately
 344 10 days.

345 Inference Regime

346 The inference process of HF-S1 (illustrated in Fig. 1a) consists of three stages: Contact Prediction, Contact
 347 Sampling, and Contact-Guided Structure Prediction and Ranking.

348 In the Contact Prediction stage, the contact prediction task is executed five times to reduce prediction
 349 variance, producing five contact probability matrices. These matrices are averaged element-wise to generate
 350 the final contact probability matrix, where each element p_{ij}^{contact} represents the predicted contact probability
 351 between tokens i and j . Only inter-chain contact probabilities are retained, with intra-chain contacts set to
 352 zero. For protein–antibody complexes, contact sampling is performed exclusively between the antigen chain
 353 and each antibody chain (heavy and light), excluding contacts between heavy and light chains.

354 In the Contact Sampling stage, inter-chain contacts are selected sequentially in descending order accord-
 355 ing to their predicted contact probabilities. Each selected contact is used as a binary constraint in the
 356 subsequent structure prediction step to generate diverse candidate structures. To improve sampling effi-
 357 ciency and avoid redundancy, two strategies are adopted: redundant contact pruning and enriched sampling
 358 of previously identified contact sets. We denote the sets of contacts extracted from previously predicted
 359 structures as C_1, C_2, \dots , where each C_k corresponds to the contacts obtained from the k -th predicted struc-
 360 ture, following the same ground-truth extraction method described earlier. During sampling, redundant
 361 contact pruning excludes any candidate contact that overlaps with contacts already present in the union
 362 of all previously sampled sets $\bigcup_{i=1}^{k-1} C_i$. Here, overlapping means the candidate contact appears in any pre-
 363 viously extracted contact set. This ensures that each newly sampled contact introduces novel constraints
 364 and helps maintain diversity among sampled structures. As sampling progresses, the predicted contact
 365 probabilities of remaining candidates naturally decrease. When these probabilities fall below a threshold
 366 (set as $0.2 \cdot \max_{i,j} p_{ij}^{contact}$), the benefit of exploring new low-confidence contacts diminishes. At this point,
 367 instead of sampling new contacts, the algorithm enriches sampling by iterating through the existing contact
 368 sets C_1, C_2, \dots in order. Contacts are drawn from these sets cyclically to further exploit high-confidence
 369 information until the total sampling budget S is reached.

In the Contact-Guided Structure Prediction and Ranking stage, each sampled contact is treated as a
 binary constraint and passed into the contact-conditioned structure prediction task to generate a candidate
 structure. A confidence score, named ranking_confidence, is computed for each structure, and the final
 prediction is selected as the one with the highest confidence among all candidates. Drawing inspiration
 from the AF3, we define the confidence score as a weighted average of the pTM and ipTM scores, with an
 additional penalty term for structural clashes. The score is computed as follows:

$$\text{ranking_confidence} = 0.2 \cdot \text{pTM} + 0.8 \cdot \text{ipTM} - 1.0 \cdot \text{has_clash},$$

370 where pTM represents the predicted TM-Score for the full complex, indicating the confidence for overall
 371 structural accuracy. ipTM represents the interface predicted TM-Score for the full complex, focusing on the
 372 accuracy of interfacial interactions. has_clash is a binary term indicating the presence of obvious clashes
 373 between polymer chains in the predicted structure. Detailed definitions of pTM, ipTM, and has_clash can
 374 be found in the AF3 paper [3].

375 We adopt consistent inference settings across structure prediction tasks, including our method and the
 376 baselines Boltz-2[31], Protenix[29], and Chai-1[30]. Each prediction is refined using 10 recycling iterations
 377 and 200 diffusion steps, where the diffusion module is run once to generate a single structure per input. In
 378 the corresponding figures, lines represent these average outcomes, while shaded areas indicate the variability
 379 between the two runs. Notably, the inference configuration for HF3 w/ AFsample adopts a more sophisticated
 380 multi-setting approach, according to the AlphaFold settings used in AFsample [16]. The complete inference
 381 specifications for HF3 with AFsample integrate three distinct hyperparameter settings as detailed in Table
 382 1. Protenix w/ S1, Chai-1 w/ S1, and Boltz-2 w/ S1 follow the same workflow as HF-S1, differing only in the
 383 backbone structure prediction module employed during the final stage. Contact probability outputs from
 384 HF-S1 are ranked from highest to lowest, and the corresponding contact pairs are incorporated as external
 385 geometric constraints to guide structure generation. It should be noted that Chai-1 w/ S1 is restricted
 386 to protein-protein/antibody datasets, as it only accepts residue-residue contact constraints as input. All
 387 methods construct MSAs using their respective built-in sequence search tools.

388 Evaluation Data

389 Evaluation sets for protein-protein, protein-ligand, protein-RNA, and protein-DNA interfaces were con-
 390 structed from all PDB entries released between May 1, 2022 and December 31, 2024, with each structure
 391 expanded to Biological Assembly 1. Interfaces were defined as pairs of entities with a minimum heavy-atom
 392 distance below 5 Å. Protein-antibody complexes were sourced from SABDab [39] within the same date
 393 range, using symmetric units instead of Biological Assembly 1.

394 For targets collected from the PDB, complexes with resolution worse than 4.5 Å or exceeding 1400 tokens
 395 under our tokenization scheme were removed. Polymer-polymer interfaces were excluded if both polymers
 396 shared more than 40% sequence identity with two chains from the same PDB entry in the training set. For
 397 protein-ligand interfaces, the following criteria were applied: (1) only ligands with CCD codes absent from
 398 the training set were retained; (2) covalently bound ligands, including those involved in glycosylation, were
 399 excluded; (3) ligands containing five or fewer atoms or occurring in ten or more PDB entries were removed;
 400 (4) only ligands with molecular weights between 100 and 900 Da were retained; (5) ligands were required
 401 to exhibit a *ranking_model_fit* score of at least 0.5, as reported in the RCSB structure validation dataset,

402 indicating above-median model quality for X-ray crystallographic structures [40]; and (6) binding pockets
403 were required to include between 5 and 100 protein residues within 5 Å of the ligand.

404 We clustered the remaining targets by grouping proteins with nine or more residues at 40% sequence
405 identity, while nucleic acids and proteins with nine or fewer residues were clustered at 100%, using MMseqs2
406 with a minimum coverage of 80% and default clustering mode. Each interface was assigned a binary, order-
407 independent cluster ID based on entity pairs—(polymer1_cluster, polymer2_cluster) for polymer–polymer
408 interfaces and (polymer_cluster, ligand_CCD-code) for protein–ligand interfaces. Evaluation was performed
409 on one representative entry per cluster.

410 Protein–antibody complexes were sourced from SAbDab [39], including only those with resolution better
411 than 9 Å, and containing antigen chains. The antigen sequences were grouped into clusters based on a 40%
412 sequence identity threshold. Subsequently, we retained only those clusters that contained no cases released
413 prior to September 30, 2021. Further filtering was applied to select cases released within the period from
414 May 1, 2022, to December 31, 2024. Ultimately, one case was chosen from each of the remaining clusters
415 to constitute the evaluation set for protein-antibody analysis. For efficiency considerations in the extensive
416 sampling test (as shown in Fig. 2c), we exclusively selected samples released between January 1, 2024, and
417 December 31, 2024, and filtered out those with over 800 residues, ultimately retaining 74 samples.

418 Evaluation Metrics

419 To evaluate structure prediction performance across different interaction types, we adopt distinct metrics
420 tailored to the characteristics of each molecular interface.

421 Protein–protein complexes, including protein–antibody interactions, are evaluated using DockQ [41],
422 which integrates interface RMSD, FNAT, and FNAS to provide a reliable summary of interface quality. For
423 protein–antibody complexes specifically, all antibody chains are treated collectively as the “ligand”, and
424 DockQ is computed over the interface between the antibody and the rest of the complex using the DockQ
425 v1 implementation.

426 Nucleic acid–protein interfaces, including both protein–RNA and protein–DNA complexes, are assessed
427 using interface LDDT (iLDDT) [42], computed over atom pairs across different chains within a 30Å inclusion
428 radius to accommodate the larger and more diffuse interaction footprints characteristic of nucleic acids.

429 Protein–ligand complexes are evaluated using pocket-aligned RMSD, which measures ligand pose accu-
430 racy after aligning the predicted structure to the binding pocket of the ground truth. The pocket is defined
431 as all heavy atoms within 10Å of any heavy atom of the ligand in the ground truth structure, restricted to
432 the primary protein chain—identified as the chain containing the most atoms within this radius. The C^α
433 atoms of this pocket are used to perform a least-squares alignment between predicted and reference struc-
434 tures. After alignment, a symmetry-corrected ligand RMSD is computed over all heavy atoms of the ligand
435 using RDKit’s Chem.rdMolAlign.CalcRMS [43], which aligns the ligands while accounting for molecular
436 symmetry before computing the final deviation.

437 The precision of the Top-K ranked structures is defined as the highest accuracy achieved among the top
438 K structures ordered by their confidence score.

439 Data Availability

440 To train HelixFold-S1, PDB can be downloaded at [https://www.rcsb.org/docs/programmatic-access/
441 file-download-services](https://www.rcsb.org/docs/programmatic-access/file-download-services) and AlphaFold Protein Structure Database as the distillation dataset can be down-
442 loaded at <https://ftp.ebi.ac.uk/pub/databases/alphafold/v2/>. The test set are filtered and clustered from
443 PDB with conditions detailed in Methods. The protein-antibody complexes for test can be downloaded
444 at <https://opig.stats.ox.ac.uk/webapps/sabdab-sabpred/sabdab/>. Detailed data processing procedures,
445 including filtering, clustering, and dataset construction, are described in the Methods section.

446 Code Availability

447 The source code, trained weights, and inference scripts for HelixFold-S1 are publicly available at [https://
448 github.com/PaddlePaddle/PaddleHelix/tree/dev/apps/protein_folding/HelixFold-S1](https://github.com/PaddlePaddle/PaddleHelix/tree/dev/apps/protein_folding/HelixFold-S1). To ensure long-term
449 availability, the version of the code at the time of acceptance has been archived on Zenodo (DOI: <https://doi.org/10.5281/zenodo.8202943>) [44].
450

451 Acknowledgments

452 This work was supported by the National Science and Technology Major Project (2023ZD0120803).

⁴⁵³ **Author Contributions**

⁴⁵⁴ X. Fang and X. Zhang led the research. L. Liu and X. Fang contributed technical ideas. L. Liu, Y. Liu, and
⁴⁵⁵ X. Ye developed the proposed method. S. Zhang, Y. Li, K. Zhu, Y. Xue, and J. Zhou developed analytics.
⁴⁵⁶ L. Liu, X. Fang, Y. Liu, and X. Ye wrote the paper.

⁴⁵⁷ **Competing Interests**

⁴⁵⁸ The authors declare no competing interests.

Settings	Templates	Dropout	Recycles	Ratio %
setting-1	Yes	Yes	3	30
setting-2	No	Yes	3	30
setting-3	No	Yes	9	40

Table 1 Inference configurations of HF3 w/ AFsample. The term *Templates* indicates whether structural templates were employed. *Dropout* denotes whether the dropout mechanism was activated. *Recycles* signifies the number of recycling operations utilized, with a default value of 3. *Ratio* represents the proportion that this particular setting occupies within the entire sampling process.

Fig. 1— **Overall framework of HelixFold-S1.** **a, b,** Conceptual comparison of sampling paradigms: **a,** Aimless exploration in conventional conformation sampling, often concentrated within limited regions of the conformational space; **b,** Guided exploration in planning-based sampling, illustrating the strategic advantage of directing sampling toward diverse conformational states. **c, d,** Two-stage inference of HelixFold-S1: **c,** stage 1 predicts a blueprint of the conformational space; **d,** stage 2 generates conformations based on this blueprint. **e,** Network architecture of HF-S1, which incorporates a contact conditioning module and a contact prediction module into the HelixFold3’s model architecture.

Fig. 2— **Structural performance and module evaluation of HelixFold-S1 across multiple complex types.** **a,** Top-5 structural precision among 25 sampled conformations for various folding models. Benchmark complexes were collected from the RCSB PDB between January 1, 2022, and December 31, 2024, including protein–antibody ($n = 221$), protein–protein ($n = 198$), protein–ligand ($n = 238$), protein–RNA ($n = 177$), and protein–DNA ($n = 254$) complexes. Statistical significance was determined using a two-sided paired t-test. On the protein–protein dataset, the difference between HF-S1 and HF3 was significant ($p \approx 0.0029$); for all other datasets, the improvements were highly significant ($p < 10^{-6}$). **b,** Relative improvements of HelixFold-S1 over HelixFold3 across different complex types in Top-1, Top-5, and Top-25 precision. **c,** Comparison of HF-S1 with various baseline models on protein–antibody complexes released in 2024 ($n=74$), showing that extensive sampling progressively improves structural accuracy. Some methods were sampled 1000 times, while others used 100 samples. **d,** Performance comparison between HF3 w/ AFsample and HF-S1 across different antibody types on the same 2024 protein–antibody complexes ($n=74$), using 1000 samples. **e,** Performance of the contact prediction module in HF-S1, evaluated using the area under the precision–recall curve (AUPRC). **f,** Effectiveness of the contact conditioning module, measured as the fraction of predicted structures satisfying contact constraints. In the box plots of **e** and **f**, the center line indicates the median; the box limits represent the 25th and 75th percentiles; and the whiskers extend to the maximum and minimum values, excluding outliers. **g,** Example structures predicted by HelixFold-S1 using two contact constraints (PDB ID: 8ozb).

Fig. 3— **Contact probability as an indicator of prediction difficulty and sampling utility.** **a,** Correlation between target-level contact probability (maximum value in the predicted matrix) and structural precision across interface types, evaluated using the Top-5 of 25 sampled conformations. In the box plots, the center line indicates the median; the box limits represent the 25th and 75th percentiles; and the whiskers extend to the minimum and maximum values (excluding outliers). **b,** Distribution of predicted contact probabilities across interface types, constructed from all benchmark targets in each category. **c,** Structural improvement from single to extensive sampling by HF-S1 across contact probability groups, measured using the best-performing conformations, showing larger gains for lower-probability interfaces. **d,** Relative Top-5 precision (of 25 sampled conformations) of HF-S1 compared with HF3 across contact probability levels, indicating stronger benefits on lower-probability targets. The sample numbers in this analysis match those in **a** and **c**. **e,** Predicted contact probabilities reflect residue physicochemical preferences.

Fig 4— **Guided sampling improves exploration of conformational space.** **a,** Target-level standard deviation of structural precision among sampled conformations for HF3 and HF-S1. Each point represents a target; higher values indicate greater conformational diversity. **b,** Fraction of targets with higher standard deviation for HF-S1, higher for HF3, or comparable, summarized across different interface types. **c,** Structural precision as a function of the number of sampled conformations. **d,** Distribution of ranking confidence and structural precision of sampled structures for two representative complexes (PDB: 8PE1 and 8HHM). HF-S1 shows a broader distribution of structural precision, whereas HF3 predictions are more concentrated. Right panels show the corresponding predicted structures by HF-S1 and HF3 compared to experimental structures.

References

- [1] Jumper, J., Evans, R., Pritzel, A., Green, T., Figurnov, M., Ronneberger, O., Tunyasuvunakool, K., Bates, R., Žídek, A., Potapenko, A., *et al.*: Highly accurate protein structure prediction with alphafold. *nature* **596**(7873), 583–589 (2021)
- [2] Evans, R., O’Neill, M., Pritzel, A., Antropova, N., Senior, A., Green, T., Žídek, A., Bates, R., Blackwell, S., Yim, J., *et al.*: Protein complex prediction with alphafold-multimer. *bioRxiv*, 2021–10 (2021)
- [3] Abramson, J., Adler, J., Dunger, J., Evans, R., Green, T., Pritzel, A., Ronneberger, O., Willmore, L., Ballard, A.J., Bambrick, J., *et al.*: Accurate structure prediction of biomolecular interactions with alphafold 3. *Nature*, 1–3 (2024)
- [4] Baek, M., DiMaio, F., Anishchenko, I., Dauparas, J., Ovchinnikov, S., Lee, G.R., Wang, J., Cong, Q., Kinch, L.N., Schaeffer, R.D., *et al.*: Accurate prediction of protein structures and interactions using a three-track neural network. *Science* **373**(6557), 871–876 (2021)
- [5] Krishna, R., Wang, J., Ahern, W., Sturmfels, P., Venkatesh, P., Kalvet, I., Lee, G.R., Morey-Burrows, F.S., Anishchenko, I., Humphreys, I.R., *et al.*: Generalized biomolecular modeling and design with rosettafold all-atom. *Science* **384**(6693), 2528 (2024)
- [6] Lin, Z., Akin, H., Rao, R., Hie, B., Zhu, Z., Lu, W., Smetanin, N., Verkuil, R., Kabeli, O., Shmueli, Y., *et al.*: Evolutionary-scale prediction of atomic-level protein structure with a language model. *Science* **379**(6637), 1123–1130 (2023)
- [7] Fang, X., Wang, F., Liu, L., He, J., Lin, D., Xiang, Y., Zhu, K., Zhang, X., Wu, H., Li, H., *et al.*: A method for multiple-sequence-alignment-free protein structure prediction using a protein language model. *Nature Machine Intelligence* **5**(10), 1087–1096 (2023)
- [8] Fang, X., Gao, J., Hu, J., Liu, L., Xue, Y., Zhang, X., Zhu, K.: Helixfold-multimer: Elevating protein complex structure prediction to new heights. *arXiv preprint arXiv:2404.10260* (2024)
- [9] Liu, L., Zhang, S., Xue, Y., Ye, X., Zhu, K., Li, Y., Liu, Y., Gao, J., Zhao, W., Yu, H., *et al.*: Technical report of helixfold3 for biomolecular structure prediction. *arXiv preprint arXiv:2408.16975* (2024)
- [10] Mirdita, M., Schütze, K., Moriwaki, Y., Heo, L., Ovchinnikov, S., Steinegger, M.: Colabfold: making protein folding accessible to all. *Nature methods* **19**(6), 679–682 (2022)
- [11] Hayes, T., Rao, R., Akin, H., Sofroniew, N.J., Oktay, D., Lin, Z., Verkuil, R., Tran, V.Q., Deaton, J., Wiggert, M., *et al.*: Simulating 500 million years of evolution with a language model. *Science*, 0018 (2025)
- [12] Bryant, P., Noé, F.: Improved protein complex prediction with alphafold-multimer by denoising the msa profile. *PLOS Computational Biology* **20**(7), 1012253 (2024)
- [13] Mirabetto, C., Wallner, B., Nystedt, B., Azinas, S., Carroni, M.: Unmasking alphafold to integrate experiments and predictions in multimeric complexes. *Nature Communications* **15**(1), 8724 (2024)
- [14] Bryant, P., Pozzati, G., Elofsson, A.: Improved prediction of protein-protein interactions using alphafold2. *Nature communications* **13**(1), 1265 (2022)
- [15] Akiyama, Y., Zhang, Z., Mirdita, M., Steinegger, M., Ovchinnikov, S.: Scaling down protein language modeling with msa pairformer. *bioRxiv*, 2025–08 (2025)
- [16] Wallner, B.: Afsample: improving multimer prediction with alphafold using massive sampling. *Bioinformatics* **39**(9), 573 (2023)
- [17] Kalakoti, Y., Wallner, B.: Afsample2 predicts multiple conformations and ensembles with alphafold2. *Communications Biology* **8**(1), 373 (2025)
- [18] Stein, R.A., Mchaourab, H.S.: Speech_af: Sampling protein ensembles and conformational heterogeneity

- 547 with alphafold2. *PLOS Computational Biology* **18**(8), 1010483 (2022)
- 548 [19] Yin, R., Pierce, B.G.: Evaluation of alphafold antibody–antigen modeling with implications for
549 improving predictive accuracy. *Protein Science* **33**(1), 4865 (2024)
- 550 [20] Xing, E., Zhang, J., Wang, S., Cheng, X.: Leveraging sequence purification for accurate prediction of
551 multiple conformational states with alphafold2. *Research Square*, 3 (2025)
- 552 [21] Silva, G., Cui, J.Y., Dalgarno, D.C., Lisi, G.P., Rubenstein, B.M.: High-throughput prediction of protein
553 conformational distributions with subsampled alphafold2. *nature communications* **15**(1), 2464 (2024)
- 554 [22] Wayment-Steele, H.K., Ojoawo, A., Otten, R., Apitz, J.M., Pitsawong, W., Hömberger, M., Ovchin-
555 nikov, S., Colwell, L., Kern, D.: Predicting multiple conformations via sequence clustering and
556 alphafold2. *Nature* **625**(7996), 832–839 (2024)
- 557 [23] Bryant, P., Noé, F.: Structure prediction of alternative protein conformations. *Nature Communications*
558 **15**(1), 7328 (2024)
- 559 [24] Stahl, K., Graziadei, A., Dau, T., Brock, O., Rappsilber, J.: Protein structure prediction with in-
560 cell photo-crosslinking mass spectrometry and deep learning. *Nature Biotechnology* **41**(12), 1810–1819
561 (2023)
- 562 [25] Stahl, K., Warneke, R., Demann, L., Bremenkamp, R., Hormes, B., Brock, O., Stülke, J., Rappsil-
563 ber, J.: Modelling protein complexes with crosslinking mass spectrometry and deep learning. *Nature*
564 *communications* **15**(1), 7866 (2024)
- 565 [26] Heo, L., Feig, M.: Multi-state modeling of g-protein coupled receptors at experimental accuracy.
566 *Proteins: Structure, Function, and Bioinformatics* **90**(11), 1873–1885 (2022)
- 567 [27] Burley, S.K., Berman, H.M., Kleywegt, G.J., Markley, J.L., Nakamura, H., Velankar, S.: Protein data
568 bank (pdb): the single global macromolecular structure archive. *Protein crystallography: methods and*
569 *protocols*, 627–641 (2017)
- 570 [28] Steinegger, M., Söding, J.: Mmseqs2 enables sensitive protein sequence searching for the analysis of
571 massive data sets. *Nature biotechnology* **35**(11), 1026–1028 (2017)
- 572 [29] Team, B.A.A., Chen, X., Zhang, Y., Lu, C., Ma, W., Guan, J., Gong, C., Yang, J., Zhang, H., Zhang,
573 K., et al.: Protenix-advancing structure prediction through a comprehensive alphafold3 reproduction.
574 *BioRxiv*, 2025–01 (2025)
- 575 [30] team, C.D., Boitreaud, J., Dent, J., McPartlon, M., Meier, J., Reis, V., Rogozhonikov, A., Wu, K.:
576 Chai-1: Decoding the molecular interactions of life. *BioRxiv*, 2024–10 (2024)
- 577 [31] Passaro, S., Corso, G., Wohlwend, J., Reveiz, M., Thaler, S., Ram Somnath, V., Getz, N., Portnoi, T.,
578 Roy, J., Stark, H., et al.: Boltz-2: Towards accurate and efficient binding affinity prediction. *BioRxiv*,
579 2025–06 (2025)
- 580 [32] Yan, C., Wu, F., Jernigan, R.L., Dobbs, D., Honavar, V.: Characterization of protein–protein interfaces.
581 *The Protein Journal* **27**(1), 59–70 (2008)
- 582 [33] Kastiris, P.L., Bonvin, A.M.J.J.: On the binding affinity of macromolecular interactions: daring to ask
583 why proteins interact. *Journal of The Royal Society Interface* **10**(79), 20120835 (2013)
- 584 [34] Zhou, H.-X., Pang, X.: Electrostatic interactions in protein structure, folding, binding, and condensa-
585 tion. *Chemical Reviews* **118**(4), 1691–1741 (2018). PMID: 29319301
- 586 [35] Seychell, B.C., Beck, T.: Molecular basis for protein–protein interactions. *Beilstein Journal of Organic*
587 *Chemistry* **17**, 1–10 (2021)
- 588 [36] Redrado-Hernández, S., Macías-León, J., Castro-López, J., Belén Sanz, A., Dolader, E., Arias, M.,
589 González-Ramírez, A.M., Sánchez-Navarro, D., Petryk, Y., Farkaš, V., et al.: Broad protection against

- 590 invasive fungal disease from a nanobody targeting the active site of fungal β -1, 3-glucanosyltransferases.
591 *Angewandte Chemie* **136**(34), 202405823 (2024)
- 592 [37] Omura, S.N., Nakagawa, R., Südfeld, C., Villegas Warren, R., Wu, W.Y., Hirano, H., Laffeber, C.,
593 Kusakizako, T., Kise, Y., Lebbink, J.H., *et al.*: Mechanistic and evolutionary insights into a type vm
594 crispr-cas effector enzyme. *Nature structural & molecular biology* **30**(8), 1172–1182 (2023)
- 595 [38] Kingma, D.P., Ba, J.: Adam: A method for stochastic optimization. arXiv preprint arXiv:1412.6980
596 (2014)
- 597 [39] Schneider, C., Raybould, M.I.J., Deane, C.M.: SAbDab in the age of biotherapeutics: updates including
598 SAbDab-nano, the nanobody structure tracker. *Nucleic Acids Research* **50**, 1368–1372
- 599 [40] Berman, H.M., Westbrook, J., Feng, Z., Gilliland, G., Bhat, T.N., Weissig, H., Shindyalov, I.N., Bourne,
600 P.E.: The Protein Data Bank. *Nucleic Acids Research* **28**(1), 235–242 (2000)
- 601 [41] Basu, S., Wallner, B.: Dockq: A quality measure for protein-protein docking models. *PLOS ONE* **11**(8),
602 1–9 (2016)
- 603 [42] Mariani, V., Biasini, M., Barbato, A., Schwede, T.: lddt: a local superposition-free score for comparing
604 protein structures and models using distance difference tests. *Bioinformatics* **29**(21), 2722–2728 (2013)
- 605 [43] Landrum, G.: RDKit: Open-source Cheminformatics. <http://www.rdkit.org>
- 606 [44] Fang, X. et al.: PaddleHelix: HelixFold-S1 Source Code and Inference Scripts. Zenodo. [https://doi.org/](https://doi.org/10.5281/zenodo.8202943)
607 [10.5281/zenodo.8202943](https://doi.org/10.5281/zenodo.8202943)

1 **Editor Summary:**

2

3 Liu and colleagues introduce HelixFold-S1, a guided sampling strategy for
4 biomolecular complex structure prediction that targets high-probability interaction
5 regions. The method achieves higher accuracy than traditional unguided methods
6 while reducing computational costs.

7

8 **Peer Review Information:**

9 *Nature Machine Intelligence* thanks Sonya M. Hanson, and the other, anonymous,
10 reviewers for their contribution to the peer review of this work.

11

12

13

14 **1. Supplementary Information:**

15 **A. PDF Files**

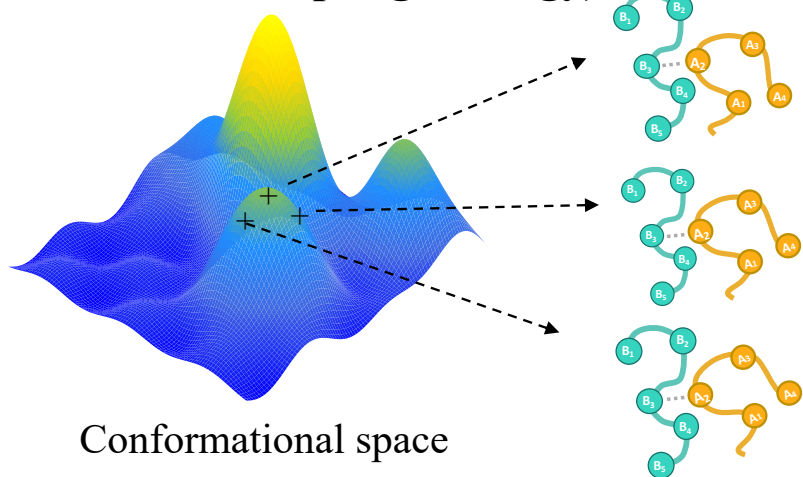
16

Item	Present?	Filename Whole original file name including extension. i.e.: Smith_SI.pdf. The extension must be .pdf	A brief, numerical description of file contents. i.e.: <i>Supplementary Figures 1-4, Supplementary Discussion, and Supplementary Tables 1-4.</i>
Supplementary Information	Yes	supplementary_information.pdf	Supplementary Fig.1-12
Reporting Summary	Yes	reporting_summary.pdf	
Peer Review Information	Yes	<i>OFFICE USE ONLY</i>	

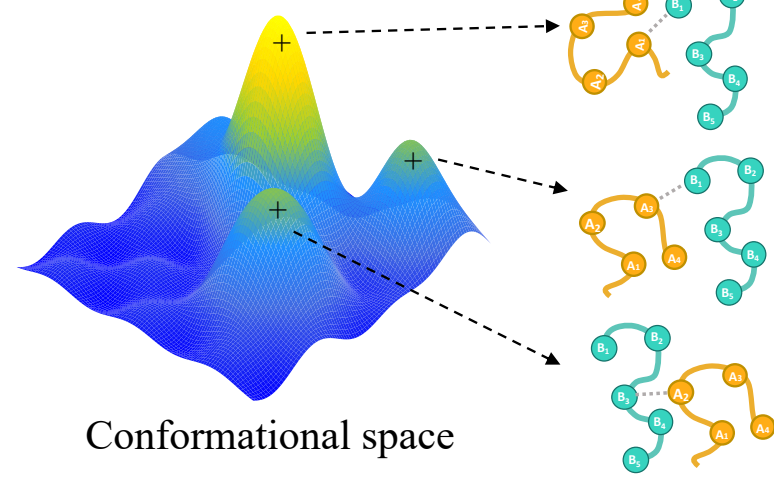
17

Conceptual Comparison

a Aimless sampling
(conventional sampling strategy)

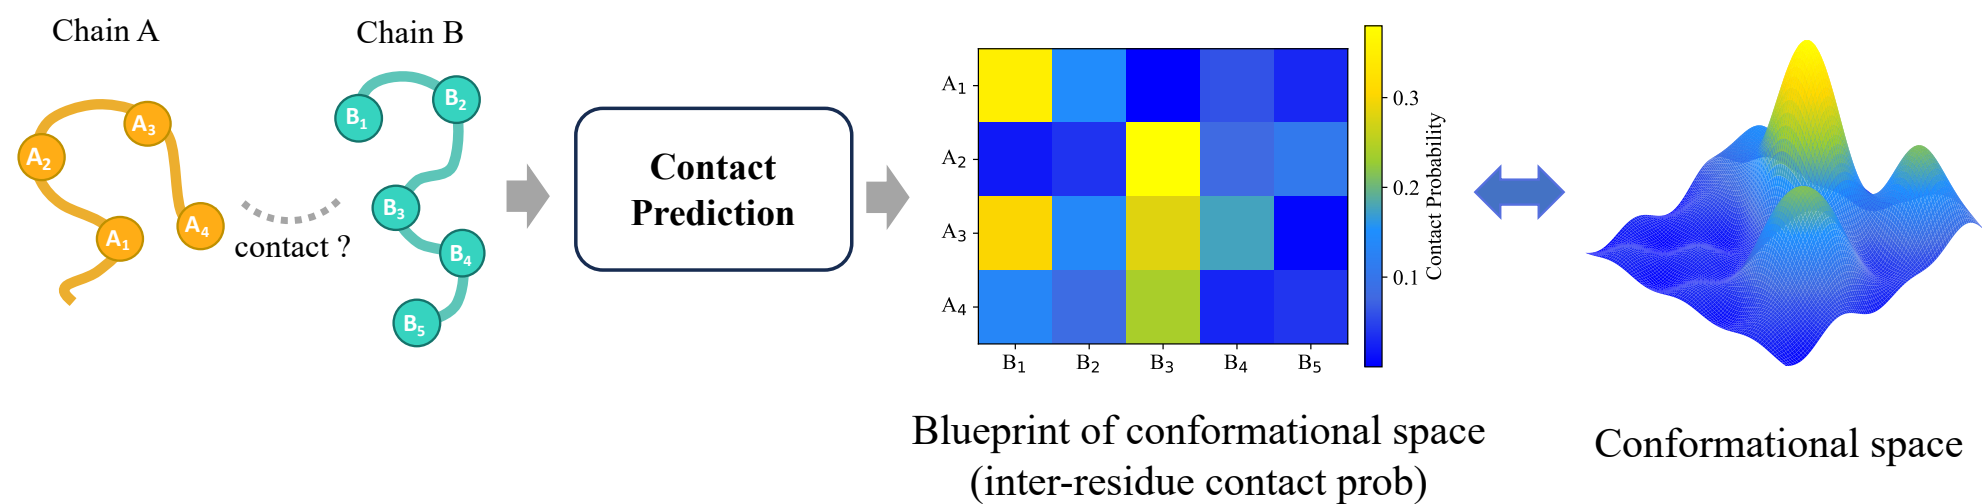


b Guided planning
(our sampling strategy)

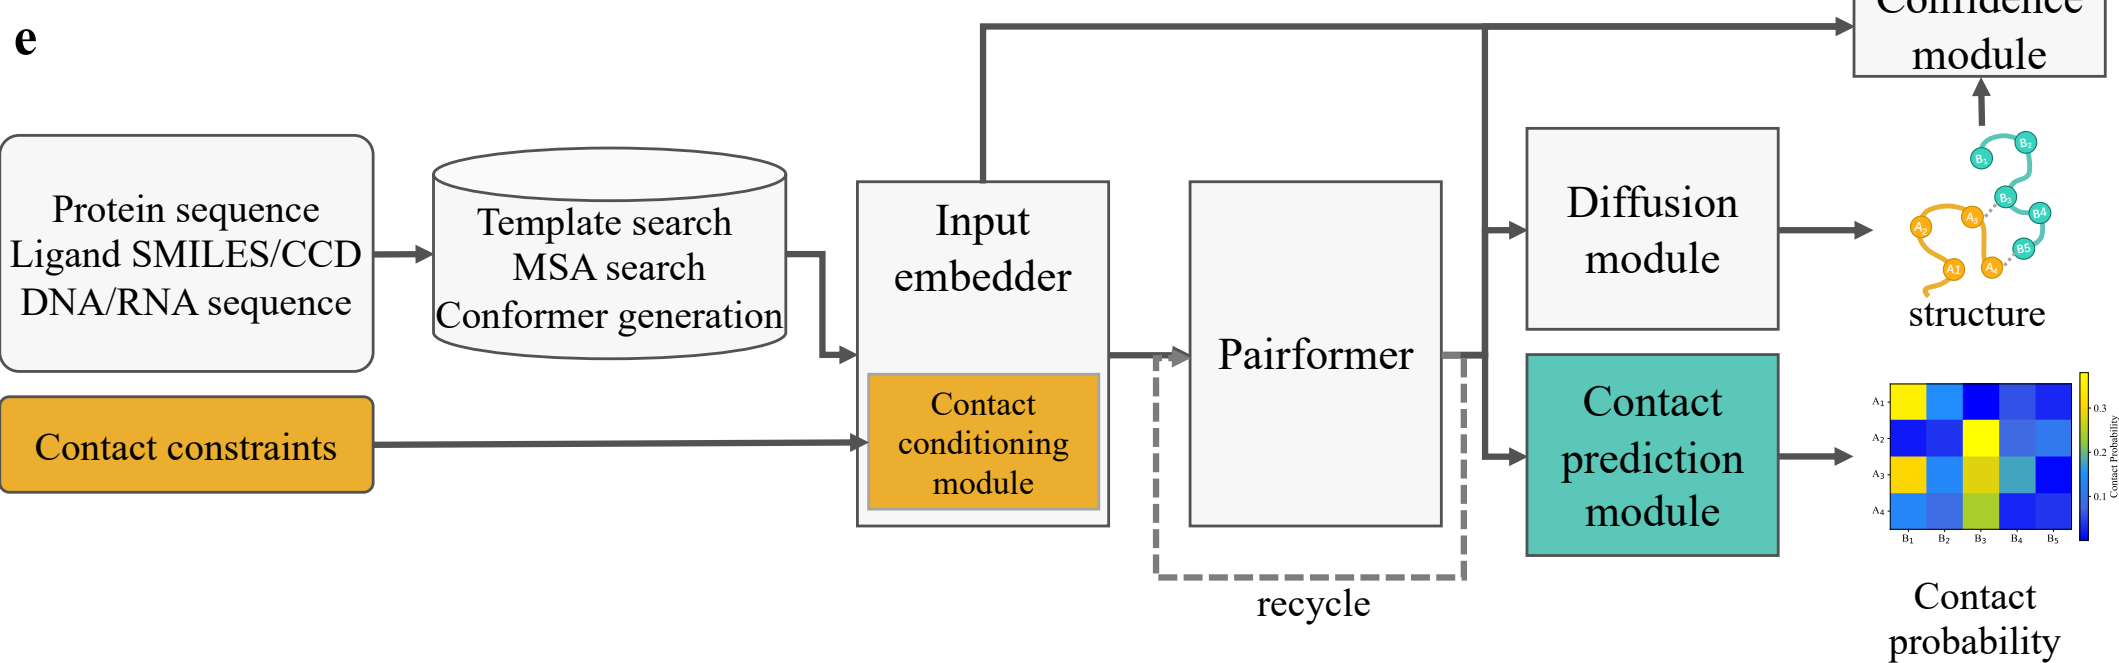
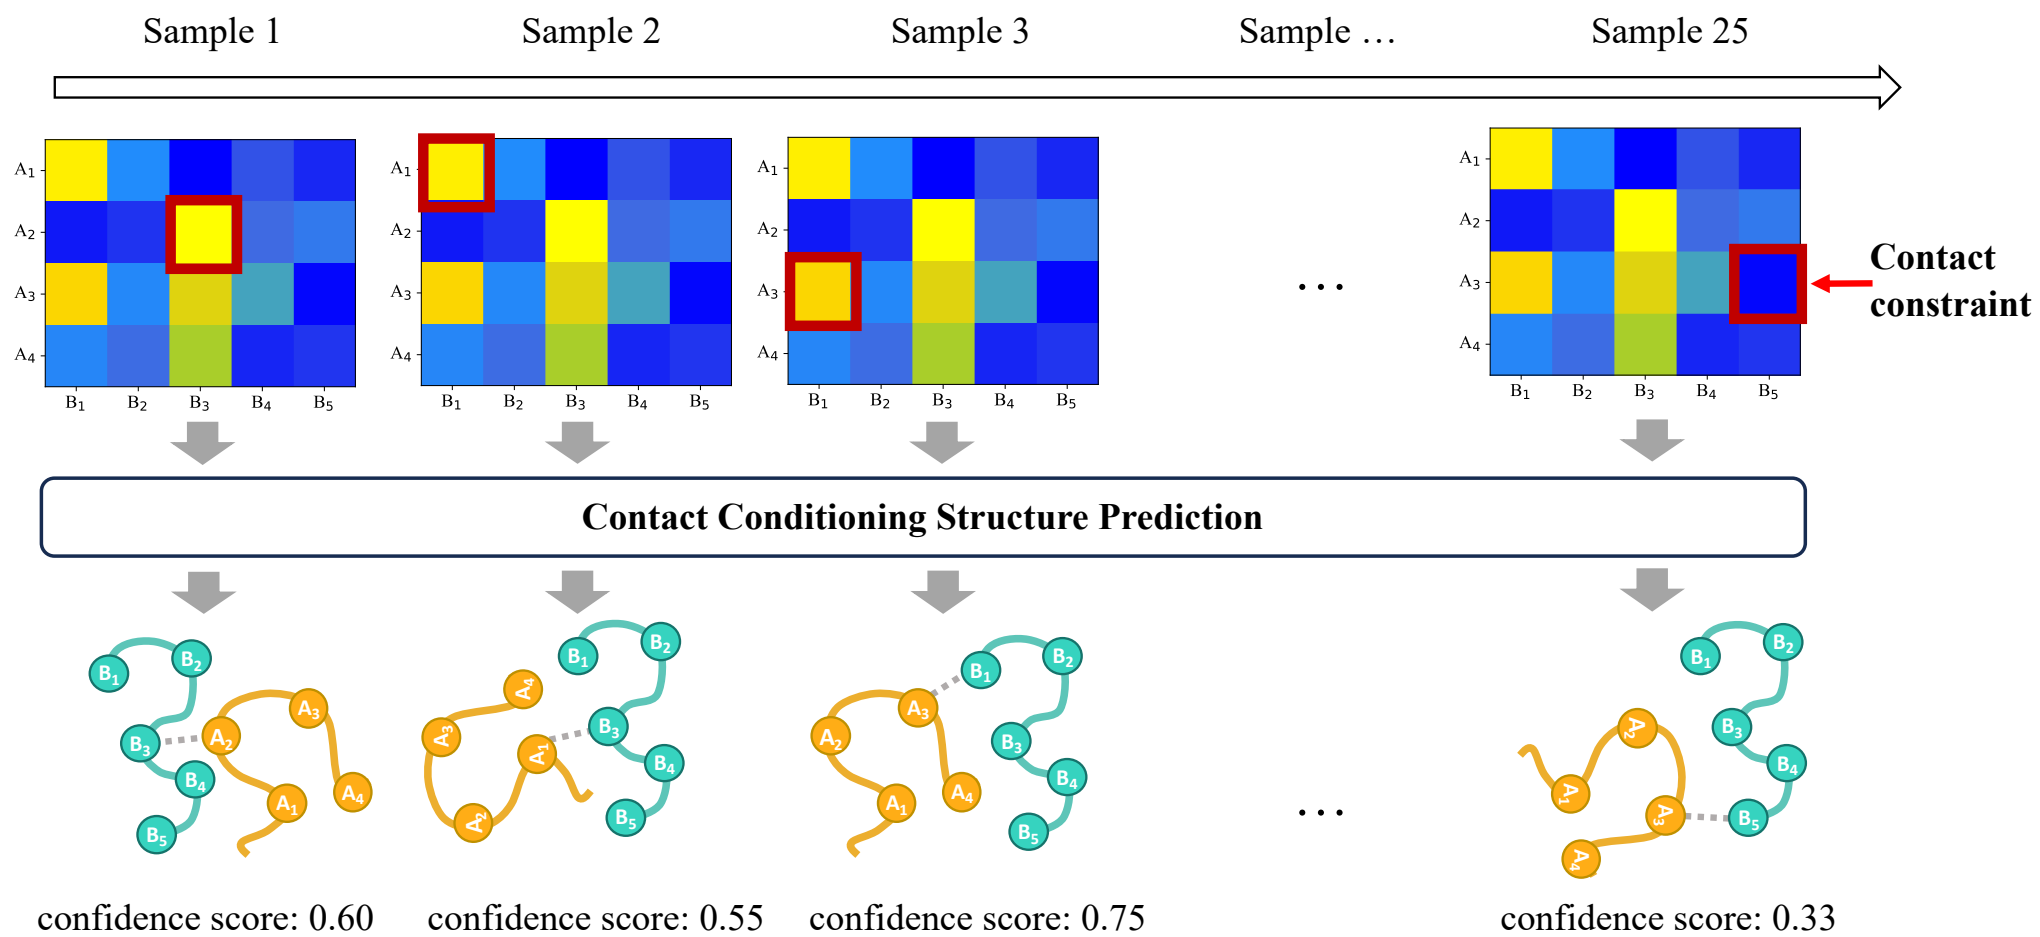


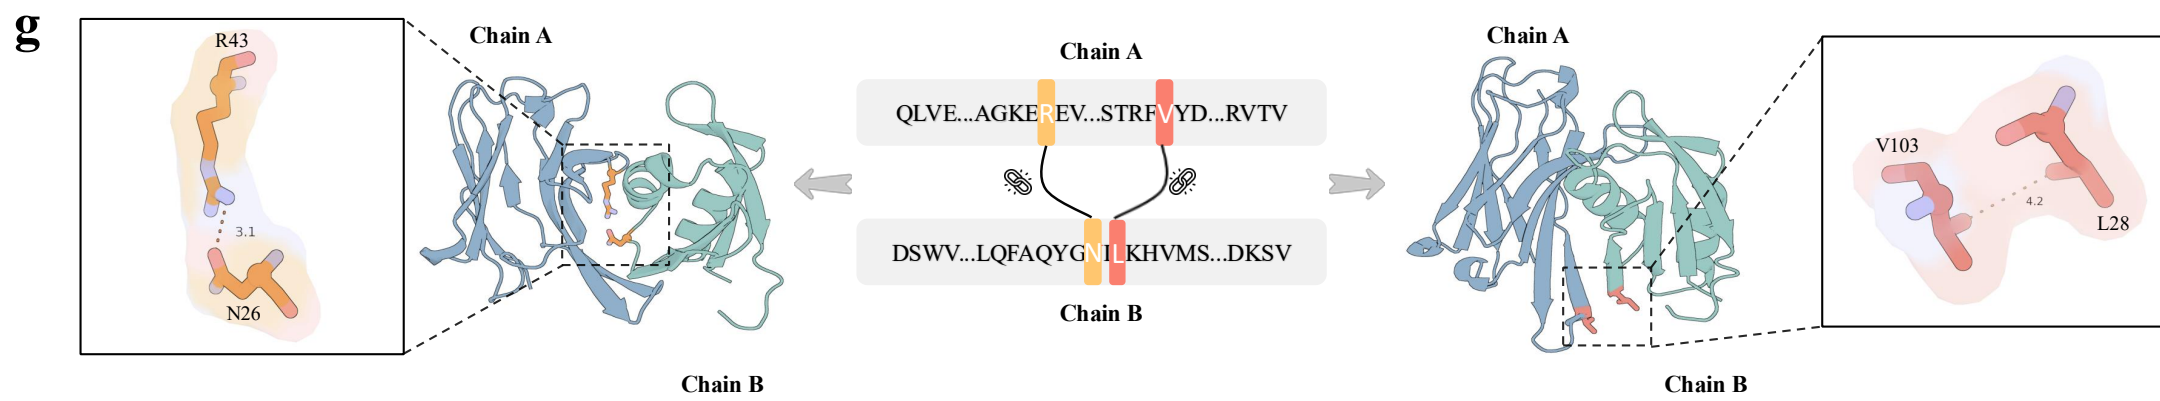
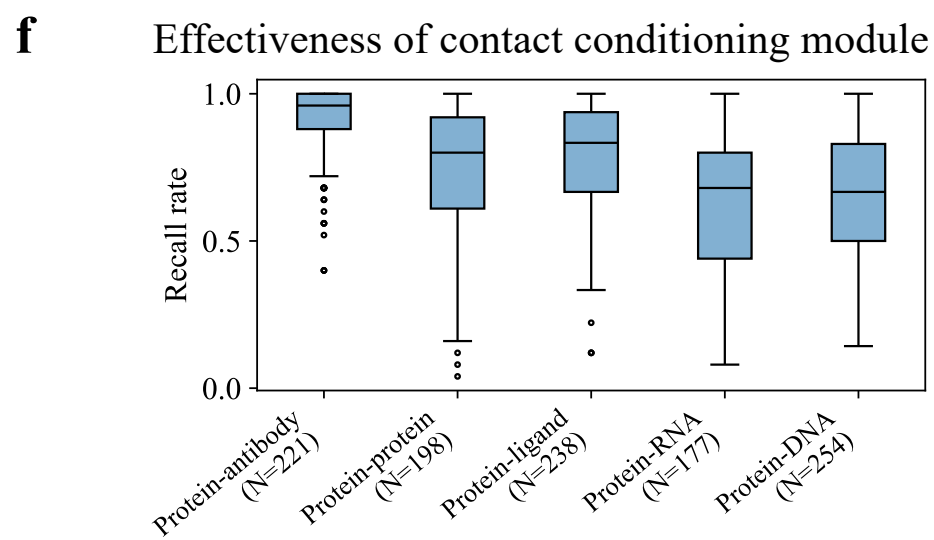
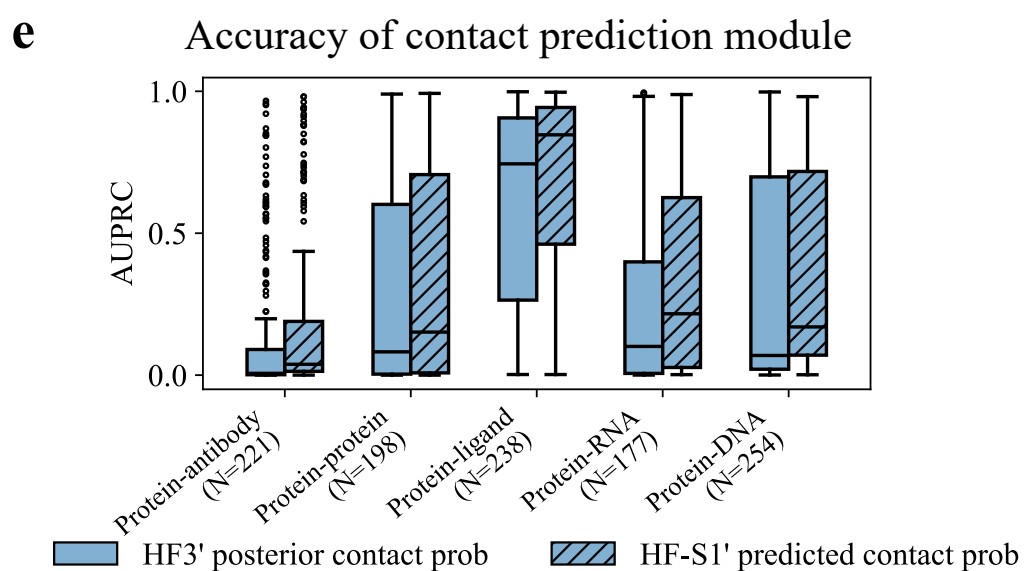
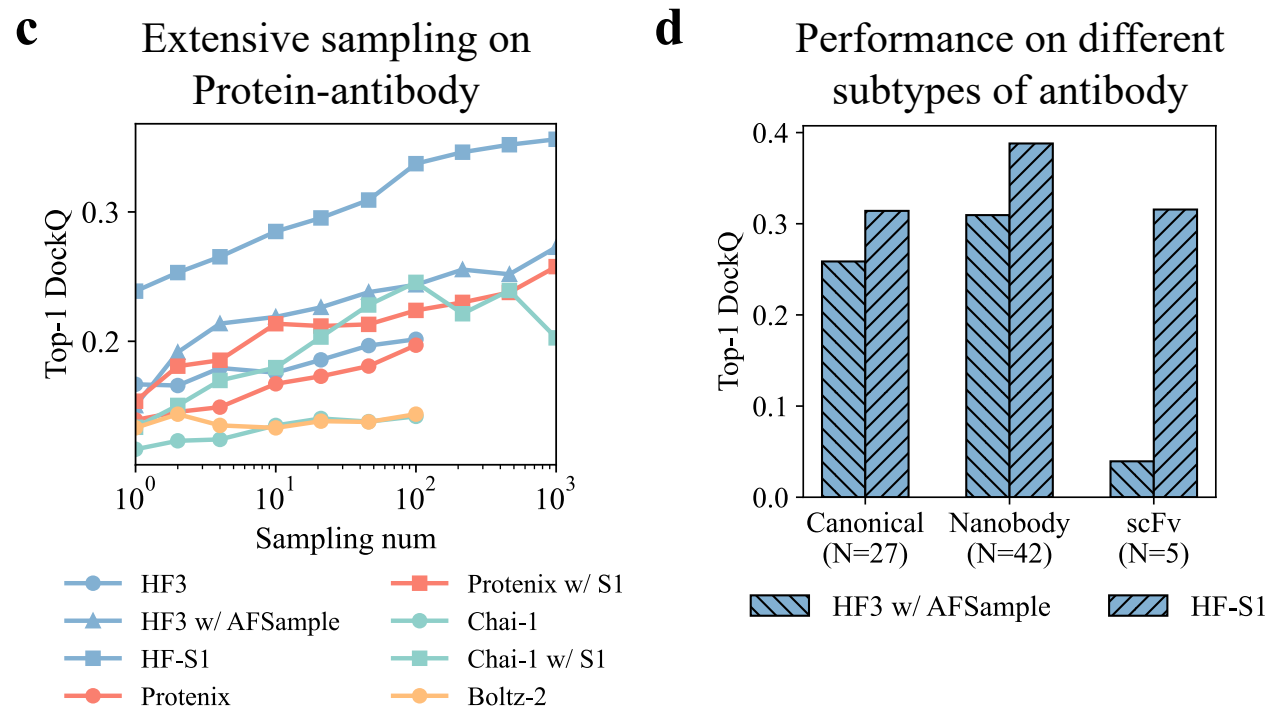
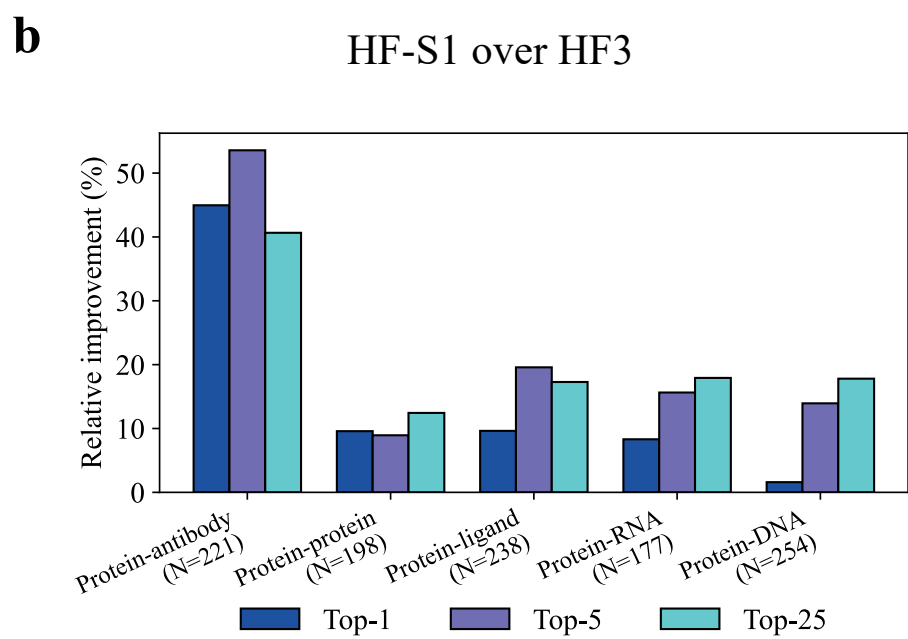
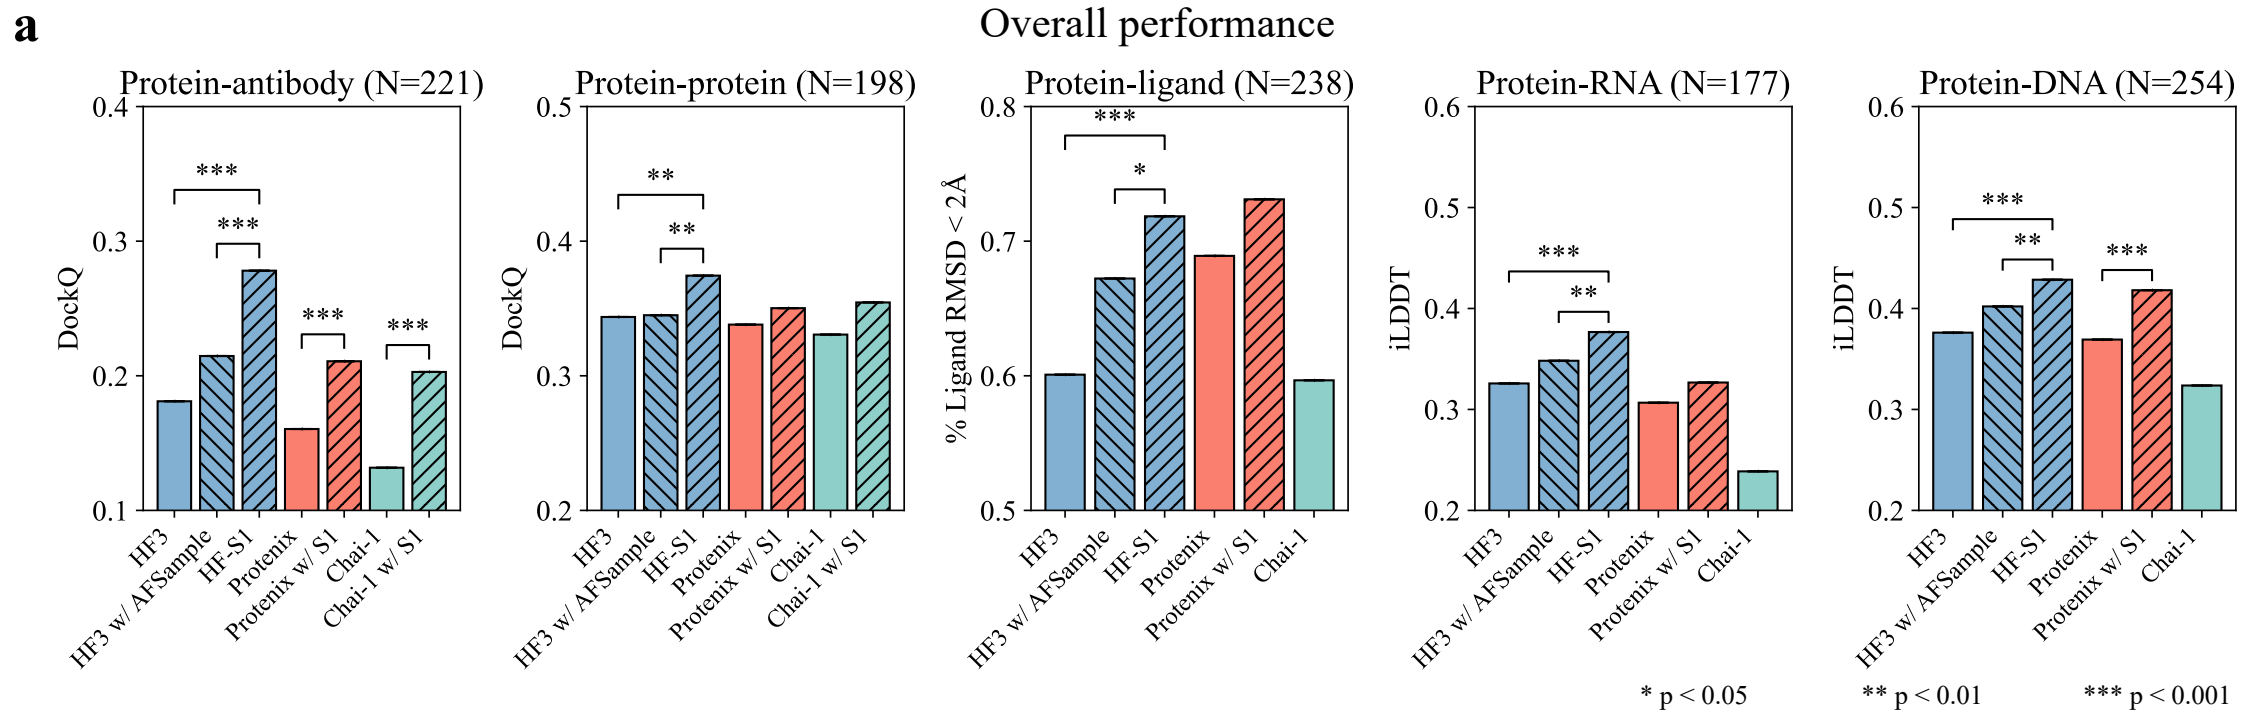
Inference Pipeline

c Stage 1: Predict blueprint of conformational space




d Stage 2: Generate conformations based on the blueprint


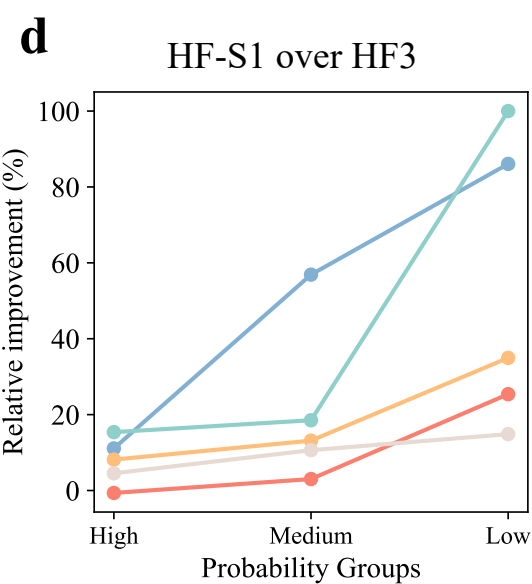






Correlation between target-level contact probability and structural precision



Improvement of multiple sampling compared with single sampling



Association between contact probability and residue interaction types

