Atomic resolution ensembles of intrinsically disordered proteins with Alphafold

Vincent Schnapka,[1] Tatiana I. Morozova,[1,2] Samiran Sen,[1] and Massimiliano Bonomi*,[1]

[1] *Institut Pasteur, Université Paris Cité, CNRS UMR 3528, Computational Structural Biology Unit, Paris, France*
[2] *Current address: CNRS, ENS de Lyon, LPENSL, UMR5672, 69342, Lyon cedex 07, France*

*E-mail: massimiliano.bonomi@pasteur.fr

Abstract

Intrinsically disordered proteins are ubiquitous in biological systems and play essential roles in a wide range of biological processes and diseases. Despite recent advances in high-resolution structural biology techniques and breakthroughs in deep learning-based protein structure prediction, accurately determining structural ensembles of IDPs at atomic resolution remains a major challenge. Here we introduce bAIes, a Bayesian framework that integrates AlphaFold2 predictions with physico-chemical molecular mechanics force fields to generate accurate atomic-resolution ensembles of IDPs. We show that bAIes produces structural ensembles that match a wide range of high- and low-resolution experimental data across diverse systems, achieving accuracy comparable to atomistic molecular dynamics simulations but at a fraction of their computational cost. Furthermore, bAIes outperforms state-of-the-art IDP models based on coarse-grained potentials as well as deep-learning approaches. Our findings pave the way for integrating structural information from modern deep-learning approaches with molecular simulations, advancing ensemble-based understanding of disordered proteins.

Introduction

Intrinsically disordered proteins (IDPs) and regions constitute approximately one-third of the human proteome[1] and play a key role in many biological systems,[2] yet they are exceptionally challenging to characterize due to their intrinsic dynamic nature. Instead of adopting a single stable structure, IDPs explore a wide continuum of conformations. The structural characterization of these systems is essential for understanding their biological functions and often requires a combination of numerous expensive experimental techniques.[3–6] For this reason, the ability to predict the conformational behavior of IDPs *in silico* is of paramount importance.

Recently, significant progress toward this objective has been achieved through the development of coarse-grained models for molecular dynamics (MD) simulations.[7–12] These models enable predicting some of the properties of IDPs with good agreement with experimental data, such as overall chain compaction distribution, and liquid-liquid phase separation propensities. As a result, they provide valuable insight into the global compaction distribution of IDPs in the human proteome.[13] Furthermore, these remarkable advancements have enabled training deep generative models capable of predicting, from sequence, IDP conformational properties,[12] intermolecular interactions involving disordered regions[14] and generating IDPs ensembles that agree with low-resolution experimental data such as small-angle X-ray scattering (SAXS) and Förster resonance energy transfer (FRET), at high throughput.[15–17]

Despite their efficiency, coarse-grained models lack the resolution needed to capture detailed local conformational properties and secondary structure propensities of IDPs, features that are crucial for understanding recognition and binding mechanisms[18] and interpreting high-resolution experiments such as Nuclear Magnetic Resonance (NMR) spectroscopy. There is also increasing demand for high-throughput generation of atomistic IDP models for drug discovery, where atomic detail is essential.[19–21] Achieving this level of detail typically re-

quires extensive high-resolution experimental data[3] as well as all-atom explicit-solvent MD simulations.[22] However, MD approaches are constrained by high computational cost and limitations in force field accuracy, despite recent advances in IDP-specific force fields[23] and efficient solvation models.[24,25] As a result, existing databases of IDP ensembles consistent with experimental data remain sparse, limiting both our understanding of IDP conformational landscapes and the development of accurate deep generative models capable of generating accurate ensembles at atomic-resolution.[26–28]

In parallel, recent advances in deep learning, particularly AlphaFold2,[29,30] have revolutionized structural biology by predicting highly accurate protein structures from amino acid sequences using data from the Protein Data Bank and co-evolutionary information from multiple sequence alignments (MSAs). Although IDPs lack strong sequence conservation, limiting MSA-based predictions, AF2 has proven valuable in this context.[31] It can predict disorder,[32] identify conditional folding via pLDDT scores,[33] and provide insights into binding mechanisms.[34–36] AF2 also outputs residue-residue distance distributions (*distograms*), which measure the local confidence in the AF2 prediction but may also reflect local dynamics and thus aid ensemble modeling. This has sparked growing interest in using AF2 for generating IDP ensembles,[37,38] including through deep learning models trained on MD simulations.[39,40] However, the reliability of AF2 for ensemble generation remains debated and requires thorough experimental validation against a wide range of both high- and low-resolution data.

In this work, we leverage the pairwise distograms from AF2 to generate atomistic ensembles of IDPs using bAIes, a Bayesian framework that integrates AF2-derived structural information with prior physico-chemical models. We applied bAIes alongside a computationally efficient random coil model that alone captures key features of disordered proteins. We benchmarked bAIes on several biologically relevant IDPs and evaluated its accuracy against diverse

NMR and SAXS data, which probe both local and global conformational properties.[41,42] Our results show that bAIes efficiently generates accurate ensembles across a spectrum of disorder, from random coils to multi-domain proteins. It is also robust to AF2 prediction errors and can match, or surpass, the accuracy of atomistic molecular dynamics simulations at a fraction of the computational cost. Moreover, bAIes outperforms IDP ensembles generated by state-of-the-art coarse-grained and deep learning-based methods. Overall, bAIes offers a promising and efficient approach for determining accurate atomic-resolution IDP ensembles by leveraging the structural information provided by AF2.

Results

We assessed the ability of bAIes to model accurate atomistic ensembles of IDPs using a benchmark set composed of 21 proteins (Fig. S1 and Tab. S1). This collection includes IDPs of varying length, amino acid composition, residual secondary structure content, compactness, and overall structural complexity. For most selected systems, NMR chemical shift data are available; in the most favorable cases, this is complemented by an extensive set of NMR observables, including Residual Dipolar Couplings (RDCs), scalar couplings, and Paramagnetic Relaxation Enhancement (PREs) data, as well as SAXS profiles (Tab. S2). This collection of both low- and high-resolution data provide a comprehensive basis for the validation of the ensembles obtained *in silico*.

An atomic-resolution random coil model for IDPs

To assess the information content provided by AF2, we first developed an efficient method to generate ensembles of IDPs to serve as a structural prior. Specifically, we created a random coil model that excludes any conformational preference or secondary structure propensity. This allows IDPs to theoretically explore the full conformational space of a polymer chain. Similar strategies have been employed to generate initial conformational pools for further

refinement using NMR and SAXS data.[43,44] Our coil model is based on a modified atomistic force field that includes only bonded interactions (bonds, angles, and dihedrals), corrected backbone dihedral distributions, and repulsive Van der Waals terms (Methods; Fig. 1a,b; Fig. S2). Ensemble generation with this model is computationally efficient, requires modest resources, and produces independent structures at a rate comparable to coarse-grained simulations with CALVADOS2[9] (Fig. S3). This random coil model serves as the structural prior in our bAIes approach and will be integrated with structural information provided by AF2 within our Bayesian framework, as detailed in the next section.

Introducing structural information from AlphaFold

In integrative structural biology, modeling approaches based on Bayesian statistics combine experimental data with prior physico-chemical knowledge through the so-called *data likelihoods* (Methods). These likelihoods ultimately translate into additional structural restraints, which are added to the standard molecular mechanics force fields (the *prior*) to enforce agreement between model and experimental data. Similarly, in bAIes, data likelihoods are modeled from the AF2 distograms, which are distance distributions between C_β atoms (or C_α for glycines). These distributions are provided as a set of equally spaced bins spanning a distance range from 0.2 to 2.2 nm, with the last bin also capturing distances beyond the upper limit. Distograms are used by the AF2 structure module to generate the final model, and as a result, the most probable C_β–C_β distance is typically close to the value observed in the resulting structure (Fig. S4).

To integrate AF2 structural information, bAIes likelihood functions were modeled from the AF2 distograms using one Gaussian distribution for each selected pair of residues (Fig. 1c,d). The selected residue pairs are associated with inter-residue contacts (see Methods). To take into account possible errors in the AF2 predictions, we considered the standard deviation of each Gaussian distribution as an unknown (*nuisance*) parameter σ_i, with minimum value

equal to the width of the corresponding AF2 distogram. A non-informative Jeffreys prior[45] was used for each σ_i to penalize large errors (Methods). Finally, the AF2 distance-dependent contributions to the force fields derived from the Gaussian likelihoods were added on-the-fly to the prior random-coil model using the open source PLUMED library[46] (Fig. 1e).

bAIes captures conformational preferences across the disorder spectrum

For the most disordered proteins with no secondary structure, AF2 is expected to contain only little structural information, as seen in the distogram average distance matrices (Fig. S5). Consequently, very few distance restraints are applied by bAIes to these very dynamic proteins. To demonstrate this point, we used our random coil model and bAIes to generate two ensembles of $A\beta40$ (Fig. 2a,b), a peptide involved in Alzheimer's disease that is known to exhibit a behavior close to random coil.[47] To assess the accuracy of our ensembles, we computed the expected NMR chemical shifts, HN-HA ^3J couplings, and NH residual dipolar couplings (Fig. 2c-e). The distribution of these NMR observables along the $A\beta40$ sequence shows that the random coil ensemble already agrees well with available experimental data and, as expected, the incorporation of AF2-derived information through bAIes does not compromise this agreement. This observation was further supported by modeling five other IDPs with no secondary structure, for which no discrepancies were observed between the chemical shifts calculated from the random coil and bAIes ensembles (Fig. 2f and Fig. S6a).

The behavior of disordered proteins and regions cannot always be fully captured by a random coil model. Varying levels of secondary structure propensity often lead to the formation of transient contacts and helical motifs, which are essential to perform specific biological functions. For example, the antitoxin PaaA2, an IDP involved in bacterial stress management, features two helical motifs that mediate toxin binding and are critical for neutralizing its toxic counterpart PaaT, thereby modulating the bacterial stress response.[48] The AF2

prediction of PaaA2, which accurately identified these motifs (Fig. 2g), is characterized by very high confidence scores (pLDDT > 90) even in the highly flexible region between the two helices. The bAIes ensemble correctly captured the flexibility of this protein (Fig. 2h), as demonstrated by the comparison with experimental NMR chemical shifts, RDCs, and SAXS data (Fig. 2i-k). This analysis revealed a clear improvement in the agreement with experimental data with respect to the pure random coil model. In particular, the two helical domains are better represented in the bAIes ensemble, as indicated by larger NMR C_α chemical shifts and positive HN RDC values in the helical domains with respect to the random coil values. In addition, the SAXS profile improved as well with respect to the random coil model, suggesting that additional short-range contacts and better representation of the secondary structure substantially improved the shape distribution of the protein. To confirm that bAIes accurately captures residual secondary structure features, we generated bAIes ensembles for five additional IDPs exhibiting varying levels of structural motifs and for which NMR chemical shift data were available. Compared to the random coil ensemble, we observed a systematic improvement in agreement with the experimental data, which became more pronounced as the amount of residual structural features increased (Fig. 2l and Fig. S6b).

bAIes is robust to AlphaFold prediction errors

In the previous section, we demonstrated that bAIes can generate accurate ensembles of IDPs across a broad range of disorder. However, AF2 can occasionally mispredict folded elements or secondary structure motifs in IDPs, particularly in regions that fold upon binding. Since AF2 is trained on stable, folded structures deposited in the PDB, the model might be biased toward predicting stable secondary and tertiary structures, even for proteins that are intrinsically disordered in solution. Additionally, because AF2 incorporates evolutionary information from multiple sequence alignments, motifs that fold upon binding to perform specific functions might have a strong sequence conservation signal that can further bias the

model toward predicting structured elements, reflecting potential binding-competent states rather than the disordered conformations sampled in solution. To evaluate the robustness of bAIes in such challenging cases, we identified several IDPs for which AF2 models have different levels of errors and we quantified the agreement of our bAIes ensembles against all the available experimental data.

We started by looking at α-synuclein, an IDP involved in Parkinson's disease. Despite the fact that this IDP is very disordered in solution,[49] AF2 predicted a 90-residue-long helix in the N-terminal domain (Fig. 3a). Surprisingly, the α-synuclein bAIes ensemble appeared in reasonable agreement with the available experimental data (Fig. 3b-e). The random coil model alone already captured very well the residue-specific conformational sampling preferences of α-synuclein as probed by the NMR chemical shifts and RDCs (Fig. 3c,e). This agreement was only marginally affected by the introduction of AF2 information in bAIes, only a very slight increase in the preference for alpha-helical conformations was observed as seen by the marginal increase in secondary chemical shifts and RDCs in the N-terminal region. Surprisingly, the agreement with SAXS data was improved by adding AF2 data. Taken together, these results indicate that bAIes is robust to errors in the AF2 predictions, thanks to the fact that uncertainties in the AF2 predictions are explicitly modeled using Gaussian likelihoods with variable width (Methods).

To test this hypothesis, we designed a modified version of bAIes (bAIes-N), where data likelihoods are derived from a Gaussian fit of the AF2 distograms, without accounting for potential errors in the distribution widths (Methods). The bAIes-N ensembles showed a poorer agreement with the NMR RDCs, J-coupling, and chemical shift data (Fig. 3f,g), in support of our hypothesis. Specifically, the chemical shift profile exhibited higher helical propensities in bAIes-N compared to both bAIes and random coil ensembles. Interestingly, these profiles resemble the experimental conditions in which α-synuclein is bound to lipids

(Fig. 3g). Despite the poorer agreement with NMR data, the bAIes-N ensemble provided a better fit to the SAXS data compared to the random coil ensemble. This highlights the complementarity of each data type and thus the need for performing a variety of different experiments to validate IDP ensembles. Finally, we did not observe a significant improvement in the agreement with PRE data (Fig. 3f), suggesting that AF2 failed to capture long-range contacts known to exist in this IDP,[50,51] which are reflected in the PRE profiles (Fig. S7).

To further validate our observations and explore the limitations of our method, we generated ensembles of five additional IDPs known to exhibit conditional folding or to be mispredicted by AF2: ACTR, NHE1, Nsp2_idr, spm_FrpC, and the SH3 domain of drk (drkN). The first three IDPs showed a similar trend in the agreement with NMR chemical shift data across random coil, bAIes, and bAIes-N ensembles (Fig. 3h and Fig. S6c), demonstrating that bAIes remains robust while a too high confidence in AF2 prediction (bAIes-N) yields worse agreement against experiments. The case of spm_FrpC proved more challenging as AF2 predicts a folded structure, whereas the protein is fully disordered in solution and only folds upon binding calcium ions.[52] A similar situation arises with drkN, which exists in a dynamic equilibrium between folded and unfolded states in solution,[53] yet is also predicted to be folded by AF2. Superimposing the NMR structures of the folded state of drkN and the Calcium-bound folded state of spm_FrpC with their respective AF2 models confirmed this observation (Fig. S8). These examples highlight a key limitation of AF2: its inability to account for environmental context or the subtle, transient nature of certain protein conformations. As a result, proteins such as drkN remain particularly challenging.

bAIes is competitive with state-of-the-art IDPs models

To further put our method to the test, we compared bAIes against state-of-the-art all-atom explicit solvent MD simulations. Specifically, we analyzed large ensembles of Aβ40, PaaA2, α-synuclein, and ACTR generated by 30μs-long MD simulations with CHARMM36m,[54]

CHARMM22s,[55] and amber99SB-disp[23] force fields, performed with the dedicated *Anton* supercomputer.[56] Overall, the accuracy of bAIes ensembles was comparable to that of atomistic MD simulations (Fig. 4a). Notably, bAIes outperformed the ensembles generated with the two CHARMM force fields across the entire benchmark set, showing superior agreement with experimental data. When compared to amber99SB-disp, which is widely considered the state-of-the-art force field for IDPs, the performance was generally comparable. Finally, for highly disordered proteins such as Aβ40 and α-synuclein, the random coil model already achieved performance comparable to the other IDP models tested, with bAIes offering marginal improvements. This highlights the strength of random coil models in accurately capturing local conformational properties (Fig. 4b-d). A detailed analysis for each protein in our benchmark set is reported in Supplementary Information (Fig. S10).

In addition to classical all-atom MD simulations, we compared bAIes ensembles with those generated using coarse-grained models, represented here by CALVADOS2,[9] and emerging deep learning-based approaches, exemplified by BioEmu.[39] These alternative methods offer substantial speed advantages over all-atom simulations. However, bAIes consistently outperformed both in terms of agreement with available experimental data (Fig. 4a). By design, CALVADOS2 is unable to capture local structural motifs or accurately reproduce dihedral angle distributions (Fig. S11). BioEmu, in turn, is limited by the quality of its input embeddings, which are derived from ColabFold predictions. To evaluate this, we compared bAIes and BioEmu ensembles for four IDPs: PaaA2, α-synuclein, spm_FrpC, and drkN (Fig. S12). While both methods are influenced by the initial AF2 predictions, bAIes, which is explicitly designed to account for prediction errors, yielded ensembles for PaaA2 and α-synuclein that are more consistent with experimental data. In contrast, BioEmu tends to produce overly rigid and overconfident ensembles based on its input AF2 structures (Fig. S12a,b), underscoring the strengths of our approach for IDPs. Nonetheless, both BioEmu and bAIes exhibited similar limitations in the cases of drkN and spm_FrpC (Fig. S12c,d).

We further analyzed each ensemble by evaluating the agreement with each individual types of experimental data. For chemical shifts, J-couplings, and RDCs, which primarily report on local conformational properties, bAIes performed on par with amber99SB-disp and outperformed all other models (Fig. 4b-d and Fig. S13). This can be attributed to the combination of secondary structure information captured by bAIes and the accurate local conformational sampling of the random coil prior. In terms of SAXS data, both CHARMM force fields were outperformed by bAIes, which achieved a level of agreement comparable to amber99SB-disp. CALVADOS2, which was explicitly optimized using SAXS data, naturally yielded the best agreement with these measurements (Fig. 4f). In contrast, bAIes ensembles were less successful at reproducing PRE data compared to amber99SB-disp (Fig. 4e). This discrepancy likely reflects the limited ability of AF2 to consistently model the transient long-range contacts visible in the PRE profiles of certain IDPs (Fig. S7).

What contributes to bAIes accuracy?

While we showed in the previous section that bAIes ensembles are in excellent agreement with various types of experimental data, it remains unclear whether this improvement over the coil models arises primarily from AF2-predicted local structural motifs or also involves long-range contacts. To better understand the origin of these improvements, we generated ensembles of Aβ40, α-synuclein, ACTR, and PaaA2 using information derived either exclusively from amino acids separated by four residues along the sequence (local contacts) or from long-range predicted contacts. The overall RMSE$_N$, which quantifies the error across all available chemical shifts, J-couplings, RDCs, PREs, and SAXS data (Methods), indicates that both local and long-range inter-residue contacts contribute to the observed agreement with experimental measurements (Fig. S9). While for Aβ40 all ensembles show comparable agreement with the data, ACTR and PaaA2 clearly demonstrate that incorporating both sources of information yields the most accurate ensembles.

From bAIes to integrative ensembles of IDPs

We have shown that bAIes can generate structural ensembles that are robust to AF2 prediction errors, achieving an accuracy comparable to state-of-the-art IDP models. We next asked whether this accuracy could be further improved by directly incorporating the experimental data previously used only for validation. To address this question, we applied a Maximum Entropy reweighting technique to refine the ensembles of ACTR, α-synuclein, Aβ40, and PaaA2 using extensive datasets of NMR observables and SAXS measurements. This approach minimally modulates the population of a prior ensemble in order to improve the agreement with solution experimental data.

In all systems studied, the reweighted bAIes ensembles outperformed those obtained by refining the CHARMM36m and CHARMM22* ensembles (Fig. 5a). When compared to amber99SB-disp, the reweighted bAIes ensembles were slightly less accurate only in the case of PaaA2. Notably, in the case of α-synuclein, where the initial bAIes ensemble was less accurate than amber99SB-disp, reweighting against a large set of experimental data resulted in superior quality compared to the atomistic force field. Furthermore, the bAIes reweighted ensembles still outperformed the reweighted ensembles generated with CALVADOS2 and BioEmu, except for the case of Aβ40, where BioEmu provided the best overall agreement with experimental data across all the IDP models tested.

Among the IDPs in our benchmark, ACTR was one that contained errors in the prediction of secondary structure, with extra helical propensities that are associated with folding upon binding mechanisms in the function of ACTR (Fig. 5b). Nevertheless, the bAIes ensemble (Fig. 5c) both before and after reweighting shows an overall agreement that competes with amber99SB-disp. Comparing the chemical shift, RDC and SAXS profiles between bAIes and amber99SB-disp confirms this observation, with very similar profiles observed (Fig. 5d-f).

bAIes ensembles of multi-domain proteins

Finally, we explored the possibility of using bAIes to model multi-domain proteins (MDPs) with folded domains connected by flexible linkers. Since bAIes maintains the integrity of well-defined folded motifs (as seen with drkN and spm_FrpC), we expect it to efficiently sample flexible linker conformations while preserving folded motifs. To test this, we generated ensembles of three MDPs: a Green Fluorescent Protein (GFP) chimera composed of two fluorescent proteins separated by a $(GS)_8$ linker (GS8),[57] and two polyubiquitin chains, bi-ubiquitin (Ubq2) and triubiquitin (Ubq3).[58] As expected, in all cases the folded motifs were well predicted by AF2 and preserved in the ensembles generated by bAIes. In addition, our approach allowed efficient sampling of the flexible linkers (Fig. 6a,b,c), showing that bAIes can generate MDPs conformational ensembles without the need to add additional restraints to conserve the folded domains.

We assessed the quality of our ensembles by quantifying the agreement with SAXS data, which provides useful low-resolution information on the overall shape of the protein in solution. The agreement with experimental SAXS profiles, while not perfect, exceeded our expectations given that bAIes only uses AF2 information (Fig. 6d,e,f). In the ensemble of GS8, the two GFP domains sampled a large conformational space, making GS8 relatively extended. Reweighting greatly improved the agreement with SAXS and resulted in more inter-domain contacts (Fig. S14). The polyubiquitins ensembles generated with bAIes showed instead preferred relative domain orientations driven by the inter-domain contacts predicted by AF2. While reweighting slightly improved the agreement with the SAXS data (Fig. S15b,e yellow bars), we wondered whether the predicted inter-domain contacts actually contributed to the match with the experimental profiles. To address this question, we generated bAIes ensembles excluding the inter-domain contacts predicted by AF2. As expected, these ensembles sampled a broader conformational space (Fig. S15a,d): for Ubq2, this had little (negative) impact on the fit with SAXS data (Fig. S15b,c), but for Ubq3, the absence

of inter-domain contacts led to poorer agreement (Fig. S15e,f). Furthermore, reweighting ensembles obtained excluding inter-domain contacts failed to reproduce the agreement with SAXS data observed in the original bAIes ensembles (Fig. S15b,e blue bars), suggesting that inter-domain contacts are indeed important for accurately modeling polyubiquitin conformational landscapes.

Although bAIes was not originally designed to model multi-domain proteins, these examples demonstrate its potential for studying this family of proteins. Despite current limitations due the AF2 limited ability to reliably predict transient long-range interactions, our approach holds great promise, especially if enhanced with more accurate physico-chemical priors that can better model inter-domain interactions and integrated with experimental data such as PRE or FRET that probe specific inter-domain contacts.

Discussions

In this work, we introduced bAIes, a Bayesian framework that leverages the structural information from AF2 to generate accurate atomistic ensembles of intrinsically disordered proteins. Across a range of IDPs with varying residual structure, bAIes matches the accuracy of atomistic molecular dynamics simulations and outperforms coarse-grained and deep-learning-based methods. We attribute this performance to at least three factors. First, the random coil prior captures key IDP properties, as shown by its excellent agreement with experimental data for highly disordered proteins. Second, bAIes effectively incorporates AF2-derived local secondary structure and long-range contact information, improving agreement with NMR and SAXS data over the prior alone. This raises a fundamental question: how does AF2 learn such features despite being trained on folded proteins? One explanation is that it generalizes from semi-ordered loops resolved in PDB structures; more intriguingly, it may have learned general principles of interatomic interactions that also apply to disor-

dered proteins.[59] Finally, bAIes explicitly models uncertainty in AF2 predictions, preventing overfitting and ensuring robustness to prediction errors.

Leveraging AI-based predictions for ensemble determination is becoming increasingly widespread.[37–39] In parallel to our work, Metainference[60] has been applied to integrate AF2-derived structural information with the CALVADOS2 coarse-grained model[9] for generating IDPs ensembles.[37] In this approach, the CALVADOS2 force field is minimally adjusted to reproduce the average inter-residue distances derived from the AF2 distograms, generating coarse-grained ensembles that show improved agreement with SAXS data. While conceptually similar to bAIes, this approach differs in several important ways. First, bAIes does not rely on filtering interatomic distances based on AF2 prediction confidence. We highlighted some cases such as PaaA2, where highly flexible regions can exhibit misleadingly high confidence scores (pLDDT > 90 and pAE < 0.5 nm); applying strong restraints to such flexible regions would artificially rigidify them. Second, bAIes is explicitly designed to capture the dynamic nature of IDPs by leveraging the full distribution of selected AF2 distograms, rather than relying solely on average distance values. Additionally, our method avoids the use of arbitrarily strong structural restraints to enforce secondary and tertiary structure formation, an important advantage when modeling multi-domain proteins. Finally, our method does not require multiple replica simulations and uses a fully atomistic representation that does not require backbone reconstruction. In addition to Alphafold-Metainference, emerging deep generative models also aim to generate protein ensembles from AF2 embeddings. BioEmu,[39] for instance, is trained on MD data and uses ColabFold[61] embeddings to emulate backbone ensembles. While promising, we demonstrated that such method is not yet as robust as bAIes in the presence of AF2 prediction errors, and currently yield IDPs ensembles that overall underperform relative to bAIes.

The limitations of our method, which are shared with AlphaFold-Metainference and BioEmu,

become evident when AF2 predicts folded states for IDPs. This typically occurs when a protein can adopt multiple states, either under exchange conditions, as in the SH3 domain of drkN, or in distinct environments, such as spm_FrpC. These mispredictions are due to the fact that AF2 was not designed to account for the effect of the environment on protein folding. Addressing this limitation would require either training structure prediction models on proteins characterized under standardized conditions or developing models that explicitly incorporate environmental parameters, both solutions currently hindered by the limited availability of experimental data. In addition, AF2 is inherently limited in its ability to capture the transient nature of certain intramolecular contacts, which are often critical and better captured by experimental techniques such as PRE and FRET.[62][?]

A potential route for improving our model is to incorporate additional energy terms into the random coil model to account for structural features not captured by AF2. One such strategy could involve parameterizing residue-specific potentials, similar to those implemented in modern IDPs coarse-grained models.[9,11] Additionally, integrating bAIes with enhanced-sampling techniques, such as metadynamics,[63] could facilitate the exploration of regions of conformational space that are rarely sampled in standard simulations. For example, in the case of drkN, bAIes simulations do not sample the unfolded states, due to the fact that AF2 predict a folded conformation. By using metadynamics to bias the sampling toward less populated, unfolded conformations, one can enhance the exploration of the conformational space. These enhanced ensembles can then be refined using reweighting approaches guided by experimental data to recover the correct equilibrium between folded and unfolded states. Another idea to tackle the existence of distinct conformational states is to exploit multiple distograms generated by running AF2 with modified multiple sequence alignments (MSAs).[64,65] One could envision a future extension of bAIes that incorporates multiple distograms, derived from different MSAs, to better capture sparsely populated folded states within the conformational landscape of IDPs. Looking ahead, the growing number of struc-

ture prediction tools like Alphafold3[30] and Boltz2[66] could further expand the range of applications of our approach. For instance, bAIes could be applied to model dynamic RNA molecules, to efficiently sample small-molecule conformations within protein binding pockets, or to investigate the effects of post-translational modifications on IDPs structural properties.

Generating accurate atomic-resolution ensembles of IDPs is becoming increasingly important. These proteins play essential roles in cellular signaling, regulation, and molecular recognition, functions that depend critically on their structural flexibility and ability to engage in dynamic interactions with multiple partners. Atomic-level insights are thus essential for elucidating the molecular mechanisms underlying these behaviors. In this context, bAIes offers a valuable and computationally efficient approach for generating accurate, atomic-resolution IDP ensembles. These ensembles not only advance fundamental research but also support a wide range of applications. First, they can provide preliminary structural descriptions of novel systems, guiding the design of future experimental studies. Second, they serve as large conformational pools that can be validated against, or refined using, diverse experimental datasets. Atomistic detail is also crucial for small-molecule docking, where scoring functions commonly employed in drug discovery rely heavily on atomic-level accuracy.[21] Finally, there is a growing need for ensembles that are consistent with both high- and low-resolution experimental data to serve as *ground truth* for training machine learning models capable of predicting structural ensembles directly from sequence.[67] In summary, bAIes provides a computationally efficient approach for generating accurate atomistic ensembles of IDPs. The method is implemented in the open-source PLUMED library (www.plumed.org),[46] making it readily accessible to the broader computational science community.

Methods

General Alphafold settings

Unless stated otherwise, Alphafold[29] 2.3.2 was run on a local cluster with default settings. MSAs were performed on Uniref90[68] and MGnify[69] and template search was performed on the pdb70[70] with HHsearch[70] and HHblits.[71] MSAs on BFD[29] and Uniref30[72] were performed with jackhmmer.[73] To maintain maximum accuracy and applicability for multido­main proteins, we systematically used the AF2 prediction with the highest global pLDDT score among the five generated models.

Development of the random coil model

The amber99SB-ILDN force field[74] was modified to create a computationally efficient random coil model. We first defined a simplified force field E_{FF} composed by all bonded terms as well as the repulsive part of the Van der Waals potential. In order to do so, these interactions were switched off at a distance cut-off set to $r_{i,j}^{cut} = 2^{\frac{1}{6}}\sigma_{r_{i,j}}$ for each pair of atoms i and j. This distance cutoff corresponds to the energy minimum of the Lennard-Jones potential:

$$E_{i,j} = 4\varepsilon_{i,j} \left[\left(\frac{\sigma_{i,j}}{r_{i,j}} \right)^{12} - \left(\frac{\sigma_{i,j}}{r_{i,j}} \right)^{6} \right]$$

where $r_{i,j}$ in the distance between atoms i and j; $\varepsilon_{i,j} = \sqrt{\varepsilon_i \varepsilon_j}$ and $\sigma_{i,j} = 0.5(\sigma_i + \sigma_j)$ are the Lennard-Jones parameters. An offset was added to $E_{i,j}$ to preserve the continuity of the energy function at the cutoff distance. The backbone dihedral distributions $P_{\mathrm{FF}}(\phi_i, \psi_i)$ obtained with E_{FF} were then corrected to account for the absence of non-bonded interactions. To do so, we computed for each residue type i a correction MAP (CMAP) term,[75] defined as:

$$E_{\mathrm{CMAP}}(\phi_i, \psi_i) = -k_B T \, \log \frac{P_{\mathrm{coil}}(\phi_i, \psi_i)}{P_{\mathrm{FF}}(\phi_i, \psi_i)} \tag{1}$$

where $P_{\mathrm{coil}}(\phi_i, \psi_i)$ is the target residue-specific random coil dihedral distribution, T the temperature of the system, and k_B the Boltzmann constant.

To obtain an estimate of the residue-specific random coil distributions $P_{\mathrm{coil}}(\phi_i, \psi_i)$, we used a well-established approach that consists in extracting dihedral distributions corresponding to residues in coil regions from structures deposited in the PDB database.[43,44,76,77] In our study, we selected a pool of protein X-ray structures deposited in the PDB database before January 1st 2024. These structures were filtered to an RFree parameter below 0.25 and sequence identity below 50%, resulting in a total of approximately 30,000 structures. The structures were then analyzed with DSSP[78] and residues belonging to alpha-helical and beta strand secondary structures were discarded to generate residue-specific (ϕ, ψ) distributions representative of random coil residues. Pre-proline residues were analyzed separately,[43] resulting in a total of 40 $P_{\mathrm{coil}}(\phi_i, \psi_i)$ distributions, which were stored on a 24x24 grid following the format adopted by the CHARMM force fields.[75]

To compute $P_{\mathrm{FF}}(\phi_i, \psi_i)$, we used well-tempered metadynamics[79] on 40 tripeptides with amino acid sequence equal to AXA or AXP, where X is one of the 20 canonical amino acids. While the pre-proline residues were considered separately, we arbitrarily chose Alanine as the neighboring residue. Since in the simplified force field E_{FF} the non-bonded interactions are turned off, we assumed that this set of tripeptides give a sufficiently accurate estimate of the residue-specific backbone dihedral angles distributions of the central aminoacid, which was later confirmed by the good performance of our random coil model. For each tripeptide sequence, the metadynamics collective variables used were the backbone dihedral angles Φ and Ψ. Gaussians with height equal to 1.2 kJ/mol and widths of 0.2 radians for both dihedrals were added every 1 ps using a bias factor of 10. Metadynamics simulations were carried out with LAMMPS 2 Aug. 2023[80] equipped with PLUMED 2.8[46] for 100 nanoseconds each. Finally, $E_{\mathrm{CMAP}}(\phi_i, \psi_i)$ were calculated from Eq. 1, stored on a grid, and added to the

simplified force field E_{FF} to obtain our final random coil model:

$$E_{\text{RC}}(\mathbf{X}) = E_{\text{FF}}(\mathbf{X}) + \sum_{i=1}^{N} E_{\text{CMAP}}(\phi_i(\mathbf{X}), \psi_i(\mathbf{X})) \tag{2}$$

where $\mathbf{X}$ are the coordinates of the system, and N the number of residues.

bAIes theory

Our bAIes approach is based on a Bayesian inference framework that estimates the probability of a model M, defined in terms of its structure $\mathbf{X}$ and other parameters, given the information available about the system, including prior physico-chemical knowledge and newly acquired data.[81] The posterior probability $p(M \mid D)$ of model M given data D and prior knowledge is expressed as:

$$p(M \mid D) \propto p(D \mid M) \cdot p(M) \tag{3}$$

where the likelihood function $p(D \mid M)$ represents the probability of observing data D given M. This function quantifies the agreement between the observed data and the model and accounts for all sources of uncertainty and error. The term $p(M)$ denotes the prior probability of model M based on existing knowledge. In the following, we specify each component of this general Bayesian framework as applied in our bAIes approach.

The data. In integrative structural biology applications, D typically represents a collection of experimental data points, which may originate from different experimental techniques, such as NMR, SAXS, or cryo-electron microscopy.[81–83] In bAIes, D represents the structural information provided by AF2, defined as a collection of distances $\{d_i\}$ between C_β atoms (or C_α in the case of glycines). Each data point d_i represents the average distance computed from the corresponding AF2 distogram. Two selection criteria were applied: (i) distances between residues separated in sequence by less than 3 aminoacids were discarded; (ii) pairs with most probable distance in the AF2 distogram greater than a residue-pair specific cutoff[84] were

discarded to retain only potential inter-residue interactions.

Data likelihood. For each pair of selected residues, a Gaussian likelihood was used:

$$p(d_i \mid \mathbf{X}, \sigma_i) = \frac{1}{\sqrt{2\pi}\sigma_i} \exp\left\{ -\frac{[d_i(\mathbf{X}) - d_i]^2}{2\sigma_i^2} \right\} \tag{4}$$

where $d_i(\mathbf{X})$ is the distance between a pair of C_β atoms (or C_α in the case of glycines) in structure $\mathbf{X}$ and σ_i an error parameter that modulates the deviation between observed and computed data. Under the assumption of independent data points, the total likelihood for a set of of M points $D = \{d_i\}$ can be factorized as:

$$p(D \mid \mathbf{X}, \{\sigma_i\}) = \prod_{i=1}^{M} p(d_i \mid \mathbf{X}, \sigma_i) \tag{5}$$

Priors. In bAIes, the error parameters σ_i are treated as unknown variables and are determined by sampling the posterior distribution. To penalize large errors, we used an informative Jeffreys prior:

$$p(\sigma_i) = \frac{1}{\sigma_i} \tag{6}$$

Additionally, the width of the corresponding distogram, calculated by fitting this distribution with a Gaussian function, was used as the minimum value $\sigma_i^{\min}$ for each uncertainty parameter. This choice allows us to handle cases where AF2 predictions are inaccurate in the same way as we would treat outlier experimental data points in other Bayesian inference approaches for structural biology.[83] The random coil model previously described was used as structural prior:

$$p(\mathbf{X}) \propto \exp\left[-\frac{E_{\mathrm{RC}}(\mathbf{X})}{k_B T} \right] \tag{7}$$

Marginalization. To avoid sampling the parameters σ_i, we marginalized the corresponding likelihood distributions along with their uncertainty priors. For each data point,

the resulting marginal data likelihood is:

$$p(d_i \mid \mathbf{X}) = \int_{\sigma_i^{min}}^{\infty} p(d_i \mid \mathbf{X}, \sigma_i)\, p(\sigma_i) = \frac{0.5}{d_i(\mathbf{X}) - d_i}\, erf\left(\frac{d_i(\mathbf{X}) - d_i}{\sqrt{2}\sigma_i}\right) \tag{8}$$

Simplified noise model. To assess the effect of our noise model in presence of AF2 prediction errors, we defined a simplified version of bAIes (bAIes-N) in which the σ_i parameters are modeled by a delta Dirac function $\delta(\sigma_i - \sigma_i^{min})$, which results in the following marginal data likelihood:

$$p(d_i \mid \mathbf{X}) = \int_{\sigma_i^{min}}^{\infty} p(d_i \mid \mathbf{X}, \sigma_i)\, \delta(\sigma_i - \sigma_i^{min}) = \frac{1}{\sqrt{2\pi}\sigma_i^{min}} \exp\left\{-\frac{[d_i(\mathbf{X}) - d_i]^2}{2(\sigma_i^{min})^2}\right\} \tag{9}$$

The bAIes hybrid energy function. After defining all the component of our approach and marginalizing the uncertainty parameters, we obtain the final bAIes posterior distribution:

$$p(\mathbf{X} \mid D) \propto \prod_{i=1}^{M} p(d_i \mid \mathbf{X}) \cdot p(\mathbf{X}) \tag{10}$$

where the marginal data likelihood is given by Eq. 9 and the structural prior by Eq. 7. To sample the posterior, we define the associated bAIes hybrid energy function as:

$$E_{\text{bAIes}}(\mathbf{X}) = -k_B T \log p(\mathbf{X} \mid D) = E_{\text{RC}}(\mathbf{X}) + E_{\text{AF2}}(\mathbf{X}) \tag{11}$$

which is composed by our random coil model E_{RC} and the spatial restraints E_{AF2} that incorporate the structural information provided by the AF2 distograms. Finally, models are sampled by Molecular Dynamics under the bAIes hybrid energy function defined in Eq. 11.

Details of the random coil model and bAIes simulations

The simulations were carried out using LAMMPS[80] version 2 Aug. 2023 and the development version of PLUMED.[46] The setup, minimization, and production steps described below are

common to random coil model and bAIes simulations.

- For each protein, the LAMMPS topology and input files are generated from GROMACS inputs prepared in vacuo with the top ranked AF2 model and using the amber99SB-ILDN force field.[74] The conversion from LAMMPS to GROMACS is performed using InterMol 0.1.0 and ParmEd 4.2.2.[85] The amber99SB-ILDN force field is then transformed into our random coil model by a custom Python script, which generates the final LAMMPS input and topology files.

- Prior to production, an energy minimization is performed using the Polak-Ribiere version of the conjugate gradient algorithm (LAMMPS default). Stopping tolerances of 10^{-4} and 10^{-6} $kcal.mol^{-1}.A^{-1}$ are used for energy and force respectively.

- Production simulations are carried out using the Bussi-Donadio-Parrinello thermostat[86] set at a temperature of 298.1 K, a damp parameter of 100 fs and a time step of 1 fs. The coordinates are saved every 100 ps, producing semi-independent conformations due to the rapid motion of the protein in this framework (Fig. S3). The final ensembles contain 20000 structures for the 18 studied IDPs and 10000 structures for the 3 multi-domain proteins (GS8, Ubq2 and Ubq3).

- For bAIes simulations, AF2 predictions are first performed. Then, a preprocessing python script reads the AF2 distograms and generates the necessary PLUMED files. The bAIes forces added by PLUMED are updated every two timesteps.

CALVADOS2 and BioEmu ensembles

CALVADOS2 simulations were carried out using openMM 8.2.0[87] following the procedure described in Tesei *et al.*[9] using a Langevin integrator at a temperature of 298 K, a friction coefficient of 0.01 ps^{-1} and a time step of 5 fs. The pH was set to 7.4 and the beads coordinates were saved every 100 ps. A total of 20,000 conformations were generated for

each protein. The all-atom model was then reconstructed with Pulchra[88] and protons were added with GROMACS 2022.5.[89]

BioEmu ensembles were generated using the protocol described in Lewis *et al.*,[39] aiming for 30,000 conformations with default settings. The resulting conformations were filtered to remove non-physical structures using the built-in BioEmu option. Finally, side chains and protons were added with hpacker[90] and GROMACS 2022.5,[89] respectively.

NMR and SAXS data back-calculation

Chemical shifts were predicted with SPARTA+.[91] The $^3J_{HNHA}$ couplings were calculated with a Karplus relationship reported by Vogeli and Bax,[92] with Karplus parameters of 7.97, -1.26 and 0.63 Hz respectively. The $^3J_{CC}$ couplings were calculated with a Karplus relationship reported by Li *et al.*[93] with Karplus parameters of 1.61, -0.93 and 0.55 Hz respectively. The PREs were calculated using a rotameric ensemble of MTSL paramagnetic side chains,[94] assuming a correlation time of 4 ns for the electron-nucleus interaction vector. The PREs were converted to intensity ratios assuming a transverse relaxation rate of 4 s^{-1} and an INEPT time of 10 ms. The RDCs were calculated with PALES[95] using a local alignment window of 15 residues. A local alignment window implicitly neglects the contributions of long-range interactions to the calculated RDC, but allows a rapid convergence of the calculated data.[96] For each ensemble, the alignment scaling factor for the calculated RDC was optimized on the calculated and experimental data points d_i^{calc} and d_i^{exp} that verify $d_i^{calc} \times d_i^{exp} > 0$. This allows us to optimize agreement with the experimental data while conserving physically sound scaling factors in cases where the experimental profile strongly deviates from the calculations.

The SAXS profiles were calculated with Pepsi-SAXS[97] using the following procedure. We first computed, for each ensemble conformer, a theoretical profile using a fixed scale factor of 1 and no offset. For the 18 IDPs, the values of the hydration shell contrast $d\rho$ and effective atomic radius r_0 were set for all conformers at $3.34 \cdot 10^{-3} \ e\text{Å}^{-3}$ and 1.68 Å, respectively, as

reported in a previous optimization procedure for IDPs.[98] For GS8 and polyubiquitins, we scanned for each conformer 31 values of $d\rho$ between $3.34 \cdot 10^{-5}$ and $3.34 \cdot 10^{-2}$ $e\mathring{A}^{-3}$, and 11 values of r_0 between 1.45 and 1.77 Å. For each choice of parameters, an ensemble-averaged profile was then computed across all conformers along with an optimal scale factor and offset to minimize the χ^2 with respect to the experimental SAXS profile. Finally, the $d\rho$ and r_0 pair that resulted in the minimum value of χ^2 across all the parameters tested was selected.[99]

Maximum Entropy reweighting

To improve agreement with ensemble-averaged experimental data in solution, we refined our *in silico* structural ensembles using the Maximum Entropy approach developed by Cesari *et al.*[100] Briefly, this refinement approach aims at minimally modifying the weights of each conformation in a prior ensemble to improve the fit with the observed data. This approach requires optimizing one or more regularization parameters that reflect the confidence in the experimental data. We used a protocol recently proposed[22] to seamlessly determine both individual regularization parameters, once for each experimental data type, as well as a global one. The protocol is based on reaching a target effective reduction of the initial ensemble, or Kish ratio, equal to 0.1. Experimental data used for reweighting include all the available chemical shifts, J-couplings, PREs, RDCs and SAXS profiles reported in this work. The optimal RDC scaling factors as well as the SAXS scaling and background parameters were determined by an iterative process of reweighting and parameter optimization.[98]

Details of ensemble validation

The quality of all ensembles generated in this study was assessed based on their agreement with the available experimental data. Five types of experimental data were used in this work: chemical shifts, J-couplings, RDCs, PREs, SAXS profiles. For chemical shifts, J-couplings and PREs, different experimental datasets are available corresponding to different nuclei, nuclei pair and paramagnetic label location in the sequence, respectively. The available data

for each protein are reported in Tab. S2.

For each experiment type and dataset composed of N data points, the Root Mean Square Error (RMSE) between the experimentally measured values d_i^{exp} and the ensemble-averaged values d_i^{calc} computed from the ensemble generated with the force field ff was defined as:

$$\text{RMSE}_{\text{t,s}}(ff) = \sqrt{\frac{1}{N} \sum_{i=1}^{N} \left[d_i^{exp} - d_i^{calc}(ff) \right]^2} \tag{12}$$

where t is the type of experiment (chemical shift, J-coupling, RDC, PRE, or SAXS), and s is the dataset associated with the type of experiment, such as C_α chemical shift or H_N-H_α J-coupling.

The RMSE for a given type of experiment was defined as the average RMSE across multiple datasets:

$$\text{RMSE}_{\text{t}}(ff) = \frac{1}{S} \sum_{s} \text{RMSE}_{\text{t,s}}(ff) \tag{13}$$

where S is the number of datasets associated with a given experiment type. Finally, to provide a global assessment of a given force field ff, we defined a global normalized Root Mean Square Error $\text{RMSE}_N(ff)$ relative to the random coil reference rc, as follows:

$$\text{RMSE}_N(ff) = \frac{1}{T} \sum_{t=1}^{T} \frac{1}{S} \sum_{s=1}^{S} \frac{\text{RMSE}_{t,s}(ff)}{\text{RMSE}_{t,s}(rc)} \tag{14}$$

where T is the total number of experimental data types considered. The normalization against the random coil model provides a unitless metric that allows us to provide an overall evaluation of the quality of an ensemble using all the types of experimental data available.

Error bars on the various RMSE metrics introduced above were obtained using blocking analysis. Each trajectory was first divided into five non overlapping blocks, and all analysis were repeated separately on these five blocks. The error bar associated with a given RMSE corresponds to the standard error of the mean across the five blocks.

Software and data availability

bAIes is implemented as a part of the PLUMED-ISDB module[101] in the GitHub master branch of PLUMED (`https://github.com/plumed/plumed2`).[46] The LAMMPS topologies and PLUMED input files used in our benchmark are available in PLUMED-NEST (`www.plumed-nest.org`), the public repository of the PLUMED consortium,[102] as plumID:25.014. Scripts to prepare bAIes simulations as well as a complete tutorial are available at `https://github.com/COSBlab/bAIes-IDP`.

Acknowledgments

The authors would like to thank Paul Robustelli and Guillaume Bouvignies for useful discussions. This project has received funding from the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation programme (Grant agreement No. 101086685 – bAIes).

Figures

Figure 1: Overview of bAIes. bAIes is a Bayesian modeling framework that combines a random coil model and structural information provided by AF2 to generate accurate ensembles of IDPs at atomic resolution. **a**, The random coil model is a computationally efficient force field with adjusted residue-specific backbone dihedral distributions. **b**, With the coil model only, the system of interest samples the whole conformational space with no structural preference, yielding appropriate prior ensembles. **c**, An Alphafold2 prediction allows the extraction of contact information using distance-based criteria (Methods). **d**, The C_β-C_β distograms reporting on contacts are then analyzed and transformed into distance restraints. **e**, These restraints are added on-the-fly to the prior model during the bAIes simulations, yielding ensembles enriched with AF2 information.

Figure 2: bAIes across the disorder spectrum: from random coil to partially ordered proteins. **a**, Alphafold2 model of Aβ40, colored based on pLDDT. **b**, bAIes ensemble of Aβ40. **c-e**, Validation of the Aβ40 bAIes (orange) and random coil (blue) ensembles with experimental data (grey): **c**, C_α chemical shift index; **d**, $^3J_{HNHA}$ scalar couplings; and **e**, NH residual dipolar couplings. **f**, Average Root Mean Square Error (RMSE) of the random coil (blue) and bAIes (orange) ensembles with respect to experimental backbone carbon chemical shift data (C_α, C_β, C) for six disordered IDPs lacking residual secondary structure. **g**, Alphafold2 model of PaaA2, colored based on pLDDT. **h**, bAIes ensemble of PaaA2. **i-k**, Validation of the PaaA2 bAIes (orange) and random coil (blue) ensembles with experimental data (grey): **i**, C_α chemical shift index; **j**, SAXS profile; and **k**, NH residual dipolar couplings. **l**, Average RMSE of the random coil (blue) and bAIes (orange) ensembles with respect to experimental backbone carbon chemical shift data (C_α, C_β, C) for six partially ordered IDPs. Error bars are obtained from blocking analysis (Methods).

Figure 3: bAIes noise model makes it robust to prediction errors. a, AF2 model of α-synuclein. **b**, bAIes ensemble of α-synuclein. **c-e**, Validation of the α-synuclein bAIes (orange) and random coil (blue) ensembles with experimental data (grey): **c**, C_α chemical shift index; **d**, SAXS profile; and **e**, NH residual dipolar couplings. **f**, Average Root Mean Square Error RMSE of the random coil (blue), bAIes (orange), and bAIes-N (dark red) ensembles with respect to RDCs, J-couplings, SAXS, and PREs. **g**, C_α chemical shift index profiles of α-synuclein in the random coil (blue) and bAIes-N (dark red) ensembles compared to the experimental chemical shifts (pink) measured on α-synuclein in presence of lipids.[103] **h**, Average RMSE of the random coil (blue), bAIes (orange), and bAIes-N (dark red) ensembles with respect to experimental backbone carbon chemical shift data (C_α, C_β, C) for six IDPs featuring errors in the Alphafold2 predictions. The IDP labeled in red is associated with erroneous predictions of a folded state of the IDP. The error bars are obtained from blocking analysis (Methods).

Figure 4: Accuracy of bAIes ensembles against state-of-the-art IDP models. **a**, Normalized Root Mean Square Error (RMSE_N) of the $A\beta 40$, PaaA2, α-synuclein, and ACTR random coil (blue), bAIes (orange), amber99SB-disp (grey), CHARMM36m (dark grey), CHARMM22* (black), BioEmu (hatched green) and CALVADOS2 (hatched blue) ensembles. **b-f**, Averaged error across our benchmark set for each data type and force field: **b**, chemical shifts; **c**, J-couplings; **d**, RDCs; **e**, PRE; and **f**, SAXS profile. The error bars are obtained from blocking analysis (Methods).

Figure 5: Maximum Entropy reweighting of IDPs ensembles. a, Normalized Root Mean Square Error (RMSE$_N$) of Aβ40, PaaA2, α-synuclein, and ACTR for the random coil (blue), bAIes (orange), amber99SB-disp (grey), CHARMM36m (darkgrey), CHARMM22* (black), BioEmu (hatched green) and CALVADOS2 (hatched blue) ensembles for the original (transparent) and reweighted (plain) ensembles. **b**, Alphafold2 model of ACTR, colored based on pLDDT. **c**, bAIes ensemble of ACTR. **d-f**, Validation of the reweighted ACTR bAIes (red) and amber99SB-disp (black) ensembles with experimental data (grey): **c**, C_α chemical shift index; **d**, SAXS profile; and **e**, ^{3}NH residual dipolar couplings.

Figure 6: bAIes ensembles of multi-domain proteins. a, **b**, **c**, bAIes ensemble of GS8, biubiquitin, and triubiquitin, respectively. The models are colored with the pLDDT associated with the AF2 predictions. **d**, **e**, **f**, Validation of the bAIes (orange) and reweighted bAIes (red) ensembles using SAXS data (grey) for (**d**) GS8, (**e**) biubiquitin, and (**f**) triubiquitin. Top row: SAXS intensity profiles in logarithmic scale; bottom row: difference between experimental and theoretical profiles (residuals).

References

(1) Ward, J. J.; Sodhi, J. S.; McGuffin, L. J.; Buxton, B. F.; Jones, D. T. Prediction and Functional Analysis of Native Disorder in Proteins from the Three Kingdoms of Life. *Journal of Molecular Biology* **2004**, *337*, 635–645.

(2) Holehouse, A. S.; Kragelund, B. B. The Molecular Basis for Cellular Function of Intrinsically Disordered Protein Regions. *Nature Reviews Molecular Cell Biology* **2024**, *25*, 187–211.

(3) Jensen, M. R.; Zweckstetter, M.; Huang, J.-r.; Blackledge, M. Exploring Free-Energy Landscapes of Intrinsically Disordered Proteins at Atomic Resolution Using NMR Spectroscopy. *Chemical Reviews* **2014**, *114*, 6632–6660.

(4) Maiti, S.; Singh, A.; Maji, T.; Saibo, N. V.; De, S. Experimental methods to study the structure and dynamics of intrinsically disordered regions in proteins. *Current Research in Structural Biology* **2024**, *7*, 100138.

(5) Gomes, G.-N. W.; Krzeminski, M.; Namini, A.; Martin, E. W.; Mittag, T.; Head-Gordon, T.; Forman-Kay, J. D.; Gradinaru, C. C. Conformational Ensembles of an Intrinsically Disordered Protein Consistent with NMR, SAXS, and Single-Molecule FRET. *Journal of the American Chemical Society* **2020**, *142*, 15697–15710.

(6) Naudi-Fabra, S.; Tengo, M.; Jensen, M. R.; Blackledge, M.; Milles, S. Quantitative Description of Intrinsically Disordered Proteins Using Single-Molecule FRET, NMR, and SAXS. *Journal of the American Chemical Society* **2021**, *143*, 20109–20121.

(7) Dignon, G. L.; Zheng, W.; Kim, Y. C.; Best, R. B.; Mittal, J. Sequence determinants of protein phase behavior from a coarse-grained model. *PLOS Computational Biology* **2018**, *14*, e1005941.

(8) Joseph, J. A.; Reinhardt, A.; Aguirre, A.; Chew, P. Y.; Russell, K. O.; Espinosa, J. R.; Garaizar, A.; Collepardo-Guevara, R. Physics-driven coarse-grained model for biomolecular phase separation with near-quantitative accuracy. *Nature Computational Science* **2021**, *1*, 732–743.

(9) Tesei, G.; Lindorff-Larsen, K. Improved predictions of phase behaviour of intrinsically disordered proteins by tuning the interaction range. *Open Research Europe* **2023**, *2*, 94.

(10) Valdes-Garcia, G.; Heo, L.; Lapidus, L. J.; Feig, M. Modeling Concentration-dependent Phase Separation Processes Involving Peptides and RNA via Residue-Based Coarse-Graining. *Journal of Chemical Theory and Computation* **2023**, *19*, 669–678.

(11) Cao, F.; von Bülow, S.; Tesei, G.; Lindorff-Larsen, K. A coarse-grained model for disordered and multi-domain proteins. *Protein Science* **2024**, *33*, e5172.

(12) Lotthammer, J. M.; Ginell, G. M.; Griffith, D.; Emenecker, R. J.; Holehouse, A. S. Direct prediction of intrinsically disordered protein conformational properties from sequence. *Nature Methods* **2024**, *21*, 465–476.

(13) Tesei, G.; Trolle, A. I.; Jonsson, N.; Betz, J.; Knudsen, F. E.; Pesce, F.; Johansson, K. E.; Lindorff-Larsen, K. Conformational ensembles of the human intrinsically disordered proteome. *Nature* **2024**, *626*, 897–904.

(14) Ginell, G. M.; Emenecker, R. J.; Lotthammer, J. M.; Keeley, A. T.; Plassmeyer, S. P.; Razo, N.; Usher, E. T.; Pelham, J. F.; Holehouse, A. S. Sequence-Based Prediction of Intermolecular Interactions Driven by Disordered Regions. *Science* **2025**, *388*, eadq8381.

(15) Janson, G.; Valdes-Garcia, G.; Heo, L.; Feig, M. Direct generation of protein conformational ensembles via machine learning. *Nature Communications* **2023**, *14*, 774.

(16) Janson, G.; Feig, M. Transferable deep generative modeling of intrinsically disordered protein conformations. *PLOS Computational Biology* **2024**, *20*, e1012144.

(17) Novak, B.; Lotthammer, J. M.; Emenecker, R. J.; Holehouse, A. S. Accurate predictions of conformational ensembles of disordered proteins with STARLING. *bioRxiv* **2025**, 10.1101/2025.02.14.638373.

(18) Mollica, L.; Bessa, L. M.; Hanoulle, X.; Jensen, M. R.; Blackledge, M.; Schneider, R. Binding Mechanisms of Intrinsically Disordered Proteins: Theory, Simulation, and Experiment. *Frontiers in Molecular Biosciences* **2016**, *3*.

(19) Heller, G. T.; Aprile, F. A.; Vendruscolo, M. Methods of probing the interactions between small molecules and disordered proteins. *Cellular and Molecular Life Sciences* **2017**, *74*, 3225–3243.

(20) Ruan, H.; Sun, Q.; Zhang, W.; Liu, Y.; Lai, L. Targeting intrinsically disordered proteins at the edge of chaos. *Drug Discovery Today* **2019**, *24*, 217–227.

(21) Dhar, A.; Sisk, T. R.; Robustelli, P. Ensemble docking for intrinsically disordered proteins. *bioRxiv* **2025**, 10.1101/2025.01.23.634614.

(22) Borthakur, K.; Sisk, T. R.; Panei, F. P.; Bonomi, M.; Robustelli, P. Determining accurate conformational ensembles of intrinsically disordered proteins at atomic resolution. *bioRxiv* **2024**, 10.1101/2024.10.04.616700.

(23) Robustelli, P.; Piana, S.; Shaw, D. E. Developing a molecular dynamics force field for both folded and disordered protein states. *Proceedings of the National Academy of Sciences* **2018**, *115*, E4758–E4766.

(24) Vitalis, A.; Pappu, R. V. ABSINTH: A new continuum solvation model for simulations of polypeptides in aqueous solutions. *Journal of Computational Chemistry* **2009**, *30*, 673–699.

(25) Choi, J.-M.; Pappu, R. V. Improvements to the ABSINTH Force Field for Proteins Based on Experimentally Derived Amino Acid Specific Backbone Conformational Statistics. *Journal of Chemical Theory and Computation* **2019**, *15*, 1367–1382.

(26) Varadi, M. et al. pE-DB: a database of structural ensembles of intrinsically disordered and of unfolded proteins. *Nucleic Acids Research* **2014**, *42*, D326–D335.

(27) Lazar, T. et al. PED in 2021: a major update of the protein ensemble database for intrinsically disordered proteins. *Nucleic Acids Research* **2021**, *49*, D404–D411.

(28) von Bülow, S.; Tesei, G.; Lindorff-Larsen, K. Machine learning methods to study sequence–ensemble–function relationships in disordered proteins. *Current Opinion in Structural Biology* **2025**, *92*, 103028.

(29) Jumper, J. et al. Highly accurate protein structure prediction with AlphaFold. *Nature* **2021**, *596*, 583–589.

(30) Abramson, J. et al. Accurate structure prediction of biomolecular interactions with AlphaFold 3. *Nature* **2024**, *630*, 493–500.

(31) Ruff, K. M.; Pappu, R. V. AlphaFold and Implications for Intrinsically Disordered Proteins. *Journal of Molecular Biology* **2021**, *433*, 167208.

(32) Tunyasuvunakool, K. et al. Highly accurate protein structure prediction for the human proteome. *Nature* **2021**, *596*, 590–596.

(33) Alderson, T. R.; Pritišanac, I.; Kolarić, ; Moses, A. M.; Forman-Kay, J. D. Systematic identification of conditionally folded intrinsically disordered regions by AlphaFold2. *Proceedings of the National Academy of Sciences* **2023**, *120*, e2304302120.

(34) Evans, R. et al. Protein complex prediction with AlphaFold-Multimer. *bioRxiv* **2021**, 10.1101/2021.10.04.463034.

(35) Omidi, A.; Møller, M. H.; Malhis, N.; Bui, J. M.; Gsponer, J. AlphaFold-Multimer accurately captures interactions and dynamics of intrinsically disordered protein regions. *Proceedings of the National Academy of Sciences* **2024**, *121*, e2406407121.

(36) Orand, T.; Jensen, M. R. Binding Mechanisms of Intrinsically Disordered Proteins: Insights from Experimental Studies and Structural Predictions. *Current Opinion in Structural Biology* **2025**, *90*, 102958.

(37) Brotzakis, Z. F.; Zhang, S.; Murtada, M. H.; Vendruscolo, M. AlphaFold prediction of structural ensembles of disordered proteins. *Nature Communications* **2025**, *16*, 1632.

(38) Pajkos, M.; Clerc, I.; Zanon, C.; Bernadó, P.; Cortés, J. AFflecto: A web server to generate conformational ensembles of flexible proteins from AlphaFold models. *Journal of Molecular Biology* **2025**, 169003.

(39) Lewis, S. et al. Scalable emulation of protein equilibrium ensembles with generative deep learning. *bioRxiv* **2025**, 10.1101/2024.12.05.626885.

(40) Liu, Y.; Yu, Z.; Lindsay, R. J.; Lin, G.; Chen, M.; Sahoo, A.; Hanson, S. M. ExEnDiff: An Experiment-Guided Diffusion Model for Protein Conformational Ensemble Generation. *PRX Life* **2025**, *3*, 023013.

(41) Sibille, N.; Bernadó, P. Structural Characterization of Intrinsically Disordered Proteins by the Combined Use of NMR and SAXS. *Biochemical Society Transactions* **2012**, *40*, 955–962.

(42) Camacho-Zarco, A. R.; Schnapka, V.; Guseva, S.; Abyzov, A.; Adamski, W.; Milles, S.; Jensen, M. R.; Zidek, L.; Salvi, N.; Blackledge, M. NMR Provides Unique Insight into the Functional Dynamics and Interactions of Intrinsically Disordered Proteins. *Chemical Reviews* **2022**, *122*, 9331–9356.

(43) Bernadó, P.; Blanchard, L.; Timmins, P.; Marion, D.; Ruigrok, R. W. H.; Blackledge, M. A structural model for unfolded proteins from residual dipolar couplings and small-angle x-ray scattering. *Proceedings of the National Academy of Sciences* **2005**, *102*, 17002–17007.

(44) Ozenne, V.; Bauer, F.; Salmon, L.; Huang, J.-r.; Jensen, M. R.; Segard, S.; Bernadó, P.; Charavay, C.; Blackledge, M. Flexible-meccano: a tool for the generation of explicit ensemble descriptions of intrinsically disordered proteins and their associated experimental observables. *Bioinformatics* **2012**, *28*, 1463–1470.

(45) Jeffreys, S. H.; Jeffreys, S. H. *The Theory of Probability*; Oxford University Press, 1998.

(46) Tribello, G. A.; Bonomi, M.; Branduardi, D.; Camilloni, C.; Bussi, G. PLUMED 2: New feathers for an old bird. *Computer Physics Communications* **2014**, *185*, 604–613.

(47) Roche, J.; Shen, Y.; Lee, J. H.; Ying, J.; Bax, A. Monomeric A1–40 and A1–42 Peptides in Solution Adopt Very Similar Ramachandran Map Distributions That Closely Resemble Random Coil. *Biochemistry* **2016**, *55*, 762–775.

(48) Sterckx, Y. G. J.; Volkov, A. N.; Vranken, W. F.; Kragelj, J.; Jensen, M. R.; Buts, L.; Garcia-Pino, A.; Jové, T.; Van Melderen, L.; Blackledge, M.; van Nuland, N. A. J.; Loris, R. Small-Angle X-Ray Scattering- and Nuclear Magnetic Resonance-Derived Conformational Ensemble of the Highly Flexible Antitoxin PaaA2. *Structure* **2014**, *22*, 854–865.

(49) Emamzadeh, F. N. Alpha-synuclein structure, functions, and interactions. *Journal of Research in Medical Sciences* **2016**, *21*, 29.

(50) Dedmon, M. M.; Lindorff-Larsen, K.; Christodoulou, J.; Vendruscolo, M.; Dobson, C. M. Mapping Long-Range Interactions in -Synuclein using Spin-Label NMR

and Ensemble Molecular Dynamics Simulations. *Journal of the American Chemical Society* **2005**, *127*, 476–477.

(51) Bertoncini, C. W.; Jung, Y.-S.; Fernandez, C. O.; Hoyer, W.; Griesinger, C.; Jovin, T. M.; Zweckstetter, M. Release of long-range tertiary interactions potentiates aggregation of natively unstructured -synuclein. *Proceedings of the National Academy of Sciences* **2005**, *102*, 1430–1435.

(52) Kuban, V.; Macek, P.; Hritz, J.; Nechvatalova, K.; Nedbalcova, K.; Faldyna, M.; Sebo, P.; Zidek, L.; Bumba, L. Structural Basis of Ca2+-Dependent Self-Processing Activity of Repeat-in-Toxin Proteins. *mBio* **2020**, *11*, 10.1128/mbio.00226–20.

(53) Bezsonova, I.; Singer, A.; Choy, W.-Y.; Tollinger, M.; Forman-Kay, J. D. Structural Comparison of the Unstable drkN SH3 Domain and a Stable Mutant,. *Biochemistry* **2005**, *44*, 15550–15560.

(54) Huang, J.; Rauscher, S.; Nawrocki, G.; Ran, T.; Feig, M.; de Groot, B. L.; Grubmüller, H.; MacKerell, A. D. CHARMM36m: an improved force field for folded and intrinsically disordered proteins. *Nature Methods* **2017**, *14*, 71–73.

(55) Piana, S.; Lindorff-Larsen, K.; Shaw, D. E. How Robust Are Protein Folding Simulations with Respect to Force Field Parameterization? *Biophysical Journal* **2011**, *100*, L47–L49.

(56) Shaw, D. E. et al. Anton, a special-purpose machine for molecular dynamics simulation. *Commun. ACM* **2008**, *51*, 91–97.

(57) Moses, D.; Guadalupe, K.; Yu, F.; Flores, E.; Perez, A. R.; McAnelly, R.; Shamoon, N. M.; Kaur, G.; Cuevas-Zepeda, E.; Merg, A. D.; Martin, E. W.; Holehouse, A. S.; Sukenik, S. Structural biases in disordered proteins are prevalent in the cell. *Nature Structural & Molecular Biology* **2024**, *31*, 283–292.

(58) Jussupow, A.; Messias, A. C.; Stehle, R.; Geerlof, A.; Solbak, S. M. ; Paissoni, C.; Bach, A.; Sattler, M.; Camilloni, C. The dynamics of linear polyubiquitin. *Science Advances* **2020**, *6*, eabc3786.

(59) Sen, S.; Hoff, S. E.; Morozova, T. I.; Schnapka, V.; Bonomi, M. Advancing in silico drug design with Bayesian refinement of AlphaFold models. *bioRxiv* **2025**, 10.1101/2025.06.25.661454.

(60) Bonomi, M.; Camilloni, C.; Cavalli, A.; Vendruscolo, M. Metainference: A Bayesian Inference Method for Heterogeneous Systems. *Science Advances* **2016**, *2*, e1501177.

(61) Mirdita, M.; Schütze, K.; Moriwaki, Y.; Heo, L.; Ovchinnikov, S.; Steinegger, M. ColabFold: Making Protein Folding Accessible to All. *Nature Methods* **2022**, *19*, 679–682.

(62) Salmon, L.; Nodet, G.; Ozenne, V.; Yin, G.; Jensen, M. R.; Zweckstetter, M.; Blackledge, M. NMR Characterization of Long-Range Order in Intrinsically Disordered Proteins. *Journal of the American Chemical Society* **2010**, *132*, 8407–8418.

(63) Bussi, G.; Laio, A. Using Metadynamics to Explore Complex Free-Energy Landscapes. *Nature Reviews Physics* **2020**, *2*, 200–212.

(64) del Alamo, D.; Sala, D.; Mchaourab, H. S.; Meiler, J. Sampling alternative conformational states of transporters and receptors with AlphaFold2. *eLife* **2022**, *11*, e75751.

(65) Wayment-Steele, H. K.; Ojoawo, A.; Otten, R.; Apitz, J. M.; Pitsawong, W.; Hömberger, M.; Ovchinnikov, S.; Colwell, L.; Kern, D. Predicting Multiple Conformations via Sequence Clustering and AlphaFold2. *Nature* **2024**, *625*, 832–839.

(66) Passaro, S.; Corso, G.; Wohlwend, J.; Reveiz, M.; Thaler, S.; Somnath, V. R.; Getz, N.; Portnoi, T.; Roy, J.; Stark, H.; Kwabi-Addo, D.; Beaini, D.; Jaakkola, T.; Barzilay, R.

Boltz-2: Towards Accurate and Efficient Binding Affinity Prediction. *bioRxiv* **2025**, 10.1101/2025.06.14.659707.

(67) Wankowicz, S. A.; Bonomi, M. From Possibility to Precision in Macromolecular Ensemble Prediction. *arXiv* **2025**, https://doi.org/https://doi.org/10.48550/arXiv.2505.01919.

(68) Suzek, B. E.; Wang, Y.; Huang, H.; McGarvey, P. B.; Wu, C. H.; the UniProt Consortium UniRef Clusters: A Comprehensive and Scalable Alternative for Improving Sequence Similarity Searches. *Bioinformatics* **2015**, *31*, 926–932.

(69) Mitchell, A. L. et al. MGnify: The Microbiome Analysis Resource in 2020. *Nucleic Acids Research* **2020**, *48*, D570–D578.

(70) Steinegger, M.; Meier, M.; Mirdita, M.; Vöhringer, H.; Haunsberger, S. J.; Söding, J. HH-suite3 for Fast Remote Homology Detection and Deep Protein Annotation. *BMC Bioinformatics* **2019**, *20*, 473.

(71) Remmert, M.; Biegert, A.; Hauser, A.; Söding, J. HHblits: Lightning-Fast Iterative Protein Sequence Searching by HMM-HMM Alignment. *Nature Methods* **2012**, *9*, 173–175.

(72) Mirdita, M.; von den Driesch, L.; Galiez, C.; Martin, M. J.; Söding, J.; Steinegger, M. Uniclust Databases of Clustered and Deeply Annotated Protein Sequences and Alignments. *Nucleic Acids Research* **2017**, *45*, D170–D176.

(73) Eddy, S. R. Accelerated Profile HMM Searches. *PLOS Computational Biology* **2011**, *7*, e1002195.

(74) Lindorff-Larsen, K.; Piana, S.; Palmo, K.; Maragakis, P.; Klepeis, J. L.; Dror, R. O.; Shaw, D. E. Improved side-chain torsion potentials for the Amber ff99SB protein force field. *Proteins: Structure, Function, and Bioinformatics* **2010**, *78*, 1950–1958.

(75) MacKerell, A. D. J.; Feig, M.; Brooks, C. L. Improved Treatment of the Protein Backbone in Empirical Force Fields. *Journal of the American Chemical Society* **2004**, *126*, 698–699.

(76) Ting, D.; Wang, G.; Shapovalov, M.; Mitra, R.; Jordan, M. I.; Jr, R. L. D. Neighbor-Dependent Ramachandran Probability Distributions of Amino Acids Developed from a Hierarchical Dirichlet Process Model. *PLOS Computational Biology* **2010**, *6*, e1000763.

(77) Mantsyzov, A. B.; Shen, Y.; Lee, J. H.; Hummer, G.; Bax, A. MERA: A Webserver for Evaluating Backbone Torsion Angle Distributions in Dynamic and Disordered Proteins from NMR Data. *Journal of Biomolecular NMR* **2015**, *63*, 85–95.

(78) Kabsch, W.; Sander, C. Dictionary of protein secondary structure: Pattern recognition of hydrogen-bonded and geometrical features. *Biopolymers* **1983**, *22*, 2577–2637.

(79) Barducci, A.; Bussi, G.; Parrinello, M. Well-Tempered Metadynamics: A Smoothly Converging and Tunable Free-Energy Method. *Physical Review Letters* **2008**, *100*, 020603.

(80) Thompson, A. P.; Aktulga, H. M.; Berger, R.; Bolintineanu, D. S.; Brown, W. M.; Crozier, P. S.; in 't Veld, P. J.; Kohlmeyer, A.; Moore, S. G.; Nguyen, T. D.; Shan, R.; Stevens, M. J.; Tranchida, J.; Trott, C.; Plimpton, S. J. LAMMPS - a flexible simulation tool for particle-based materials modeling at the atomic, meso, and continuum scales. *Computer Physics Communications* **2022**, *271*, 108171.

(81) Rieping, W.; Habeck, M.; Nilges, M. Inferential Structure Determination. *Science* **2005**, *309*, 303–306.

(82) Paissoni, C.; Jussupow, A.; Camilloni, C. Determination of Protein Structural Ensembles by Hybrid-Resolution SAXS Restrained Molecular Dynamics. *Journal of Chemical Theory and Computation* **2020**, *16*, 2825–2834.

(83) Hoff, S. E.; Thomasen, F. E.; Lindorff-Larsen, K.; Bonomi, M. Accurate Model and Ensemble Refinement Using Cryo-Electron Microscopy Maps and Bayesian Inference. *PLOS Computational Biology* **2024**, *20*, e1012180.

(84) Kamisetty, H.; Ovchinnikov, S.; Baker, D. Assessing the utility of coevolution-based residue–residue contact predictions in a sequence- and structure-rich era. *Proceedings of the National Academy of Sciences* **2013**, *110*, 15674–15679.

(85) Shirts, M. R.; Klein, C.; Swails, J. M.; Yin, J.; Gilson, M. K.; Mobley, D. L.; Case, D. A.; Zhong, E. D. Lessons learned from comparing molecular dynamics engines on the SAMPL5 dataset. *Journal of Computer-Aided Molecular Design* **2017**, *31*, 147–161.

(86) Bussi, G.; Donadio, D.; Parrinello, M. Canonical sampling through velocity rescaling. *The Journal of Chemical Physics* **2007**, *126*, 014101.

(87) Eastman, P. et al. OpenMM 8: Molecular Dynamics Simulation with Machine Learning Potentials. *The Journal of Physical Chemistry B* **2024**, *128*, 109–116.

(88) Rotkiewicz, P.; Skolnick, J. Fast procedure for reconstruction of full-atom protein models from reduced representations. *Journal of Computational Chemistry* **2008**, *29*, 1460–1465.

(89) Abraham, M. J.; Murtola, T.; Schulz, R.; Páll, S.; Smith, J. C.; Hess, B.; Lindahl, E. GROMACS: High performance molecular simulations through multi-level parallelism from laptops to supercomputers. *SoftwareX* **2015**, *1-2*, 19–25.

(90) Visani, G. M.; Galvin, W.; Pun, M.; Nourmohammad, A. H-Packer: Holographic Rotationally Equivariant Convolutional Neural Network for Protein Side-Chain Packing. Proceedings of the 18th Machine Learning in Computational Biology meeting. 2024; pp 230–249.

(91) Shen, Y.; Bax, A. SPARTA+: a modest improvement in empirical NMR chemical shift prediction by means of an artificial neural network. *Journal of Biomolecular NMR* **2010**, *48*, 13–22.

(92) Vögeli, B.; Ying, J.; Grishaev, A.; Bax, A. Limits on Variations in Protein Backbone Dynamics from Precise Measurements of Scalar Couplings. *Journal of the American Chemical Society* **2007**, *129*, 9377–9385.

(93) Li, F.; Lee, J. H.; Grishaev, A.; Ying, J.; Bax, A. High Accuracy of Karplus Equations for Relating Three-Bond J Couplings to Protein Backbone Torsion Angles. *ChemPhysChem* **2015**, *16*, 572–578.

(94) Polyhach, Y.; Bordignon, E.; Jeschke, G. Rotamer libraries of spin labelled cysteines for protein studies. *Physical Chemistry Chemical Physics* **2011**, *13*, 2356–2366.

(95) Zweckstetter, M. NMR: prediction of molecular alignment from structure using the PALES software. *Nature Protocols* **2008**, *3*, 679–690.

(96) Marsh, J. A.; Baker, J. M. R.; Tollinger, M.; Forman-Kay, J. D. Calculation of Residual Dipolar Couplings from Disordered State Ensembles Using Local Alignment. *Journal of the American Chemical Society* **2008**, *130*, 7804–7805.

(97) Grudinin, S.; Garkavenko, M.; Kazennov, A. Pepsi-SAXS: an adaptive method for rapid and accurate computation of small-angle X-ray scattering profiles. *Acta Crystallographica Section D: Structural Biology* **2017**, *73*, 449–464.

(98) Pesce, F.; Lindorff-Larsen, K. Refining conformational ensembles of flexible proteins against small-angle x-ray scattering data. *Biophysical Journal* **2021**, *120*, 5124–5135.

(99) Larsen, A. H.; Wang, Y.; Bottaro, S.; Grudinin, S.; Arleth, L.; Lindorff-Larsen, K. Combining molecular dynamics simulations with small-angle X-ray and neutron scat-

tering data to study multi-domain proteins in solution. *PLOS Computational Biology* **2020**, *16*, e1007870.

(100) Cesari, A.; Reißer, S.; Bussi, G. Using the Maximum Entropy Principle to Combine Simulations and Solution Experiments. *Computation* **2018**, *6*, 15.

(101) Bonomi, M.; Camilloni, C. Integrative structural and dynamical biology with PLUMED-ISDB. *Bioinformatics* **2017**, *33*, 3999–4000.

(102) Bonomi, M. et al. Promoting Transparency and Reproducibility in Enhanced Molecular Simulations. *Nature Methods* **2019**, *16*, 670–673.

(103) Rao, J. N.; Kim, Y. E.; Park, L. S.; Ulmer, T. S. Effect of Pseudorepeat Rearrangement on -Synuclein Misfolding, Vesicle Binding, and Micelle Binding. *Journal of Molecular Biology* **2009**, *390*, 516–529.