nature computational science

ARTICLES

https://doi.org/10.1038/s43588-021-00155-3

Physics-driven coarse-grained model for biomolecular phase separation with near-quantitative accuracy

Jerelle A. Joseph[1,2,3,4] ✉, Aleks Reinhardt[1,4] ✉, Anne Aguirre[1], Pin Yu Chew[1], Kieran O. Russell[1], Jorge R. Espinosa[2], Adiran Garaizar[2] and Rosana Collepardo-Guevara[1,2,3] ✉

Various physics- and data-driven sequence-dependent protein coarse-grained models have been developed to study biomolecular phase separation and elucidate the dominant physicochemical driving forces. Here we present Mpipi, a multiscale coarse-grained model that describes almost quantitatively the change in protein critical temperatures as a function of amino acid sequence. The model is parameterized from both atomistic simulations and bioinformatics data and accounts for the dominant role of π–π and hybrid cation–π/π–π interactions and the much stronger attractive contacts established by arginines than lysines. We provide a comprehensive set of benchmarks for Mpipi and seven other residue-level coarse-grained models against experimental radii of gyration and quantitative in vitro phase diagrams, demonstrating that Mpipi predictions agree well with experiments on both fronts. Moreover, Mpipi can account for protein–RNA interactions, correctly predicts the multiphase behavior of a charge-matched poly-arginine/poly-lysine/RNA system, and recapitulates experimental liquid–liquid phase separation trends for sequence mutations on FUS, DDX4 and LAF-1 proteins.

Under certain conditions, macromolecules within cells demix into membraneless organelles. These organelles, often termed biomolecular condensates, are sustained by the physicochemical process of liquid–liquid phase separation (LLPS)[1,2]. The ensuing condensates play important roles in cellular function as well as dysfunction[3], so delineating the mechanisms of intracellular LLPS is now an active area of research. Intracellular LLPS is principally driven by biomolecular multivalency, that is, the ability of multidomain proteins, intrinsically disordered proteins/regions (IDPs/IDRs), ribonucleic acids (RNA) and chromatin to engage multiple interaction partners simultaneously. This multivalency is, in turn, predominantly encoded in the chemical makeup of the macromolecules in question. It is well established that both hydrophobic and electrostatic interactions are important drivers of biomolecular LLPS, including charge–charge, π–π, cation–π, dipole–dipole and non-polar interactions. Additionally, there is strong evidence that certain chemical building blocks have a bigger stake in biomolecular LLPS than others[4–7]. Biophysical models for studying LLPS must therefore be able to capture the correct balance between these myriad driving forces. In this Article we aim to achieve precisely this.

Together, the stickers-and-spacers framework of Pappu and colleagues[8,9] and the quantitative experimental phase diagrams of Mittag and colleagues position aromatic residues as the chief drivers of biomolecular phase separation[4,10]. Moreover, it is evident that even within the subset of aromatic residues, tyrosine, for example, is a stronger contributor than phenylalanine to LLPS stability[7,10–12], perhaps because it has more side-chain binding modes than phenylalanine. In addition to forming aromatic π–π contacts, tyrosine can form strong hydrogen bonds via its phenol group.

The dominant role of π–π interactions in LLPS was also suggested by Vernon et al.[13], who, via a comprehensive survey of the Protein Data Bank (PDB), identified an abundance of planar π–π contacts involving not only aromatic but also non-aromatic residues in protein structures. Additionally, in recent work, π–π interactions have emerged as a major driver of LLPS at both low and high salt concentration[7]. Specifically, our atomistic simulations revealed that, for proteins, the strongest pairwise interactions arise when the two amino acids in question both possess π electrons in their side chains, including both aromatic or non-aromatic residues with sp^2-hybridized groups[7].

The role of hydrophobic π–π contacts in LLPS is even more obvious when we consider differences in the strengths of cation–π interactions. Notably, cationic residues, namely arginine and lysine, have been shown to act as unequal contributors to LLPS[6,7,14–16], with arginine establishing appreciably stronger cation–π and charge–charge interactions due to the presence of the guanidinium group[6,7,13,16–18]. Furthermore, Pappu and colleagues have recently demonstrated that the free energy of hydration of arginine is considerably less favorable than that of lysine; thus, although they both carry the same charge, arginine is substantially more hydrophobic than lysine[16]. Given that π-based contacts play such a dominant role in biomolecular LLPS, achieving the correct balance of these interactions is essential for making quantitative predictions.

Complementing experimental and theoretical work, computer simulations have provided a unique lens for probing biomolecular LLPS. Because LLPS is a collective phenomenon, coarse-graining is essential to reduce the system dimensionality while retaining essential physicochemical information and allowing sufficient sampling of the phase space in computationally tractable timescales. There are numerous possible approaches for parameterizing biomolecular coarse-grained models[19], from bottom-up strategies that rely on higher-resolution models[9,10,20–22], to knowledge-driven approaches

[1]Yusuf Hamied Department of Chemistry, University of Cambridge, Cambridge, UK. [2]Department of Physics, University of Cambridge, Cambridge, UK. [3]Department of Genetics, University of Cambridge, Cambridge, UK. [4]These authors contributed equally: Jerelle A. Joseph, Aleks Reinhardt.
✉e-mail: jaj52@cam.ac.uk; ar732@cam.ac.uk; rc597@cam.ac.uk

that aim to reproduce experimental properties using a data-based parameterization[23–25], top-down strategies that account for emergent behavior by approximating fundamental physical forces[26,27] or to combinations of these[28,29]. Coarse-grained models can also be broadly classed as system-specific, bottom-up parameterizations focusing on finding an optimum representation for a particular system using fine-grained simulations as a reference, often derived in a systematic way using, for instance, iterative Boltzmann inversion[30,31] or force matching[32,33], and transferable, either bottom-up or top-down parameterizations, aiming to achieve a generally applicable potential.

Developing a coarse-grained model involves invoking multiple approximations and making design decisions (for example, the type of model, resolution, bead characteristics and types of interaction) that are more or less appropriate depending on the question being investigated. As discussed by Choi et al.[9], there is no unequivocal reason that makes one scheme intrinsically superior to others—each approach has its advantages and drawbacks. For instance, systematic multiscale coarse-graining from higher-resolution models[21,34] by construction results in an excellent description of the system under investigation and allows us to work out precisely what underlying building blocks have been coarse-grained. However, system-specific coarse-graining does not generally result in a unique solution[35] and requires sufficiently long simulations of the entire system of interest to be run with an expensive high-resolution potential. Indeed, bulk phase behavior can be substantially different between a machine-learned potential and the underlying quantum-mechanical potential-energy surface, even for systems much simpler than biomolecules[36], illustrating the formidable challenge of this approach. Similarly, transferable data-driven or machine-learning-based approaches can give excellent agreement with the data from which they were parameterized, but, because in such high-dimensional problems many solutions are similarly good, careful curation is required to obtain statistically meaningful results for a specific system, and even these may still not be transferable to similar molecules[37]. On the other hand, a transferable physics-based approach to interaction parameters provides us with a simple way of rationalizing complex behavior based on relatively simple interactions. However, it risks introducing our biases regarding which interactions are important into the predictions of the model, and the predictions and rationalizations of observed behaviors with such models can therefore be to some extent self-fulfilling.

Various coarse-graining strategies have now been applied to gain insights into the problem of biomolecular LLPS. For example, the mean-field stickers-and-spacers model can be parameterized to reach quantitative agreement with experimental phase diagrams of specific proteins, providing a tool to dissect the driving forces behind the observations[4,8–10]. A different approach, pioneered by Mittal and colleagues[26], combines residue-level coarse-grained models with direct-coexistence simulations[38], offering a transferable method to predict protein phase diagrams and augmenting our ability to link molecular sequences to their experimental phase behavior.

Inspired by previous computational work and guided by the accumulated knowledge of the LLPS interaction landscape, we set out to design a chemically accurate coarse-grained model for predicting biomolecular LLPS. Specifically, the model aims to achieve optimal strengths of protein–protein and protein–RNA interactions. We demonstrate that our model can accurately predict biomolecular phase separation while achieving quantitative agreement with experiments, recapitulating post-translational modification effects and even capturing more complex features such as sequence-dependent multiphasic compartmentalization. Simulations using our model are particularly simple to set up because all its components are already implemented in open-source software.

In what follows, we first describe the design of our multiscale π–π model (termed 'Mpipi') for probing the phase separation of biomolecules. We outline the use of atomistic potential-of-mean-force (PMF) calculations coupled with bioinformatics data to yield a chemically accurate interaction scale for coarse-grained simulations (Fig. 1). We then assess the balance of key interactions in the Mpipi model alongside other commonly used residue-based coarse-grained models. Finally, we present benchmarks for several LLPS systems and directly compare our predictions with other models and quantitative experimental phase diagrams and other experiments.

Results

Designing coarse-grained models for biomolecular LLPS. We have designed a residue-level coarse-grained model for predicting biomolecular phase behavior (Fig. 1a–c) that captures the fundamental van der Waals and electrostatic interactions of a top-down approach and the interaction strengths obtained from bottom-up atomistic simulations and bioinformatics data. In the Mpipi model, each amino (or nucleic) acid is mapped onto a unique bead (Fig. 1b) based on simulation and experimental data. Following Dignon et al.[26], the potential energy of molecules is computed as the sum of a harmonic bond energy, Debye–Hückel and short-ranged energy terms (Methods), which account for π–π, cation–π and other non-ionic interactions. The main differences between the Mpipi model and other sequence-based coarse-grained models for LLPS are (1) the functional form of short-ranged terms, (2) the parameterization of short-ranged interactions and (3) the relative contribution of long-ranged electrostatics and short-ranged terms to the total energy. Specifically, for short-range interactions, we use the recently developed Wang–Frenkel[39] pair potential (Fig. 1b and Methods), which accounts for the key physical interactions, namely a short-ranged excluded-volume repulsion and a longer-ranged attraction that gradually decays to zero. The Wang–Frenkel potential has several advantages[39] over the Lennard-Jones-like potentials that are commonly adopted in molecular simulations. We outline these and how our model is fitted within the Wang–Frenkel framework in Methods.

When deciding on the energy scale for short-ranged interactions, our main objective is to achieve the correct balance of π–π and non-π-based contacts. To this end, we combine bioinformatics data and atomistic short-ranged free-energy estimates (Fig. 2a–d). In Methods, we explain our parameterization of short-ranged pairwise contacts and long-ranged charge–charge interactions (Fig. 2e).

Cation–π, π–π and non-π interactions in residue-level models. To validate our model parameters, we first compared the Mpipi model with seven other residue-level coarse-grained models for LLPS, namely the KH (Kim–Hummer)[26], HPS-KR (hydrophobicity scale)[26], FB-HPS[24] and HPS-Urry[27] models, as well as the HPS model with augmented cation–π interactions (schemes (i) and (ii))[40]. We also included analyses for TSCL-M2, the M2 parameter set of Tesei et al.[25] proposed while our work was under review. Das et al.[40] recently provided a thorough comparison of the KH, HPS-KR, HPS+cation–π(i) and HPS+cation–π(ii) models. In the following, we briefly discuss the key features of all models and then evaluate them in terms of the balance of π–π, cation–π and non-π-based interactions.

A key difference between Mpipi and other residue-level models is the parameterization of short-ranged interactions. In the KH model[26,41], short-ranged interactions (ε_{ij}) are based on the residue contact statistics of folded proteins. The energy scale of the KH model was tuned to approximately reproduce the radii of gyration of selected unfolded proteins/IDPs[26]; here, we utilize parameter set D[26] for interactions involving disordered proteins. The KH model has been used successfully to predict LLPS propensities for variants of the N-terminal domain (NTD) of the DEAD-Box helicase 4 (DDX4) protein[40] and to describe qualitatively the phase behavior of the proteins fused in sarcoma (FUS) and lethal-and-feminizing-1 (LAF-1)[26].

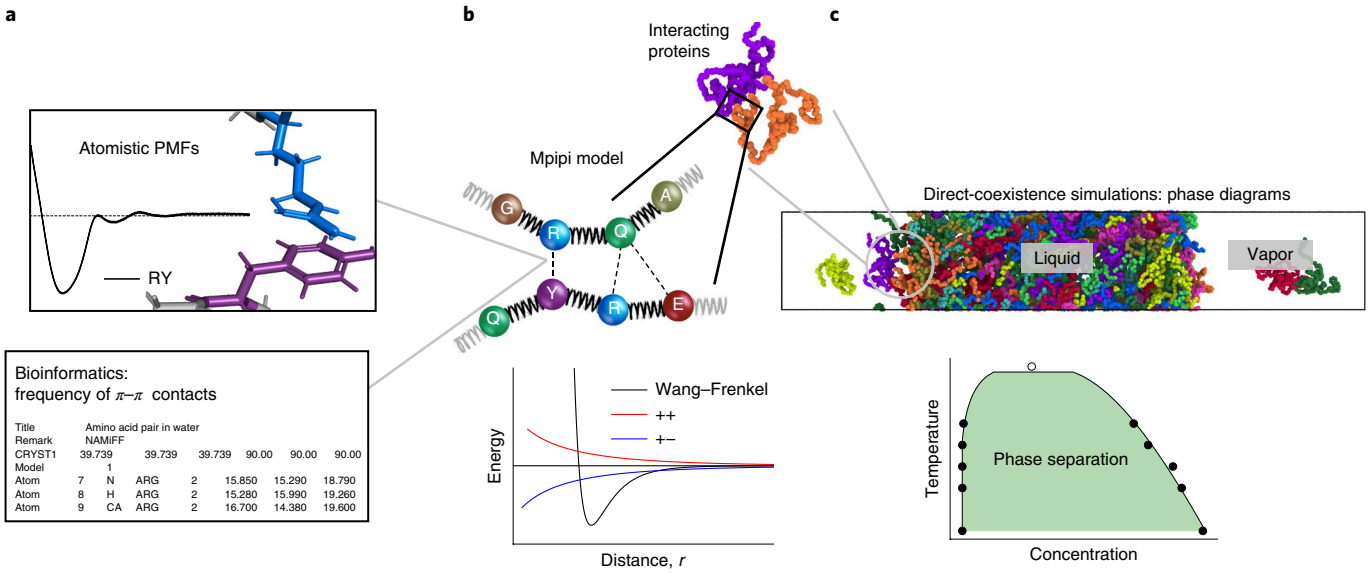

Fig. 1 | Designing a coarse-grained model for LLPS from PMF calculations and bioinformatics data. a, Top: PMFs of selected amino acid (or nucleic acid) pairs are computed in all-atom simulations with explicit solvent and ions. The computed curves provide a free energy of interaction for the pair in question. Bottom: the frequencies of π–π contacts for amino acids are obtained from bioinformatics work[13]. Together, these data are used to parameterize the pairwise interactions in the Mpipi model. **b**, Top: in the Mpipi model, each amino acid (or nucleic acid) is represented by a unique bead. The potential energy is computed as a sum of short-ranged pairwise terms, electrostatic interactions and bonded interactions modeled as harmonic springs. Bottom: short-ranged pairwise and electrostatic interactions are computed via the Wang–Frenkel potential[39] (black curve) and a Coulomb term with Debye–Hückel screening (red and blue curves), respectively. **c**, To study biomolecular phase behavior, we use direct-coexistence molecular dynamics simulations[38] and compute phase diagrams in the temperature–concentration (or density) space.

The next model, HPS-KR[26], is perhaps the most widely used sequence-based continuum model for studying biomolecular LLPS. In this model, short-ranged interactions are based on the hydrophobicity scale of Kapcha and Rossky[42], and each amino acid is assigned a λ_i value that accounts for its hydrophobicity, and residue–residue contacts (λ_{ij}) are determined by the arithmetic mean of the λ_i values of each residue[43]. Additionally, the absolute energy scale of the model was optimized to reproduce the experimental radii of gyration (R_g) of an IDP subset. However, as previously noted[40], the HPS-KR model is inconsistent with experimental data when accounting for the balance between Arg and Lys interactions. An improved version of the HPS model, HPS-Urry[27], was recently parameterized, which employs instead the hydrophobicity scale of Urry et al.[44] to determine λ_{ij}. Moreover, two free parameters (μ and Δ) are introduced to scale and shift the λ_{ij} values; these are optimized to reproduce the experimental R_g (ref. [27]). Recently, Dannenhoffer-Lafage and Best[24] also reparameterized the short-ranged interactions in the HPS-KR model using machine-learning techniques. Their model, FB-HPS, was optimized against the experimental R_g of unfolded, phase-separating and intrinsically disordered proteins.

Before these studies, Das et al.[40] augmented the HPS-KR model so as to better account for the cation–π interactions. They presented two schemes: scheme (i), where Arg/Lys–π interactions are scaled uniformly, and scheme (ii), where they vary. Notably, the authors commented that, despite these changes, the augmented models fail to capture fully the experimental LLPS propensities of their test set of proteins[40]. In another study, it was demonstrated that the HPS+cation–π(i) model can reasonably reproduce the experimental trends of selected RNA binding proteins[45]. Here, we have considered these augmented models to achieve a more complete view of how cation–π interactions contribute to biomolecular LLPS.

Recently, Tesei et al.[25] used experimental data to reparameterize the hydrophobicity scale of HPS-KR via a Bayesian parameter-learning procedure. The M2 parameter set predicted well both the single-molecule and collective behaviors of the tested IDPs, so we have included benchmarks for this parameter set in our work.

Figure 3 summarizes the relative contributions of selected π–π and non-π-based interactions for the residue-level coarse-grained models assessed in this work (Supplementary Figs. 5 and 6 present the results for the HPS+cation–π(i) and TSCL-M2 models). The relative interaction strengths are obtained by computing the integral of the curves of the short-ranged potential (with consistent limits of σ–3σ). In the Mpipi, KH, FB-HPS and TSCL-M2 models, aromatic residue pairs (Fig. 3a,b,e, magenta bars and Supplementary Fig. 6a) are generally considerably stronger than residue pairs not involving π contacts (Fig. 3a,b,e, dark yellow bars, and Supplementary Fig. 6a). Hence, consistent with the stickers-and-spacers framework, YY and FF are expected to act as stickers, while AA, SS and PP should behave as spacers in these models. Interestingly, in the FB-HPS model, glycine (fig. 4 of ref. [24]), which is normally classified as a spacer, has an interaction strength that is stronger than even the aromatic residues. Although this result is attributed to glycine forming strong backbone π–π contacts[24], mutational studies have consistently found that replacing sticker-like residues with Gly substantially suppresses biomolecular LLPS[10,11]. The stronger contacts for Gly arising from the machine-learning algorithm optimization[24] may be a result of how commonly occurring glycine is in many proteins, particularly those for which experimental radii of gyration are available.

With regard to Tyr versus Phe, a survey of the PDB[13], our atomistic PMF calculations (Fig. 2d) and experiments[10–12] all suggest that Phe–Phe contacts are weaker than Tyr–Tyr ones. By contrast, the KH, HPS-KR, FB-HPS, HPS+cation–π(i) and HPS+cation–π(ii) models all predict stronger Phe–Phe contacts than Tyr–Tyr interactions (Fig. 3 and Supplementary Fig. 5). The trend for the KH model is particularly striking, with the weighted interaction energy of FF predicted to be about twice that of YY (Fig. 3b). Taken together, we do not expect these models to reproduce LLPS propensities faithfully as far as Tyr versus Phe mutations are concerned.

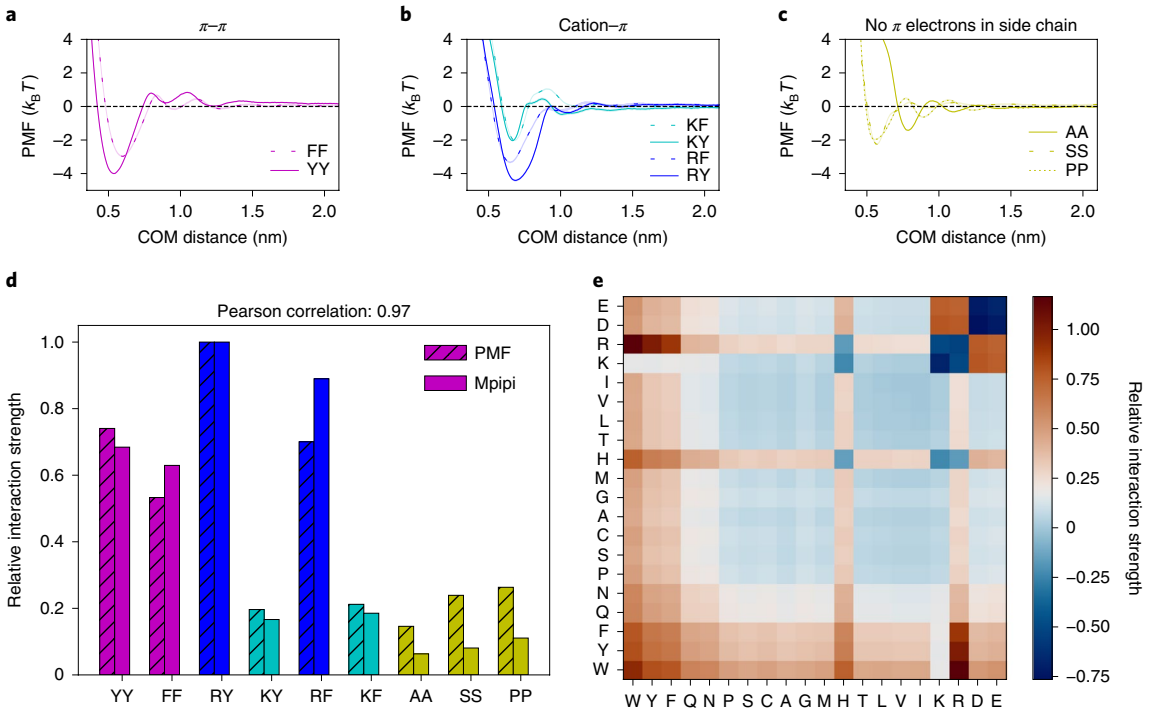

Fig. 2 | Obtaining the correct balance of π–π and non-π-based interactions in the Mpipi model. a–**c**, PMF calculations at 150 mM NaCl salt concentration for π–π (**a**), cation–π (**b**) and non-π-based (**c**) interactions, respectively, as a function of the center-of-mass (COM) distance. Statistical errors (mean ± s.d.) are given as error bands, and are only just larger than the line width. These were computed via Bayesian bootstrapping of three independent simulations. Each pair is labeled using one-letter amino acid codes (Supplementary Table 1). **d**, Comparison of the relative interaction strengths of selected residue pairs (Supplementary Table 1) from the PMF calculations with those implemented in the Mpipi model, relative to the Arg–Tyr (RY) interaction. Values are computed by taking the integral of the curves in **a**–**c** and the integral of the Wang–Frenkel potential only (between σ and 3σ) for the PMF and Mpipi sets, respectively. For the PMF data, only the leftmost well is considered. These correspond to mean energies in the high-temperature limit. **e**, Summary of relative interaction strengths in the Mpipi model. These relative interaction strengths include electrostatic interactions and are computed by numerically integrating equation (8) and normalizing the result by the RY interaction strength.

We also examine the relative strengths of cation–π interactions in the coarse-grained models, again focusing on the contributions of cation–π contacts versus aromatic π–π contacts and Arg–π contacts compared to Lys–π ones. The HPS+cation–π(ii) (Fig. 3d) and TSCL-M2 (Supplementary Fig. 6a) models are most similar to the Mpipi model in terms of the relative contributions of Arg–π and Lys–π contacts. Although in the KH model, Arg–π interactions are also stronger than Lys–π ones, the overall strength of these is low, which makes Arg–π interactions closer in strength to spacer-type interactions (Fig. 3b); accordingly, the dominant role of these interactions may not be properly accounted for in the KH model. The HPS-Urry model also captures the overall trend for Arg–π versus Lys–π contacts (Fig. 3f); in addition, the weights of these interactions are more similar than those encoded in Mpipi, HPS+cation–π(ii) and TSCL-M2. The opposite trend is found in both the HPS-KR and the FB-HPS models, where Lys–π interactions are now stronger than Arg–π interactions. Moreover, in the HPS-KR model, non-π-based interactions are comparable to (or even stronger than) cation–π interactions (Fig. 3c). We therefore speculate that the HPS-KR model might represent an upper bound in terms of predicting the LLPS propensities of proteins. Strikingly, in the HPS+cation–π(i) model, cation–π interactions convincingly dominate all other types of interaction (Supplementary Fig. 5). LLPS systems driven by cation–π interactions are thus likely to be over-stabilized with this model[40].

Estimating single-molecule radii of gyration. Certain single-molecule properties of proteins, such as R_g in the context of coil-to-globule transitions, are often governed by similar driving forces as bulk LLPS[46–48], and coiling transitions are sometimes used as a proxy for the upper critical solution temperature (T_c)[22]. Importantly, the strong correlation between single-molecule dimensions and T_c has been used as a target for optimizing coarse-grained LLPS models. It is often assumed that models that correctly reproduce the experimental R_g of single proteins (at infinite dilution) should accurately predict homotypic LLPS propensities.

Accordingly, we tested the ability of Mpipi to recapitulate the experimental R_g of IDPs (Fig. 4a and Supplementary Section 2a and Supplementary Table 3). The set of IDPs has a good distribution of net charge—from $-44e$ for ProTα to $+16e$ for Ash1, where e is the elementary charge. These proteins therefore provide an indirect measure of how well electrostatic and short-ranged pairwise interactions are balanced in the coarse-grained models. Most of the proteins in our test set largely comprise neutral residues that lack π electrons in their side chains. Proteins amenable to single-molecule experiments are likely to have a high content of these neutral residues that lack π electrons, because these residues form weaker contacts and so the resulting proteins are less prone to aggregation and self-assembly. Notwithstanding the dominance of this class, the test set of proteins does exhibit appreciable variation in protein composition.

Figure 4b–g compares simulated R_g values with experimental R_g values for each coarse-grained model we have considered. Each protein is colored according to its dominant residue class, using the same coloring code as in Fig. 4a, but ignoring the neutral class. The Mpipi, FB-HPS (Fig. 4f), HPS-Urry (Fig. 4g) and TSCL-M2 (Supplementary Fig. 6b) models achieve the closest match with experiment. This result is not unexpected for these three models,

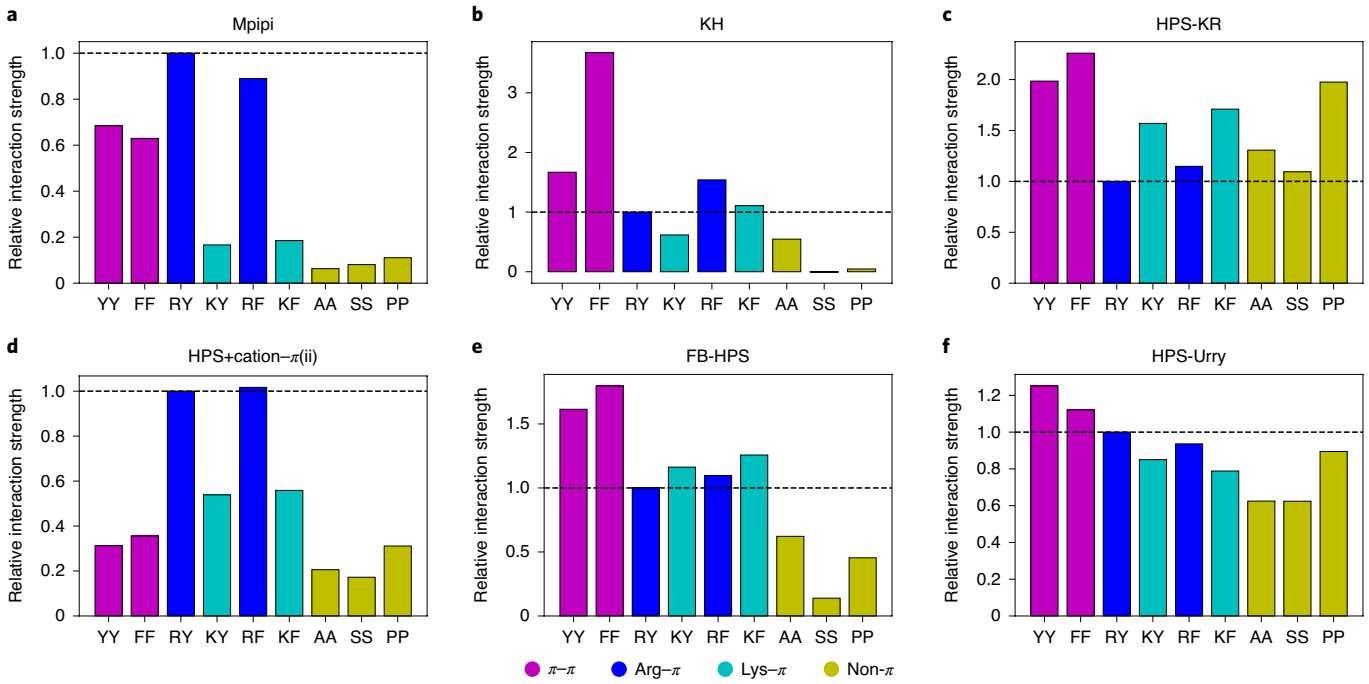

Fig. 3 | Relative contributions of π–π, cation–π and non-π-based interactions in different residue-level models. a–**f**, Relative interaction strengths (equation (8)) for selected residue pairs (Supplementary Table 1 presents the one-letter amino acid codes) in the Mpipi (**a**), KH (**b**), HPS-KR (**c**), HPS+cation–π(ii) (**d**), FB-HPS (**e**) and HPS-Urry (**f**) models. For each model, the dataset is normalized relative to the corresponding Arg–Tyr (RY) interaction. In each plot, a horizontal dashed line at the RY interaction strength is provided for comparison purposes. Aromatic π–π interactions are colored in magenta, Arg–π in blue, Lys–π in cyan and non-π-based interactions in dark yellow.

as they were all optimized to reproduce experimental R_g values, and several proteins in the current study were used either to train or to validate the respective model parameters. Importantly, compared to HPS-KR, HPS-Urry and TSCL-M2 both perform better at predicting single-molecule radii of gyration. Thus, in this regard, these models fulfill their goal of offering an improvement over their common predecessor.

Notably, Mpipi (Fig. 4b), whose parameters were not optimized on R_g data but rather on physicochemical information, is able to predict the R_g values to within a root mean squared deviation D of 0.3 nm for the tested IDPs. Fits to the bioinformatics data and atomistic PMFs therefore appear to be physically sound, at least with respect to capturing sequence-dependent single-molecule chain dimensions.

Although the HPS+cation–π(ii) (Fig. 4e), HPS-KR (Fig. 4d) and HPS+cation–π(i) (Supplementary Fig. 5) models yield reasonable agreement with experiment, all generally predict more compact proteins than experiments. Interestingly, the KH model (Fig. 4c) gives the poorest agreement with experiment, perhaps because short-ranged pairwise interactions in the KH model were obtained from residue–residue contacts in folded proteins and may thus overestimate the relative strengths of such interactions.

Recapitulating the phase behavior of A1-LCD variants. To ascertain the extent to which the Mpipi potential is able to capture the bulk properties of protein solutions, we computed the critical solution temperatures for a series of variants of the LCD of the heterogeneous nuclear ribonucleoprotein A1 (hnRNPA1), referred to here as A1-LCD, whose experimental phase diagrams were recently determined[10].

We estimated the experimental critical temperatures (Fig. 5a) of a range of A1-LCD variants, using the fitting procedure described in Supplementary Section 2b. The variants' sequences are given in Supplementary Section 2b, following the nomenclature of Bremer and colleagues[10]. For each model, we show the computed phase diagrams in Fig. 5b–g, and the correlation between simulation and experimental values in Fig. 5h–m. The corresponding data for the HPS+cation–π(i) and TSCL-M2 models are provided in Supplementary Fig. 5d,e and Supplementary Fig. 6d,e, respectively.

Although the fitted linear regression has a positive gradient for the eight models considered, indicating that, broadly speaking, all the models capture some of the underlying physics, the Pearson correlation coefficient varies considerably across the models. Mpipi achieves a Pearson correlation coefficient (r) of 0.97 (Fig. 5h), indicating that the parameterization works well not only for single-molecule properties (Fig. 4), but also accounts for bulk behavior. The HPS-Urry (Fig. 5m; $r=0.91$), TSCL-M2 (Supplementary Fig. 6e; $r=0.80$) and HPS+cation–π(ii) (Fig. 5k; $r=0.79$) models also achieve high correlation with experiments. By contrast, the FB-HPS model, which was parameterized principally on R_g data, performs quite well when predicting the radius of gyration, but the improvement of its predictions of the phase behavior of the A1-LCD variants relative to the underlying HPS-KR potential (Fig. 5j; $r=0.35$) is only marginal (Fig. 5l; $r=0.37$). Although the R_g properties do correlate with phase behavior in experiments, there are evidently several parameterizations of coarse-grained models that are able to capture one property but not the other.

Coarse-graining necessarily entails integrating out some degrees of freedom, and interaction 'energies' are therefore approximate free energies. It is thus not likely that such models could capture faithfully the behavior of protein systems far from the temperature range in which they were parameterized. Nevertheless, given that all the models were parameterized to reproduce protein behavior close to room temperature, we may also consider the agreement between the experimental and simulation temperature scales. To this end, we can compare the deviation of the experimental and simulation critical temperatures. For the A1-LCD wild type, this can be

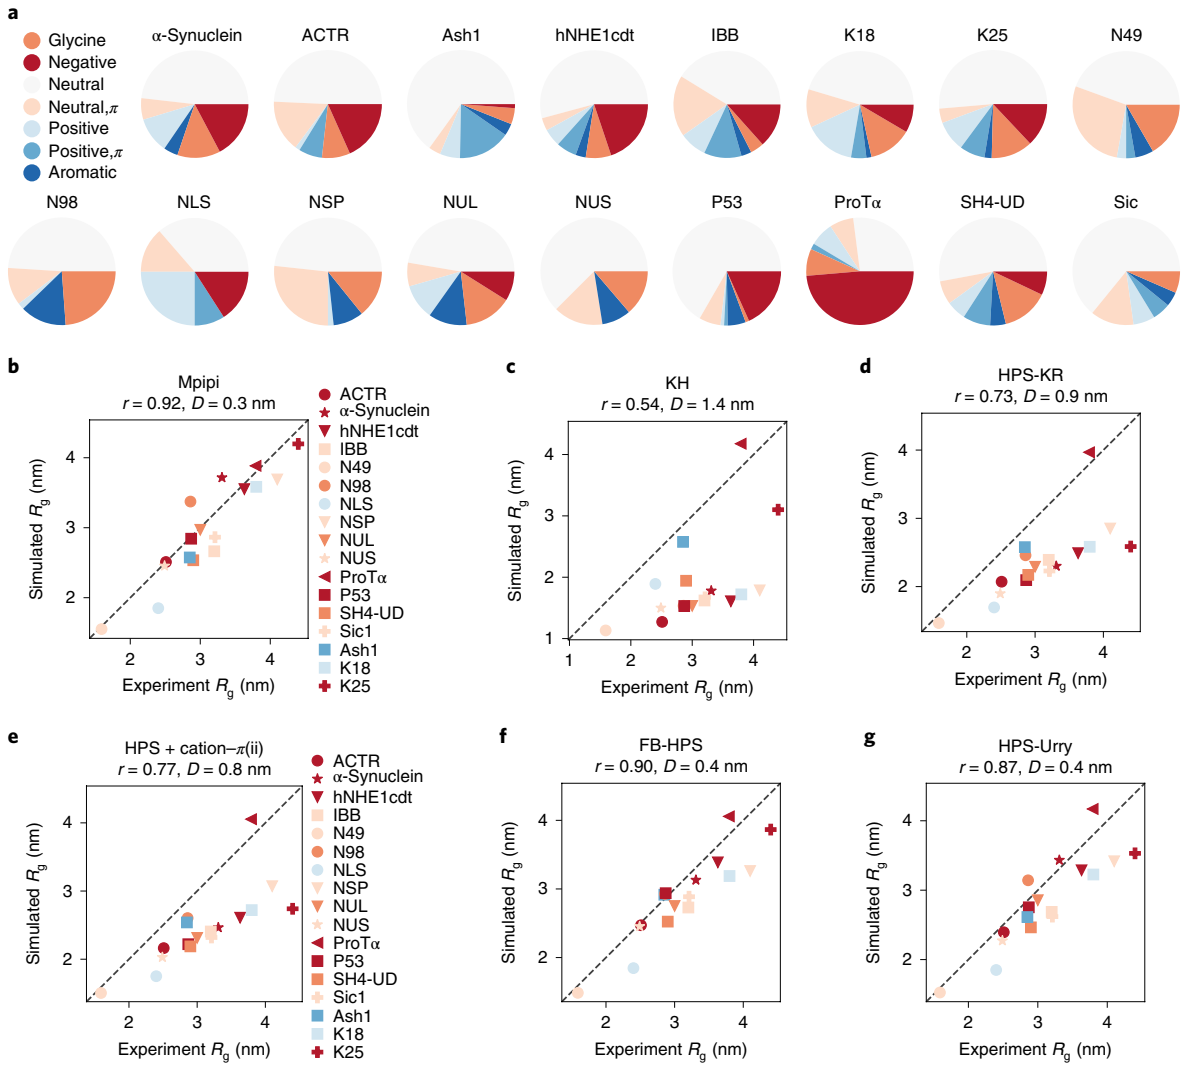

Fig. 4 | Comparison of single-molecule radii of gyration with experiment. a, The composition of simulated IDPs. We selected 17 IDPs for which experimental radii of gyration (R_g) were available (Supplementary Section 2a and Supplementary Table 3) and assessed the composition of the IDPs in terms of the proportion of glycine (orange), negative (red: D, E), neutral (white; no net charge at pH 7 and no π electrons in the side chain: A, C, I, L, M, P, S, T, V), neutral with π (salmon; no net charge at pH 7 with π electrons in the side chain: N, Q), positive (light blue; without π electrons in the side chain: K), positive with π (blue; with π electrons in the side chain: H, R) and aromatic (dark blue: F, W, Y) residues. **b**–**g**, Comparison of simulated and experimental R_g for the Mpipi (**b**), KH (**c**), HPS-KR (**d**), HPS+cation–π(ii) (**e**), FB-HPS (**f**) and HPS-Urry (**g**) models. The R_g values were computed at 300 K in each model. Each protein is colored based on its dominant residue class (as categorized in **a** and excluding the 'neutral' class). The dashed lines represent the 'perfect fit' lines. For each model, r and D are shown.

visualized by comparing the solid horizontal black lines in Fig. 5b–g, representing the experimental critical solution temperature of the A1-LCD wild type, to the maximum temperature of the binodals in black, the simulation results for the same system. In addition, a quantification of the deviation between the experimental and simulation results is given by the difference in slope between the black linear fits shown in Fig. 5h–m and the $y=x$ lines shown in red, as well as by the root mean squared deviation between the experimental and simulation critical temperatures (D).

These comparisons demonstrate that, at least within the range of experimental data available, Mpipi ($D=9$ K), HPS+cation–π(ii) ($D=18$ K) and TSCL-M2 ($D=19$ K) give good quantitative predictions for A1-LCD (that is, high r and small D values). Finally, although HPS-Urry achieves good agreement with experiment when considering the Pearson coefficients, its range of predicted critical temperatures is smaller than in experiments, which results in a relatively large D.

LLPS of other proteins and multiphasic compartmentalization. To probe the model's transferability, we tested its performance for other well-studied IDRs of RNA-binding proteins. Specifically, we computed phase diagrams for the prion-like domain (PLD) of the FUS protein, three variants of the arginine/glycine-rich (RGG) domain of LAF-1 (Fig. 6a) and four variants of the DDX4 NTD (Fig. 6b). We also computed phase diagrams for five variants of the full-length FUS protein (Supplementary Section 2 and Fig. 6c). For all these systems, Mpipi achieves good qualitative agreement with experiment and in some cases achieves quantitative accuracy as well. A full account of these results is provided in Supplementary Section 7.

Finally, beyond predicting critical temperatures, achieving the correct balance of interactions is essential to recapitulate more complex condensate behaviors. Inside cells, condensates are multicomponent systems that can have complex molecular architectures that are meaningful to their functions (for example, the nucleolus), and

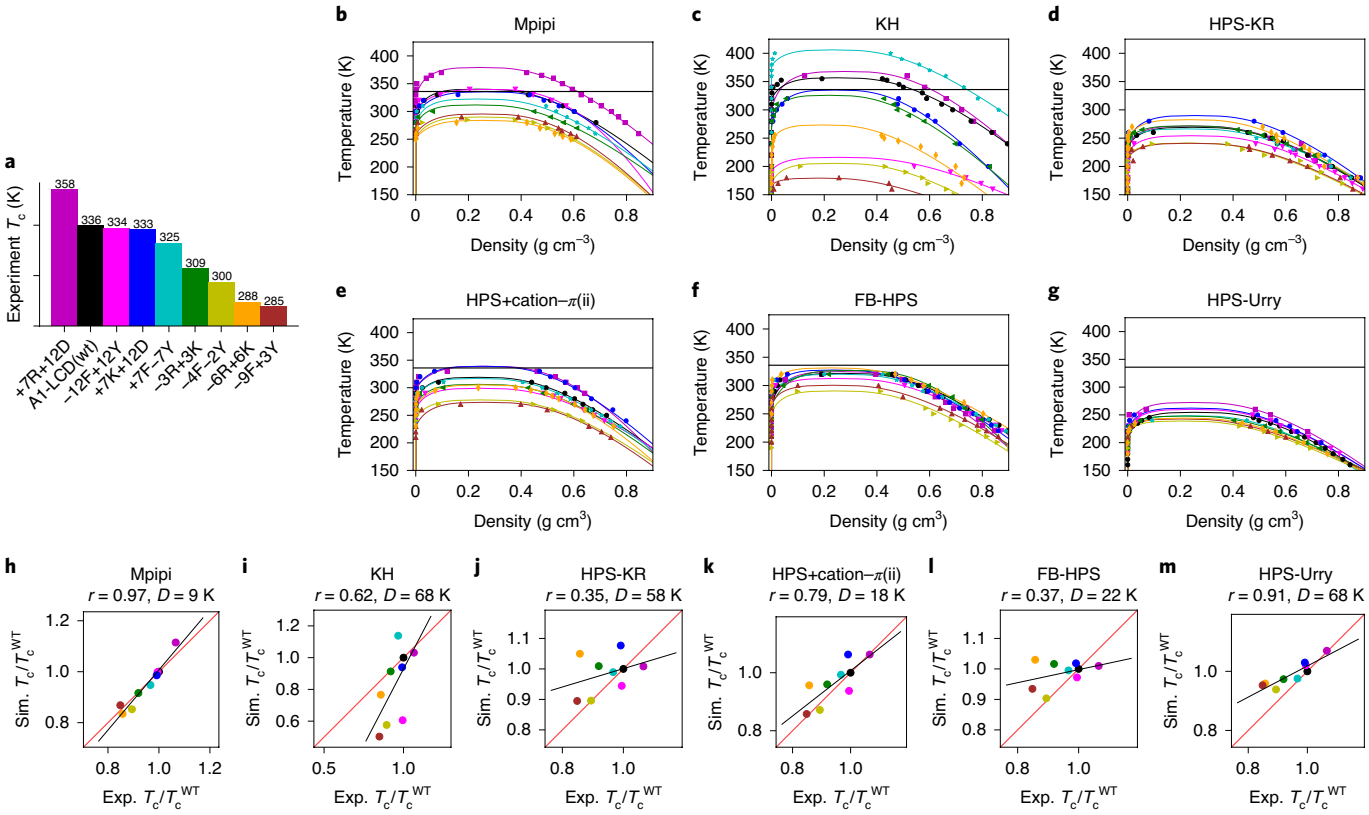

Fig. 5 | Recapitulating the phase behavior of A1-LCD variants. a, Nine variants of A1-LCD (including the wild type) were studied in this work. The variants were prepared following ref. [10]. The experimental critical temperatures were estimated as described in Supplementary Section 2b. The colors used in **a** for each variant are used in all other panels. **b**–**g**, Phase diagrams for A1-LCD variants obtained via direct-coexistence simulations using the Mpipi (**b**), KH (**c**), HPS-KR (**d**), HPS+cation–π(ii) (**e**), FB-HPS (**f**) and HPS-Urry (**g**) models, respectively. Estimation of the critical points of the simulated phase diagrams is described in Methods. The curves are derived from empirical fits of the data to equations (6) and (7). Typical errors are discussed in Supplementary Section 8d. The horizontal black lines represent the experimental critical temperature of the wild type. **h**–**m**, Simulated critical temperature T_c relative to the critical temperature of the wild type (T_c^{WT}) shown against the experimental analog for Mpipi (**h**), KH (**i**), HPS-KR (**j**), HPS+cation–π(ii) (**k**), FB-HPS (**l**) and HPS-Urry (**m**) models, respectively. In **h**–**m**, r and D are provided above each graph. The red lines correspond to a perfect fit to the experimental data, while the black lines represent the linear regression fit.

the variance in the chemical makeup of biological phase-separating systems can give rise to multilayered architectures[49]. For example, Fisher and Elbaum-Garfinkle[6] recently demonstrated that charge-matched mixtures of poly-arginine (polyR), poly-lysine (polyK) and polynucleotides form multiphasic droplets in which arginine is positioned towards the center of the condensates and lysine is concentrated at the interface.

To investigate this behavior in simulations, we extended Mpipi to include parameters for RNA (Methods and Supplementary Fig. 4) and studied the phase-separation behavior of a mixture of polyK, polyR and polyU RNA. Consistent with experiments, our simulations recapitulate multiphase droplet architectures (Fig. 6d). Interestingly, we find that the density of the Arg-rich region at the droplet core is considerably higher than the Lys-rich phase density towards the interface; these results also agree well with experiments[6]. For comparison, we also simulated this mixture using the extended HPS-KR model, which includes parameters for RNA[50]; however, using this model, multiphase droplets are not stabilized (Supplementary Fig. 8a). This result is not especially surprising because, as noted above, the balance between Arg and Lys interactions in HPS-KR is incorrect[27,40].

More broadly, we speculate that it is generally difficult to account for multiphasic behavior if the standard arithmetic mixing rules for interaction strengths are utilized. Collectively, our work suggests that it is not sufficient to obtain the correct trends for short-ranged pairwise terms: one must also achieve the right balance of these interaction parameters to yield quantitative accuracy.

Discussion

Although the balance of interactions in our current model is in agreement with several experimental and computational studies, there is conflicting evidence on the exact ordering of certain key interactions. For example, Bremer et al.[10] report weaker Arg–π contacts than the corresponding π–π ones, while our work and the work of Wang et al.[11] appear to favor the view that Arg–aromatic interactions are stronger than analogous aromatic–aromatic ones. Furthermore, the data suggest that, in some cases, the precise ordering of these interactions within coarse-grained models is fundamental to recapitulate the observed behaviors, while in other cases an approximate ordering could suffice. For example, in FUS protein we find that the LLPS propensities of the full-length protein versus its PLD domain are highly sensitive to the relative ordering of these interactions[7], whereas our current benchmarks reveal that for the A1-LCD variants, all models that favor Arg–π and π–π contacts over the non-π-based contacts achieve a high Pearson correlation, regardless of the precise ordering of Arg–π and π–π contacts. Hence, we postulate that the ordering of these and other interaction strengths is likely to be context-specific, and a system-specific coarse-graining strategy may be necessary to achieve good agreement with experiment in some cases. Consequently, one set of

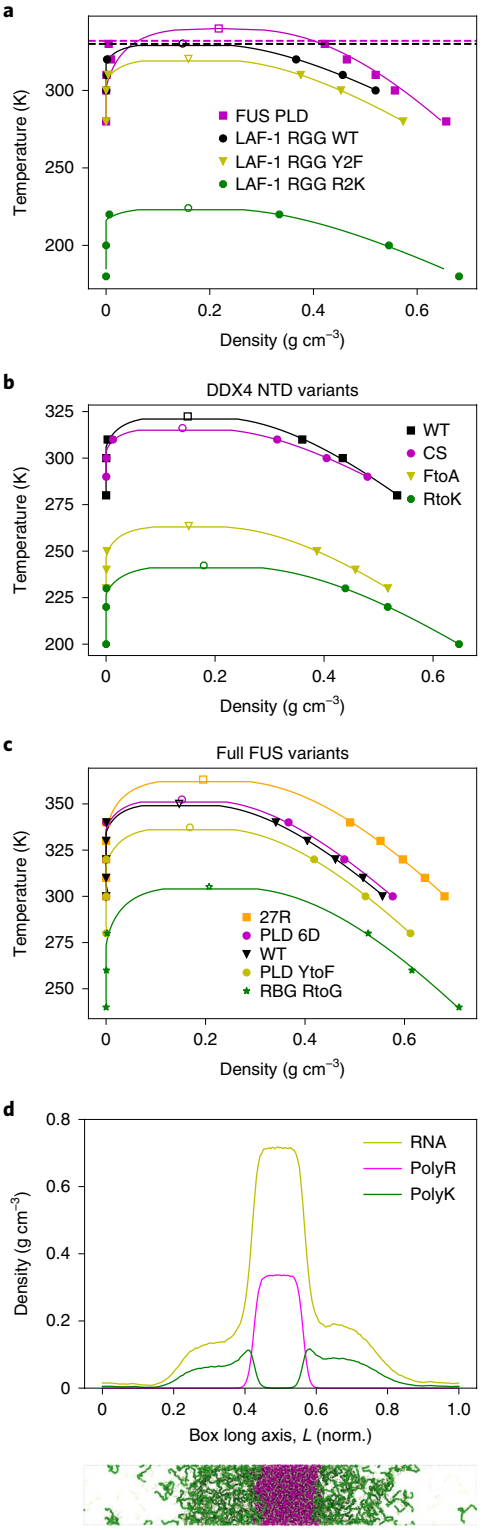

Fig. 6 | Predicting the LLPS propensities of other proteins and multiphasic compartmentalization. a, Temperature–density phase diagrams for FUS PLD, LAF-1 RGG (wild type, WT) and two other variants of LAF-1 RGG for the Mpipi model. Filled symbols represent simulation data and open symbols depict estimated simulation critical points (Methods). The horizontal dashed lines represent estimated T_θ (the temperature of the coil-to-globule transition) for FUS PLD (magenta) and LAF-1 RGG (WT) (black) obtained with the ABSINTH potential. Solid curves are fits to equations (6) and (7). **b**,**c**, As in **a**, but for four DDX4 variants (**b**) and full FUS variants (**c**). **d**, We simulated a mixture of polyK (50 residues; 128 chains), polyR (50 residues; 128 chains) and RNA (10 residues; 1,280 chains) with an extended Mpipi model (Methods and Supplementary Fig. 4). The density profile along the simulation box's long axis (L, normalized) is given for each mixture component. A simulation snapshot is provided below the density plot. The color code in the snapshot is consistent with that used in the density plot. The mixture is simulated at $T/T_c \approx 0.8$, where T_c is the critical temperature for liquid–vapor phase separation.

measurements, be it experiments or simulations, will be unlikely to yield the complete picture.

A key assumption in our work is that differences in the LLPS propensities of biomolecules can be captured via pairwise amino acid interactions. This approximation allows us to construct a transferable coarse-grained model that can capture several qualitative and quantitative trends for phase-separating systems, especially for those characterized in vitro. However, in crowded intracellular environments, three- and higher-body energy terms may become important and, accordingly, cooperative interactions can reshape the phase boundaries of LLPS systems[5,51]. It is therefore important to consider carefully the contribution of such cooperative interactions in intracellular LLPS systems.

As we discussed above, interaction energies in coarse-grained potentials are in fact effective free energies, and they should, in principle, depend on temperature. In particular, because we have not considered explicit protein–solvent interactions, the solubility of all proteins studied increases with increasing temperature, even when other effects, such as hydrogen bonding, could result in very different phase behavior as the temperature is lowered. This can even result in complete mixing at low temperatures, leading to a lower critical solution temperature or re-entrant phase behavior, especially in multicomponent systems[52–54].

Capturing such effects within computational models can further extend our ability to elucidate the driving forces for intracellular LLPS and to probe the ensuing material properties. The current parameterization of the Mpipi potential is not able to account for such phase behavior; as an extension of the current work, an approach similar to that of Dignon et al.[55] for the HPS-KR model, involving an explicit temperature dependence of the interaction strengths, could be undertaken to enable successful simulations of a broader range of proteins to be performed.

This work highlights the promise that multiscale coarse-grained models can prove robust in delineating the link between chemical changes in biomolecules and their emergent collective behavior. In particular, the ability of Mpipi to predict quantitatively both single-molecule radii of gyration (which are computationally inexpensive to determine) and the collective behavior of proteins and RNA in solution (which is computationally more expensive) makes it a prime candidate for efficiently assisting the design of experiments and for gaining physical insight into LLPS at the microscopic scale. Our approach therefore augments the set of rigorous tools that are narrowing down the gap towards achieving a predictive quantitative description of the influence of amino acid sequence in biological phase behavior. Alongside experimental advances, theoretical work and other computational approaches, the Mpipi model has the potential to help discover the molecular mechanisms underpinning phase separation and to provide biophysical understanding of how biomolecular condensates are formed, sustained and regulated.

Methods

Atomistic PMF calculations. To quantify the relative contributions of different types of interaction at physiological salt levels, we performed atomistic PMF computations for a subset of residue pairs, namely WW, YY, FF, RY, RF, KY, KF, AA, SS, PP, RE, RD, KE and KD. All residue pairs from our previous work[7] were recomputed at 150 mM

salt concentration, and we have included additional pairs. We also performed PMF calculations for four RNA dimer pairs (Supplementary Section 4).

Preparation of structures. Amino acids and nucleic acids were modeled using the AMBER ff03ws force field[56]. This force field is well-suited for probing protein–protein interactions. For modeling the solvent (water) and ions, we used the JC-SPC/E-ion/TIP4P/2005 force field[57], as in our previous work[7]. The N- and C-terminal ends of each amino acid were capped with acetyl and N-methyl capping groups, respectively. Pairs of amino acids were orientated with their side chains facing each other, based on the most common arrangements observed in protein structures. In cases where the interaction preference is uncertain, multiple arrangements were tested to determine the strongest interaction mode.

Each dimer was then immersed in a cubic box containing TIP4P/2005 water molecules (~960–11,020 molecules) with a minimum distance of 1 nm between the dimer and the edge of the box. Na^+ and Cl^- ions were added to achieve a salt concentration of ~150 mM, as well as to produce charge-neutral systems. The resulting systems were then minimized (force tolerance 500 J mol^{-1} pm^{-1}), with positional restraints of 200 J mol^{-1} pm^{-2} applied in each dimension to all heavy atoms.

Umbrella sampling. The interaction between each dimer was probed with umbrella sampling. For production runs, positional restraints of 1 J mol^{-1} pm^{-2} in directions perpendicular to the pulling direction were used to constrain heavy atoms. The COM distance between interacting pairs was restrained with a harmonic umbrella potential (pulling force constant of 6 J mol^{-1} pm^{-2}). All bonds with hydrogens were constrained using the LINCS algorithm, permitting an integration time step of 2 fs. Periodic boundary conditions were used during molecular dynamics (MD) simulations. Electrostatics were computed using particle-mesh Ewald summations with a Coulomb cutoff of 0.9 nm. For each umbrella sampling run, ~40 windows, spaced at 50 pm from 0.1 nm to 2 nm, were used per pair. Each window was simulated for 10 ns. Three independent simulations were conducted for each umbrella sampling window (that is, an aggregate simulation time of 30 ns per window). Umbrella sampling data were analyzed using the weighted histogram analysis method (WHAM). The first 1 ns of simulations was used for equilibration and was not included in the WHAM analysis. Error analysis was performed using the Bayesian bootstrap method. All atomistic simulations and analyses were carried out using the GROMACS simulation package.

Although we focused on COM distances for fixed molecular orientations in the PMF calculations, we ultimately mapped these to the Cα–Cα distances of the coarse-grained potential. The effective free energy as expressed with different order parameters may not be the same and depends on the Jacobian determinant of the transformation. However, our choice of order parameter cannot affect the observable properties of the system, and the two distances were related by a simple linear relationship for a fixed molecular orientation. Provided we use the PMFs in a self-consistent manner, the resulting ratios of interaction strengths should not depend on this choice of order parameter.

Cation–π charge refitting. Cation–π interactions involve the polarization of π electron clouds of aromatic side chains in the proximity of cationic side chains (that is, arginine and lysine), especially at physiological salt conditions. There have been many efforts to correctly capture cation–π interactions in atomistic force fields, both with fixed-charge and polarizable force fields (see the discussion by Liu et al.[58]). Recently, Paloni et al. demonstrated that the fixed-charge AMBER 99SB-disp force field was able to account for Arg/Lys–π interactions for the DDX4 NTD[59]. In another study, Liu et al. used quantum-mechanical calculations to reparameterize the Lennard-Jones parameters in the CHARMM36 force field to model cation–π pairs[58]. Their modified parameters led to improved descriptions of the selected folded proteins[58], achieving a closer match to experimental crystal structures.

In this Article, to model cation–π interactions atomistically, we followed our previous approach[7] and first refitted the charges on tyrosine and phenylalanine side chains. Specifically, the dimers (Arg/Lys–Phe/Tyr) were first optimized using constrained geometry optimizations at the MP2/6-31G(d) level of theory, where the backbone and capping-group heavy atoms were frozen. The electrostatic surface potential (ESP) was then computed for the respective optimized pairs at the HF/6-31G(d) level. These calculations were carried out using the Gaussian 09 code. Finally, the restrained electrostatic potential method in AMBER was used to refit the side-chain charges of Tyr and Phe to the ESPs from the quantum-mechanical calculations (the charge symmetry of the rings was maintained during the refitting procedure). The refitted charges were then used when probing the pairwise interaction strengths via umbrella sampling, as described above.

Mpipi model. In the Mpipi model, each amino acid or nucleic acid is represented by a single bead, with corresponding mass, molecular diameter (σ), charge (q) and an energy scale reflecting the relative planar π–π contact frequency (ε). We broadly followed the approach of Dignon et al.[26] to compute the potential energy of a given protein or RNA molecule as

$$E_{\text{Mpipi}} = E_{\text{bond}} + E_{\text{elec}} + E_{\text{pair}}. \quad (1)$$

The bond energy was computed by using a harmonic bond potential

$$E_{\text{bond}} = \sum_{\text{bonds } i} \frac{1}{2} k(r_i - r_{i,\text{ref}})^2, \quad (2)$$

where the spring constant k was set to 8.03 J mol^{-1} pm^{-2} and r_i is the bond length (reference bond lengths, $r_{i,\text{ref}}$, of 381 pm and 500 pm were used when bond i connected two protein and two RNA beads, respectively). The electrostatic contribution to the potential energy was computed using a Coulomb term with Debye–Hückel electrostatic screening

$$E_{\text{elec}} = \sum_{i,j} \frac{q_i q_j}{4\pi\varepsilon_r\varepsilon_0 r_{ij}} \exp(-\kappa r_{ij}), \quad (3)$$

where $\varepsilon_r = 80$ is the relative dielectric constant of water, ε_0 is the electric constant and $\kappa^{-1} = 795$ pm is the Debye screening length, corresponding to a monovalent salt concentration of 0.15 M to be consistent with the PMF calculations. We used a Coulomb cutoff of 3.5 nm. The dielectric constant and the Debye length control the range of ionic interactions and determine the relative importance of charges relative to all other interactions. A more careful treatment of electrostatics, perhaps in the spirit of Wessén et al.[60], would be an important next step to consider in the development of more accurate potentials.

Finally, the non-bonded interactions between protein/RNA beads were modeled via the Wang–Frenkel (WF) potential[59]. The WF potential between two beads of types i and j a distance r apart is given by

$$\phi_{ij}(r) = \varepsilon_{ij}\alpha_{ij}\left[\left(\frac{\sigma_{ij}}{r}\right)^{2\mu_{ij}} - 1\right]\left[\left(\frac{R_{ij}}{r}\right)^{2\mu_{ij}} - 1\right]^{2\nu_{ij}}, \quad (4)$$

where

$$\alpha_{ij} = 2\nu_{ij}\left(\frac{R_{ij}}{\sigma_{ij}}\right)^{2\mu_{ij}}\left[\frac{2\nu_{ij}+1}{2\nu_{ij}\left(\left\{\frac{R_{ij}}{\sigma_{ij}}\right\}^{2\mu_{ij}} - 1\right)}\right]^{2\nu_{ij}+1}, \quad (5)$$

and σ_{ij}, ε_{ij} and μ_{ij} are parameters specified for each pair of interacting beads. We used $\nu_{ij} = 1$ and $R_{ij} = 3\sigma_{ij}$. The total pairwise energy E_{pair} was then taken as the sum over all pairs of beads evaluated within their respective interaction ranges (that is R_{ij}, at which ϕ_{ij} vanishes).

Most importantly, the WF potential is finite-ranged, vanishing quadratically to zero at the user-specified cutoff distance, and so obviates the need for truncating and shifting the potential. This key feature makes the WF potential better suited for numerical calculations and removes any ambiguities or inconsistencies that may arise from one implementation to the next. For example, Lennard-Jones-based potentials can exhibit undesirable finite-size effects as a function of the cutoff distance and subsequent tail corrections[61]. The computational performance of the WF potential is comparable to the Lennard-Jones potential for the same cutoff; although we have not done this here; if one wished to simulate particularly large systems, the WF potential's more flexible functional form affords an opportunity for optimizing the distance at which the potential vanishes, which could enable an appreciable computational boost without degrading performance. Moreover, although from its scaling properties the Lennard-Jones potential appears at first glance to account for London dispersion interactions, in reality this is not the case in solution, where the potential accounts for many interactions in a coarse-grained way. Indeed, a further advantage of the WF potential is that it removes this misleading appearance of physicality.

To obtain the parameters that appear in the WF parameterization, we first determined the relative planar π–π contact frequencies of the amino acids from the work of Vernon et al.[13], determined Ashbaugh–Hatch-style Lennard-Jones interactions following Dignon et al.[26], and from these obtained the initial WF parameters. The steepness of the repulsive region of the potential and the width of the attractions can easily be modulated in this framework by allowing the μ parameter to take values larger than unity. We next adjusted the values of ε_{ij} by a suitable multiplicative factor so that the integrals of the well depths of the PMF curves of residue pairs i and j approximate their WF analogs, including any screened charge–charge interaction (Supplementary Fig. 3), if relevant, to ensure that the overall interaction energy was correctly taken into account[62]. We provide a full parameter listing in Supplementary Table 11 and a LAMMPS implementation in the Supplementary data. Although it has been suggested[24,43] that simple arithmetic combination rules are often sufficient, unlike in previous models, the pairwise interactions for those residue pairs that dominate the phase behavior are explicitly specified, giving the model greater flexibility. There is no a priori reason to assume that coarse-grained interactions between unlike species will be well described by an arithmetic mean of homotypic interactions and, in particular, we find that the heterotypic interactions of arginine and lysine can be rather different from the mixing-rule prediction. Parameters for the nucleic acids were determined directly by fitting the respective PMF well depths and widths to the WF framework (see below). Both disordered proteins/regions and RNA were modeled as fully flexible polymers.

The validation simulations used various previously reported models. Mostly, these were based on the functional form introduced in the work of Dignon et al.[26]. The bonded and electrostatic contributions to the potential were given by the same functional form in each case, although with slightly different constants (Supplementary Table 10).

Background on the parameterization of Mpipi. To parameterize the non-bonded short-ranged terms described via the WF potential, we first determined the relative π–π contact frequencies for amino acids from the work of Vernon et al.[13] (Supplementary Table 1), who predicted the planar π–π contact frequencies from a survey of ~6,000 high-resolution structures in the PDB. We utilized these contact frequencies as an initial energy scale for short-ranged interactions in our model.

We then refined this initial energy scale using atomistic PMF calculations, focusing on aromatic π–π (Fig. 2a), cation–π (Fig. 2b) and a subset of non-π-based (Fig. 2c) interactions. Pappu, Mittag and colleagues position aromatic 'stickers' as the chief drivers of biomolecular LLPS[4,8–10]. Our recent findings are also consistent with the stickers-and-spacers model, where we predict that aromatic π–π interactions constitute dominant forces in LLPS, even at extremely high salt concentrations[7].

In this Article we compute the PMF between YY, FF (Fig. 2a) and WW (Supplementary Fig. 2) at physiological salt concentration (Methods) and find that, in agreement with the bioinformatics data[13] and experiments[10–12], the relative strength of aromatic π–π interactions increases in the order FF < YY < WW (Fig. 2d, magenta bars and Supplementary Fig. 2). Importantly, we find that aromatic π–π interactions are at least twice as strong as non-π-based interactions (dark yellow bars, Fig. 2d). The latter interactions include non-polar, polar and special residues (for example, Pro) and are commonly categorized as spacers[4,8,11]. Interestingly, Bremer et al. predicted that the disparity in spacer–spacer and sticker–sticker residue interaction strengths can be as high as 1:8 (ref. [10]). Our fitted spacer-type interactions represent a compromise between the predictions of Bremer et al.[10] and those suggested by our PMF calculations (Fig. 2e).

We next concentrated on the interactions between basic residues (Arg and Lys in particular) and aromatics. These cation–π interactions also make substantial contributions to the LLPS of biomolecules. In an early bioinformatics survey, Gallivan and Dougherty[63] revealed that Trp was most likely to form cation–π interactions, followed by Tyr and then Phe. Song et al.[64] subsequently used experiments and simulations to demonstrate much higher binding strengths for RW interactions compared to RY/F ones, with RY slightly stronger than RF.

Furthermore, recent work by Wang et al.[11] suggests that Arg–Tyr interactions may be stronger drivers of LLPS than Tyr–Tyr contacts. Our PMF calculations agree that Arg–Tyr interactions are stronger than Tyr–Tyr. However, whether cation–π interactions are indeed stronger contributors to protein LLPS than π–π interactions remains contested: the work of Bremer et al.[10] and single-residue solubility measurements[65] suggest instead that Tyr–Tyr interactions are stronger than Arg–Tyr contacts. A potential source of error in the relative ordering of interactions in our work might come from the approximate nature of atomistic PMFs simulations and the use of a pairwise energy to describe an interaction that is probably affected by cooperative effects. Reassuringly, despite the differences, in all cases, both cation–π and π–π interactions are considerable.

A further important consideration is the balance of Arg–π and Lys–π interactions. The differences between Arg–π and Lys–π contacts were highlighted by Gallivan and Dougherty[63], who reported a higher percentage of Arg–π contacts in protein structures. We recently proposed that Arg–π interactions are best described as hybrid cation–π/π–π, whereas Lys–π contacts represent 'purer' cation–π interactions[7]. This distinction arises largely from the presence of π electrons in the Arg side chain[6,7,13,16–18], which enable Arg residues to interact much more strongly with π-binding partners than Lys can[6,7,15,16]. The dominance of Arg over Lys in these interactions is also consistent with the less favorable hydration free energy recently reported for Arg versus Lys[16]. An earlier study by Kumar et al. revealed that, whereas Lys–π interactions are more favorable in the gas phase, in solution Arg establishes stronger interactions with aromatic rings than Lys, the latter being dominated by electrostatics and therefore weakened by the surrounding dielectric medium[66]. Collectively, our PMF calculations and previous studies all suggest a preference for Arg–π interactions over Lys–π contacts in biomolecular systems. Accordingly, we reparameterized cation–π interactions so that the relative weights in our model more closely matched those suggested by the atomistic simulations (Fig. 2d). A summary of the relative interaction strengths between amino acid pairs is provided in Fig. 2e. These interaction strengths correspond to the average interaction energy in the high-temperature limit relative to a fixed (albeit arbitrary) energy of zero obtained by numerically integrating equation (8). The integration over the energy well in this high-temperature limit enables the relative interaction strength to account, at least approximately, for both enthalpic and entropic contributions.

Direct-coexistence simulations. Proteins and RNA were represented via the Mpipi model (or other residue-level coarse-grained model) and direct-coexistence[38] simulations were used to compute their phase diagrams. In such simulations, the high- and low-density fluid phases were simulated in the same simulation box, delimited by an interface.

The target number of copies of the protein were placed in an elongated box, which was initially simulated at high temperature and then cooled to the desired temperature. Canonical-ensemble simulations were then run at temperatures below the estimated critical temperatures for each system. A relaxation time of 5 ps was typically used for the Langevin thermostat and an integration time step of 10 fs was used for all coarse-grained simulations. Calculations were carried out using LAMMPS. Finite-size effects[9,67,68] are described below.

Estimation of critical points on phase diagrams. Critical temperatures were estimated using the law of coexistence densities

$$(\rho_{\mathrm{high}}(T) - \rho_{\mathrm{low}}(T))^{3.06} = d\,(1 - T/T_{\mathrm{c}})\,, \tag{6}$$

and critical densities were computed by assuming that the law of rectilinear diameters holds, namely

$$\rho_{\mathrm{high}}(T) + \rho_{\mathrm{low}}(T) = 2\rho_{\mathrm{c}} + 2A(T - T_{\mathrm{c}}), \tag{7}$$

where $\rho_{\mathrm{high}}(T)$, $\rho_{\mathrm{low}}(T)$ and ρ_{c} are the densities of the high-density and low-density phases and the critical density, respectively, T_{c} is the critical temperature, and d and A are fitting parameters.

Data analysis. When comparing relative interaction strengths, we computed the average energy in the high-temperature limit, that is, assuming that the well depth is much smaller than $k_{\mathrm{B}}T$ for each individual interaction. For each pair of amino acid residues, we computed

$$\mathcal{E}_{\mathrm{avg}} = \int_{\sigma}^{3\sigma} \phi(r)\,\mathrm{d}r + \int_{\sigma}^{3.5\ \mathrm{nm}} E_{\mathrm{elec}}(r)\,\mathrm{d}r, \tag{8}$$

where $\phi(r)$ is the WF potential (equation (4)) and $E_{\mathrm{elec}}(r)$ the Coulomb energy (equation (3)), if relevant. We then normalized the result by the interaction strength of the RY (Arg–Tyr) pair. We computed the integral in equation (8) in one-dimensional space, that is, not including the $4\pi r^2$ volume element, because we wished to compare these strengths to PMF calculations, where a constrained approach was used. However, including a spherical-polar volume element does not substantially affect the appearance of Fig. 2e.

To assess the agreement between simulation results and experiment for a given observable X (where X is either the critical temperature or the radius of gyration), we computed both the Pearson correlation coefficient (r), which is a measure of deviation from a linear fit to the data and is a good measure of the quality of the ordering of the predictions, and the root mean squared deviation D, the square of which we define as

$$D^2 = \frac{1}{n}\sum_{i=1}^{n}[X_i(\mathrm{experiment}) - X_i(\mathrm{simulation})]^2, \tag{9}$$

where n is the number of data points. D is a measure of the absolute deviation from the experimental results.

Radii of gyration computation. We computed single-molecule radii of gyration (R_{g}) for the protein sequences presented in Supplementary Section 2a. Each protein was simulated in a large cubic box (~60 nm × 60 nm × 60 nm). Canonical-ensemble simulations were then performed at 300 K for 5 µs, with a time step of 10 fs. A Langevin thermostat was used, with a relaxation time of 50 ps. R_{g} measurements were made every 100,000 time steps (that is, 1 ns). The first 1,000 fs of the simulation were not used in the estimation of R_{g} values. Simulations were run using LAMMPS.

Estimation of the coil–globule transition temperature. We estimated the coil–globule transition temperature T_θ with ABSINTH[69], a continuum solvation all-atom model of proteins. To determine this temperature, we simulated a single protein in a spherical cell with explicit ions, and determined the temperature at which there was a sudden change in the radius of gyration as a function of temperature[22].

Finite-size scaling. Finite-size effects can play an important role in direct-coexistence simulations[67]. To ascertain that phase-diagram calculations with direct-coexistence simulations yield robust results, we first confirmed that reducing the system size by ~30% yielded the same phase diagrams, within error bars, for the hnRNPA1 variants and FUS-LCD as reported above. Because there was no difference in the predicted critical temperatures, we hypothesized that finite-size effects were not dominant in our simulations. To test this hypothesis more carefully, we investigated the finite-size scaling behavior of the FUS-LCD system systematically. In Supplementary Fig. 7, we show two sets of results for this system. We first tested the effect of the size of the cross-sectional area of the interface, starting from a particularly small area of 4 nm × 4 nm, that is, with box dimensions only just larger than the largest cutoff distance in the interparticle potential (3.5 nm). Supplementary Fig. 7a shows a considerable spread in

individual values; with the possible exception of the smallest system size, there is no significant difference in the mean density computed across the entire high-density portion of the density profile, which is needed for the phase diagram. In other words, a cross-sectional area of ~10 nm × 10 nm, as used in the majority of phase-diagram calculations, appears to be more than sufficient to avoid finite-size effects.

Next, we tested the finite-size scaling of the bulk of the system by maintaining a constant cross-sectional area and increasing the length of the long box dimension at a constant density, increasing the number of chains in the system. We show these results in Supplementary Fig. 7b. For the very smallest system size considered here, with a z-axis dimension of 11 nm, the density profile is close, but not exactly consistent with that of the larger systems. This is not especially surprising, because the long box dimension is only marginally longer than the other two, and the interface is considerably more fluxional as a result. However, from the 22 nm simulation onwards, the high-density profile has a flat region that changes in width, but not the mean density, suggesting that finite-size effects are negligible beyond this point. The system sizes used in our phase-diagram calculations are shown in Supplementary Table 9. All sizes are well beyond the point where finite-size effects dominate the system's behavior.

One caveat here is that the results for FUS-LCD that we show in Supplementary Fig. 7 correspond to a temperature just above 80% of the critical temperature. The interface naturally becomes less well defined as the critical point is approached, and data points very close to the critical temperature are not usually very robust. However, it is not necessary to obtain such data points to estimate the critical temperature from a fit to equations (6) and (7).

Implementation of parameters for RNA. We have parameterized an initial set of RNA nucleotide parameters that is compatible with the Mpipi model for proteins. Nucleotide–nucleotide interaction strengths were derived by first performing atomistic PMF calculations for homo-dimer pairs (Supplementary Fig. 4a,b). We used dimers instead of monomers because it is more straightforward to study homodimers than single nucleic acid monomers in standard protein/RNA force fields. From these simulations we extracted the relative weights of the RNA nucleotide–nucleotide interactions. Specifically, we computed the base–base binding free energies, which encode the short-range pairwise terms in the Mpipi model.

To map the PMFs to our Mpipi model parameters, we first fitted the weighted atomistic interaction energies (that is, the integral of the PMF curves at 298.15 K; equation (8)) to the corresponding WF weighted interaction energies at the same temperature. This procedure involved performing a linear fit between known WF interaction energies in our model and their atomistic counterparts (that is, for the set in Supplementary Fig. 2a). The fit parameters were then used to determine the corresponding weights for the RNA nucleotides. Next, using an iterative procedure, we determined the WF parameters (that is, ε and μ) for each RNA bead that yielded the target binding strengths.

Each RNA bead was then described by a unique set of WF parameters and a charge of $-0.75e$. Finally, we reduced the ε in the WF part of the potential for the RNA beads until self-assembly for polyA/polyG RNA (50 beads, 64 chains; that is, nucleotides with stronger base-stacking propensities) was sufficiently destabilized at 200 K. In subsequent work, we aim to refine our RNA parameters (including short-ranged binding strengths, bond constants and angular constants).

Data availability

All relevant supporting data are available in the figshare data repository at https://doi.org/10.6084/m9.figshare.16772812[70]. The data for this study were generated with the simulation codes and algorithms outlined in Supplementary Table 14, using the supporting code[70], alongside standard command-line tools. Source data are provided with this paper.

Code availability

LAMMPS input scripts and parameter files are available in the figshare data repository at https://doi.org/10.6084/m9.figshare.16772812[70].

Received: 15 May 2021; Accepted: 8 October 2021;
Published online: 22 November 2021

References

1. Hyman, A. A. & Simons, K. Beyond oil and water-phase transitions in cells. *Science* **337**, 1047–1049 (2012).
2. Li, P. et al. Phase transitions in the assembly of multivalent signalling proteins. *Nature* **483**, 336–340 (2012).
3. Alberti, S. & Dormann, D. Liquid-liquid phase separation in disease. *Annu. Rev. Genet.* **53**, 171–194 (2019).
4. Martin, E. W. et al. Valence and patterning of aromatic residues determine the phase behavior of prion-like domains. *Science* **367**, 694–699 (2020).
5. Choi, J.-M., Holehouse, A. S. & Pappu, R. V. Physical principles underlying the complex biology of intracellular phase transitions. *Annu. Rev. Biophys.* **49**, 107–133 (2020).
6. Fisher, R. S. & Elbaum-Garfinkle, S. Tunable multiphase dynamics of arginine and lysine liquid condensates. *Nat. Commun.* **11**, 4628 (2020).
7. Krainer, G. et al. Reentrant liquid condensate phase of proteins is stabilized by hydrophobic and non-ionic interactions. *Nat. Commun.* **12**, 1085 (2021).
8. Harmon, T. S., Holehouse, A. S., Rosen, M. K. & Pappu, R. V. Intrinsically disordered linkers determine the interplay between phase separation and gelation in multivalent proteins. *eLife* **6**, e30294 (2017).
9. Choi, J. M., Dar, F. & Pappu, R. V. LASSI: a lattice model for simulating phase transitions of multivalent proteins. *PLoS Comput. Biol.* **15**, e1007028 (2019).
10. Bremer, A. et al. Deciphering how naturally occurring sequence features impact the phase behaviors of disordered prion-like domains. Preprint at *bioRxiv* https://doi.org/10.1101/2021.01.01.425046 (2021).
11. Wang, J. et al. A molecular grammar governing the driving forces for phase separation of prion-like RNA binding proteins. *Cell* **174**, 688–699 (2018).
12. Qamar, S. et al. FUS phase separation is modulated by a molecular chaperone and methylation of arginine cation-π interactions. *Cell* **173**, 720–734 (2018).
13. Vernon, R. M. et al. π-π contacts are an overlooked protein feature relevant to phase separation. *eLife* **7**, e31486 (2018).
14. Brady, J. P. et al. Structural and hydrodynamic properties of an intrinsically disordered region of a germ cell-specific protein on phase separation. *Proc. Natl Acad. Sci. USA* **114**, E8194–E8203 (2017).
15. Dubreuil, B., Matalon, O. & Levy, E. D. Protein abundance biases the amino acid composition of disordered regions to minimize non-functional interactions. *J. Mol. Biol.* **431**, 4978–4992 (2019).
16. Fossat, M. J., Zeng, X. & Pappu, R. V. Uncovering differences in hydration free energies and structures for model compound mimics of charged side chains of amino acids. *J. Phys. Chem. B* **125**, 4148–4161 (2021).
17. Dyson, H. J., Wright, P. E. & Scheraga, H. A. The role of hydrophobic interactions in initiation and propagation of protein folding. *Proc. Natl Acad. Sci. USA* **103**, 13057–13061 (2006).
18. Andrew, C. D. et al. Stabilizing interactions between aromatic and basic side chains in α-helical peptides and proteins. Tyrosine effects on helix circular dichroism. *J. Am. Chem. Soc.* **124**, 12706–12714 (2002).
19. Noid, W. G. Perspective: coarse-grained models for biomolecular systems. *J. Chem. Phys.* **139**, 090901 (2013).
20. Hills, R. D., Lu, L. & Voth, G. A. Multiscale coarse-graining of the protein energy landscape. *PLoS Comput. Biol.* **6**, e1000827 (2010).
21. Ruff, K. M., Harmon, T. S. & Pappu, R. V. CAMELOT: a machine learning approach for coarse-grained simulations of aggregation of block-copolymeric protein sequences. *J. Chem. Phys.* **143**, 243123 (2015).
22. Zeng, X., Holehouse, A. S., Chilkoti, A., Mittag, T. & Pappu, R. V. Connecting coil-to-globule transitions to full phase diagrams for intrinsically disordered proteins. *Biophys. J.* **119**, 402–418 (2020).
23. Latham, A. P. & Zhang, B. Consistent force field captures homologue-resolved HP1 phase separation. *J. Chem. Theory Comput.* **17**, 3134–3144 (2021).
24. Dannenhoffer-Lafage, T. & Best, R. B. A data-driven hydrophobicity scale for predicting liquid-liquid phase separation of proteins. *J. Phys. Chem. B* **125**, 4046–4056 (2021).
25. Tesei, G., Schulze, T. K., Crehuet, R. & Lindorff-Larsen, K. Accurate model of liquid–liquid phase behavior of intrinsically disordered proteins from optimization of single-chain properties. *Proc. Natl Acad. Sci. USA* **118**, e2111696118 (2021).
26. Dignon, G. L., Zheng, W. W., Kim, Y. C., Best, R. B. & Mittal, J. Sequence determinants of protein phase behavior from a coarse-grained model. *PLoS Comput. Biol.* **14**, e1005941 (2018).
27. Regy, R. M., Thompson, J., Kim, Y. C. & Mittal, J. Improved coarse-grained model for studying sequence dependent phase separation of disordered proteins. *Protein Sci.* **30**, 1371–1379 (2021).
28. Souza, P. C. T. et al. Martini 3: a general purpose force field for coarse-grained molecular dynamics. *Nat. Methods* **18**, 382–388 (2021).
29. Benayad, Z., von Bülow, S., Stelzl, L. S. & Hummer, G. Simulation of FUS protein condensates with an adapted coarse-grained model. *J. Chem. Theory Comput.* **17**, 525–537 (2021).
30. Reith, D., Pütz, M. & Müller-Plathe, F. Deriving effective mesoscale potentials from atomistic simulations. *J. Comput. Chem.* **24**, 1624–1636 (2003).
31. van Hoof, B., Markvoort, A. J., van Santen, R. A. & Hilbers, P. A. A novel method for coarse graining of atomistic simulations using Boltzmann inversion. *Biophys. J.* **100**, 309a (2011).
32. Ercolessi, F. & Adams, J. B. Interatomic potentials from first-principles calculations: the force-matching method. *Europhys. Lett.* **26**, 583–588 (1994).
33. Lu, L., Dama, J. F. & Voth, G. A. Fitting coarse-grained distribution functions through an iterative force-matching method. *J. Chem. Phys.* **139**, 121906 (2013).
34. Izvekov, S. & Voth, G. A. A multiscale coarse-graining method for biomolecular systems. *J. Phys. Chem. B* **109**, 2469–2473 (2005).
35. Johnson, M. E., Head-Gordon, T. & Louis, A. A. Representability problems for coarse-grained water potentials. *J. Chem. Phys.* **126**, 144509 (2007).
36. Reinhardt, A. & Cheng, B. Quantum-mechanical exploration of the phase diagram of water. *Nat. Commun.* **12**, 588 (2021).

37. Wang, J. et al. Machine learning of coarse-grained molecular dynamics force fields. *ACS Cent. Sci.* **5**, 755–767 (2019).
38. Opitz, A. Molecular dynamics investigation of a free surface of liquid argon. *Phys. Lett. A* **47**, 439–440 (1974).
39. Wang, X., Ramírez-Hinestrosa, S., Dobnikar, J. & Frenkel, D. The Lennard-Jones potential: when (not) to use it. *Phys. Chem. Chem. Phys.* **22**, 10624–10633 (2020).
40. Das, S., Lin, Y.-H., Vernon, R. M., Forman-Kay, J. D. & Chan, H. S. Comparative roles of charge, π, and hydrophobic interactions in sequence-dependent phase separation of intrinsically disordered proteins. *Proc. Natl Acad. Sci. USA* **117**, 28795–28805 (2020).
41. Kim, Y. C. & Hummer, G. Coarse-grained models for simulations of multiprotein complexes: application to ubiquitin binding. *J. Mol. Biol.* **375**, 1416–1433 (2008).
42. Kapcha, L. H. & Rossky, P. J. A simple atomic-level hydrophobicity scale reveals protein interfacial structure. *J. Mol. Biol.* **426**, 484–498 (2014).
43. Li, H., Tang, C. & Wingreen, N. S. Nature of driving force for protein folding: a result from analyzing the statistical potential. *Phys. Rev. Lett.* **79**, 765–768 (1997).
44. Urry, D. W. et al. Hydrophobicity scale for proteins based on inverse temperature transitions. *Biopolymers* **32**, 1243–1250 (1992).
45. Tejedor, A. R., Garaizar, A., Ramírez, J. & Espinosa, J. R. Dual RNA modulation of protein mobility and stability within phase-separated condensates. Preprint at *bioRxiv* https://doi.org/10.1101/2021.03.05.434111 (2021).
46. Lin, Y.-H. & Chan, H. S. Phase separation and single-chain compactness of charged disordered proteins are strongly correlated. *Biophys. J.* **112**, 2043–2046 (2017).
47. Riback, J. A. et al. Stress-triggered phase separation is an adaptive, evolutionarily tuned response. *Cell* **168**, 1028–1040 (2017).
48. Dignon, G. L., Zheng, W., Best, R. B., Kim, Y. C. & Mittal, J. Relation between single-molecule properties and phase behavior of intrinsically disordered proteins. *Proc. Natl Acad. Sci. USA* **115**, 9929–9934 (2018).
49. Fare, C. M., Villani, A., Drake, L. E. & Shorter, J. Higher-order organization of biomolecular condensates. *Open Biol.* **11**, 210137 (2021).
50. Regy, R. M., Dignon, G. L., Zheng, W., Kim, Y. C. & Mittal, J. Sequence dependent phase separation of protein-polynucleotide mixtures elucidated using molecular simulations. *Nucleic Acids Res.* **48**, 12593–12603 (2020).
51. Choi, J.-M., Hyman, A. A. & Pappu, R. V. Generalized models for bond percolation transitions of associative polymers. *Phys. Rev. E* **102**, 042403 (2020).
52. Zeng, X. et al. Design of intrinsically disordered proteins that undergo phase transitions with lower critical solution temperatures. *APL Mater.* **9**, 021119 (2021).
53. Banerjee, P. R., Milin, A. N., Moosa, M. M., Onuchic, P. L. & Deniz, A. A. Reentrant phase transition drives dynamic substructure formation in ribonucleoprotein droplets. *Angew. Chem. Int. Ed.* **56**, 11354–11359 (2017).
54. Alshareedah, I. et al. Interplay between short-range attraction and long-range repulsion controls reentrant liquid condensation of ribonucleoprotein-RNA complexes. *J. Am. Chem. Soc.* **141**, 14593–14602 (2019).
55. Dignon, G. L., Zheng, W., Kim, Y. C. & Mittal, J. Temperature-controlled liquid-liquid phase separation of disordered proteins. *ACS Cent. Sci.* **5**, 821–830 (2019).
56. Best, R. B., Zheng, W. & Mittal, J. Balanced protein-water interactions improve properties of disordered proteins and non-specific protein association. *J. Chem. Theory Comput.* **10**, 5113–5124 (2014).
57. Benavides, A. L., Aragones, J. L. & Vega, C. Consensus on the solubility of NaCl in water from computer simulations using the chemical potential route. *J. Chem. Phys.* **144**, 124504 (2016).
58. Liu, H., Fu, H., Shao, X., Cai, W. & Chipot, C. Accurate description of cation-π interactions in proteins with a nonpolarizable force field at no additional cost. *J. Chem. Theory Comput.* **16**, 6397–6407 (2020).
59. Paloni, M., Bailly, R., Ciandrini, L. & Barducci, A. Unraveling molecular interactions in liquid-liquid phase separation of disordered proteins by atomistic simulations. *J. Phys. Chem. B* **124**, 9009–9016 (2020).
60. Wessén, J., Pal, T., Das, S., Lin, Y.-H. & Chan, H. S. A simple explicit-solvent model of polyampholyte phase behaviors and its ramifications for dielectric effects in biomolecular condensates. *J. Phys. Chem. B* **125**, 4337–4358 (2021).
61. Holcomb, C. D., Clancy, P. & Zollweg, J. A. A critical study of the simulation of the liquid-vapour interface of a Lennard-Jones fluid. *Mol. Phys.* **78**, 437–459 (1993).
62. Reinhardt, A. Phase behavior of empirical potentials of titanium dioxide. *J. Chem. Phys.* **151**, 064505 (2019).
63. Gallivan, J. P. & Dougherty, D. A. Cation-π interactions in structural biology. *Proc. Natl Acad. Sci. USA* **96**, 9459–9464 (1999).
64. Song, J., Ng, S. C., Tompa, P., Lee, K. A. W. & Chan, H. S. Polycation-π interactions are a driving force for molecular recognition by an intrinsically disordered oncoprotein family. *PLoS Comput. Biol.* **9**, e1003239 (2013).
65. Auton, M. & Bolen, D. W. Application of the transfer model to understand how naturally occurring osmolytes affect protein stability. *Methods Enzymol.* **428**, 397–418 (2007).
66. Kumar, K. et al. Cation-π interactions in protein-ligand binding: theory and data-mining reveal different roles for lysine and arginine. *Chem. Sci.* **9**, 2655–2665 (2018).
67. Chapela, G. A., Saville, G., Thompson, S. M. & Rowlinson, J. S. Computer simulation of a gas-liquid surface. Part 1. *J. Chem. Soc. Faraday Trans. 2* **73**, 1133–1144 (1977).
68. Nilsson, D. & Irbäck, A. Finite-size scaling analysis of protein droplet formation. *Phys. Rev. E* **101**, 022413 (2020).
69. Vitalis, A. & Pappu, R. V. ABSINTH: a new continuum solvation model for simulations of polypeptides in aqueous solutions. *J. Comput. Chem.* **30**, 673–699 (2009).
70. Joseph, J. A. et al. Code and data for 'Physics-driven coarse-grained model for biomolecular phase separation with near-quantitative accuracy'. *figshare* https://doi.org/10.6084/m9.figshare.16772812 (2021).

Acknowledgements

We thank D. Frenkel for useful comments on the manuscript, J. Mittal and G. L. Dignon for helping us implement the HPS-KR potential in LAMMPS and G. Tesei and K. Lindorff-Larsen for helping us debug our implementation of their potential. This project has received funding from the European Research Council under the European Union's Horizon 2020 research and innovation programme (grant no. 803326; R.C.-G.). J.A.J. is a Junior Research Fellow at King's College. R.C.-G. is an Advanced Fellow of the Winton Programme for the Physics of Sustainability. J.R.E. acknowledges funding from the Oppenheimer Fellowship of the University of Cambridge and the Roger Ekins Fellowship from Emmanuel College. A.G. is funded by the EPSRC (Doctoral Training Partnership, grant no. EP/N509620/1) and the Winton Programme for the Physics of Sustainability. P.Y.C. is funded by the University of Cambridge Ernest Oppenheimer Fund and the Winton Programme for the Physics of Sustainability. K.O.R. is funded by the EPSRC (Doctoral Training Partnership, grant no. EP/T517847/1). This work was performed using resources provided by the Cambridge Tier-2 system operated by the University of Cambridge Research Computing Service funded by EPSRC Tier-2 capital grant no. EP/P020259/1 (R.C.-G., J.A.J. and A.R.). The funders had no role in study design, data collection and analysis, decision to publish or preparation of the manuscript.

Author contributions

J.A.J. and R.C.-G. conceived the project. J.A.J., A.R. and R.C.-G. designed the model and benchmarking framework. J.A.J. and A.R. implemented and optimized the model. J.A.J., A.R., A.A., P.Y.C., K.O.R., J.R.E. and A.G. validated the model and analyzed the data. J.A.J. and A.R. wrote the manuscript with help from R.C.-G. All authors reviewed the manuscript. J.A.J., A.R. and R.C.-G. acquired funding. J.A.J., A.R. and R.C.-G. supervised the research.

Competing interests

The authors declare no competing interests.

Additional information

Supplementary information The online version contains supplementary material available at https://doi.org/10.1038/s43588-021-00155-3.

Correspondence and requests for materials should be addressed to Jerelle A. Joseph, Aleks Reinhardt or Rosana Collepardo-Guevara.

Peer review information *Nature Computational Science* thanks Hue Sun Chan and the other, anonymous, reviewer(s) for their contribution to the peer review of this work. Peer reviewer reports are available. Handling editor: Jie Pan, in collaboration with the *Nature Computational Science* team.