A study of intrinsically disordered proteins using molecular dynamics simulations

Autor: Pablo Herrera Nieto

TESI DOCTORAL UPF / year L’any de la tesi: 2020

THESIS SUPERVISOR
Dr. Gianni De Fabritiis
Department Departament de Ciències Experimentals i de la Salut

Nothing ever ends.

--- Watchmen ---

Acknowledgements

I would like to thank my thesis supervisor, labmates, friends, family, partner, and, specially, Ministerio de Ciencia, for the funding, and GPUGRID volunteers, for their time.

Abstract

Over the last decades, molecular dynamics simulations have been successfully applied to relevant biological problems such as protein-ligand and protein-protein binding, as well as protein folding. A perfect challenge for molecular simulations is the study of intrinsically disordered proteins as they present faster timescales than structured proteins, and, therefore, can be explored more exhaustively. The main objectives of this work include the exploration of the conformational space of p53 by revealing the presence of many partially ordered states, the reconstruction of the coupled folding and binding of a disordered protein and its folded partner by applying novel reinforcement learning inspired sampling algorithms, and the performance of free-ligand binding assays to address the potential druggability of disordered proteins. The compendium of work presented here contributes to the understanding of such intrinsically disordered proteins at an atomic level, highlighting key aspects of their behavior in isolation, binding mechanisms, and external modulation.

Resumen

A lo largo de las últimas décadas las simulaciones de dinámica molecular han sido aplicadas con éxito en problemas biológicos como la interación proteína-ligando o proteína-proteína, así como en su plegamiento. Un desafío idóneo para las simulaciones es el estudio de proteína desordenadas, ya que presentan escalas de tiempo mas rápidas que la proteínas plegadas, lo que permite una exploración mas exhaustiva de las mismas. Entre los principales objetivos del presente trabajo figuran la exploración del paisaje conformacional de p53 revelando la presencia de estados parcialmente ordenados, la reconstrucción del acoplamiento de unión y plegamiento de una proteína desordenada a su pareja, aplicando novedosos algoritmos de muestreo inspirados en aprendizaje reforzado, y la realización de ensayos de unión de ligando libre para abordar la potencial drogabilidad de las proteínas desordenadas. El compendio de trabajo presentado aquí contribuye a entender dichas proteínas a nivel atómico, destacando aspectos clave de su comportamiento en aislamiento, de sus mecanismos de unión, y de su modulación externa.

Contents

List of Figures

Chapter 1

INTRODUCTION

1.1. Intrinsically Disordered Proteins

1.1.1. Breaking the paradigm

Up to the end of the twentieth century, the common understanding of protein behavior can be summarized with the *sequence-structure-function paradigm*, which establishes that sequence, structure, and function of a protein are intimately related, being the latter the consequences of the former. The paradigm had to be adapted with the discovery of a group of proteins able to perform biologically relevant tasks despite lacking a stable three-dimensional structure, the intrinsically disordered proteins or IDPs (1; 2).

The analysis of genomic data showed that a significant fraction of proteomes, up to 30% in the case of the human, corresponds to fully disordered proteins or disordered regions within folded proteins (3; 4). Comparing predicted levels of proteome disorder across species served to illustrate the asymmetrical distribution of disorder between eukaryotic and prokaryotic proteomes, with a remarkable enrichment of IDPs in more evolved forms of life, especially in cellular signaling and regulatory-related pro-

cesses (1; 5; 6). Hence, IDPs serve as hubs in complex protein networks controlling transcription, translation, and cell-cycle routines. At the same time, IDPs themselves are also subject to several layers of regulation. On the one hand, their synthesis and degradation levels are finely tuned, affecting the location and time-lapse time-frame they will be found (7). On the other hand, additional mechanisms such as alternative splicing and post-translational modifications (PTMs) control the presence or absence of particular functional sections that, in turn, control their dynamic behavior (8; 9; 10).

IDPs composition, sequence-wise, is biased towards hydrophilic and charged residues, and prolines, while depleted of bulkier hydrophobic amino acids (4). In this manner, this depletion prevents the formation of a hydrophobic core that could promote folding. Consequently, IDPs exhibit a relatively flat energy surface composed of many local minima allocating different fast interchanging conformations, in opposition to the funnel-like surface of folded proteins, which has a well defined global minima region, as shown in Figure 1.1.

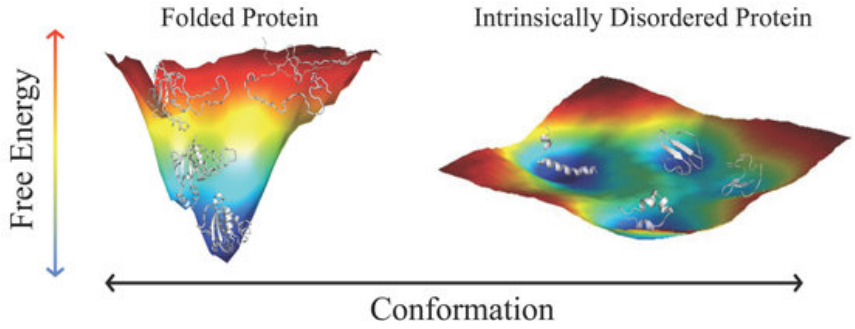

Figure 1.1: **Free energy surface of proteins.** The free energy surface of folded proteins shows a global minima corresponding to folded conformations. IDPs, on the other hands, display a flat energy surface, including many isoenergetic conformations. Extracted from (11).

However, some IDPs do not interconvert between random coil-like structures endlessly, since their flexibility can be modified by interacting with other proteins. There are many examples of IDPs that fold upon binding to their targets, acquiring secondary structural elements (12; 13; 14; 15; 16; 17; 18). IDPs interactions are not restricted to one partner only, and they bind promiscuously. A single IDP can interact with many targets, exhibiting alternative conformations. For instance, the C-terminal of p53, as illustrated in Figure 1.2, binds up to four targets. To the question of how folding couples with binding two extreme solutions have been proposed: *conformational selection*, by which pre-folded conformations of the IDP bind to is partners, or *induced fit*, where folding follows binding. Other IDPs, contrarily, remain disordered upon binding, especially polyampholytes chains, which mainly interact by electrostatic complementarity reaching affinities in the order of picomolar (19).

From the kinetic and thermodynamic point of view, binding of disordered proteins takes place in a high-specificity and low-affinity fashion (21). In the first place, specificity arises from their amino acid sequence. IDPs bind their partners through short linear interactive motifs (SLIMs) (22) of 10 to 20 residues where 3 or 4 positions take a central role in the interaction. The lack of structure is an advantage from an evolutionary perspective, as it allows for the quick evolution of SLIMs. On the other hand, affinity is regulated by binding (k_{on}) and unbinding (k_{off}) rates. k_{on} of IDPs is as fast, and sometimes even faster than for the folded counterparts. A typical IDP binds with a k_{on} of 10^6-10^7 $M^{-1}s^{-1}$, in line, for example, with ligand binding to GPCRs. On the other hand, k_{off} ranges from 1 to 1000 s^{-1} (23) including both fast and slow unbind. Low affinities are due to the higher entropic penalty of binding for IDPs compared to folded proteins, as the associated folding will reduce the potential number of conformations likely to be explored. Thus, the higher the rigidity of the bound conformation, the higher this penalty would be. However, the enthalpic term can compensate it (24) thanks to the creation of favorable interactions, as revealed by protein-IDP complexes, where the driving

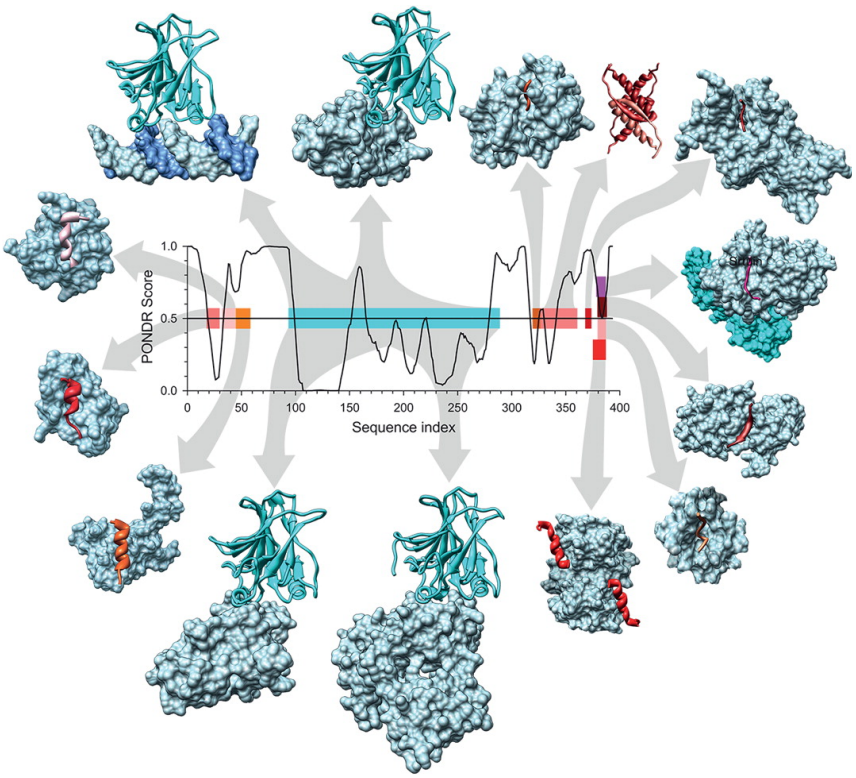

Figure 1.2: **Binding promiscuity of p53.** By residue disorder prediction for p53. A single segment of an IDP is able to bind multiple partners with unique conformations. p53 highlight this feature: its C-terminal binds to at least five different partners, while other regions can also interact with many other proteins. Extracted from (20).

forces for the interaction are lead by the burial of hydrophobic residues on the disordered partner and the formation of electrostatic interactions (25; 15). Altogether, these two features allow IDPs to exert their function as molecular switches and rheostats, shaping signaling networks through

transient albeit specific contacts.

Finally, one field of study of IDP biology that has gathered more attention in recent years is liquid-liquid phase separation (LLPS). LLPS is a reversible process by which a homogeneous liquid segregates in a diluted and a condensed phase (26; 27). LLPS has vast implications in the cytosolic organization and results in the formation of dynamic membrane-less organelles mediated by RNA molecules and proteins, especially IDPs. It has been suggested that these bodies host machinery to speed up biochemical reactions (28) and could also promote pathological protein fibrillization events such as the formation of tau aggregates.

As previously mentioned, IDPs participate in many regulatory processes, and their misregulation has implications in neurological disorders and cancer. However, they are generally not considered as feasible drug targets, given that traditional approaches focus on the development of small molecules able to block well defined 3D pockets, which are not found in IDPs. Other approximations have been studied, including the development of peptidomimetics resembling bound conformation of IDPs (i.e., a set of molecules called *nutlins* which mimic the conformation of p53 bound to Mdm2 (29)), as well as small molecules directly interacting with disordered targets such as the androgen receptor (by the EPII molecule (30; 31)), the *nuclear protein 1* (32), the cell-cycle regulator $p27^{kip1}$ (33; 34), and others (35; 36).

To sum up, IDPs represent an additional facet of protein functionality, and while providing a set of advantageous features from the biological perspective, their study has also proven experimentally complicated.

1.1.2. Experimental methodologies on structural biology

The methodologies employed in structural biology, either experimental or computational, are limited in both resolution (from atoms to complex molecular assemblies) and accessible timescales (ranging from picoseconds to minutes) as represented in Figure 1.3. Notably, the study of IDPs presents the additional problem of how to deal with their inherent flexibility, as most approaches are better suited for the analysis of folded ones.

X-ray crystallography is one of the leading examples. It is the primary source of protein structures in the PDB (currently around 83%), and it allows deriving a 3D electron density map of a protein crystal from an X-ray diffraction pattern, thus providing a static image of the molecule. Despite not being applicable for the general study of IDPs, it was especially useful at the initial stages of the field: the problems when crystallizing several proteins and the presence of missing sections hinted about the notion of disorder within proteins. Additionally, both X-ray crystallography and Nuclear Magnetic Resonance spectroscopy (NMR) were able to resolve the firsts folded protein—IDP complexes, such as p53 bond to Mdm2 (12), $p27^{kip1}$ to the cyclin A-Cdk2 complex (25), KIX—KID (13), and KIX—c-Myb (15), showing how they could fold upon interaction with their partners.

Small Angle X-ray Scattering (SAXS) (38) is another biophysical method based on X-rays that provides quantitative information about protein shape and size feasible to study, which is to study flexible systems such as IDPs. This technique is based on the scattering profile of an X-ray beam irradiating a protein sample in solution, rather than a crystal, which offers low-resolution information on protein conformations. The primary SAXS-derived parameter is the radius of gyration (R_g), an informative measurement of the average root mean square deviation from the center of mass of the molecule, that is, its level of compaction.

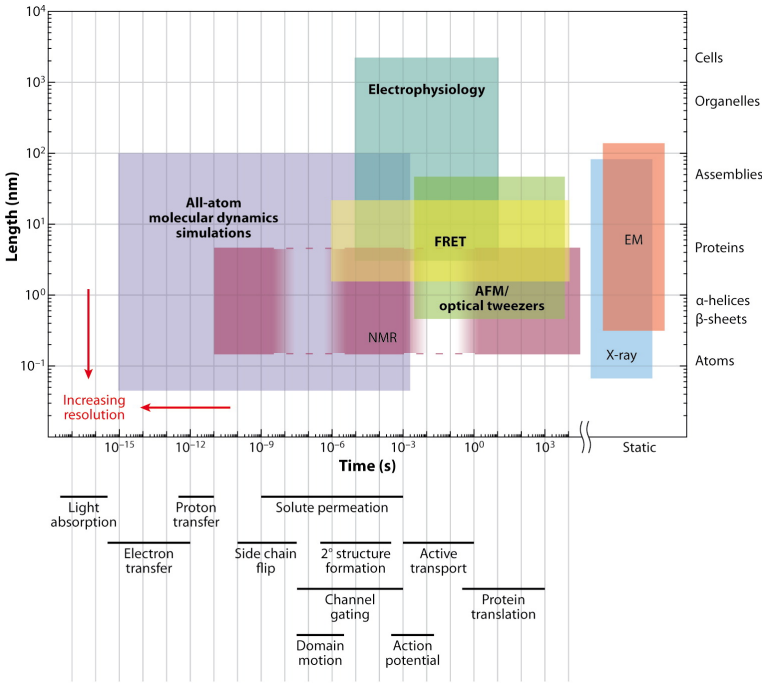

Figure 1.3: **Research tools employed on structural biology.** Each tool is limited by its spatio-temporal resolution. Computational techniques, such as MD simulations, are able to provide atomistic resolutions while accessing long times scales. Similarly, NMR resonance also have access to a wide range of temporal scales, while others (X-ray and EM) only provides a static image. mExtracted from (37).

Nevertheless, the most capable experimental methodology for the study of IDPs is NMR spectroscopy (39; 40), as previously mentioned. It is based on the NMR phenomenon, by which nuclei under a homogeneous external magnetic field absorb energy at a given frequency depending on their type and environment. The main advantage of NMR over X-ray diffraction is

that, by using samples in solution, protein crystals are no longer needed. In a broad sense, the frequency of absorption is referred as chemical shift (as summarized in Figure 1.4), and they represent one of the most descriptive parameters of protein dynamics. Most of the studies focus on the chemical shifts of ^{1}H and ^{15}N atoms, but for peptides and IDPs ^{1}H dispersion is limited and in the recent years the detection of ^{13}C has became an alternative for IDP investigation (39; 41).

The power of NMR resides in its potential to characterize dynamic behaviors, such as population shifts, exchange rates between conformers, and structural backbone constraints (including secondary structure assignment), that altogether shape the ensemble of explored conformations. NMR is also relevant in the study of how IDPs ensembles are affected by PTMs (42), by other proteins (15; 43), and even small molecules (33; 34). In this way, NMR is able to discriminate binding regions, to characterize intermediate states, and to deepen into the processes of aggregation and phase separation, amongst others.

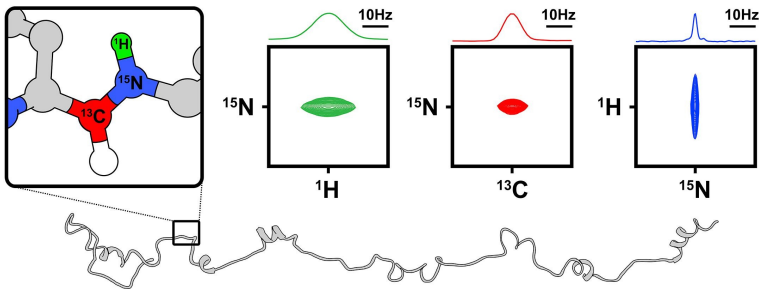

Figure 1.4: **Chemical shift assignment**. Backbone chemical shift assignments is performed with ^{1}H, ^{15}N, and ^{13}C atoms. Extracted from (41).

Other computational techniques for the study of IDPs and extensively employed throughout the thesis are described in more detail in *Section 1.2*.

1.2. From simulations to sampling: one step at a time

Since the appearance of the first protein crystal structure in 1958 (44), one of the main objectives in the structural biology field was to provide these static structures with dynamic behavior that could advance the knowledge of how protein machinery works. First Molecular Dynamics (MD) simulations were performed in the 1970s, reaching the modest amount from today's perspective of 8.8 ps (45). Since then and after multiple technological upgrades, MD has consolidated as one of the most relevant computational tools in the exploration of protein functionality.

Current technology allows generating massive amounts of aggregate simulation time, reaching up to hundreds of microseconds or even the millisecond scale. Some reviews (46) have estimated that the frontier of one second on simulation could be reached by 2022. All this data is used to build probabilistic models known as Markov States Models (MSMs) (47; 48).

In order to explore even longer timescales or to perform a better sampling of the configurational space, new algorithms have been developed under the name of Adaptive Sampling techniques (49). Their primary purpose is to guide sampling, starting new rounds of simulations based on the most promising conformations previously sampled.

This stack of methodologies defined by MD simulations for data generation, MSMs for analysis, and Adaptive Sampling for boosting exploration and performance, compose the fundamental tools used throughout the present thesis for the study of IDPs.

1.2.1. Molecular Dynamics Simulations

MD is a computational technique for simulating time evolution of atomic models in phase space. From an elementary perspective, MD could also be understood as a time-lapse movie where every frame represents the current coordinates of all atoms in the system. Additionally, the time separation between consecutive frames is in the order of picoseconds. Many elements intervene in the production of a simulation trajectory. Each choice made in the building and running process implies a trade-off between performance and accuracy. Some of the most relevant ones will be covered next.

In the first place, a molecular system is built comprising solvent, protein and other molecules. For the proper representation of the solvent, there are two main choices, either implicit or explicit. Implicit solvent (50) treats water as a continuous medium, whereas more costly explicit solvent models include all atoms for each molecule. Explicit water models diverge in the number of points used to represent each water molecule (three or four), in their geometry, (planar or non-planar), and the rigidity of the bonds, The water model employed throughout the thesis is the widely used rigid three-points TIP3P (51).

The next step in the process supposes one of the most critical decisions as it determines the accuracy of the results obtained: the choice of the force field. Force fields are designed to fit experimental or quantum mechanics (QM) derived data and include all necessary equations and parameters to describe the attributes of the atoms, their bonds, and interactions. Most used families of force fields in academia include CHARMM (52) and AMBER (53; 54). Each one encompassed several force fields adapted to different macromolecules such as nucleic acids (55; 56) or proteins.

Simulations propagate the evolution of the system through time by solving Newton's equation of motion. At each step, it is necessary to calculate the forces acting over every atom (58). This forces are generally divided into two groups: *bonded terms* —including bond, angle, dihedral,

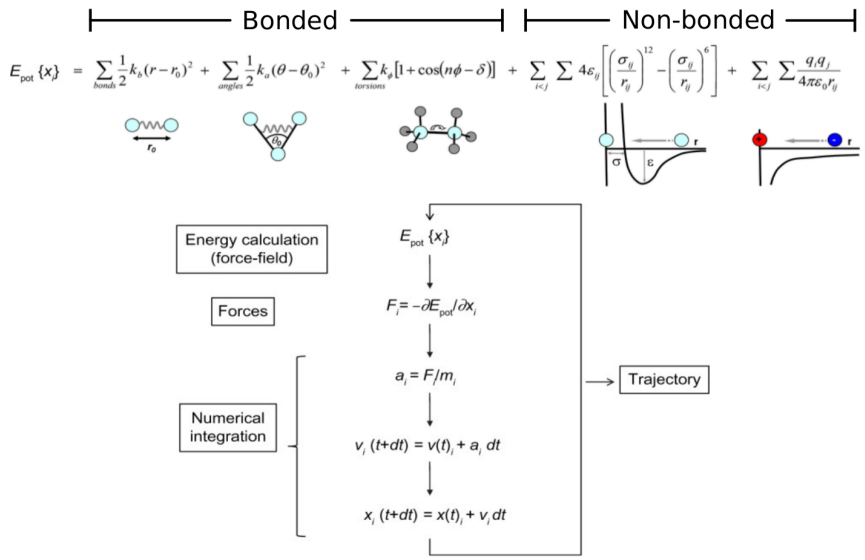

Figure 1.5: **Molecular Dynamics simulations summary.** *top)* MD potential energy function. Classical MD represents atomic systems as a collection of *balls* (atoms) and *springs* (bonds) that are related by bonded terms (bond, bond angles, torsions), and non-bonded terms (van de Waals and Coulumbic interactions). Adapted from (57). *bottom)* Basic steps performed by the MD engine. Extracted from (58).

and improper terms— and *non-bonded terms* —with electrostatic and van der Waals terms—, as depicted in scheme 1.5. In order to maintain the integrity of the system and speed up calculations, several techniques are used. Some of them include the introduction of cutoffs for the calculation of non-bonded terms, the implementation of constraints of certain types of atoms to increment the time step, and the definition of periodic boundary conditions to deal with the finite size of the system and related surface effects. The numerical solution for the equation of motion to obtain new atomic po-

sitions and velocities after each step is performed with integration schemes are such as the Verlet algorithm. The time step separating frames in a trajectory needs to be shorter than the fastest processes observed, to ensure the stability of the system. Hence, it ranges on the femtoseconds ($10^{-15}s$ scale, faster than bond vibrations).

Simulations of classical MD represent the system with a *"ball and springs"* model, for atoms and bonds, respectively. However, other types of simulations choose a different representation. Coarse-grained simulations (59), for instance, define a broader representation of atomic systems by grouping atoms into single entities. This approach simplifies the representation of the system, allowing it to increment the time step and, therefore, making it possible to explore longer times scales at the cost of reduced accuracy. On the other hand, systems can also be represented in greater detail than classical MD. Quantum mechanics, for instance, also describes the electronic behavior. Therefore this technique is suitable for the study of reactivity (bond rupture and formation). However, it is limited to small systems due to its high computational cost.

GPU-based MD

Biologically relevant timescales are orders of magnitude higher than the time step used in MD simulations, for instances, side chain rotation is in the range of nanoseconds (10^{-9} s), secondary structure element formation can take up to microseconds (10^{-6} s), and binding/unbinding of proteins with their ligands or protein folding can extend up to the second. This means that it is necessary to perform millions or billions of steps to reach those timescales, thus making of MD simulation a computationally intensive technique. Therefore, its performance over the years has been intimately linked to the development of more capable software and hardware.

For a long time, MD was run almost exclusively in CPU and CPU-based clusters, and the main advances included the development of spe-

cialized hardware, such as ANTON supercomputers (60), which were able to reach microsecond-length simulations. To mitigate the performance problems, many biased MD techniques appeared, including metadynamics (61; 62) or umbrella sampling (63), which applied external forces in the system to bias sampling towards the events of interest. Fortunately, MD is a highly parallelizable process, and it greatly benefited from the introduction of novel GPU architectures like CUDA. With the appearance of the firsts GPU-based engines, such as ACEMD (64), MD performance was greatly boosted, making it possible to routinely reach microsecond timescales with more accessible infrastructures, and even higher with distributed systems such as GPUGRID (65), the one used during this thesis.

Force Fields for IDPs

As stated before, the accuracy of a simulation is determined by the chosen force field. Using a generic protein force field has revealed to be imprecise for the simulation of IDPs. This should not be surprising, given that their parameters are developed to match those observed in folded proteins.

There are two main problems with the application of general force fields to IDPs (66): *1)* the unbalance in the sampling of secondary structural elements, and *2)* the propensity to generate over collapsed conformations, being close to maximally compact globules. Tuning backbone torsion parameters addressed the first issue. For the second, several solutions have been proposed, including the strengthening of the protein water-interactions (as done in the Amber ff03ws forcefield (67)), and the development of new water models, such as TIP4P-D (68). The implementation of additional approaches to boost accuracy is still an active area of research, with periodic releases modified force fields (69; 70; 66; 71).

Over this thesis, a modified version of the CHARMM force field called CHARMM22* (72) has been used. It was developed by Piana et al. to solve the overstabilization of helices and salt bridges.

1.2.2. Markov State Models

Initial conditions of an MD run, velocities —set at random— and coordinates might have a great impact on the output trajectory. Thus, conclusions drew from a single trajectory, despite its length, will likely be biased. Multiple trajectories can be run simultaneously to improve sampling and reduce this bias, and with the help of high-performance infrastructures such as GPU-based clusters, thousands of trajectories can be generated within weeks or months.

Markov states models (MSMs) were introduced to address the challenge of how to condense vast amounts of high dimensional data (N atoms $\times$ 3 spatial coordinates), split across many short trajectories, into a single model that discriminates between fast and slow processes on long timescales (47). MSMs describe time series data as stochastic processes defined by a set of discrete N states and an N-N transition matrix with the probabilities of jumping between them. Building an MSM with a discrete number of kinetically relevant states starting from continuous MD trajectories involves a complex pipeline summarized in Figure 1.6 and composed by the following steps:

1. **Data projection** is generally the first step. It reduces dimensionality by computing a set of features from the trajectories. In particular, the most effective metrics that are generally employed include distances or contacts between a set of atoms (i.e., C_α distances to track folding or protein-protein binding), and dihedrals angles (i.e., to monitor IDP rigidity).

2. **Time independent component analysis (TICA)** (73) is then used to reduce the number of dimensions. In principle, it is similar to other techniques, such as PCA (74), but instead of capturing the variance of the data, TICA finds the slowest processes.

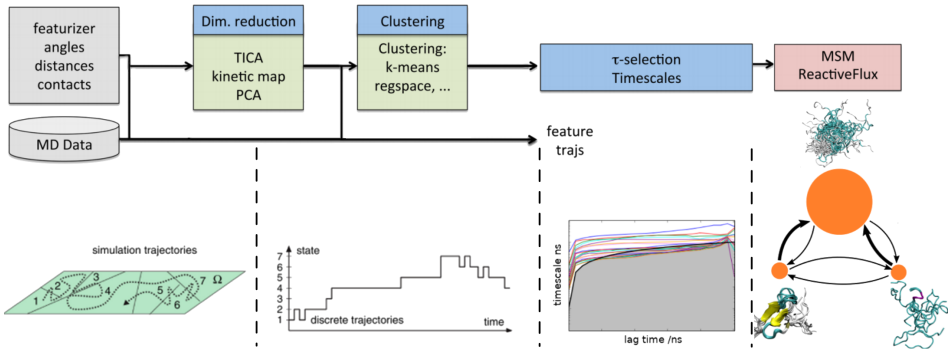

Figure 1.6: **Markov state model building.** Schematic and graphical diagram representing the mains steps in the construction of a MSM. Original MD data is featurized by computing distances, contacts, or angles. TICA is used to fatherly reduce the number of dimensions. Clustering provides the space discretization needed to finally create a MSM. Adapted from (75; 48).

3. **Clustering** can be applied now in the reduced number of dimensions obtained in the previous step. It provides a basic discretization of the continuous space of atomic coordinates found in a trajectory into a succession of transitions from one state to another. One of the most widespread algorithms to perform this task is KMeans (76), which can handle thousands or millions of data points. Determining the optimal number of clusters is complicated, and testing is generally required.

4. **MSM building** constitutes the last step in the pipeline. It is achieved by generating the transition matrix with the jumps between every cluster at a given lag time. Finally, to get a reduced number of macrostates (generally less than 10 is enough), the fast interconverting microstates are merged with the Perron-cluster cluster analysis

(PCCA) algorithm (77).

MSMs can be used to estimate diverse parameters of the system and have been successfully applied in protein folding (78; 79) and protein-ligand (80; 81) binding studies. The transition matrix allows estimating the equilibrium probabilities of the states, which are, in turn, intimately related to the free energy associated with the transition from one state to another. Kinetic and energetic parameters of interest such as the association (k_{on}) and dissociation (k_{off}) rates, mean first passages times ($mpft_{on}$ and $mpft_{off}$), and free energies can also be derived from the MSM. Similarly, transition path theory (82) can be used to understand how the flux evolves from one source state *A* to a sink state *B*, what intermediates states are visited, and quantitatively assess their participation in the transition.

1.2.3. Adaptive Sampling

High throughput MD simulations are performed with a goal in mind: to create a kinetic model as complete as possible of a particular phenomenon, generally binding or folding. As stated before, the space explored by a single trajectory will correspond to a limited region of the configurational landscape. Ensuring the exhaustive exploration of such a surface is not trivial, as challenging tasks such as traversing high energy barriers in funneled landscapes are not likely to happen at the first try. This *"luck"* factor can be outperformed by launching many simulations in parallel from the same point and trusting that rare events will be sampled sooner or later. Taking advantage of MSMs analysis at this point pushes the approach to a new stage. After a set of simulations is finished, an MSM is constructed with the available data to select the most promising conformations for spawning a new run of simulations (49). This iterative scheme, where rounds of simulations are run cyclically, and a policy is applied to select a set of new generators, is known as adaptive sampling (Figure 1.7).

Adaptive sampling methodologies represent a natural evolution of the increased power of the analytical tools and the capabilities of running multiple parallel simulations. They can be designed to provide a more efficient exploration of the configurational space, finding new states, or reaching the desired one faster. They have proved their usefulness in many problems such as protein-ligand binding, with a speedup of one order of magnitude (49; 83), as well as in protein-protein binding (84).

The deciding point in the pipeline is the policy implemented for selecting starting conformations for the new set of simulations. Initial approaches can be defined as *heuristic*. They rank conformations based on various criteria such as the equilibrium probability or numbers of counts (49), which will consequently prioritize sampling from less explored states. Other approaches focus on spawning new simulations from states which

contribute the most to errors in the mean first passage time (85) or eigenvalues and eigenvectors (86). Finally, there are also explicit rankings, defined by the so-called *goal functions*, which assigns a score to every structure based on external knowledge (83), i.e., the similarity with an experimental structure.

However, when performing adaptive sampling, fully exploratory approaches are also disadvantageous, as they can explore high energy regions or rare conformations of low interest without performing sufficient sampling of biologically relevant events. Then, one of the fundamental decisions that an adaptive sampling policy has to encode is how to face the *exploration vs. exploitation* dilemma: the decision of how to trade between exploring new states and sampling from already known ones, or in other words, how to choose between the actions with the maximum estimated value or non-optimal actions to update their value-estimates.

The *exploration vs. exploitation* dilemma has been extensively studied in other research areas, such as reinforcement learning (RL) (87). RL is a subarea of machine learning concerned with the development of optimal behaviors given an environment, a set of actions that can be performed, and a notion of performance or fitness that wants to be maximized based on the actions taken. Framing adaptive sampling in an RL context is the latest step towards the development of a more evolved sampling policy.

In particular, MD sampling can be described as a multi-armed bandit problem, a subproblem of RL where an agent has to choose actions from a limited action set, identified here with the states of an MSM, each one with an associated numerical reward coming from an unknown stationary probability distribution (87). The objective is to maximize the cumulative reward over a finite number of choices. The main challenge to optimally solve the problem is learning which actions are the most valuable and exploit them.

In this context, the novel AdaptiveBandit was recently developed (88). AdaptiveBandit defines adaptive sampling under a multi-armed bandit framework precisely to tackle the exploration-exploitation dilemma, faced

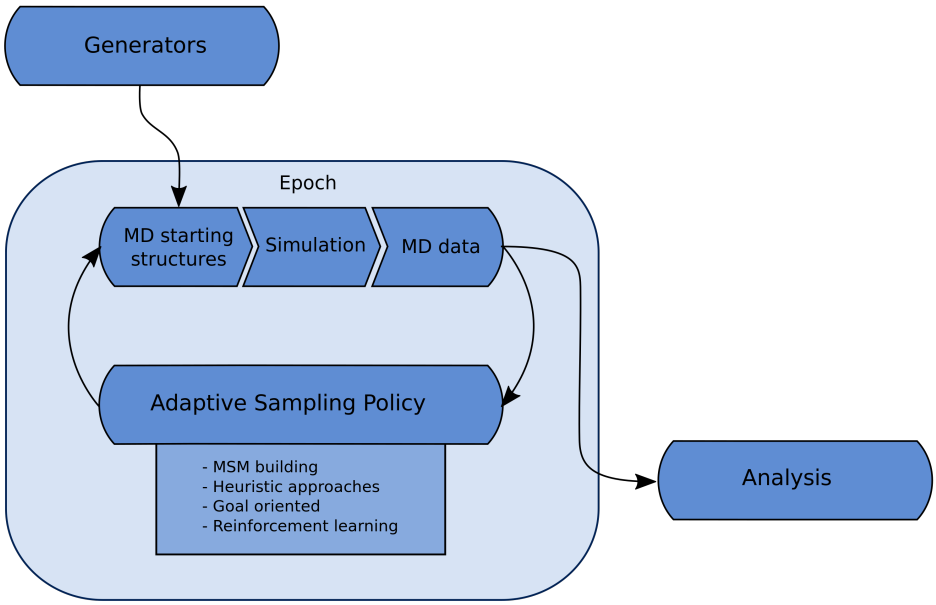

Figure 1.7: **Adaptive Sampling** is an iterative process where every epoch a new set of simulations is run, and the accumulated data is evaluated by a policy that chooses starting conformations for another round.

at each new round of simulations, to handle the exploration of new conformations while sampling from the most rewarding known states. AdaptiveBandit makes use of the MSM's discretized conformational space to define the action set. As for the reward definition, different choices depending on the objective, i.e. if the interest is sampling a bound metastable state, actions should be rewarded based on the stability of conformations using MSM estimations of equilibrium probabilities for each state. Lastly, at each round, a random conformation is picked from the selected states/actions to respawn new simulations.

To handle the exploration-exploitation dilemma, AdaptiveBandit relies

on the UCB1 algorithm (89), defining an upper confidence bound for each action-value estimate based on the number of times an action has been picked and the total amount of actions taken:

$$a_t = \underset{a \in \mathcal{A}}{\operatorname{argmax}} \left[Q_t(a) + c\sqrt{\frac{\ln t}{N_t(a)}} \right] \tag{1.1}$$

where t denotes the total number of actions taken, $Q_t(a)$ is the action-value estimation, $N_t(a)$ is the number of times action a has been selected (prior to time t) and c is a constant controlling the degree of exploration. UCB1 follows the principle of *"optimism in the face of uncertainty"*, prioritizing actions with uncertain value-estimations, even if those are not the greatest.

In summary, Adaptive Sampling methodologies appeared to tackle the sampling problem faced by MD simulations at a time where high throughput simulations were possible. Since its origins in multiple parallel simulations, sampling techniques have greatly evolved, moving from heuristic methods to novel RL-based approaches, which represent the most viable solution towards optimal policies.

1.2.4. HTMD

In contrast to the theoretical complexity of all the processes and methodologies previously exposed, their practice is outstandingly simple from a technical point of view thanks to the accessibility provided by several software packages, such as HTMD (90). HTMD provide capabilities to build, run and analyze molecules simulations, including MSM (derived from PyEMMA (75)) and visualization (VMD (91)), allowing to write complex building or analysis pipelines that integrate with the rest of data-science related tools available in Python.

These features dramatically contribute to smooth the way of newcomers as well as for ensuring standardization of protocols, result reproducibil-

ity and, in general, good practices.

1.3. Systems Overview

Amongst the many topics related to IDPs and briefly discussed along *Section 1.1*, the thesis will focus on three of them: the behavior of IDPs in isolation, binding mechanism of IDPs to folded proteins, and druggability of IDPs. Each study revolves around a particular system: the N-terminal region of p53, the interaction between KIX and c-Myb, and the modulation of $p27^{kip1}$ with the small molecule SJ403.

1.3.1. p53

p53 is a protein of great biological relevance, and its fundamental role as a signaling hub in cell replication has granted it the nickname of *"genome keeper"*. Hence, its malfunction is a shared feature amongst many cancer types.

From a structural point of view, p53 includes a folded domain responsible for its tetramerization and located in the center of the protein. Disordered regions flank it at both ends. These disordered regions can interact with many other proteins, as shown in Figure 1.2. In particular, its disordered N-terminal region is arguably the most widely investigated IDP. Several studies have focused in both describing the structure of the complex formed with Mdm2 (Figure 1.8) and exploring their binding mechanism, either experimentally (NMR studies of the complex and the behavior in isolation of the disordered region, (92)) or computationally (generally using it as a benchmark for testing novel MD related technologies, (93; 94; 95; 96)), and even for the development of drugs (29). Finally, similar studies based on NMR and SAXS focusing on the N-terminal behavior have shown a similar helicity profile between the isolated state and the bound conformation (97; 92).

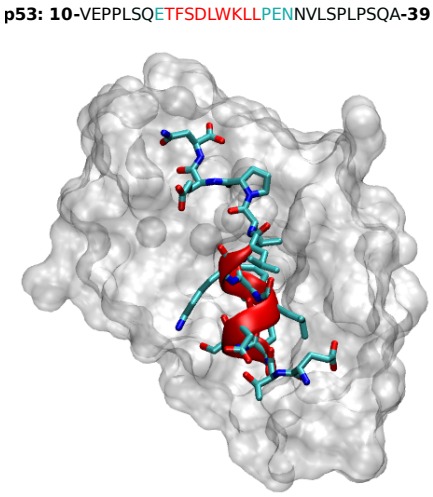

Figure 1.8: **p53 bound to Mdm2**. MDM2 protein is shown as a white surface. p53 is depicted as cyan sticks and the helical region between residues 18 and 26 as red cartoon, (PDB code: *1YCR* (12)). Above, the sequence of p53 used for the simulation is displayed: in red the helical section, in cyan the rest of the peptide found in the PDB structure, and in black the extended sequence.

1.3.2. KIX — c-Myb

Similarly to p53, the system constituted by KIX — c-Myb was also one of the pioneer structures including both a folded and a disordered protein (15). Since then, it has constituted a widely studied system in the context of protein-IDP interaction, especially in the determination of their binding mechanisms (98; 99; 100).

The folded component, the KIX domain of the CBP protein, consists of 87 amino acids arranged in a three α-helical bundle that, in the same line as p53, participates as a signaling hub in transcriptional regulation.

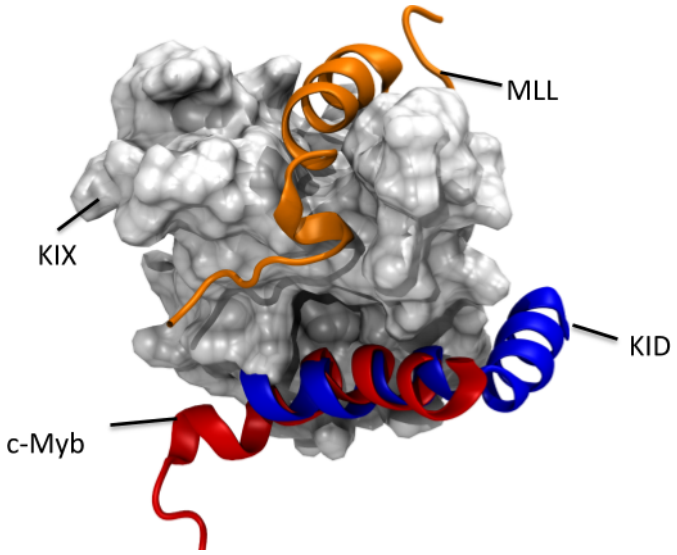

Figure 1.9: **Overview of KIX bound to several IDPs** including MLL in orange, c-Myb in red, and pKID in blue. Image is as result of the superimposition of PDB structures *2AGH* (16) and *1KDX* (13).

This hub-like behavior is due to the capabilities of KIX to accommodate a high number of ligands, in various conformations and interfaces, even several of them at the same time (Figure 1.9). With such a complex network of interaction, KIX is a fantastic system to study the main features of protein-protein interaction, such as allosterism or discrimination of binding mechanisms. For example, extensive work has been done towards describing the binding mechanism underlying the pKID—KIX interaction (43; 101; 102). PTMs modulate binding of KID, in particular, the phosphorylation in position Ser133 increases the binding affinity (102). The binding mechanism for this pair of proteins follows the induced-fit scheme where initial encounter complexes, mediated mostly by hydrophobic contacts, become increasingly rigid upon the formation of additional contacts

(43). Additionally, allosteric effects have also been shown in the system. In particular, the binding of mixed-lineage leukemia (MLL) transcription factor to a secondary interface on KIX (16) has a positive cooperative effect on the binding to other ligands to its primary binding site.

c-Myb is a proto-oncogene that works as a transcriptional activator playing an essential role in the control of proliferation and differentiation of hematopoietic progenitor cells. It includes three DNA binding domains in its N-terminal region, followed by a transactivation domain (TAD), and a regulatory C-terminal region. Upon binding to KIX, c-Myb's TAD folds in an extended α-helical conformation laying in a hydrophobic groove defined by the interface of two α-helices of KIX (Figure 1.9, (15)). Additionally, despite binding to the same interface, there are not apparent sequence similarities between pKID and c-Myb, highlighting underlying differences in how binding takes place.

Kinetics analysis of the KIX — c-Myb interaction revealed fast association kinetics (99) with a $k_{on} = (3.48 \pm 0.43) \cdot 10^7 M^{-1} s^{-1}$. Finally, there is some controversy in the literature on whether binding between this pair of proteins occurs via induced-fit, as supported by some experimental reports (103; 98; 104), or if, on the other hand, c-Myb's helicity influence the process towards conformational selection, as supported by the mixed-mechanism proposed by Arai et al. (100).

1.3.3. p27 — SJ403

Rational drug design is one of the greatest challenges faced in structural biology in general, and in the IDP field in particular. As previously stated, up to 30% of eukaryotic proteomes are disordered, and they play critical roles in disease-related processes. Hence, adapting traditional pipelines for folded proteins to IDPs should be a desirable objective, provided the amount of potential disordered drug targets. However, given the inherited difficulty of dealing with flexible proteins, not many examples of small molecule — IDP interaction has been reported in the literature

(35; 36; 31; 32). In those studies, several experimental approaches, including SAXS or circular dichroism (35), have been assessed to screen protein-ligand interactions. Finally, NMR has also been successfully employed with this purpose, as shown by the discovery of several molecules able to modulate the disordered $p27^{kip1}$ (33). On the other hand, computational studies on the topic have been used to perform virtual screenings based on ensemble conformations and simple MD studies that lead to the hypothesis of *"ligand clouds around protein clouds"*, which describes IDP-ligand binding through the formation of fuzzy complexes (105).

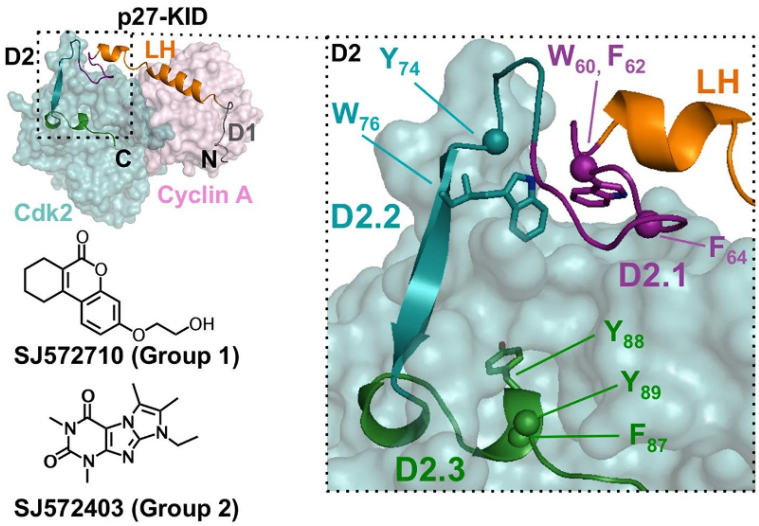

Figure 1.10: **p27 structure and ligands**. **a) p27 bound to the Cdk2—Cyclin A complex** (PDB ID *1JSU* (25)). **b) and c) Chemical structures** of two small molecules, SJ572710 (SJ710) (b) and SJ572403 (SJ403) (c), that bound to p27-KID. Modified from (33).

p27 is a cell cycle regulator, whose interaction with the cyclin A — Cdk2 complex induces cell cycle arrest. Several NMR-based studies per-

formed with the IDP and its ligands reported both the molecular determinants (33) for the interaction and a glimpse of the underlying mechanism of binding with one of the ligands (SJ403, (34)). In this binding scheme, there are some clusters of hydrophobic residues of p27 in contact with each other in the isolated state, but the addition of the ligand displaces one of them and occupies its position.

Some of the features exhibited by the system composed by the disordered $p27^{kip1}$ and the molecule SJ403 make it feasible to be studied by high-throughput MD. The main two are the relatively short length of the protein, with $\sim$50 residues, and the simplicity of SJ403, with only one dihedral angle and a planar chemical structure. Of all the approaches found within the MD toolkit, free ligand biding studies (80; 49; 106) represent a widely used and useful one that could be easily tested in this scenario. In these experiments, extensive MD simulations are run starting from the unbound state, with protein and ligand placed apart, and are let to freely diffuse, sampling various binding poses. Such approaches have never been tested in the context of IDPs.

Chapter 2

OBJECTIVES

The main objective of the present thesis is to explore new prospects in the application of high-throughput unbiased MD simulations in the study of several key aspects of the biology of IDPs. In particular, the research focus on IDP behavior in isolation through the formation of partially ordered states, how folding couples with binding in IDP—folded partner interaction, and the potential of the druggability of IDPs and the impact on the energy surface of small molecules.

2.1. Behavior of IDPs in isolation

The first objective is exemplified in the study of the N-terminal domain of p53. This short peptide folds in an α-helix upon binding to Mdm2. It was submitted to extensive MD simulations and analyzed using MSMs to study its behavior in isolation. The main goal was to create a detailed map of the partially ordered states populating the phase space as well as to understand the relationship amongst them.

2.2. IDP — protein interaction

The main focus on this problem is to reconstruct the coupled folding and binding mechanism of the protoncogene c-Myb upon interacting with the KIX domain of the CREB protein. Experimental information regarding kinetic parameters and secondary binding modes were used to validate the resulting MSMs. Over the last years, the task proved to be especially challenging and many unsuccessful attempts with a variety of methodologies we carried out. Finally, the development of the latest MD sampling algorithms helped to overcome part of the problems.

2.3. IDP — small molecule interaction

The purpose of the last project is to perform the first free ligand MD binding study between an IDP ($p27^{kip1}$) and a small molecule (SJ403). The main objective addressed with these experiments was to determine the nature of the interaction. Primarily, whether the IDP remained or not disordered upon binding. Secondly, whether the addition of small molecules either expanded or reduced the conformational space explored by the IDP.

Chapter 3

RESULTS

3.1. Characterization of partially ordered states in the p53 intrinsically disordered N-terminal domain using millisecond molecular dynamics simulations

Pablo Herrera-Nieto, Adrià Pérez, Gianni De Fabritiis. Submitted to *Scientific Reports*

Summary

The present study tackles the emergence of partially ordered states within the conformational landscape of an exemplar IDP, p53. IDPs define relatively flat energy surfaces compared to the funneled profiles observed in folded proteins. However, they are not entirely composed by random coiled conformations. Experimental research of these events is limited, as the information they provide represents population averages, and are not capable of isolating such potential ordered states. Despite the existence of

many computational tools devoted to the characterization of IDP ensembles, high throughput MD and MSMs have not been applied yet in this task. Here, we make use of these powerful tools in the context of isolated IDPs and show various partially ordered states within the N-terminal region of p53.

Characterization of partially ordered states in the intrinsically disordered N-terminal domain of p53 using millisecond molecular dynamics simulations

Pablo Herrera-Nieto[a], Adrià Pérez[a], and Gianni De Fabritiis[a,b,*]

[a]Computational Science Laboratory,Barcelona biomedical research park (PRBB), Universitat Pompeu Fabra, C Dr Aiguader 88, Barcelona 08003, Spain
[b]Institució Catalana de Recerca i Estudis Avançats (ICREA), Passeig Lluis Companys 23, 08010 Barcelona, Spain
[*]gianni.defabritiis@upf.edu

ABSTRACT

The exploration of intrinsically disordered proteins in isolation is a crucial step to understand their complex dynamical behavior. In particular, the emergence of partially ordered states has not been explored in depth. The experimental characterization of such partially ordered states remains elusive due to their transient nature. Molecular dynamics mitigates this limitation thanks to its capability to explore biologically relevant timescales while retaining atomistic resolution. Here, millisecond unbiased molecular dynamics simulations were performed in the exemplar N-terminal region of p53. In combination with state-of-the-art Markov state models, simulations revealed the existence of several partially ordered states accounting for ∼40% of the equilibrium population. Some of the most relevant states feature helical conformations similar to the bound structure of p53 to Mdm2, as well as novel β-sheet elements. This highlights the potential complexity underlying the energy surface of intrinsically disordered proteins.

Introduction

Over the last decades, the understanding of protein function was summarized by the *sequence-structure-function* triumvirate: protein sequences encode folds able to perform specific tasks. Intrinsically disordered proteins (IDPs) defy this principle by mediating their biological functions despite lacking a stable three-dimensional structure[1–3]. Such behavior configures a relatively flat energy surface where many isoenergetic conformations coexist[4]. This surface can be modified to a certain extent, as revealed by the shift towards certain subpopulations observed in the formation of protein-IDP[5–8] or molecule-IDP complexes[9,10]. Similarly, kinetic parameters governing the conversions amongst subpopulations can also be modified by post-translational modifications[11,12]. Thus, the energy surface of IDPs if far from being constituted exclusively by random coiled conformations, and pieces of evidence support the existence of partially ordered states[11]. The characterization of such partially ordered states is crucial to understand IDPs' function, their mechanisms of action, and their potential modulation.

The structural description of IDPs summarizes their structural heterogeneity as a collection or ensemble of conformations. They can be resolved experimentally by using nuclear magnetic resonance (NMR) or small-angle X-ray scattering data. The main limitation of IDP ensembles resolved in that way is that they focus on global averages rather than diving in particular atomic coordinates[13]. There are also many computational approximations to address this task[14]. They generally involve an initial step of conformer generation followed by a refinement step that minimizes differences between the generated library and experimental data. However, many computationally resolved ensembles can match the same experimental observations.

Molecular dynamics simulations (MD) have been extensively used over the years to navigate complex energy surfaces in other biological problems, i.e. in folding[15], protein-protein binding[16,17] and, modulation of IDPs by post-translational modifications[11] or by interacting with their folded partners[8]. In the context of IDP ensembles, MD simulations have been primarily applied as a tool for conformational generation. Nevertheless, the main goal of MD in this area would be to define reliable ensembles without the need for biasing or reweighting procedures. In this line, recent studies have employed enhanced sampling methods such as Hamiltonian replica exchange MD to define IDP ensembles matching the available experimental information[18]. In terms of aggregated time, the study run for ∼10 μs, while others have performed more extensive simulations, ∼200 μs but in a single trajectory[19].

Current technologies allow MD simulations to reach aggregated times in the order of milliseconds[20], thus making this tool a valuable one for the exploration of biological systems at increasingly longer time scales[16]. The potential offered by high throughput MD simulations coupled with Markov State Models (MSMs)[21] analysis for the exploration of conformational land-

scapes of IDP remains mostly untested. The main advantage offered by this tandem is the possibility to address subpopulations directly within ensembles and to study the kinetics controlling them, rather than working with population averages. By focusing on the most relevant subpopulations and their kinetic properties, it is possible to gain insight into the emergence of partially ordered states in atomistic detail.

Here we make use of extensive, unbiased full-atom MD simulations and state-of-the-art MSMs to explore the structural variability of the N-terminal region of p53 in isolation. p53 is a widely studied protein, given its relation to oncogenic processes. This protein includes disordered sections at both N and C terminals that are known to interact with various partners[22]. In particular, the interaction between the N-terminal region and the Mdm2 has gathered substantial interest over the years, as it pioneered the coupled binding-and-folding phenomena[5], The p53-Mdm2 complex has served as a template for the development of peptidomimetics drugs[23] and as the preferred benchmark for several MD studies aiming to reconstruct the binding process[16,24–26]. Finally, NMR studies of the N-terminal region in isolation revealed a helicity profile similar to the one observed in the p53-Mdm2 complex, implying that bound conformations might also be sampled prior to binding[27].

The main results show the existence of many kinetically relevant states, accounting for ∼40% of the equilibrium population, including high levels of secondary structural elements. In particular, simulations show the presence of an α-helix enriched states similar to the folded pose found in complex with Mdm2, as well as, a tangled interplay between β-strands formation leading to novel β-sheet enriched structures. Altogether, this illustrates the complexity of partially ordered states within the conformational space of an exemplar IDP, such as the N-terminal region of p53.

Results and discussion

Identification of secondary structure enriched states

The simulation time of the MD run totaled ∼1.4 ms. Initially, the secondary structure of the aggregated MD was analyzed. Data showed the coexistence of both α-helix and β-strand, each one peaking at ~ 20 % in the central region of the protein (Fig. 1.b). The helicity profile follows a bell-shaped distribution, while β-strand is more sparsely scattered in three groups in the proximity of residues S15, K24, and V31.

MD data was used to create an MSM based on $backbone_{C\alpha} + sidechain_{O,N}$ self distance matrix that splits the space into 12 different sets of kinetically related conformations referred to as macrostates (labeled as *M1-12*). MSM subpopulations successfully separate metastable sets of conformations enriched in each secondary structure type (Fig. 1.c,d), implying that these structural elements do appear in a concerted way, rather than being the average of residue independent structural propensities.

The helicity profile displayed by the helix-enriched state matches the bound conformation of p53 when interacting with Mdm2 (Fig. 1.a). It spams from residue T18 to L26, and maximum levels of helicity being found in W23 —an essential amino acid for that interaction. Similar profiles arise from NMR studies[27]. Besides this state, many others also display various degrees of helicity (Supplementary Fig. 1). The tendency of IDPs to acquire secondary structure profiles resembling their folded conformation has also been observed in other IDPs, and it has also been related to the binding mechanism to their partner[28] and their signaling properties[29].

For β-strand, segregation of secondary structural elements into their own states becomes especially evident in *M4*, where three β-strands —namely $\beta 1$, $\beta 2$, and $\beta 3$ from N to C terminal— are organized in an anti-parallel double-sheet (Fig. 1.c), defining the partially ordered state with the most significant level of structure. Besides the secondary structure enriched macrostates aforementioned, many other states also exhibit different profiles and combinations of β-strand or α-helix (Supplementary Figs. 1 and 2). This includes a number of states displaying different β-sheet arrangements, featuring only one strand, either $\beta 1$-$\beta 2$ (*M7*) or $\beta 2$-$\beta 3$ (*M3*). In those states, the potential location of the third strand is occupied by helical conformations (Supplementary Fig. 1). Altogether, this highlights the variety of possible configurations found across the conformational landscape of p53.

Kinetic characterization of the conformational landscape of p53

Population wise, partially ordered states account for a significant proportion of the equilibrium population (∼40%, Fig. 2.a). The triple-stranded macrostates, the most folded ones, have low populations (<1%), in contrast to double-stranded states like *M6,11*, which reach ∼20% at equilibrium. However, the most populated state —*M12*; with ∼60% of the population— is structurally heterogeneous, and lacks any secondary structural element or long-range contacts. Hence, partially ordered states are not energetically favored compared to the most extended configurations, and their free energies range from 0.5 to -2.5 $kcalM^{-1}$ (from less to more equilibrium probability, Fig. 2.d). This can be visualized in more detail in the entire energy surface of p53 (Supplementary Fig. 3.a). There are two well-defined minima separated by a small energy barrier. One of them is covered with the extended and the helical states (*M8,10,12*), and the other by $\beta 2$-$\beta 3$ conformations (*M11*). High energetic areas are occupied by the most structured states, like *M4*. Such profile, with many energetically similar states, fits the description generally provided to explain the behavior of IDPs in isolation.

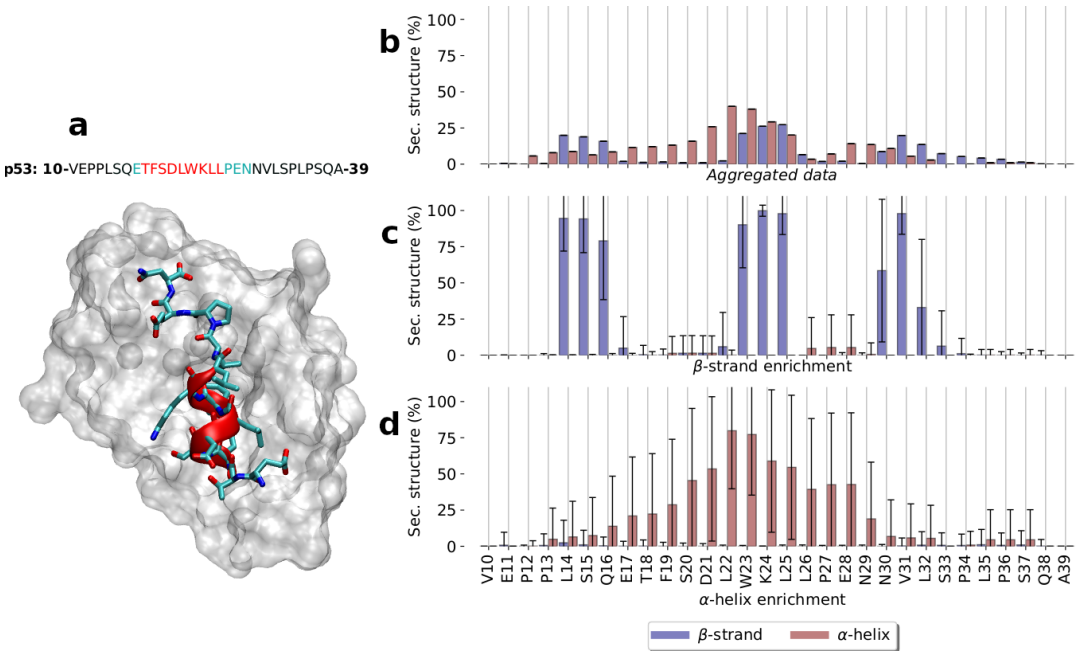

Figure 1. p53 secondary structure propensities a) p53-MDM2 complex: MDM2 protein is shown as a white surface. p53 is depicted as cyan sticks and the helical region between residues 18 and 26 as red cartoon, (PDB code: *1YCR*). Above, the sequence of p53 used for the simulation is displayed: in red the helical section, in cyan the rest of the peptide found in the PDB structure, and in black the extended sequence. **Secondary structure profiles** derived from MD data: β-strand and α-helix profiles for the **b)** aggregated data and for those macrostates of the MSM enriched in either **c)** β-sheet and **d)** α-helix.

The distinctive population profiles previously observed have an impact on the kinetic behavior of the macrostates. There is an approximately two orders of magnitude difference between maximum and minimum $k_{on}/mfpt_{on}$ estimations, which separate *faster/more populated (M7-11)* macrostates from *slower/less populated (M1-6)* macrostates (Fig. 2.b,e). *Faster* macrostates comprise the helical conformation and several double-stranded states. They feature k_{on} values between $5-10\cdot10^6 M^{-1}s^{-1}$, which, for reference, are as fast as binding between p53 and Mdm2 ($9.2\cdot10^6 M^{-1}s^{-1}$[30]). *Slower* states, on the other hand, have k_{on} values around $1-5\cdot10^5 M^{-1}s^{-1}$ and include triple stranded states and other low populated states.

We employed transition path theory[31,32] to study the most relevant pathways and fluxes for macrostate interconversion. In particular, the focus was to elucidate the folding process leading from the less structured state to the triple-stranded conformation. There are three main paths (Fig. 3 — central panel) involved in this process. The least transited one accounts for ∼15% of the total flow, and directly reaches the folded conformation from the extended one. On the other hand, the most transited paths involve the participation of double-stranded intermediates, with the β1-β2 structure taking ∼45% and the β2-β3 conformation being responsible for the remaining ∼30% of the flux. Additionally, other β-enriched states, such as the extended β2-β3 sheet found in *M11*, are disconnected from this network and can be directly reached from *M12* without the need of intermediates. It is interesting to point out that some conformations (such as *M6,11*), despite their structural similarity, show a ∼30 slowdown that explains the differences in stability aforementioned.

In summary, partially ordered states populating the conformational landscape are structural and kinetically diverse. States coexist at different timescales, even if they are structurally similar, such as the case of *M6,11*. These two states feature a short and an extended β2-β3 sheet but have k_{on} values of $5\cdot10^5 M^{-1}s^{-1}$ and $1\cdot10^7 M^{-1}s^{-1}$ respectively.

Comparison with NMR data

In order to ensure and validate MD observations, simulation data was compared against experimentally determined backbone chemical shifts (CS) for the N-terminal region of p53[33]. CS allows inferring by-residue secondary structure tendencies on folded and disordered proteins. Calculations were performed using two softwares, SPARTA+[34] and SHIFTX2[35], on a set of 2 000 structures selected at random accordingly to the macrostate equilibrium probabilities. Calculated CS with both programs

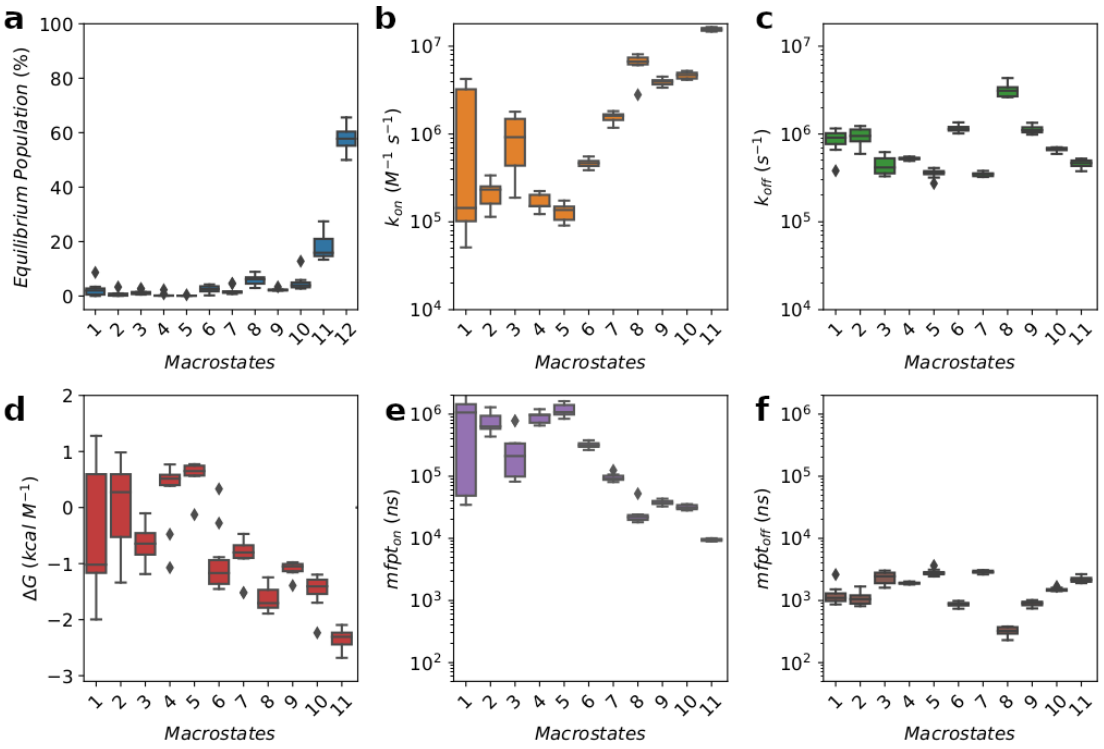

Figure 2. **Estimation** of **a)** *equilibrium population*, **b)** k_{on}, **c)** k_{off}, **d)** ΔG, **e)** $mfpt_{on}$, **f)** $mfpt_{off}$, for each macrostate after 10 rounds of bootstrap. Kinetics parameters for *Macro12* (structurally heterogeneous) are not shown as it was used as the source state for calculations.

yielded similar results. Overall, there is a high correlation between experimental and MD calculated CS values for Cα, Cβ, and N with R^2 values of 0.98, 0.99, and 0.88 (Supplementary Fig. 4). Differences between experimental and calculated CS remains within the intrinsic estimations error of each software (~ 1 $p.p.m$ for SPARTA+ and 0.4 and 0.5 for Cα and Cβ, respectively, for SHIFTX2) thus, indicating that structural rearrangements observed in MD data are in line with those determined by NMR experiments.

Conclusion

The characterization of the p53 conformational landscape using unbiased MD simulations revealed a high number of transient partially ordered states accounting for ∼40% of the equilibrium populations. Partial order arises from the formation of both α-helix and β-strand structural elements. The helical state resembles the structure acquired by p53 upon interaction with Mdm2. The MSM also revealed the presence of several β-enriched states, not described before, that established long-range contacts through the arrangement of either one or two β-sheets. These processes are kinetically different, and some of the faster states are quickly accessible from the random-coiled macrostate and highly populated at equilibrium. Thus, it would be possible for some of them to play biologically relevant roles and could even provide novel strategies for the modulation of IDPs.

The current study provides a structural and kinetically detailed description of the conformational landscape of an IDP using MD simulations in combination with MSMs. Given the high number of short linear motifs within the human proteome, a similar pipeline could, in principle, be more extensively applied in order to investigate whether other IDPs may also share such complex behaviors. However, reaching millisecond simulation time may not prove scalable in more extensive studies with multiple targets. Novel adaptive sampling techniques[36,37], which perform a more intelligent exploration of surfaces, might mitigate this problem by reducing the computational time needed to achieve similar results. One final consideration of the challenge of the task is the computational effort to simulate longer peptides. Here we use a relatively short 30 amino acids

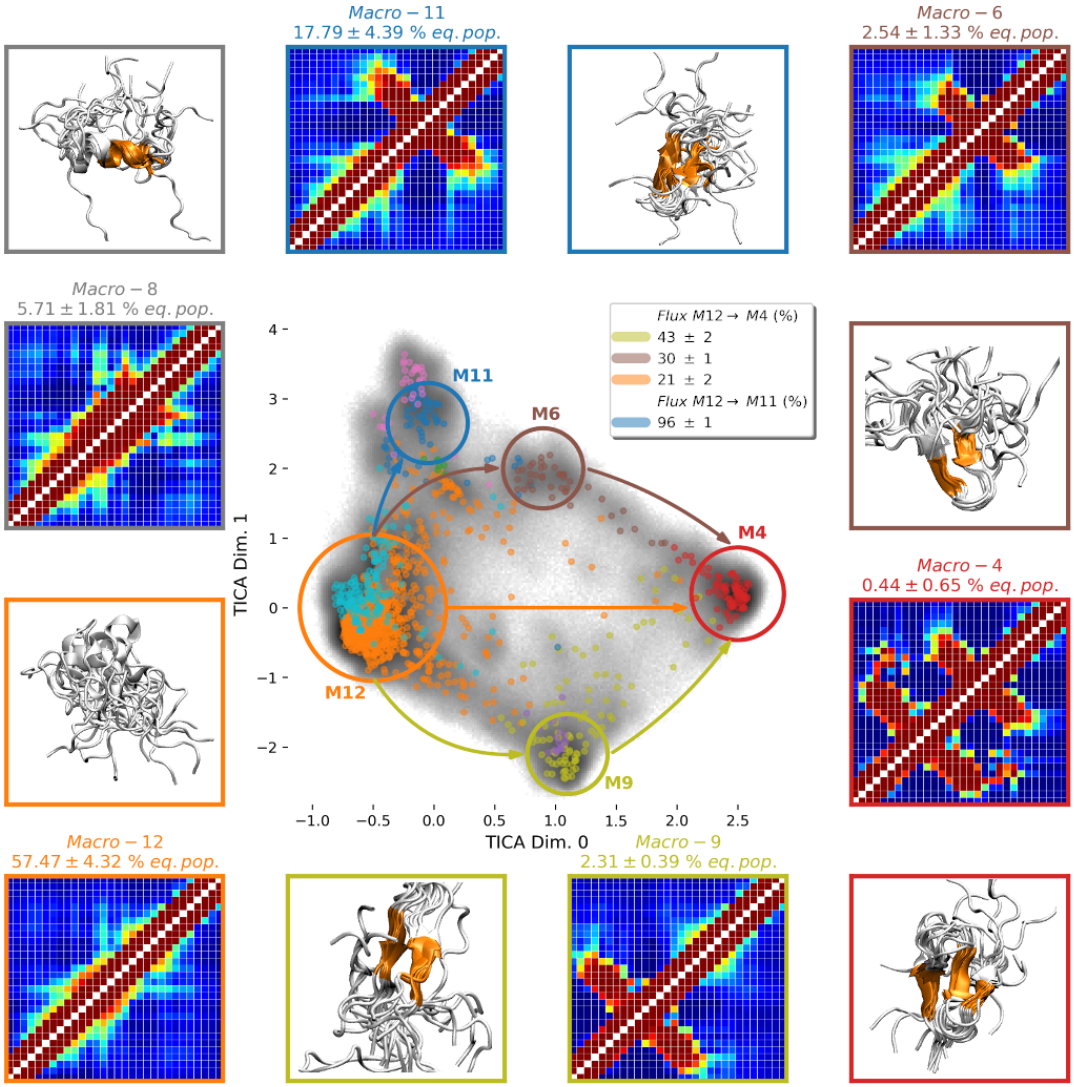

Figure 3. **p53 conformational landscape. Central panel** illustrates the first two dimensions of the TICA space. In grey, a 2D histogram 200x200 bins represents the frame count of the aggregated MD data. MSM microstates are distributed accordingly to their centers and colored with respect to their corresponding macrostate. Arrows represent the main pathways leading from the most extended macrostate (*Macro-12*) to the most β-sheet enriched one (*Macro-4*). Correspondence between macrostate location in the central panel and side panels is color mapped. **Side panels** describe macrostates in terms of residue-residue contacts maps. Protein visualization is performed by superimposing 20 structures using residues highlighted in orange for structural alignment.

section of p53, while completely disordered domains may spam hundreds of residues.

Methods

Molecular dynamics simulation set up

In order to perform the exploration of the conformational space of p53, extensive parallel simulations were run. The selected region of p53 spammed from residue 10 to 39.

A set of 110 structures was used as initial conformations for the MD run (Supplementary Fig. 5.c). All systems were built with VMD[38], solvated with TIP3P[39] (each system included ∼8 200 water molecules, resulting in a final protein concentration of ∼6.8 mM), with a final NaCl concentration of 0.05 M. A Langevin integrator with a damping constant of 0.1 ps^{-1} was used. The integration step was set to 4 fs, with heavy hydrogen atoms scaled up to four times their natural mass. Electrostatic were computed using PME with a cutoff distance of 9 Å and grid spacing of 1 Å. Equilibration was performed at 300 K, firstly undergoing 250 steps of energy minimization followed by 0.1 ns simulations in an NVE ensemble (pressure was kept at 1 atm by using the berendesen barostat) and 2 ns in an NPT ensemble.

After equilibration, no proline *cis* isomers were detected. Production run was executed in the distributed computing platform GPUGrid.net[40] using the ACEMD engine[41] and the CHARMM22* forcefield[42].

Markov State Model analysis

Production runs generated a total of 1.337 trajectories (each equilibrated system was used at least 10 times) of 1 μs each, similarly to[16]. Thus, the production runs accounted for an aggregated simulation time of ∼1.4 ms in order to maximize the exploration of the conformational space. All MD data analyses presented here were performed using HTMD[43].

MD data analysis was performed by featurizing atomic coordinates as the self-distance matrix between C_α and side chains nitrogen and oxygen atoms (54 out of 237 a total of heavy atoms), thus creating a sparse representation of the protein. Next, time independent component analysis method (TICA,[44]) furtherly reduced data dimensionality, projecting the first 5 TICA dimensions at a lag time of 20 frames. Then, TICA projected data was clustered into 1.500 states using the MiniBatchKMeans algorithm[45]. Microstates were fused at a lag time of 120 ns, following the implied time scales plot (Supplementary Fig. 5.a) into 12 macrostates (using the PCCA+ algorithm[46] and based on the discretization of the TICA space shown in Supplementary Fig. 5.e). Finally, transition path theory[31,32] was used to calculate fluxes between states.

For every measure, the error was estimated by creating 10 independent MSMs using a random set containing 80% of the simulated trajectories.

Chemical shift calculations

Calculations of MD derived chemical shifts were performed using 2.000 frames distributed amongst macrostates based on their equilibrium probability. The biological magnetic resonance data bank entry *17760*[33] was used to obtain experimental chemical shift data for the N-terminal region of p53. Two different softwares were used: SPARTA+[34] and SHIFTX2[35].

References

1. Wright, P. E. & Dyson, H. J. Intrinsically unstructured proteins: re-assessing the protein structure-function paradigm. *J. molecular biology* **293**, 321–331 (1999).
2. Dyson, H. J. & Wright, P. E. Intrinsically unstructured proteins and their functions. *Nat. reviews Mol. cell biology* **6**, 197 (2005).
3. Van Der Lee, R. *et al.* Classification of intrinsically disordered regions and proteins. *Chem. reviews* **114**, 6589–6631 (2014).
4. Burger, V., Gurry, T. & Stultz, C. Intrinsically disordered proteins: where computation meets experiment. *Polymers* **6**, 2684–2719 (2014).
5. Kussie, P. H. *et al.* Structure of the MDM2 oncoprotein bound to the p53 tumor suppressor transactivation domain. *Science* **274**, 948–953 (1996).
6. Russo, A. A., Jeffrey, P. D., Patten, A. K., Massagué, J. & Pavletich, N. P. Crystal structure of the p27Kip1 cyclin-dependent-kinase inibitor bound to the cyclin A–Cdk2 complex. *Nature* **382**, 325 (1996).
7. Zor, T., De Guzman, R. N., Dyson, H. J. & Wright, P. E. Solution structure of the KIX domain of CBP bound to the transactivation domain of c-Myb. *J. molecular biology* **337**, 521–534 (2004).
8. Chong, S.-H., Im, H. & Ham, S. Explicit characterization of the free energy landscape of pkid–kix coupled folding and binding. *ACS Cent. Sci.* **5**, 1342–1351 (2019).
9. Iconaru, L. I. *et al.* Discovery of small molecules that inhibit the disordered protein, p27 kip1. *Sci. reports* **5**, 15686 (2015).

10. Ban, D., Iconaru, L. I., Ramanathan, A., Zuo, J. & Kriwacki, R. W. A small molecule causes a population shift in the conformational landscape of an intrinsically disordered protein. *J. Am. Chem. Soc.* **139**, 13692–13700 (2017).

11. Stanley, N., Esteban-Martín, S. & De Fabritiis, G. Kinetic modulation of a disordered protein domain by phosphorylation. *Nat. communications* **5**, 5272 (2014).

12. Bah, A. & Forman-Kay, J. D. Modulation of intrinsically disordered protein function by post-translational modifications. *J. Biol. Chem.* **291**, 6696–6705 (2016).

13. Camilloni, C., De Simone, A., Vranken, W. F. & Vendruscolo, M. Determination of secondary structure populations in disordered states of proteins using nuclear magnetic resonance chemical shifts. *Biochemistry* **51**, 2224–2231 (2012).

14. Fisher, C. K. & Stultz, C. M. Constructing ensembles for intrinsically disordered proteins. *Curr. opinion structural biology* **21**, 426–431 (2011).

15. Lindorff-Larsen, K., Piana, S., Dror, R. O. & Shaw, D. E. How fast-folding proteins fold. *Science* **334**, 517–520 (2011).

16. Paul, F. *et al.* Protein-peptide association kinetics beyond the seconds timescale from atomistic simulations. *Nat. communications* **8**, 1095 (2017).

17. Plattner, N., Doerr, S., De Fabritiis, G. & Noé, F. Complete protein–protein association kinetics in atomic detail revealed by molecular dynamics simulations and markov modelling. *Nat. chemistry* **9**, 1005 (2017).

18. Shrestha, U. R. *et al.* Generation of the configurational ensemble of an intrinsically disordered protein from unbiased molecular dynamics simulation. *Proc. Natl. Acad. Sci.* **116**, 20446–20452 (2019).

19. Lindorff-Larsen, K., Trbovic, N., Maragakis, P., Piana, S. & Shaw, D. E. Structure and dynamics of an unfolded protein examined by molecular dynamics simulation. *J. Am. Chem. Soc.* **134**, 3787–3791 (2012).

20. Martinez-Rosell, G., Giorgino, T., Harvey, M. J. & de Fabritiis, G. Drug discovery and molecular dynamics: methods, applications and perspective beyond the second timescale. *Curr. topics medicinal chemistry* **17**, 2617–2625 (2017).

21. Prinz, J.-H. *et al.* Markov models of molecular kinetics: Generation and validation. *The J. chemical physics* **134**, 174105 (2011).

22. Uversky, V. N. & Dunker, A. K. Understanding protein non-folding. *Biochimica et Biophys. Acta (BBA)-Proteins Proteomics* **1804**, 1231–1264 (2010).

23. Vassilev, L. T. *et al.* In vivo activation of the p53 pathway by small-molecule antagonists of mdm2. *Science* **303**, 844–848 (2004).

24. Zwier, M. C. *et al.* Efficient atomistic simulation of pathways and calculation of rate constants for a protein–peptide binding process: application to the MDM2 protein and an intrinsically disordered p53 peptide. *The journal physical chemistry letters* **7**, 3440–3445 (2016).

25. Morrone, J. A., Perez, A., MacCallum, J. & Dill, K. A. Computed binding of peptides to proteins with meld-accelerated molecular dynamics. *J. chemical theory computation* **13**, 870–876 (2017).

26. Zhou, G., Pantelopulos, G. A., Mukherjee, S. & Voelz, V. A. Bridging microscopic and macroscopic mechanisms of p53-MDM2 binding with kinetic network models. *Biophys. journal* **113**, 785–793 (2017).

27. Wells, M. *et al.* Structure of tumor suppressor p53 and its intrinsically disordered n-terminal transactivation domain. *Proc. Natl. academy Sci.* **105**, 5762–5767 (2008).

28. Arai, M., Sugase, K., Dyson, H. J. & Wright, P. E. Conformational propensities of intrinsically disordered proteins influence the mechanism of binding and folding. *Proc. Natl. Acad. Sci.* **112**, 9614–9619 (2015).

29. Borcherds, W. *et al.* Disorder and residual helicity alter p53-mdm2 binding affinity and signaling in cells. *Nat. chemical biology* **10**, 1000–1002 (2014).

30. Schon, O., Friedler, A., Bycroft, M., Freund, S. M. & Fersht, A. R. Molecular mechanism of the interaction between mdm2 and p53. *J. molecular biology* **323**, 491–501 (2002).

31. Weinan, E. & Vanden-Eijnden, E. Towards a theory of transition paths. *J. statistical physics* **123**, 503 (2006).

32. Noé, F., Schütte, C., Vanden-Eijnden, E., Reich, L. & Weikl, T. R. Constructing the equilibrium ensemble of folding pathways from short off-equilibrium simulations. *Proc. Natl. Acad. Sci.* **106**, 19011–19016 (2009).

33. Wong, T. S. *et al.* Biophysical characterizations of human mitochondrial transcription factor a and its binding to tumor suppressor p53. *Nucleic acids research* **37**, 6765–6783 (2009).

34. Shen, Y. & Bax, A. Sparta+: a modest improvement in empirical nmr chemical shift prediction by means of an artificial neural network. *J. biomolecular NMR* **48**, 13–22 (2010).

35. Han, B., Liu, Y., Ginzinger, S. W. & Wishart, D. S. Shiftx2: significantly improved protein chemical shift prediction. *J. biomolecular NMR* **50**, 43 (2011).

36. Doerr, S. & De Fabritiis, G. On-the-fly learning and sampling of ligand binding by high-throughput molecular simulations. *J. chemical theory computation* **10**, 2064–2069 (2014).

37. Zimmerman, M. I. & Bowman, G. R. Fast conformational searches by balancing exploration/exploitation trade-offs. *J. chemical theory computation* **11**, 5747–5757 (2015).

38. Humphrey, W., Dalke, A. & Schulten, K. VMD: visual molecular dynamics. *J. molecular graphics* **14**, 33–38 (1996).

39. Jorgensen, W. L., Chandrasekhar, J., Madura, J. D., Impey, R. W. & Klein, M. L. Comparison of simple potential functions for simulating liquid water. *The J. chemical physics* **79**, 926–935 (1983).

40. Buch, I., Harvey, M. J., Giorgino, T., Anderson, D. P. & De Fabritiis, G. High-throughput all-atom molecular dynamics simulations using distributed computing. *J. chemical information modeling* **50**, 397–403 (2010).

41. Harvey, M. J., Giupponi, G. & Fabritiis, G. D. ACEMD: accelerating biomolecular dynamics in the microsecond time scale. *J. chemical theory computation* **5**, 1632–1639 (2009).

42. Piana, S., Lindorff-Larsen, K. & Shaw, D. E. How robust are protein folding simulations with respect to force field parameterization? *Biophys. journal* **100**, L47–L49 (2011).

43. Doerr, S., Harvey, M., Noé, F. & De Fabritiis, G. HTMD: high-throughput molecular dynamics for molecular discovery. *J. chemical theory computation* **12**, 1845–1852 (2016).

44. Pérez-Hernández, G., Paul, F., Giorgino, T., De Fabritiis, G. & Noé, F. Identification of slow molecular order parameters for Markov model construction. *The J. chemical physics* **139**, 07B604_1 (2013).

45. Pedregosa, F. *et al.* Scikit-learn: Machine learning in Python. *J. machine learning research* **12**, 2825–2830 (2011).

46. Röblitz, S. & Weber, M. Fuzzy spectral clustering by pcca+: application to markov state models and data classification. *Adv. Data Analysis Classif.* **7**, 147–179 (2013).

Acknowledgements

The authors thank volunteers at GPUGRID.net for contributing with computational resources and Acellera for funding. G.D.F. acknowledges support from MINECO (Unidad de Excelencia María de Maeztu MDM-2014-0370 and BIO2017-82628-P) and FEDER. This project received funding from the European Union's Horizon 2020 Research and Innovation Programme under Grant Agreement 823712 (CompBioMed2 Project).

Author contributions statement

G.D.F. conceived the experiments, P.H.N and A.P. analyzed the results. P.H.N wrote the manuscript. All authors reviewed the manuscript.

Additional information

Competing interests The author(s) declare no competing interests.

The corresponding author is responsible for submitting a competing interests statement on behalf of all authors of the paper. This statement must be included in the submitted article file.

Supplementary Information: Characterization of partially ordered states in the p53 intrinsically disordered N-terminal domain using millisecond molecular dynamics simulations

Pablo Herrera-Nieto[a], Adrià Pérez[a], and Gianni De Fabritiis[a,b]

[a]Computational Science Laboratory,Barcelona biomedical research park (PRBB), Universitat Pompeu Fabra, C Dr Aiguader 88, Barcelona 08003, Spain
[b]Institució Catalana de Recerca i Estudis Avançats (ICREA), Passeig Lluis Companys 23, 08010 Barcelona, Spain

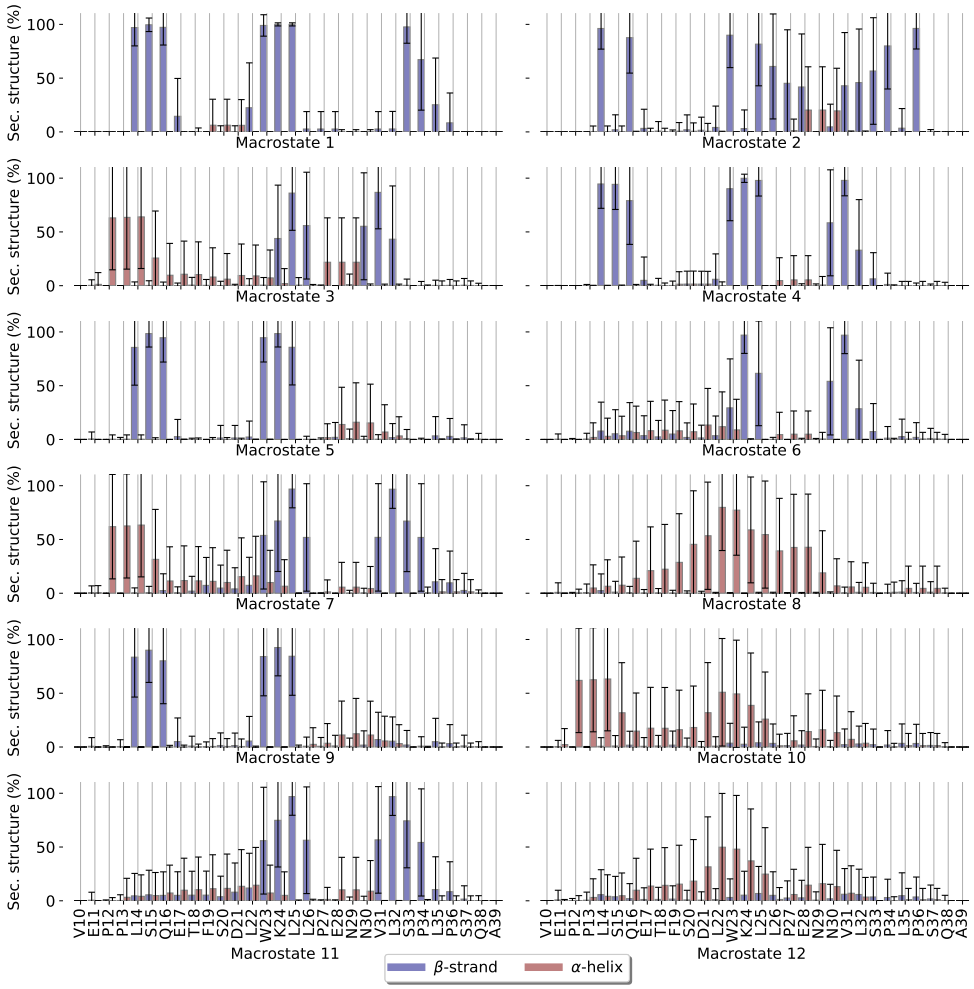

Figure 1: **Secondary structure profile** for all MSM macrostates.

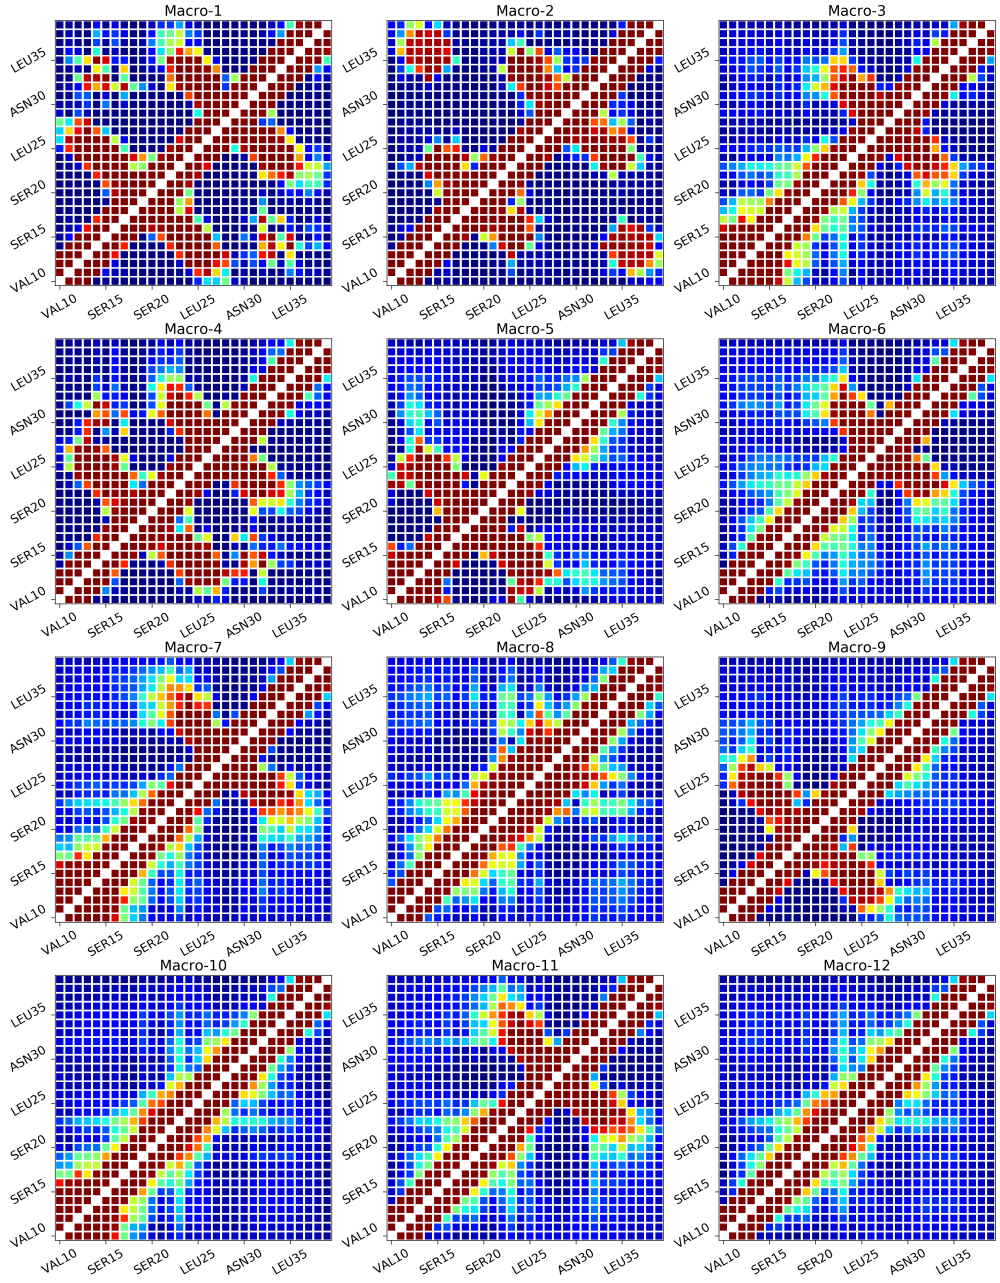

Figure 2: **Residue-residue contacts** for all MSM macrostates.

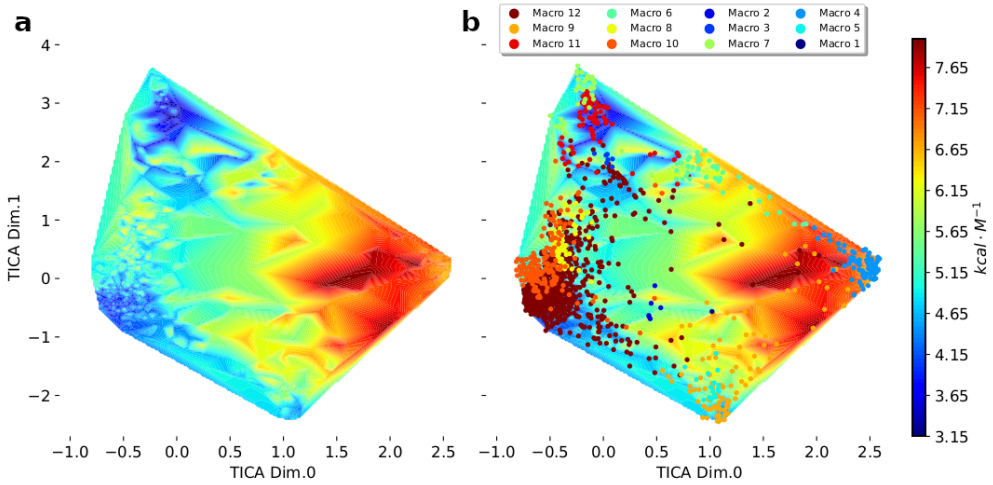

Figure 3: **Free Energy surface** of p53 in **a)** p53, and microstates centers in **b)**. Microstates are color mapped according to their macrostates.

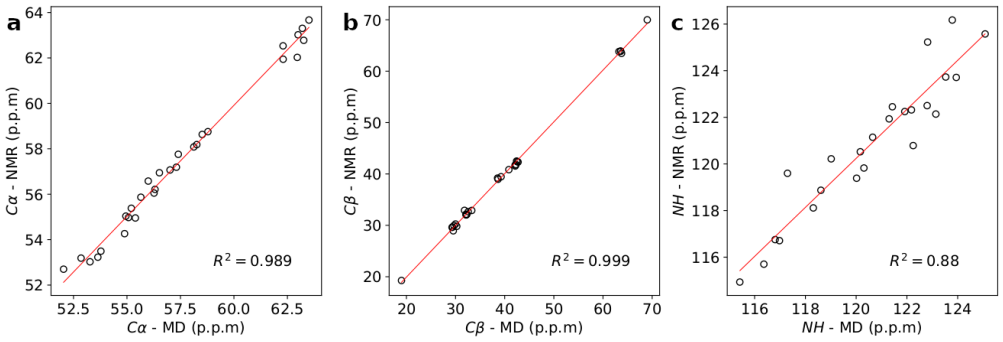

Figure 4: **Comparison against NMR data.** Chemical shift (CS) difference between MD-derived calculations for **a)** $C\alpha$, **b)** $C\beta$, and **c)** NH atoms performed with SPARTA+ and the experimentally measured for the N-terminal of p53 (Biological Magnetic Resonance Data Bank entry number 17760).

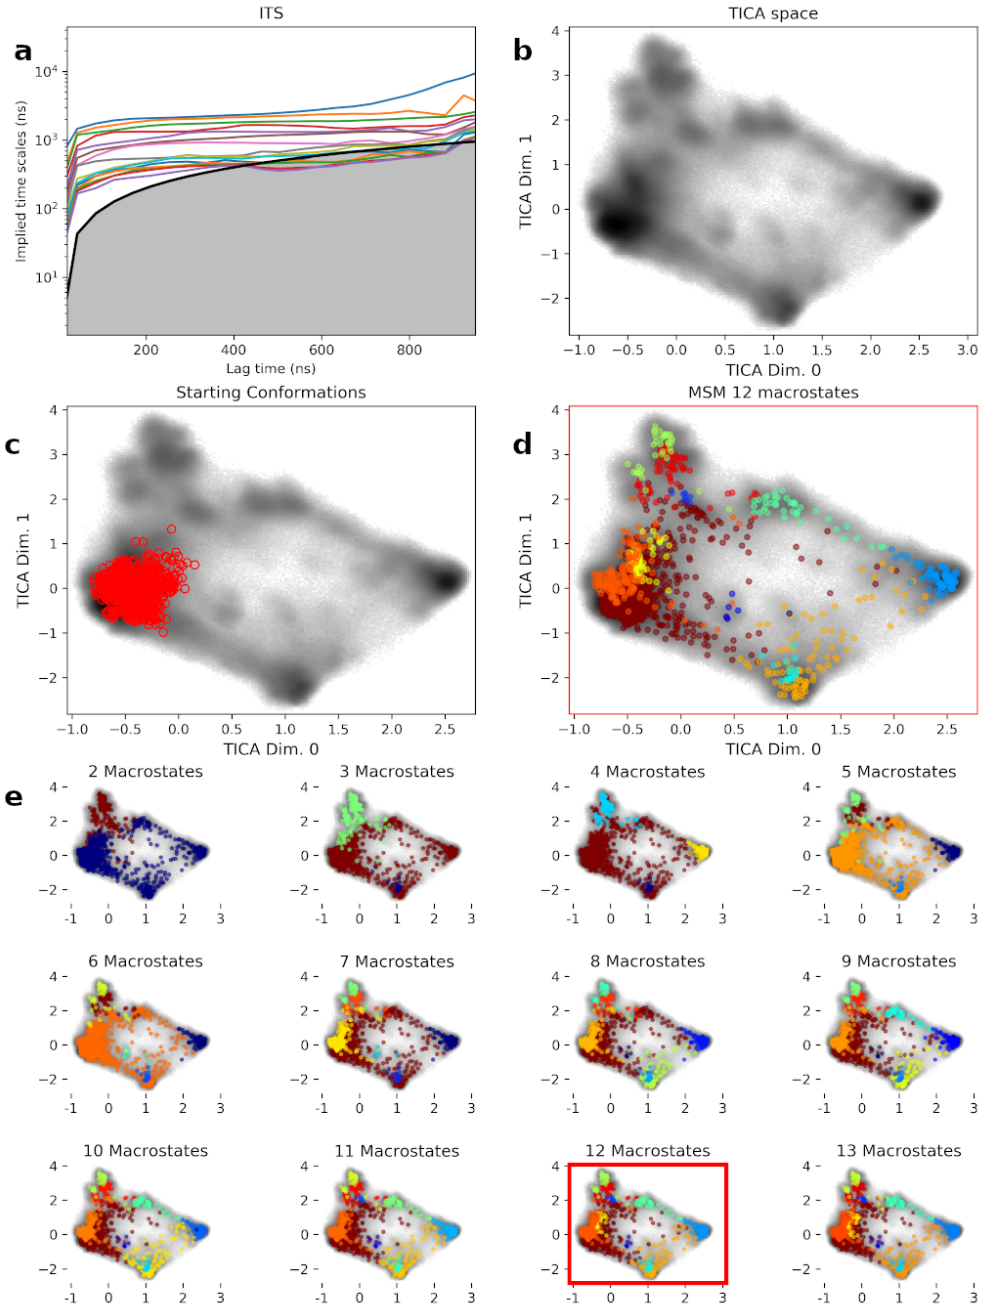

Figure 5: **Markov State Model** parametrization. **a) Implied times scales**. A final lag time of 120 ns was used to create the MSM. **b) TICA space** for the first and the second TICA dimensions. **c) Starting conformations** location on the TICA space. **d) Macrostate distribution** on the TICA space. Microstates are located based on their centers and color mapped following their macrostates assignment. **e) Discretization of the TICA space** by incrementing the number of macrostates. A final number of 12 macrostates was chosen.

3.2. Binding-and-folding recognition of an intrinsically disordered protein using adaptive molecular dynamics

Pablo Herrera-Nieto, Adrià Pérez, Gianni De Fabritiis. *Manuscript in preparation*

Summary

In this paper, we study the topic of coupled folding and binding between a disordered protein and its folded partner. Since the discovery of the potential of IDPs to undergo disorder-to-order transitions, the molecular details explaining these processes have been widely studied both experimentally and computationally. Here we summarize the main findings on such events in the system composed by the folded KIX and the disordered cMyb system. The interaction between this pair of proteins has proved challenging for us over the years. However, novel sampling techniques allowed us to partially reconstruct the NMR structure, to describe secondary binding modes, and to provide an atomically detailed overview of the binding mechanism.

Binding-and-folding recognition of an intrinsically disordered protein using adaptive molecular dynamics

Pablo Herrera-Nieto[a,1], **Adrià Pérez**[a,1], **and Gianni De Fabritiis**[a,b,2]

[a]Computational Science Laboratory,Barcelona biomedical research park (PRBB), Universitat Pompeu Fabra, C Dr Aiguader 88, Barcelona 08003, Spain; [b]Institució Catalana de Recerca i Estudis Avançats (ICREA), Passeig Lluis Companys 23, 08010 Barcelona, Spain

This manuscript was compiled on April 14, 2020

Intrinsically disordered proteins participate in many biological processes by folding upon binding with other proteins. However, coupled folding and binding processes are not well understood from a microscopic point of view. One of the main questions is whether folding occurs prior or after binding. Here we use a novel unbiased high-throughput adaptive sampling approach to reconstruct the binding and folding between the disordered transactivation domain of c-Myb and the KIX domain of the CREB-binding protein. The reconstructed long term dynamical process highlights the binding of a short stretch of amino acids on c-Myb as a pre-folded α-helix. Leucine residues, specially Leu298 to Leu302, establish initial native contacts that prime the binding and folding of the rest of the peptide, thus mixing conformational selection on the N-terminal region with the induced fit of the C-terminal.

conformational selection | coupled folding binding | intrinsically disordered proteins | molecular dynamics simulations | protein-protein interactions

Intrinsically disordered proteins (IDPs) participate in many biological functions despite lacking a stable tertiary structure (1). Initial clues for the function of IDPs were revealed by structural studies (2, 3), showing that proteins that were disordered in isolation became folded upon interacting with their partners, opening to question how folding couples with binding.

Recently, molecular dynamics (MD) simulations have been successfully applied to reconstruct biological dynamic events in problems such as protein-ligand (4) and protein-protein (5, 6) binding, as well as protein folding (7, 8). MD has also been applied in the field of IDPs (9–12). In particular, the Mdm2 protein and the disordered 12-residue N-terminal region of p53 were studied using implicit solvent simulations(9), parallel full-atom simulations totalling 831 μs (11), biased free-energy-based sampling (10), and both biased and unbiased simulations in order to estimate kinetics on the second timescale (12). For another system, KIX-pKID, a single event of binding (13) has been sampled at all-atom resolution.

The KIX—c-Myb folding-and-binding mechanism has been extensively studied experimentally as an exemplar case of protein-IDP interaction (14–20). The KIX domain of CBP is a short 87-aa region composed of three α-helices (designated as α-1, α-2 and α-3, from N-terminal to C-terminal) forming a compact bundle (3). KIX represents a paradigm of binding promiscuity: it binds many IDPs, including the proto-oncogene c-Myb (3) (Figure 1.a), with multiple binding conformations (14). However, the system composed by KIX—c-Myb remained outside of the scope of all-atom molecular simulations due to the extension of the IDP (it doubles the lengths of p53) and the existence of multiple binding modes between them (14). In particular, it is unclear whether the interaction takes place by conformational selection, i.e. c-Myb needs to be folded before binding to its partner or by induced-fit, where binding not only happens independently of c-Myb's secondary structure but also triggers its folding as shown for other IDPs (KIX-pKID) (13, 21). Another important factor is c-Myb's high helicity in isolation and the consequences it might exert on the final complex structure, which features an extended α-helical c-Myb bound to KIX. Some reports support the induced-fit approach based on kinetics and mutagenesis studies (15, 16, 20), while others advocate for a mixed mechanism (14); yet not a detailed model for the binding process is available.

In this paper, we study the binding mechanism between KIX and c-Myb. We take advantage of a novel algorithm which frames the MD sampling problem from a reinforcement learning perspective to reconstruct multiple binding modes between c-Myb and KIX. This new sampling algorithm was key for us to reconstruct the binding process, as previous attempts over the years using other state-of-the-art adaptive sampling methods (22, 23) were not successful, always failing to recover the NMR bound structure. Results provide insights in the binding mechanism between these two proteins, supporting

Significance Statement

Many intrinsically disordered proteins fold upon interacting with their protein partners. This process, known as coupled folding and binding, has been extensively studied to discriminate whether pre-folded conformations are selected or if binding takes place through induced fit. Molecular dynamic simulations have been extensively applied in the challenging task of recreating experimentally determined structures. Here, we take advantage of a novel reinforcement-learning-based molecular dynamics sampling algorithms to reconstruct the binding and folding between the disordered transactivation domain of c-Myb and the KIX domain of the CREB-binding protein. We report the reconstruction of binding of c-Myb in several interfaces of KIX, how pre-folded conformations of a short region of c-Myb are selected in the initial steps of binding, and the various pathways underlying the process.

Author contributions: P.H.N. and A.P. generated and analyzed the data; P.H.N., A.P. and G.D.F. wrote the paper; G.D.F. designed research.

[1]These authors contributed equally.
The authors declare no conflict of interest.

[2]To whom correspondence should be addressed. E-mail: gianni.defabritiis@upf.edu

DRAFT

a mixed model that combines both conformational selection and induced-fit.

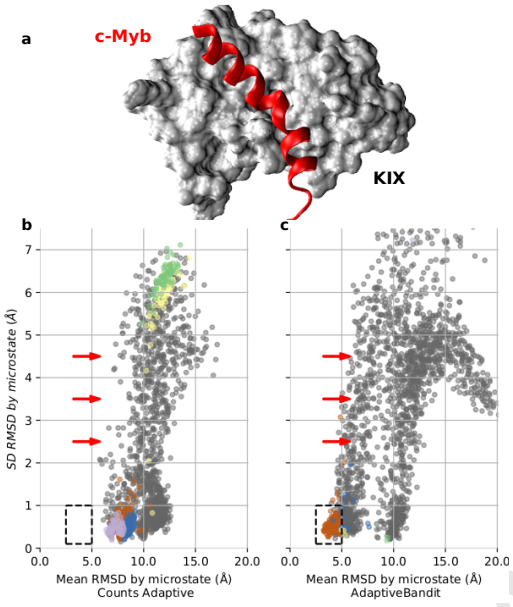

Fig. 1. Exploration performance a) KIX—cMyb NMR structure. KIX domain is shown as grey surface and bound c-Myb as a red helix (PDB code *1SB0*). **Exploration performance** by **b) counts Adaptive** ($\sim$ 480 μs) , and **c) AdaptiveBandit** ($\sim$ 225 μs) is shown by plotting mean RMSD (on the x axis) and standard deviation (on the y-axis) for each microstates of the MSM. Microstates are color mapped accordingly to their macrostate assignment. Dashed square indicates the *bound zone*, placed in the region corresponding to low RMSD mean and deviation. Red arrows are placed for reference, to highlight differences in the exploration performed by the algorithms.

Results

Adaptive sampling the conformational landscape of KIX—c-Myb binding. Simulations to reconstruct the KIX—c-Myb binding mode were performed following an adaptive sampling strategy. In adaptive sampling, successive rounds of simulations are performed in an iterative step-wise manner, where an acquisition function over the currently sampled conformation is defined. Initially, we compare two acquisition functions: a count-based one and another inspired by reinforcement learning. Standard low counts adaptive sampling (22) (called Counts Adaptive) can be shown to be optimal in pure exploration conditions (23). Counts are computed over clusters of conformations; this method is, however, noisy as clusters can be poorly populated. Therefore, in the implementation available in (23), counts are computed over a smaller subset by grouping clusters (microstates) into macrostates, constructing a Markov State Model (MSM) (24) with the available data at each round. The acquisition function is given by proportionally choosing macrostates as $1/c$, where c represents macrostate counts, and by randomly selecting conformations within them.

The new AdaptiveBandit method (25) is framed into a simplified reinforcement learning problem, i.e a multi-armed bandit problem (see *Materials & Methods*). We use the upper confidence bound (UCB) algorithm (26) to optimize an action picking policy in order to maximize future rewards, optimally balancing the exploration of new higher rewarding actions with the exploitation of the most known rewarding ones. The reward function, which associates the action to the reward given by the system, defines what we want to optimize. In this work, we choose the reward to be minus the free energy of each configuration visited in the trajectory spawn from a given action (see Eq.2 in *Materials & Methods*), where the free energy of a conformation is given by the corresponding MSM microstate computed with the data available at the current sampling epoch.

A comparison using Adaptive and AdaptiveBandit is provided in Figure 1.b,c. The batch based on Adaptive (48 epochs) failed to connect microstates similar to the NMR structure in over $\sim$ 480 μs, reaching at best an RMSD around 7 Å, indicating excessive exploration. For us, it was impossible to build an MSM with the bound state with previous methods, and novel approaches were needed to reconstruct the binding process between KIX and c-Myb successfully. In average, $\sim$ 14 μs are needed to sample the bound state, as measured by taking the concentration of the simulation and the experimental k_{on} rate. The mean first passage time for unbinding is instead 41 ms. AdaptiveBandit finds bound like microstates already in the first epochs, while it provides converged estimates of kinetics and thermodynamics after just 150 μs of sampling (see Figure 3)

Identification of an bound state & kinetics estimation. The full data set of the AdaptiveBandit run accounted for a total simulation time of $\sim$450 μs, split across 40 epochs, and was the one used to study the molecular features of KIX—c-Myb binding. An MSM was built based on all-pair $C_\alpha + C_\beta$ distances between KIX and c-Myb and self distances between C_α of c-Myb, tracking both protein-protein contacts and IDP folding. The MSM included three kinetically similar sets of conformations, referred as macrostates (Figure 2.a and Supplementary Figure 1.b): a highly populated state including an heterogeneous mixture of conformations (*unbound*), a well defined bound state located on a secondary pocket (*misbound*), and, finally, c-Myb bound to the primary interface (*bound*). Representative structures of all states can be found in Figure 2.b.

The *bound* state shares up to 40% of the fraction of native inter-molecular contacts (Q_{int}) between c-Myb and KIX with the original NMR structure (PDB ID: *1SB0*), as shown in Figure 2.c. These contacts mainly involve the interaction of c-Myb residues Leu298 and Leu302 with residues across the primary binding interface: Leu302 contacts Leu603, Leu653, and specially Leu607 of KIX, which is buried down in the pocket, whereas Leu298 establishes additional native contacts with Ala610, Ile657, and Tyr658. Q_{int} reaches up to 75% in those microstates exhibiting the tightest bound conformations, and, in addition to the leucine binding, they feature most of the contacts between the C-terminal half of c-Myb and KIX, which are not that prevalent across the bound macrostate (Supplementary Figure 2). The main contacts missing account for the electrostatic interactions established between Arg294 and the C-terminal region on α-3.

Secondary structure profile for MD derived states matches

DRAFT

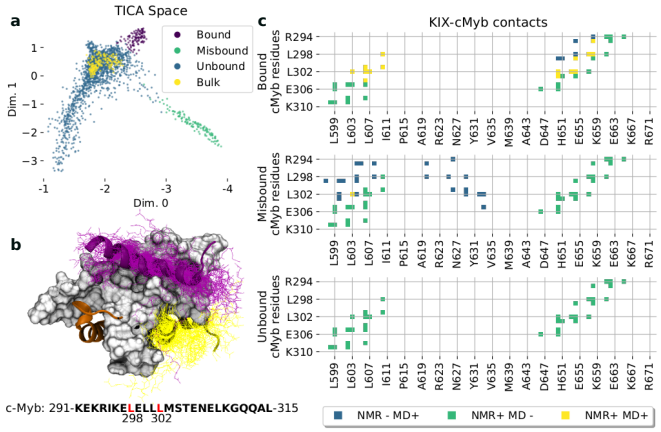

Fig. 2. KIX—c-Myb binding model. a) States distribution across the TICA space: microstates are represented as dots and are colored following their macrostate assignment. **b) Representative structures:** PDB structure *2AGH* is depicted with KIX as gray surface, c-Myb bound to the primary interface as a purple ribbon, and MLL bound to the secondary interface as an orange ribbon. c-Myb backbones for 50 representatives MD structures of *bound* and *misbound* states are displayed with purple and yellow lines, respectively. **c) Macrostate contact fingerprint:** profile of contacts established between c-Myb and KIX in each macrostate in at least 50% of the structures. Blue color represents contacts present in the state but not in the original NMR structure; green indicates original NMR contacts not found in the MSM state; and yellow squares represent contact matches, found in both NMR and MD structures.

the experimental description of c-Myb (14, 19), as shown in Supplementary Figure 3: the 25 residues are separated in two halves by residues Met303 and Ser304. The N-terminal half shows a high helical tendency, around 20-30% for residues in positions 297 to 302 with c-Myb in isolation, being maximal in bound states. Experimentally, this N-terminal half in isolation reaches even higher helicity levels (∼70%) when using an extended construct of c-Myb (14). On the other hand, the C-terminal section exhibits lower helical propensity, regardless of the presence or absence of KIX, and its folded conformation is mainly appreciated in those microstates with the tightest bound conformation.

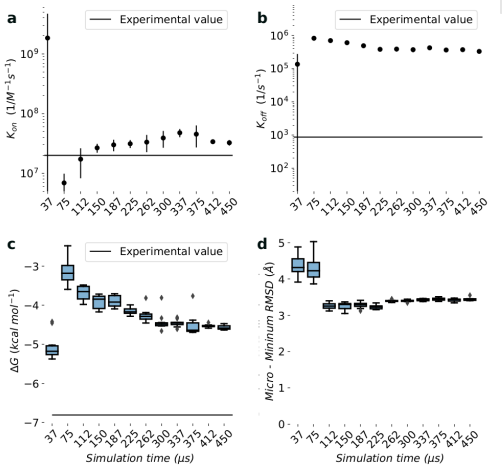

Fig. 3. Evolution through the MD run of **a)** k_{on}, **b)** k_{off}, **c)** free energy and **d)** microstate minimum RMSD. Each data point was calculated with 10 MSMs built by bootstraping 80% of the simulations.

The kinetic parameters derived from the model were compared with available information (17). Experimental values from Shammas et al. were calculated at temperatures ranging from 278 to 298 K, while simulations were executed at physiological temperature (310 K). k_{on} values display a temperature independent tendency, whereas temperature dependent k_{off} and k_d values at 310 K were, therefore, extrapolated (Supplementary Figure 4). Hence, reference values for k_{off} resulted in 866 s^{-1} and -6.81 $kcal\ mol^{-1}$ for the free energy (obtained from k_d).

MSM estimation of k_{on} is $(3.2 \pm 0.3) \cdot 10^7 M^{-1}s^{-1}$, in agreement with the experimental value $(2.2 \pm 0.2) \cdot 10^7 M^{-1}s^{-1}$ (17). However, off rates are overestimated by ∼500 fold (Figure 3.b), with calculated values of $(3.3 \pm 0.2) \cdot 10^5\ s^{-1}$. Similarly, calculations for the free energy of binding were $-4.41 \pm 0.39\ kcal\ mol^{-1}$, with the experimental being $-6.81\ kcal\ mol^{-1}$. These parameters were estimated by building and bootstrapping MSMs using incrementally more trajectories in order to ensure convergence of the model (Figure 3 and Supplementary Figure 5). Convergence is reached at ∼750 trajectories, equivalent to half of the aggregated data (Figure 3). Initially, estimations for k_{on} values are in the range of the diffusion limit, up to $10^{10} M^{-1}s^{-1}$, as a well defined bound macrostate is not found until half of the adaptive sampling run. These estimations agree with the exploration profile followed for the AdaptiveBandit algorithm (Supplementary Figure 1.c), which shows how the exploration of TICA space related to the *bound* state took place during the first part of the run.

The discrepancy between reference and computed k_{off} points out structural differences between the MD-bound state and the NMR, as highlighted by the remarkably high k_{off} values ($\sim 10^6\ s^{-1}$ compared with the reference $\sim 866\ s^{-1}$). In order to explore potential divergences responsible for these kinetic deviations, additional long trajectories (8 replicas of 2 μs each) were run starting from bound NMR and MD derived conformations (summarized in Supplementary Figure 6). The comparison between the long runs indicates an evident separation between most sampled areas between both groups, with the NMR one occupying closer areas to the experimental structure. The comparison between structural averages revealed that the most diverging regions map to the C-terminal of c-Myb, which retains its helical conformation in the NMR

DRAFT

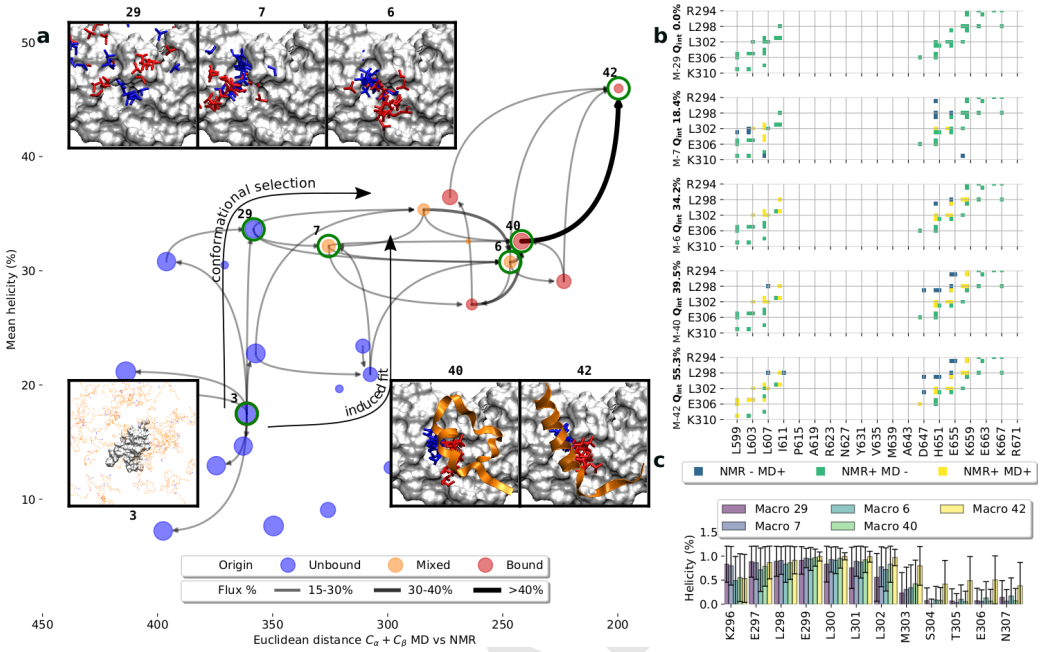

Fig. 4. Complete binding process of c-Myb to KIX. a) Main pathways leading from *"bulk"* macrostate to *"NMR-like"*. Nodes are placed according to the structural similarity respect to the NMR model —calculated as the euclidean distance between the C_α + C_β pair distances of KIX to c-Myb of MD macrostates respect the NMR structure— on the *x axis*, and mean helicity on the *y axis*. Node size is proportional to their equilibrium distribution. Edges represent the connection between macrostates, and their color and thickness, the percentage of the total flux traversing them. Nodes are colored blue, red, or green depending on whether all their microstates originally belonged to the unbound or bound state, or if they have a mixed composition. **Representative structures:** KIX is represented with a gray surface. c-Myb backbone for 10 MD superposed structures is shown in orange, with side chains of residues Leu298 and Leu302 highlighted in red and blue, respectively. Evolution of the most relevant states involved at different stages of the binding process. For most relevant macrostates the **b)** KIX—c-Myb contacts profile and **c)** mean helicity are shown.

data set while being mostly disordered in the MD derived runs. Therefore, the instability of the C-terminal portion of cMyb might be key to the overall volatility of the complex.

Binding follows both induced-fit and conformational selection. The MSM used so far has an inherent structural diversity on each macrostate that complicates the monitoring of conformational changes along the binding process just by tracking macro-to-macro transitions. At the same time, the number of microstates conforming the MSM is too high to obtain a simplified binding pathway. In order to gain structural insight of the binding process, it is necessary to create a new set of macrostates as structurally homogeneous as possible. To achieve this goal we followed the process described in the *Kinetic pathway identification* section of *Materials & Methods* to recluster microstates based on their structural similarity, which yield a final MSM of 123 macrostates (labelled as M1-M123).

Once the new model was generated, we applied transition path theory (27, 28) to calculate fluxes leading from the purely bulk state to the bound conformations. Out of the 123 macrostates of the MSM, only a reduced set of 25 is sufficient to explain binding to the primary interface, as the rest participate in secondary binding modes. The network generated by the flux interchanges between macrostates (Figure 4.a and Supplementary Figure 7) can be separated into two main events: the establishment of the initial contacts and reaching the bound conformation.

Initial binding encompasses a high number of macrostates, formerly belonging to the *unbound* state, presenting low Q_{int} values and various degrees of helicity that conform a tangled network of fast inter-converting states but low transited paths. However, helicity remains stable amongst those macrostates carrying most of the flux in these initial steps (macrostates M29, M53, and M71; Figure 4.b). They all are in the range of 30% and include conformations with a helical N-terminal of c-Myb, especially on residues 297-302 (Figure 4.c). The most important state at this stage is M7, which centralizes most of the incoming flux (Figure 4.a), serving as the main intermediate to reach the *bound* state. M127 features the first native contacts found across the KIX—c-Myb binding pathway, which involves residues Leu302 of c-Myb. The role of Leu302 as the main driving force for the interaction has already been described (3), and is due in part to the kink in the helix created by neighbours residues Met303 and Ser304, which exposes Leu302 allowing for a deep penetration inside the binding pocket. Besides, on the KIX residues contacted at this statge is Leu603, which is one the most exposed residues in the hydrophobic pocket later occupied by Leu302.

It is interesting to note the presence of a set of macrostates

(M72, M93, M114, and others) including similar contacts involving Leu302 and KIX, but featuring remarkably low levels of helicity (Supplementary Figure 7). The contribution of these pathways to the overall flux is lower, and they are channelled through M114 and M122, with higher levels of helicity, to the main path. Therefore, binding before folding is also observed, but it is not the main pathway followed along the process, and helical conformations are preferred for the formation of initial contacts.

Once M7 has been reached, the flux converges to a reduced number of highly transited paths, with M5, M6, and especially M40 collecting most of the flux leading to M5, the closest to the NMR structure. From M7, Q_{int} increases due to the establishment of the contacts involving Leu298 and Leu302. The last step, jumping to M42, involves the folding of the C-terminal region of c-Myb and the formation of its contacts with KIX. A full reconstruction of the progressive acquisition of contacts between KIX and c-Myb is shown in Figure 4.c.

In summary, initial steps are greatly benefited from pre-folded helical structures of c-Myb (Figure 4.d). The *binding before folding* pathway is also observed, but it accounts for a relatively smaller fraction of the total flux. Binding of helical conformations dominates the initial steps of the interaction, but for the interaction of the C-terminal tail, folding follows binding. According to our model, transitioning between states takes place by fast interchanges ranging from 10^6 to 10^8 $M^{-1}s^{-1}$. No limiting steps in the binding process are observed; hence no possible transition states can be defined, as pointed by experimental reports (17).

Secondary binding Modes. The existence of alternative binding poses between c-Myb and KIX has also been reported (14). The MSM shows the presence of a secondary binding mode (referred as *misbound*) apart from the interaction with the primary binding interface. It occupies a novel interface, located between α-1 and α-2 (Figure 1.b and Supplementary Figure 9). The interaction of the *misbound* state resembles the *bound* binding mode: the N-terminal half is folded in the typical α-helix, while the C-terminal section remains mostly unstructured. Kinetically, there is a 4-fold difference in the mean first passage time for binding between both sites —$(9.28 \pm 1.11) \cdot 10^3$ ns for binding to *bound* site and $(4.06 \pm 1.81) \cdot 10^4$ ns for the *misbound* site— that may account for the preferential binding of c-Myb to the primary interface. Finally, exploration of other slow processes on the MSM also shows the binding of c-Myb to the mixed-lineage leukemia (MLL) (29) site (located between helices α-2 and α-3).

Conclusion. The analysis presented here provides a detailed molecular description of binding of c-Myb to the primary interface of KIX, summarized as a two-step process, where initially the N-terminal region of c-Myb binds with a preferred helical conformation, allowing the formation of native contacts and, in the last step, folding and binding of the C-terminal. Study of the fluxes derived from the MSM shows the relevance of residue Leu302 not only in the final bound structure but also as the responsible of establishing the first contacts and serving as an anchoring point between c-Myb and KIX.

The model also indicates how conformational selection would only affect residues 298 to 302 and not the whole length of the peptide. Additionally, binding before folding is also observed, but it accounts for a relatively smaller fraction of the total flux, thus establishing a induced-fit only pathway that coexists with conformational selection.

Finally, the novel MD sampling approach used in this work, AdaptiveBandit, had a crucial role in resolving this type of folding and binding process. The method is implemented and available in the HTMD python package (23). However, more algorithms can be derived within the same adaptive bandit framework. While here we choose the reward to be minus the free energy, other choices could optimize different costs, for example, improving the precision of the off-rate or optimizing sampling in the context of structure prediction. These aspects are left for future work.

Materials & Methods

Molecular dynamics simulations. In order to generate initial conformations for c-Myb (residues 291 to 315), we ran multiple parallel simulations. The peptide was solvated in a cubic water box of 64 Å side with a NaCl concentration of 0.05 M. First, the peptide was simulated at 500 K for 120 ns to unfold the initial structure.

Then, 200 systems were built by placing one random unstructured c-Myb conformation in conjunction with KIX in opposite corners of a 64 Å side cubic water box with a NaCl concentration of 0.05 M, resulting in final protein concentration of ∼3.2 mM.

All systems were built using VMD (30) and simulated with ACEMD (31), the CHARMM22* force field (32) and TIP3P water model (33). A Langevin integrator was used with a damping constant of 0.1 ps^{-1}. The integration time step was set to 4 fs, with heavy hydrogen atoms (scaled up to four times the hydrogen mass) and holonomic constraints on all hydrogen-heavy atom bond terms. Electrostatics were computed using PME with a cutoff distance of 9 Å and grid spacing of 1 Å. After energy minimization, equilibration for all systems was done in the NPT ensemble at 303 K, 1 atm, with heavy atoms constrained at 1 $kcal\ mol^{-1}\ Å^2$. Energy minimization was run for 500 steps and equilibrated for 2 ns.

Production runs of 250 ns were performed at 310 K using the distributed computing project GPUGrid (34), following an adaptive sampling strategy.

AdaptiveBandit sampling. The multi-armed bandit problem is defined by $\langle \mathcal{A}, \mathcal{R}, \gamma \rangle$, where an action $a_t \in \mathcal{A}$ and $\mathcal{R}^a$ is a (stochastic) reward function. We choose $\gamma = 0$ for totally discounted rewards. The optimal policy $\pi_a \sim \mathbb{P}[a]$ selects actions a_t in order to maximize the cumulative future rewards. The construction of an optimal selection strategy requires handling the exploration-exploitation problem. AdaptiveBandit relies on the UCB1 algorithm (26), defining an upper confidence bound for each action-value estimate based on the number of times an action has been picked and the total amount of actions taken

$$a_t = \underset{a \in \mathcal{A}}{\operatorname{argmax}} \left[Q_t(a) + c \sqrt{\frac{\ln t}{N_t(a)}} \right], \qquad [1]$$

where t denotes the total number of actions taken, $Q_t(a) = \mathbb{E}_\pi[r|a]$ is the action-value estimation, $N_t(a)$ is the number of times action a has been selected (prior to time t) and c is a constant controlling the degree of exploration. As for the reward definition, there are different choices depending on the objective, e.g. here, the interest is sampling the bound metastable state, hence, we rewarded actions based on the stability of conformations using MSM estimations of the free energy for each state

$$\mathcal{R}_a = < k_B T \log(\mu(x)) >_{(a, x_1, \ldots, x_\tau)}, \qquad [2]$$

where $\mu(x)$ is the equilibrium distribution estimated by the MSM with the current available data and the average is performed over the frames in the trajectory starting from a. AdaptiveBandit uses the MSM discretized conformational space to define the action set and at each round acquires a random conformation from the selected states to respawn new simulations. A more formal description of the bandit framework and AdaptiveBandit in the context of adaptive sampling as well as analysis in simpler, analytical potentials are available at (25). The AdaptiveBandit sampling algorithm is made available in the HTMD (23) Python package.

DRAFT

Adaptive Sampling parameters. For both the AdaptiveBandit and the count Adaptive runs, the construction of MSMs at each epoch was done using the residue-residue contacts between KIX and c-Myb measured as the minimum contacts between residues at a threshold of 5 Å, and the backbone dihedral angles of c-Myb. Time independent component analysis (TICA) (35) was used for dimensionality reduction using a lag time of $\tau = 20$ frames and keeping the 3 first dimensions, which were later clustered with a k-centers algorithm. AdaptiveBandit was performed during 40 epochs with a c value of 0.01.

Long trajectories. Additionally, a set of long MD runs were performed starting from bound structures. Four models of the NMR determined structure and four random bound conformations were selected and equilibrated as previously described. A total of 8 long trajectories of ~ 2 μs each were generated. The analysis was performed by comparing the aggregated data from NMR and MD runs.

Markov state model analysis. The projected space used for building the MSM included all pair C_α + C_β atoms distances between KIX and c-Myb (to account for the interaction between the two proteins), and self-distances between every C_α of c-Myb (to monitor its conformation). TICA was used at a lag time $\tau = 25$ ns (implied timescales are shown in Supplementary Figure 1.a), taking the 4 most relevant components. An extra dimension including the minimum distances between KIX and c-Myb was also added at this step to enhance a clearer separation of the bulk conformations. and discretizing them into 2,000 clusters using the mini batch k-means algorithm (36). The microstates defined in the MSM were coarse-grained into three larger meta-stable macrostates by using PCCA++ (37). For the estimation of kinetic values, the original MSM was modified by creating an additional macrostate with those microstates exhibiting the minimum amount of contacts between KIX and c-Myb. This process raised a 4^{th} metastable state considered as the *bulk* state for all subsequent calculations to obtain the kinetics of binding. For every measure, the error was estimated by creating 10 independent MSMs using a random set containing 80% of the simulation data. All analysis were performed with HTMD (23).

Kinetic pathway identification. In order to create a new set of macrostates as structurally homogeneous as possible, microstates belonging to *unbound* and *bound* macrostates (see *Results*) were re-clustered. Data used for clustering included the mean KIX—c-Myb by-residue contacts (at a threshold of 5 Å) and c-Myb dihedrals for each microstate. The AffinityPropagation algorithm (38) within scikit-learn (36) was chosen for this task since it is not necessary to provide an initial number of clusters. As a result of this process, the *unbound* macrostate was split into 114 new macrostates, *bound* states into 5, and 4 macrostates included microstates originally belonging to both *unbound* and *bound*, raising a final MSM of 123 macrostates. The transition matrix between macrostates involved in binding to the primary interface is shown in Supplementary Figure 10. Finally, fluxes between macros were estimated using transition path theory (27, 28).

1. HJ Dyson, PE Wright, Intrinsically unstructured proteins and their functions. *Nat. reviews Mol. cell biology* **6**, 197 (2005).
2. PH Kussie, et al., Structure of the MDM2 oncoprotein bound to the p53 tumor suppressor transactivation domain. *Science* **274**, 948–953 (1996).
3. T Zor, RN De Guzman, HJ Dyson, PE Wright, Solution structure of the KIX domain of CBP bound to the transactivation domain of c-Myb. *J. molecular biology* **337**, 521–534 (2004).
4. I Buch, T Giorgino, G De Fabritiis, Complete reconstruction of an enzyme-inhibitor binding process by molecular dynamics simulations. *Proc. Natl. Acad. Sci.* **108**, 10184–10189 (2011).
5. N Plattner, S Doerr, G De Fabritiis, F Noé, Complete protein–protein association kinetics in atomic detail revealed by molecular dynamics simulations and markov modelling. *Nat. chemistry* **9**, 1005 (2017).
6. A Borgia, et al., Extreme disorder in an ultrahigh-affinity protein complex. *Nature* **555**, 61 (2018).
7. K Lindorff-Larsen, S Piana, RO Dror, DE Shaw, How fast-folding proteins fold. *Science* **334**, 517–520 (2011).
8. S Piana, K Lindorff-Larsen, DE Shaw, Atomistic description of the folding of a dimeric protein. *The J. Phys. Chem. B* **117**, 12935–12942 (2013).
9. MC Zwier, et al., Efficient atomistic simulation of pathways and calculation of rate constants for a protein–peptide binding process: application to the MDM2 protein and an intrinsically disordered p53 peptide. *The journal physical chemistry letters* **7**, 3440–3445 (2016).
10. JA Morrone, A Perez, J MacCallum, KA Dill, Computed binding of peptides to proteins with meld-accelerated molecular dynamics. *J. chemical theory computation* **13**, 870–876 (2017).
11. G Zhou, GA Pantelopulos, S Mukherjee, VA Voelz, Bridging microscopic and macroscopic mechanisms of p53-MDM2 binding with kinetic network models. *Biophys. journal* **113**, 785–793 (2017).
12. F Paul, et al., Protein-peptide association kinetics beyond the seconds timescale from atomistic simulations. *Nat. communications* **8**, 1095 (2017).
13. SH Chong, H Im, S Ham, Explicit characterization of the free energy landscape of pkid–kix coupled folding and binding. *ACS Cent. Sci.* **5**, 1342–1351 (2019).
14. M Arai, K Sugase, HJ Dyson, PE Wright, Conformational propensities of intrinsically disordered proteins influence the mechanism of binding and folding. *Proc. Natl. Acad. Sci.* **112**, 9614–9619 (2015).
15. R Giri, A Morrone, A Toto, M Brunori, S Gianni, Structure of the transition state for the binding of c-Myb and KIX highlights an unexpected order for a disordered system. *Proc. Natl. Acad. Sci.* **110**, 14942–14947 (2013).
16. S Gianni, A Morrone, R Giri, M Brunori, A folding-after-binding mechanism describes the recognition between the transactivation domain of c-Myb and the KIX domain of the CREB-binding protein. *Biochem. biophysical research communications* **428**, 205–209 (2012).
17. SL Shammas, AJ Travis, J Clarke, Remarkably fast coupled folding and binding of the intrinsically disordered transactivation domain of cMyb to CBP KIX. *The journal physical chemistry B* **117**, 13346–13356 (2013).
18. A Toto, et al., Molecular recognition by templated folding of an intrinsically disordered protein. *Sci. reports* **6**, 21994 (2016).
19. A Poosapati, E Gregory, WM Borcherds, LB Chemes, GW Daughdrill, Uncoupling the folding and binding of an intrinsically disordered protein. *J. molecular biology* **430**, 2389–2402 (2018).
20. SL Shammas, AJ Travis, J Clarke, Allostery within a transcription coactivator is predominantly mediated through dissociation rate constants. *Proc. Natl. Acad. Sci.* **111**, 12055–12060 (2014).
21. K Sugase, HJ Dyson, PE Wright, Mechanism of coupled folding and binding of an intrinsically disordered protein. *Nature* **447**, 1021 (2007).
22. S Doerr, G De Fabritiis, On-the-fly learning and sampling of ligand binding by high-throughput molecular simulations. *J. chemical theory computation* **10**, 2064–2069 (2014).
23. S Doerr, M Harvey, F Noé, G De Fabritiis, HTMD: high-throughput molecular dynamics for molecular discovery. *J. chemical theory computation* **12**, 1845–1852 (2016).
24. JH Prinz, et al., Markov models of molecular kinetics: Generation and validation. *The J. chemical physics* **134**, 174105 (2011).
25. A Pérez, P Herrera-Nieto, S Doerr, GD Fabritiis, Adaptivebandit: A multi-armed bandit framework for adaptive sampling in molecular simulations (2020).
26. P Auer, Using confidence bounds for exploitation-exploration trade-offs. *J. Mach. Learn. Res.* **3**, 397–422 (2002).
27. E Weinan, E Vanden-Eijnden, Towards a theory of transition paths. *J. statistical physics* **123**, 503 (2006).
28. F Noé, C Schütte, E Vanden-Eijnden, L Reich, TR Weikl, Constructing the equilibrium ensemble of folding pathways from short off-equilibrium simulations. *Proc. Natl. Acad. Sci.* **106**, 19011–19016 (2009).
29. RN De Guzman, NK Goto, HJ Dyson, PE Wright, Structural basis for cooperative transcription factor binding to the CBP coactivator. *J. molecular biology* **355**, 1005–1013 (2006).
30. W Humphrey, A Dalke, K Schulten, VMD: visual molecular dynamics. *J. molecular graphics* **14**, 33–38 (1996).
31. MJ Harvey, G Giupponi, GD Fabritiis, ACEMD: accelerating biomolecular dynamics in the microsecond time scale. *J. chemical theory computation* **5**, 1632–1639 (2009).
32. S Piana, K Lindorff-Larsen, DE Shaw, How robust are protein folding simulations with respect to force field parameterization? *Biophys. journal* **100**, L47–L49 (2011).
33. WL Jorgensen, J Chandrasekhar, JD Madura, RW Impey, ML Klein, Comparison of simple potential functions for simulating liquid water. *The J. chemical physics* **79**, 926–935 (1983).
34. I Buch, MJ Harvey, T Giorgino, DP Anderson, G De Fabritiis, High-throughput all-atom molecular dynamics simulations using distributed computing. *J. chemical information modeling* **50**, 397–403 (2010).
35. G Pérez-Hernández, F Paul, T Giorgino, G De Fabritiis, F Noé, Identification of slow molecular order parameters for Markov model construction. *The J. chemical physics* **139**, 07B604_1 (2013).
36. F Pedregosa, et al., Scikit-learn: Machine learning in Python. *J. machine learning research* **12**, 2825–2830 (2011).
37. S Röblitz, M Weber, Fuzzy spectral clustering by pcca+: application to markov state models and data classification. *Adv. Data Analysis Classif.* **7**, 147–179 (2013).
38. BJ Frey, D Dueck, Clustering by passing messages between data points. *Science* **315**, 972–976 (2007).

DRAFT

Supplementary Information for

Reconstruction of the binding pathway of an intrinsically disordered protein using Molecular Dynamics simulations

Pablo Herrera-Nieto, Adrià Pérez and Gianni De Fabritiis

Gianni De Fabritiis.
E-mail: gianni.defabritiis@upf.edu

This PDF file includes:

Figs. S1 to S10

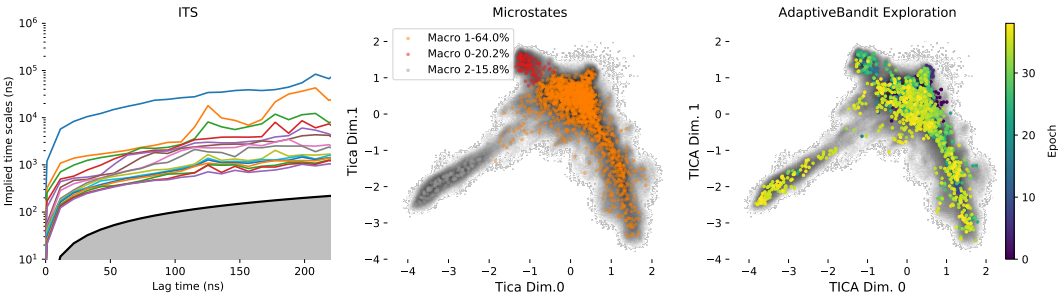

Fig. S1. Markov State Model. a) Implied time scales of the MD data. **b)** State distribution across the TICA space. **c)** AdapticeBandit exploration of the TICA space.

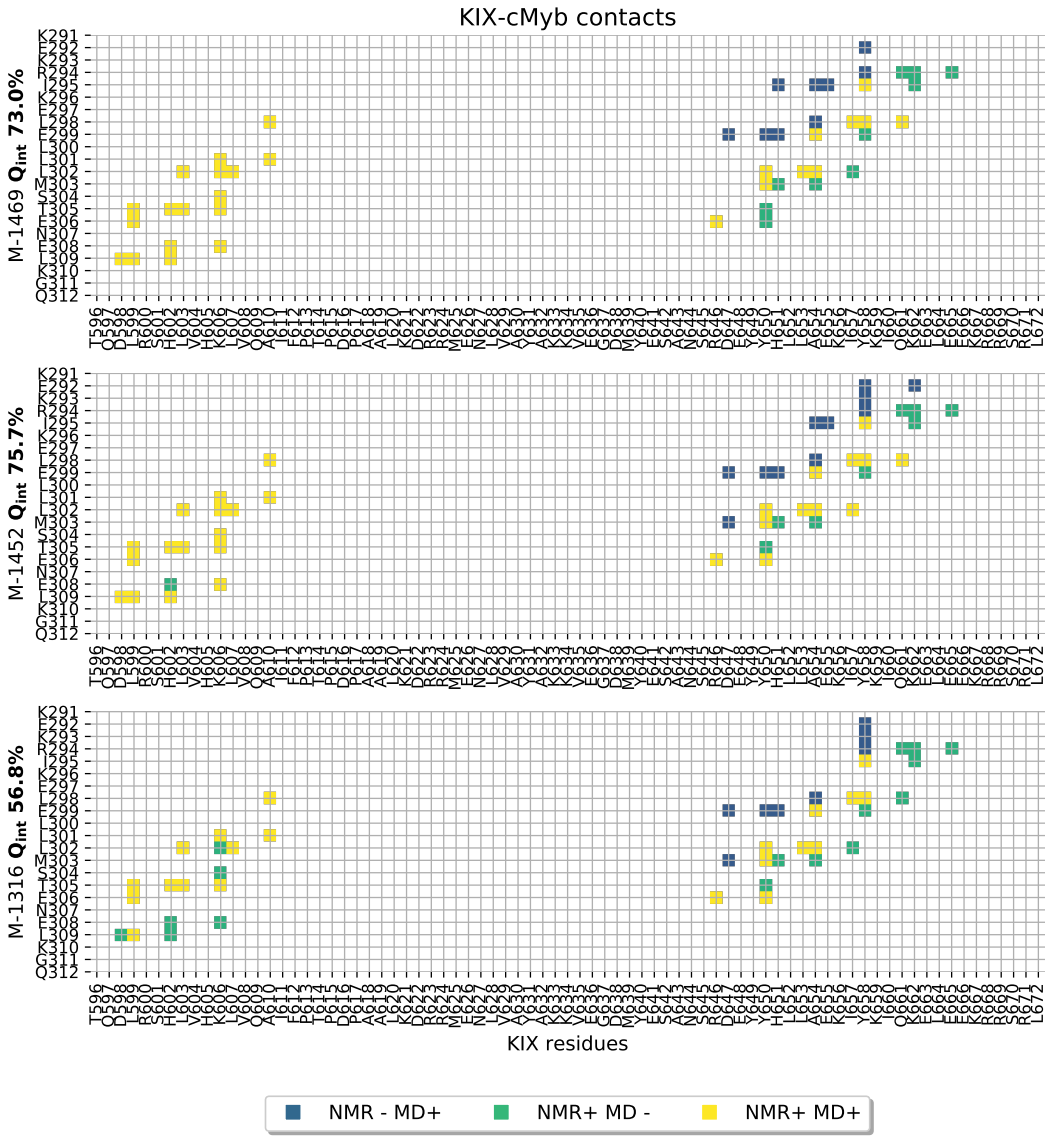

Fig. S2. Microstate contact fingerprint: profile of contacts established between c-Myb and KIX in microstates with minimum distances respect the NMR structure. Blue color represents contacts present in the state but not in the original NMR conformation; green indicates native contacts not found in the MSM state; and yellow squares represent a match on a contact, found in both the NMR model and MD structures.

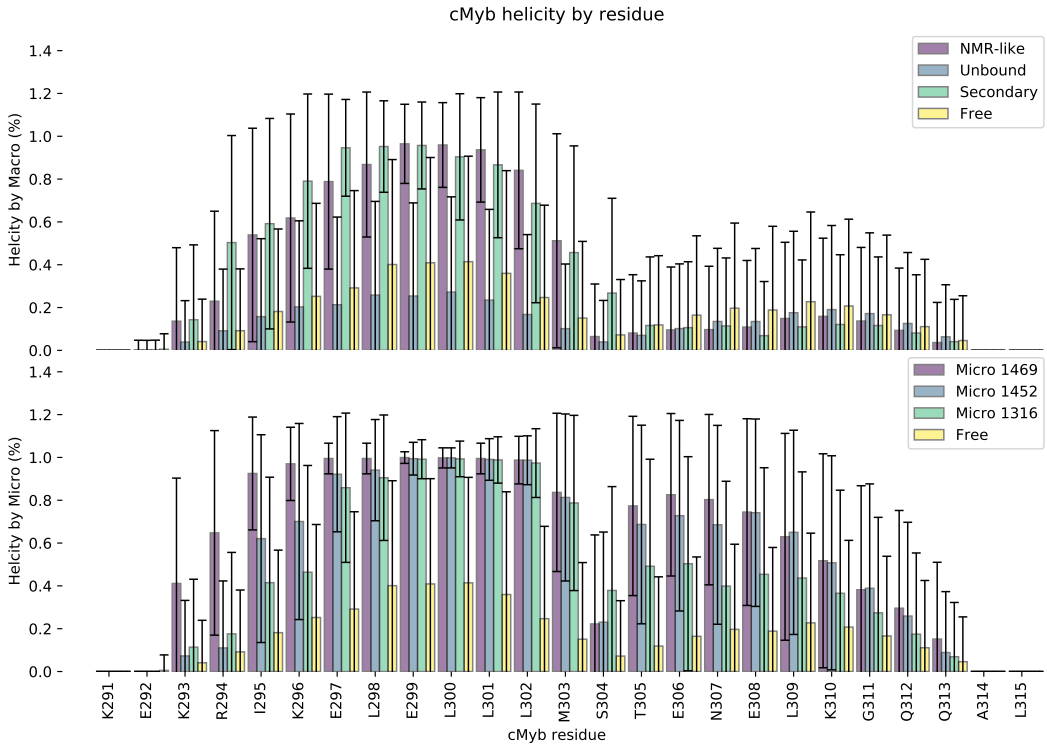

Fig. S3. c-Myb helicity. Comparison of the by-residue helicity fraction of c-Myb between **a)** the three MSM macrostates, and **b)** the three microstates with maximum Q_{int}. Additionally, the helicity profile for the peptide in isolation is depicted in yellow.

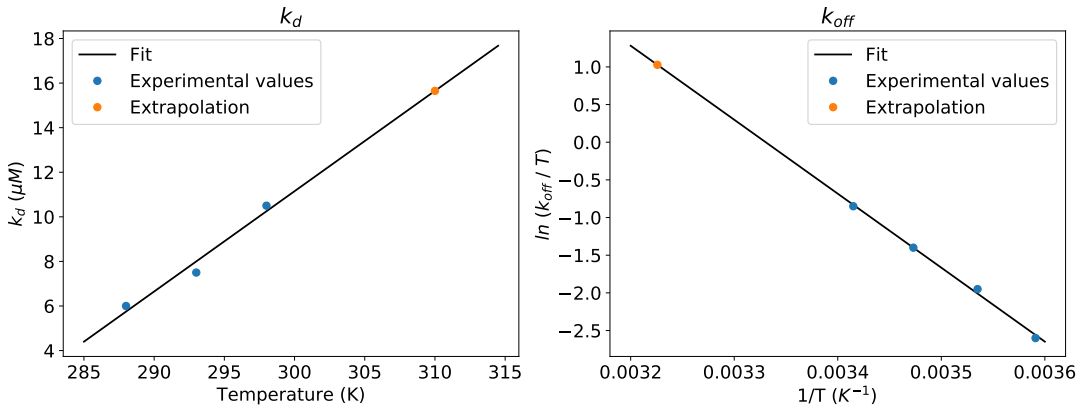

Fig. S4. Extrapolations of **a)** k_d and **b)** k_{off} values from experimental data.

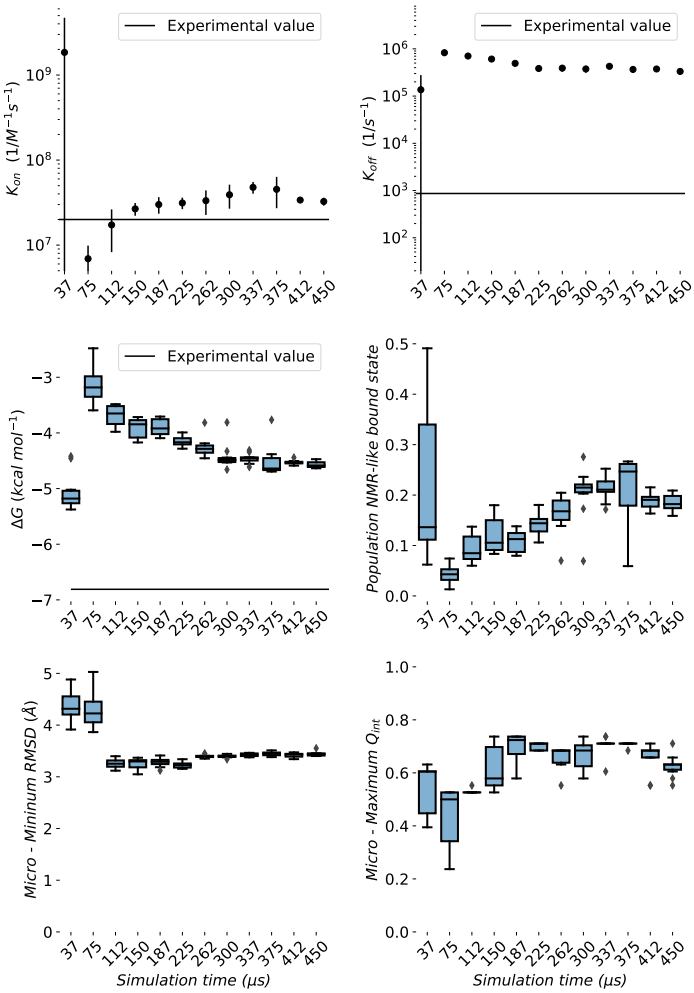

Fig. S5. Observables evolution through the MD run. Each data point was calculated with 10 MSMs built by bootstrapping 80% of the trajectories. The estimated values included here are: **a)** K_{on}, **b)** K_{off}, **c)** free energy, **d)** equilibrium probability of the NMR-like states, **e)** microstate minimum C_α RMSD; and **f)** microstate maximum Q_{int}.

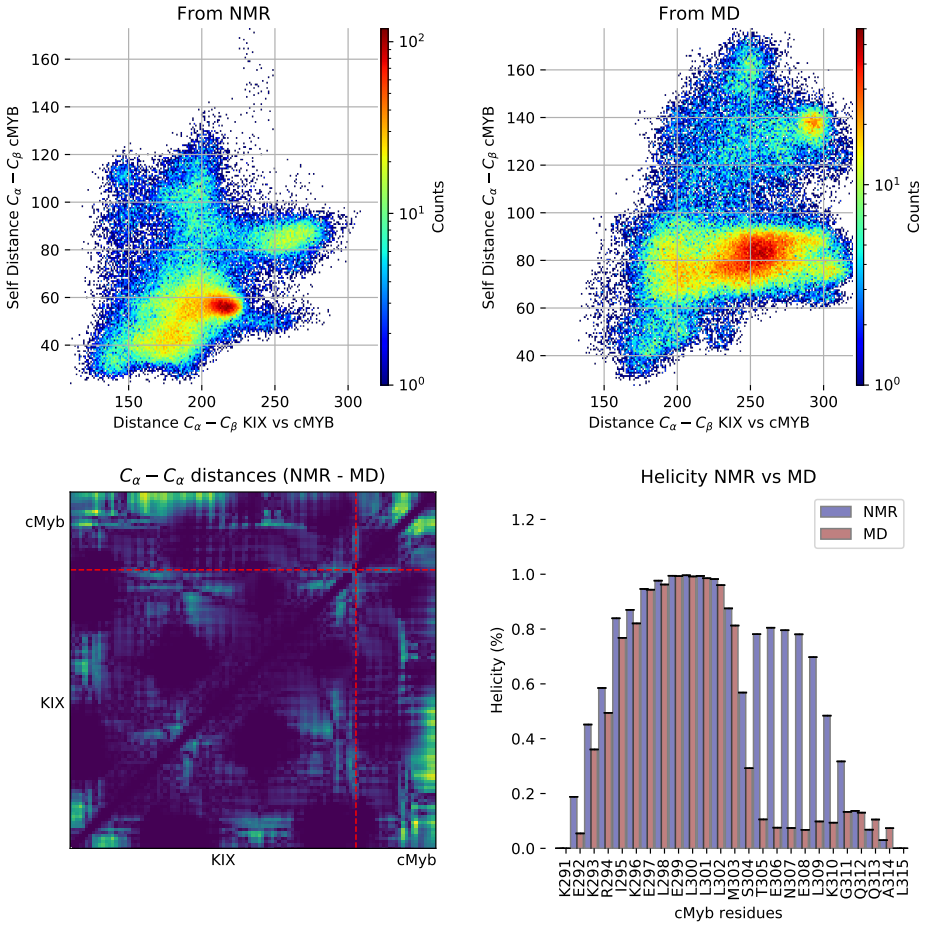

Fig. S6. Comparison between long trajectories starting from NMR and bound MD conformations. a), b) By frame distribution according to similarity between KIX and c-Myb interaction (x axis) and c-Myb conformation (y axis). **c)** $C_\alpha - C_\alpha$ distance differences between aggregated MD and NMR data. **d)** Comparison of the helicity profile for the MD and the NMR data sets.

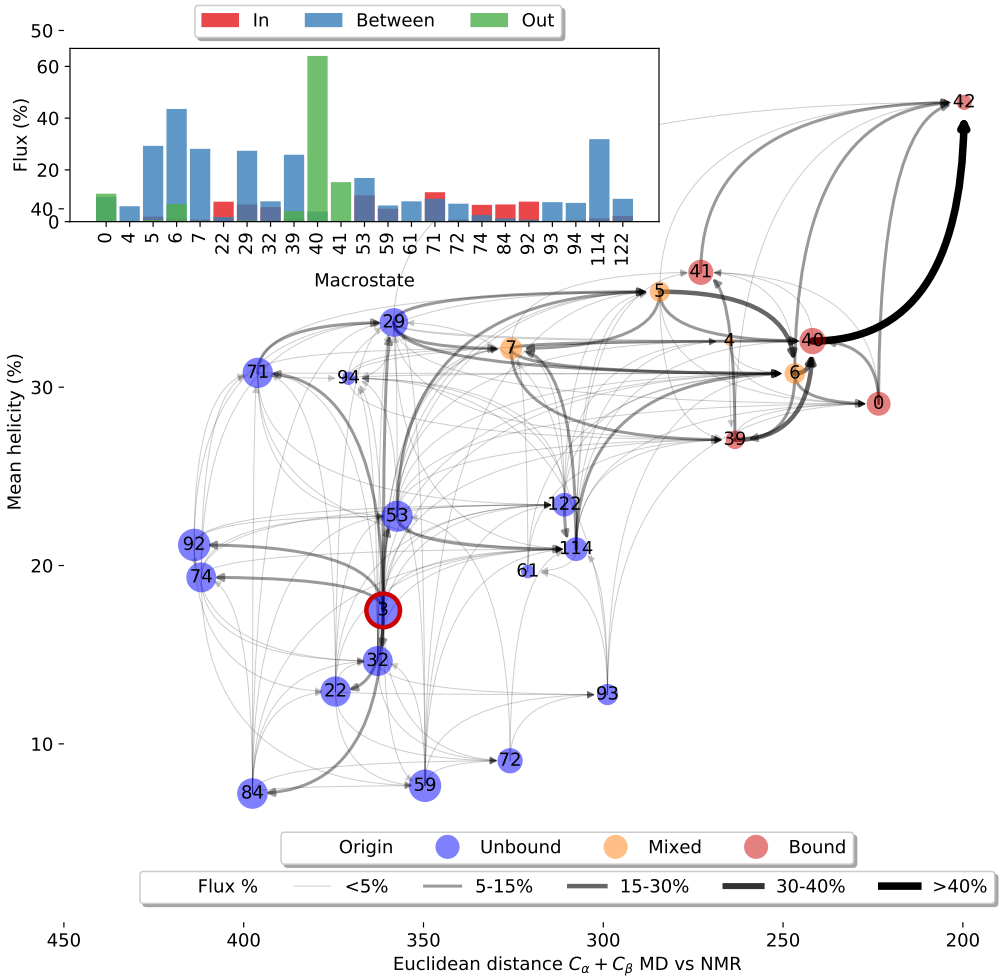

Fig. S7. Complete binding process of c-Myb to KIX. a) Flux profile by macrostate. *"In"* rates indicate the percentage of the total flux coming directly from the bulk state; *"out"* rates represent the percentage of flux directly leading to the *NMR-like* state; and *"between"*, the difference between the total flux that each node receives, and the *in* and *out* rates. **b) Main pathways leading from *"bulk"* macrostate to *"bound"*.** Nodes are placed according to the structural similarity respect to the NMR model —calculated as the euclidean distance between the C_α + C_β pair distances of KIX to c-Myb of MD macrostates respect the NMR structure— on the *x axis*, and mean helicity on the *y axis*. Node size is proportional to their equilibrium distribution. Edges represent the connection between macrostates, and their color and thickness, the percentage of the total flux traversing them. Nodes are colored blue, orange, or red depending on whether all their microstates belonged to the unbound or bound state in the 4-macro MSM, or if they have a mixed composition. Bulk states is circled in red.

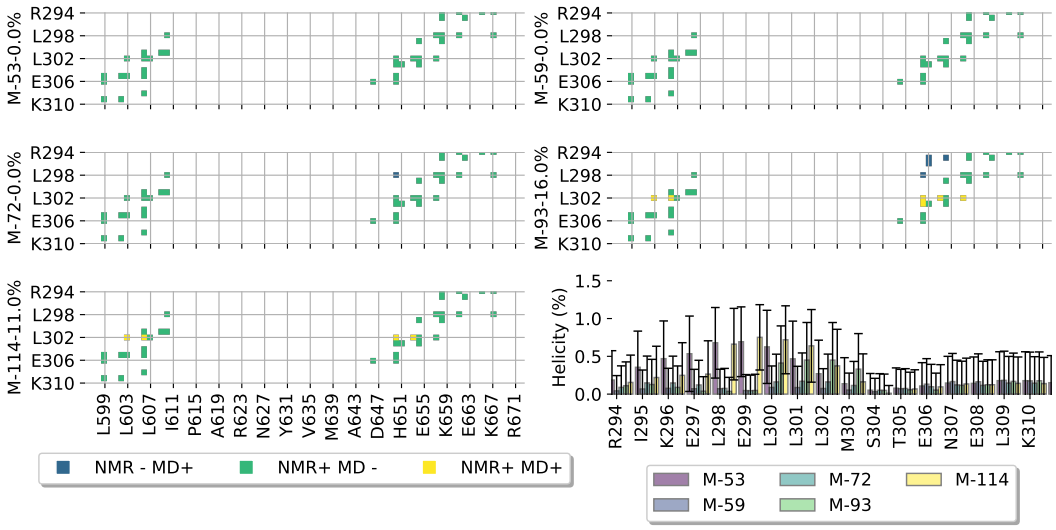

Fig. S8. Secondary binding path of KIX and c-Myb. Study of the states involved in the secondary binding pathway. For selected macrostates the **a)** KIX—c-Myb contacts profile and **b)** mean helicity is shown. Percentages on **a)** indicate the Q_{int}.

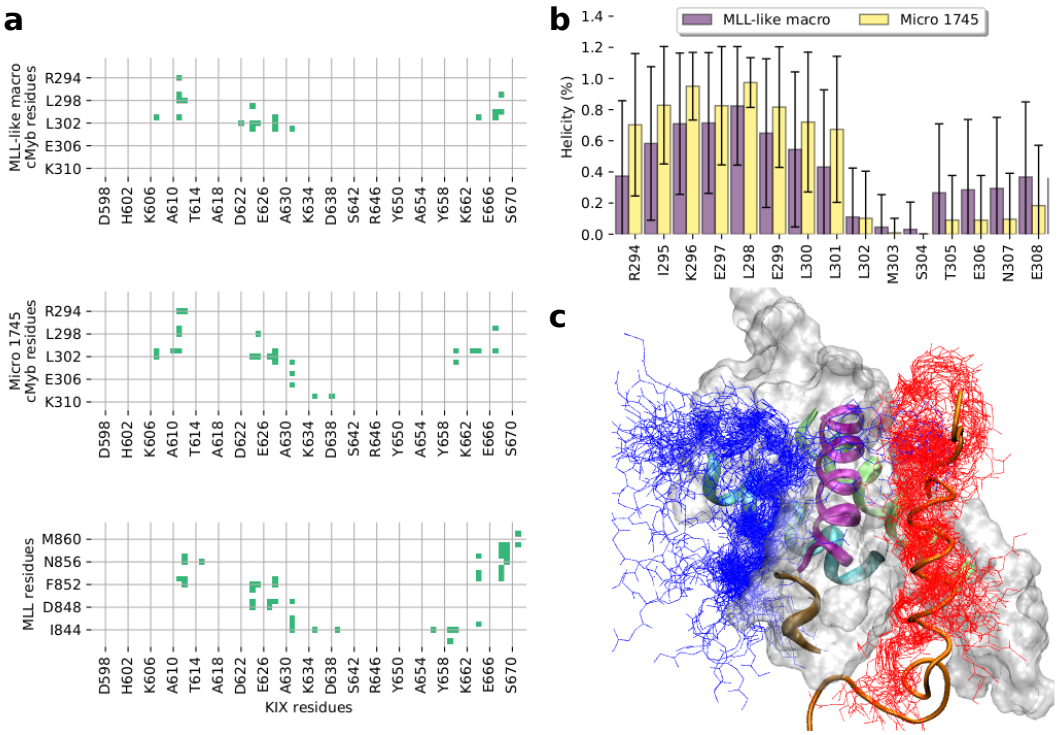

Fig. S9. MLL-like binding conformation of c-Myb to KIX. a) KIX—c-Myb contact profile for MLL-like macro and MLL-like microstate 1745 are shown. KIX—MLL contact profile taken form PDB *2AGH* in bottom panel. **b)** Mean by-residue helicity of c-Myb for MLL-like macrostates and for microstate 1745, which maximizes c-Myb's helicity. **c) Representative structures:** PDB structure *2AGH* is depicted with KIX as gray surface. Cartoon representations show α-1 in cyan, α-2 in purple, α-3 in green, and G2 (residues 617-621) in ochre. MLL is displayed as an orange ribbon. cMyb backbone for 50 MD superposed structures taken from the MLL-like 1745 microstates is shown in red, and from Misbound state (microstate 466) in blue.

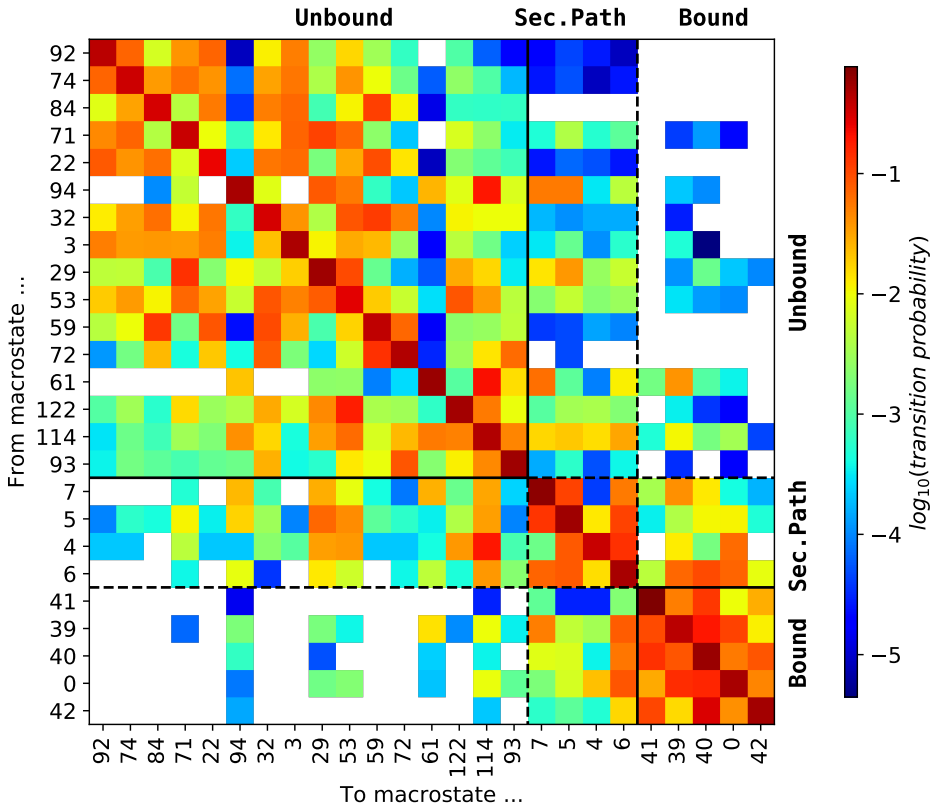

Fig. S10. Transition matrix for all macrostates involved in the binding process. *Y-axis* and *x-axis* indicate the starting and destination macrostates, respectively. Macrostates are ordered according, from top to bottom and left to right, to an increasing similarity to the experimental structure. Values represent the log10 of the probability of crossing from one macrostate to another.

3.3. Small molecule modulation of intrinsically disordered proteins using molecular dynamics simulations

Pablo Herrera-Nieto, Adrià Pérez, Gianni De Fabritiis. Submitted to *Journal of Chemical Information and Modeling*

Summary

Technological leaps in the field of MD in the last years allowed to reconstruct protein-ligand and protein-protein binding processes by letting them freely diffuse in multiple parallel simulations. However, given the limited amount of examples, similar studies have not been performed before in the context of IDPs. Here, we present the first free IDP-ligand binding study, using the system composed by the disordered protein p27 and the small molecule SJ403.

Small molecule modulation of intrinsically disordered proteins using molecular dynamics simulations

Pablo Herrera-Nieto,[†] Adrià Pérez ,[†] and Gianni De Fabritiis[*,†,‡,¶]

[†]*Computational Science Laboratory, Universitat Pompeu Fabra, Barcelona, Spain*

[‡]*Acellera Ltd., Barcelona, Spain*

[¶]*Institució Catalana de Recerca i Estudis Avançats, Barcelona, Spain*

E-mail: gianni.defabritiis@upf.edu

Abstract

The extreme dynamic behavior of intrinsically disordered proteins hinders the development of drug-like compounds capable of modulating them. There are several examples of small molecules that specifically interact with disordered peptides. However, their mechanisms of action are still not well understood. Here, we use extensive molecular dynamics simulations combined with adaptive sampling algorithms to perform free ligand binding studies in the context of intrinsically disordered proteins. We tested this approach in the system composed by the disordered protein p27 and the small molecule SJ403. The results show several protein-ligand bound states characterized by the establishment of a loosely oriented interaction mediated by a limited number of contacts between the ligand and critical residues of p27. Finally, protein conformations in the bound state are likely to be explored by the isolated protein too, therefore supporting a model where the addition of the small molecule restricts the available conformational space.

Introduction

The modulation of protein function via small molecules constitutes the main strategy for the development of novel pharmacological therapies. Amongst the many computational tools devoted to the subject, molecular dynamics simulations (MD) is one of the most powerful. In the last decade, MD boosted its performance thanks to dedicated supercomputers (ANTON[1]), and GPU-based technologies.[2] Together with more efficient sampling methodologies,[3–5] the exploration of increasingly longer timescales has become more accessible.[6,7] Such upgrades have progressively translated into more computationally demanding experiments like free ligand binding studies. In free ligand binding, protein and ligand are placed apart and, by launching multiple parallel simulations, spontaneous binding events are studied. This approximation allows investigating protein-ligand binding determinants and their mechanisms of action, as well as their kinetic parameters. Since the appearance of pioneer studies,[8,9] the approach has been employed in other challenges, such as protein-ligand binding,[10] protein-protein binding[11] and even in fragment screening.[12] Nowadays, it is routinely used for the exploration of complex energy landscapes needless of a priori biasing knowledge.

Similar studies have not been applied to the same extent with intrinsically disordered proteins (IDPs). IDPs are widely present in eukaryotic proteomes, where up to 30% of the proteins contain at least one disordered region.[13] They performed their biological functions despite lacking an stable three dimensional conformation[14,15] and they participate as crucial intermediates in complex signaling networks. Therefore, their malfunction may result in oncogenic and neurodegenerative pathologies.[16] Nevertheless, they are not generally considered as feasible drugs targets, as the most widely used methodologies are better suited for working with folded proteins.

Thus, the number of examples regarding IDP-small molecule modulation is limited, despite raising over the last few years. Several experimental approximations have been successfully used to identify small molecules interacting with disordered regions of proteins such as cMyc,[17] RNA helicase A,[18] NUPR1,[19] androgen receptor[20] and PTP1B.[21] Some computa-

tional studies in the field have proposed the so-called *"ligand clouds around protein clouds"* model,[22] which states that interaction takes place by fuzzy complexes rather than by rigid conformations.

Concerning this matter, on the most recent examples is the discovery, using NMR screening, of several small molecules (specially one referred as *SJ403*) that are able to interact with the $p27^{Kip1}$ protein[23] (Figure 1.a,b). p27 is an eukaryotic cell-cycle controller which blocks the function of the cycling-dependent kinase Cdk2—Cyclin A complex.[24] The binding model derived from experimental data proposed that the addition of a small molecule induces a population shift, instead of expanding the conformational space.[25] In particular, SJ403 is shown to disrupt the interaction between sets of hydrophobic residues, with some of them also being critical participants in binding.[25]

The amount of NMR data derived from this system, as well as the short length of the IDP (∼50 amino acids) and the chemical features of SJ403 (planarity and rigidity,), provide an ideal system for testing successful MD-based protein-ligand binding studies in the context of IDPs (Figure 1.a, b). In more detail, we study the binding between $p27^{kip1}$ and the small molecule SJ403 using full-atom unbiased MD simulations in combination with Markov state models (MSM) and novel adaptive sampling algorithms.[26] The identified bound conformations exhibit a mixture of binding-induced rigid regions with others retaining their fuzziness. Additionally, the comparison between protein-ligand bound conformations against isolated p27 reveals a reduction of the conformational space exploration and the disruption of intra-protein contacts, in line with experimental observations. Altogether, this exemplifies the unexplored potential of MD to understand the impact of small molecules in the energy landscape of IDPs.

Methodology

Molecular dynamics simulation set up

Initially, the N-terminal domain of p27 (residues 54 to 101) was built with PyRosetta in an extended conformation (release 197[27]). Next, a short implicit solvent simulation with Amber99sb[28] using OpenMM[29] was performed (100 000 steps at 550 K followed by 800 000 steps at 350 K, with the OBC water model) in order to compact the structure. Then, the peptide was simulated in explicit solvent simulations with the TIP3P water model[30] using CHARMM22*[31] for 60 ns at 500 K setting the dielectric to 80 to generate an initial pool of unfolded conformations of p27.

In parallel, the small molecule SJ403 was built with Maestro —version 10.7.2015[32]— and parameterized with Parameterize[33] (Supplementary Figure 3).

We performed two different types of experiments: one with the protein in solution (*p27* data set) and a second one with the protein and the ligand (*p27-SJ403* data set). For each of those, ten systems were constructed by placing one random, unstructured protein conformation —in conjunction with one ligand molecule, if needed— at a salt concentration of 5 mM. All systems were built using HTMD.[34] A Langevin integrator was used with a damping constant of 0.1 ps^{-1}. The integration time step was set to 4 fs, with heavy hydrogen atoms and holonomic constraints on all hydrogen-heavy atom bond terms. Electrostatics were computed using PME with a cutoff distance of 9 Å and grid spacing of 1 Å. Energy minimization was run for 500 steps. After energy minimization, equilibration for all systems was done in an NPT ensemble at 310 K, 1 atm, with heavy atoms constrained at 1 $kcal\ mol^{-1}\ \mathring{A}^2$ during 30 ns.

After equilibration, no proline *cis* isomers were detected. Production runs of 80 *ns* were performed at 310 *K* using the distributed computing project GPUGrid[35] running the ACEMD[2] simulation engine, following an adaptive sampling strategy.

AdaptiveBandit sampling

Adaptive sampling is designed to efficiently explore the conformational space by spawning new rounds of simulations based on the already generated data. After one epoch, MD data is analyzed, and a policy selects those structures of greater interest as initial conformations for a new set of simulations. The strategy employed here is called AdaptiveBandit (included in HTMD[34]), which frames MD sampling as a multi-armed bandit problem.[26] In the protein-ligand binding case studied here, we rewarded actions based on the stability of conformations using MSM estimations of the free energy for each state

$$\mathcal{R}_a =< k_B T \log(\mu(x)) >_{(a,x_1,\ldots,x_\tau)}, \tag{1}$$

where $\mu(x)$ is the equilibrium distribution estimated by the MSM with the current available data and the average is performed over the frames in the trajectory starting from a.

For both the *p27-SJ403* and *p27* experiments, AdaptiveBandit was initiated with 100 systems (every system was used ten times).

An MSM is constructed at each epoch using the available data. In the *p27-SJ403* run, the MSM was built on top of contacts between protein heavy atom vs. ligand heavy atoms at a threshold of 4 Å, in order to enhance the exploration of the conformations established between the two molecules. For the *p27* run, backbone atoms (Cα+Cβ) self distance contacts were employed instead, to optimize the conformational sampling of the IDP. For both runs, time-independent component analysis (TICA)[36] was used to reduce the number of dimensions to 5. Once the computation finished, the *p27-SJ403* data set included 6688 simulations, resulting in an aggregated simulation time of 535 μs. Finally, the *p27* data set was composed of 8900 simulations, corresponding to 712 μs.

Markov State Model analysis

MD data analysis was performed using HTMD.[34] Data featurization for the *p27-SJ403* data set was carried out using the distance between protein heavy atoms and SJ403 heavy atoms. For the *p27* data set, distances between C_α and side chain nitrogen and oxygen atoms were used. After that, a similar process was followed for building an MSM with each set. The featurized data was projected into a lower space of 3 dimensions by using TICA at a lag time of 20 frames. Afterwards, TICA derived data was clustered into 1.500 and 800 states using MiniBatchKMeans algorithm,[37] for the *p27-SJ403* and the *p27* data sets, respectively. Finally, microstates were fused into 3 macrostates (using the PCCA+ algorithm[38]) at a lag time of 20 ns for both sets, following the implied timescales plot (Supplementary Figure 4.a, c).

For every measure, errors were estimated by creating 10 independent MSMs using a random set containing 80% of the simulated trajectories.

Comparing MD runs

Once all simulations finished, we compared the distinctive exploration of the conformational landscape observed in the presence or absence of the small molecule. The approach followed to achieve it consisted on featurizing the MD trajectories from both data sets together, and then reducing the number of dimensions. The output combines all the data into a single space. In this space, it is possible to identify which regions belong to each data set, and allow to measure the overlap between them.

In particular, the scheme was applied several times, using two different metrics (dihedrals and protein contacts), and two different dimensionality reduction methods (TICA[36] and PCA[39]). The purpose of having all these combinations (4 in total) is to avoid possible biases induced either by a metric of by the dimensionality reduction method.

Results and Discussion

Identification and structural analysis of bound states

In an IDP-ligand binding scenario, both protein and protein-ligand conformations are highly variable. In essence, protein-ligand interactions are unlikely to be very stable, nor IDPs would become totally rigid upon binding. Therefore, both protein rigidity and protein-ligand orientation have to be assessed. In order to tackle this question, MD data derived from a protein-ligand adaptive run (*p27-SJ403* data set) were used to build an MSM based on distances between protein and ligand heavy atoms. The resulting MSM includes three long-lived (metastable) states: two of them (*M1* and *M2*) feature stable contacts, as shown in Figure 1.c. At the same time, the third one remains largely unfolded, without forming any stable protein-ligand contact.

In more detail, main residues involved in the interactions are located in the N-terminal half of p27 and include hydrophobic residues *W60*, *L70*, and *W76* (Figure 1.c), pointed as the main driving force in the interaction by experimental data.[23] They create an hydrophobic environment where the small molecule stacks between the aromatic rings or hydrophobic side chains of those residues (Figure 1.g). While residue-ligand contact provides a broad description of the interaction, it is not informative about their relative orientation. We computed protein-ligand spatial orientation as the contact probability between each residue and every heavy atom in the ligand (Figure 1.d). Results show a distinctive pattern between the two binding profiles. In *M1*, there are many highly probable atom-residue contacts, while in *M2*, protein residues interact more sparsely with ligand heavy atoms. Thus, in *M2*, residue-ligand atom contacts are not preserved and a looser p27-SJ403 conformation prevails.

On the other hand, we tracked macrostate flexibility as the deviation of backbone dihedral angles (Figure 1.f). All states share a reduced group of low-variation dihedral angles corresponding to section *90-RPPRPP-95* of p27, given its inherently rigid nature. Interactions with the small molecule increase protein rigidity, and the most rigid microstates integrate

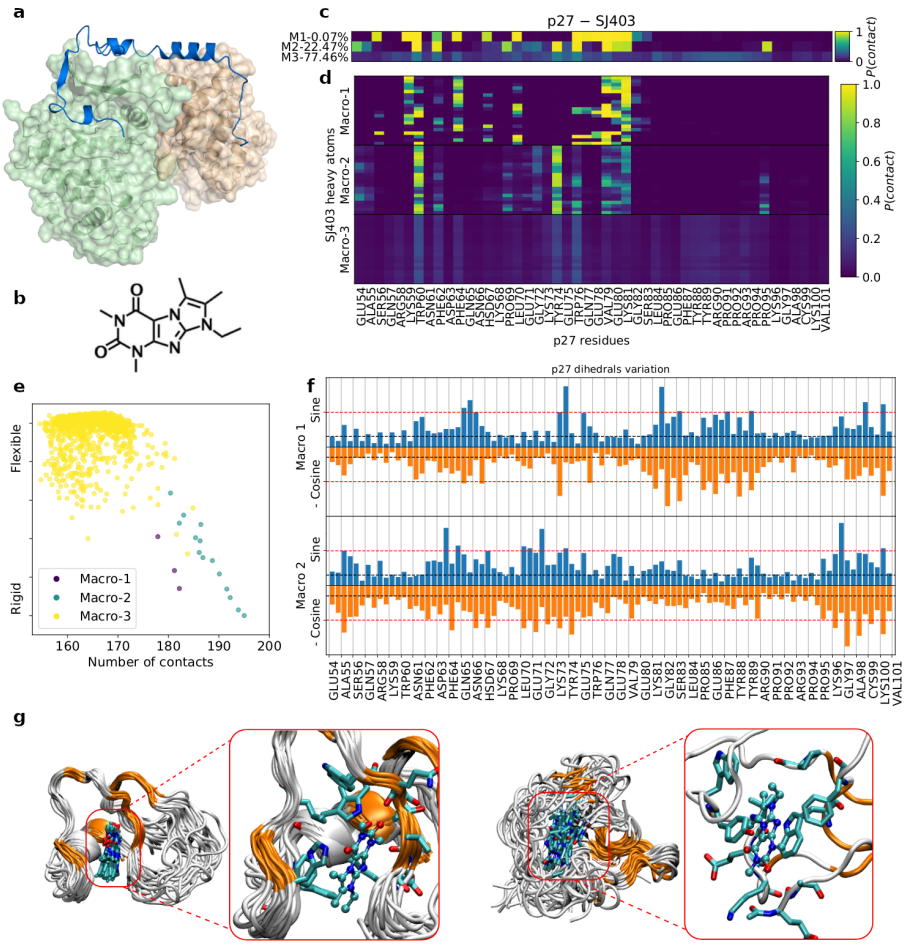

Figure 1: **System overview. a)** Kinase Cdk2, Cyclin A2, and p27 in complex. (PDB code: *1JSU*[24]). p27 is displayed as a blue ribbon. **b)** Chemical structure of SJ572403 (SJ403), a molecule able to bind to p27. **c) Residue-molecule** contact map. On the *y-axis* label the equilibrium population (%) for each state is also shown. **d) Residue-heavy atom** contact map. **e)** Microstate distribution according to the number of contacts between heavy atoms of the two molecules (*x-axis*), and their flexibility in the *y-axis*, measured as the sum of the deviation of backbone angles, and colored by macrostates assignment. **f) Backbone dihedral angle standard deviation** by state. Red and black dashed lines are placed at 0.5 and 0.1 respectively, for reference. **g) Representative structures** for macrostates 1 and 2. 20 random structures are displayed. Most rigid residues are colored in orange and their Cα were used for structural alignment.

the protein-ligand interacting states (Figure1.e). Additionally, the distinctive interaction profiles between *M1* and *M2* are also reflected in the induced backbone stability exerted by the ligand. Residues involved in the interaction become more rigid upon interaction in *M1* compared to *M2*, where most rigid residues do not participate in the interaction. Overall, contact probability and dihedral deviation reveal a more rigid bound conformation in *M1*, while *M2* features a more relaxed binding mode.

The measurement of equilibrium populations show that bound *M1* state is only marginally populated (0.06 ± 0.03 %), corresponding with an estimation for the k_{eq} of 11.93 ± 15.32 M. Such a low population and high affinity k_{eq} value are expected, given the high entropic penalty incurred by the extensive rigidity observed in this bound conformation. On the other hand, *M2* encompasses a significant portion of the population 24.23 ± 2.79 %, which corresponds with a 8.3 ± 1.3 mM in line with the experimental value (2.2 ± 0.3 mM).[23]

p27 in solution

An additional MD run was performed with the disordered p27 protein in solution (*p27* data set). The derived MSM included two macrostates featuring clusters of contacts between hydrophobic residues (*W60-Y88* and *W76-Y88*). They configure partially ordered states within the conformational surface of p27, and account for ∼2% of the equilibrium population (Figure 2). Interactions between these residues were also reported experimentally,[25] as part of a more intricate network of contacts between the triplet of hydrophobic residues *W60*, *W76*, and *Y88* that defines the ground state of p27 according to NMR data.[25] These two macrostates also encompass the most rigid structures (Supplementary Figure 1). However, in this case, there is not such pronounced segregation with respect to the most populated state as in the *p27-SJ403* data set (Figure 1.e). Furthermore, residues 69 and 90-95 keep a rigid conformation across all macrostates (Figure 2.b), similarly to *p27-SJ403* data set. Finally, a third highly populated and structurally heterogeneous state is also found.

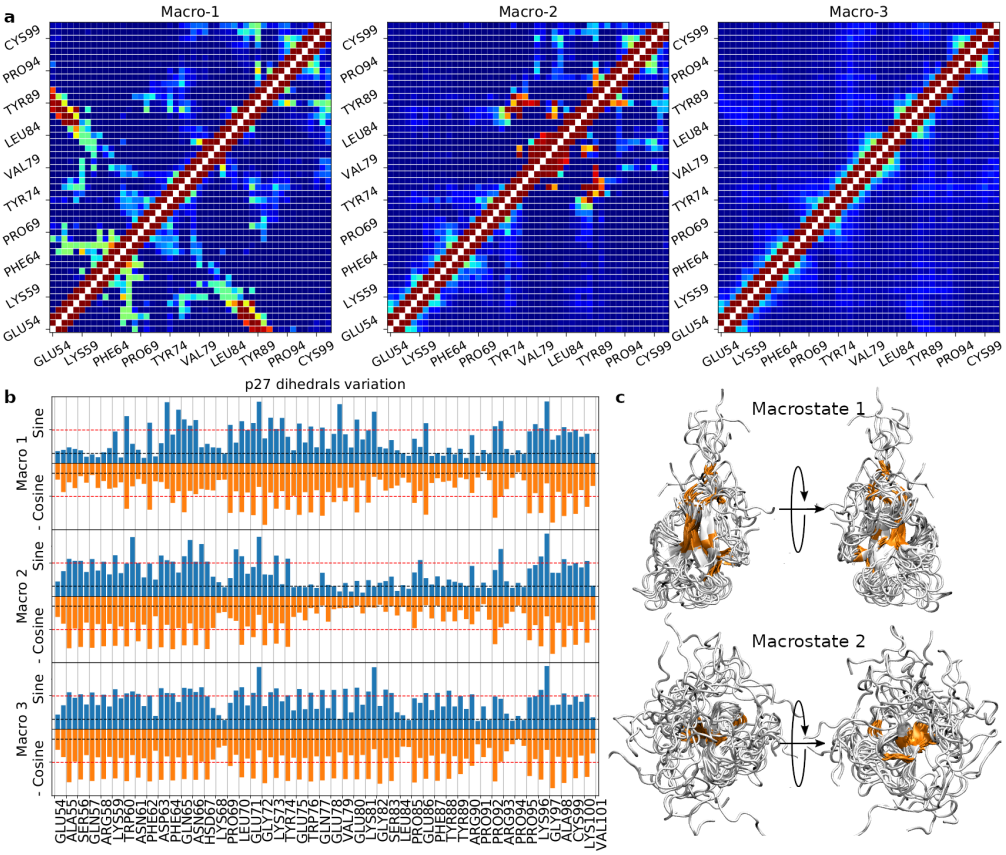

Figure 2: **Isolated system description. a) Residue-residue contact map**. **b) Backbone dihedral angle standard deviation** by macrostate. Red and black dashed lines are placed at 0.5 and 0.1 respectively, for reference. **c) Representative structures** for macrostates 1 and 2. 20 random structures are shown. Most rigid residues are colored in orange and their Cα were used for structural alignment.

Influence of small molecule addition in the conformational landscape

One of the main questions regarding IDP-ligand binding is how the addition of a molecule would impact the conformational landscape of p27. Two opposed scenarios are possible. In the first one, the molecule expands the pool of available structures, by promoting the exploration of unknown regions.[40] In the second one, the molecule restricts the populated areas of the conformational space, leading to a population shift between states.[25] For the

p27-SJ403 system, the experimental data support the latter;[25] it was observed that the addition of SJ403 disrupted the formation of intra-protein contacts, especially between the central and the C-terminal sections.

This question was also addressed using the available MD data. More specifically, the impact of SJ043 on the landscape was measured as the difference between the conformational exploration observed in the *p27-SJ403* data set and the *p27* one, as summarized in the *Comparison of MD run* section of the *Methodology* (Figure 3 and Supplementary Figure 2). We compared the exploration profiles between each data set. Firstly, we detect a high overlap between the space explored by each data set, indicating that they explore the same region. This overlap is expected, as only a small fraction of the data participates in either protein-protein or protein-ligand contacts. Secondly, *p27-SJ403* data points are contained within the *p27* space. The differences in the setup of each adaptive run can explain it. The *p27* run was designed to make a thorough exploration of the conformational space whereas the *p27-SJ403* run was done to incentivize protein-ligand conformations (a more detailed explanation can be found in the *AdaptiveBandit sampling* section on *Methodology*).

Moreover, data points of bound conformations are contained within the explored space by p27 in solution. Altogether, these observations align with the space exploration restriction scheme aforementioned: bound conformations are confined within regions also explored by the protein in isolation. The extension of theses areas is generally small compared to the entire space, consistent with their overall structural rigidity (Figure 1.e).

Finally, the comparison between regions corresponding to p27-SJ403 bound macrostates and partially ordered states of p27 in solution yields an almost complete separation profile. This implies that indeed they are structurally different and that the addition of the small molecule disrupts the formation of intra-protein contacts between those specific residues, as pointed by NMR data.[25]

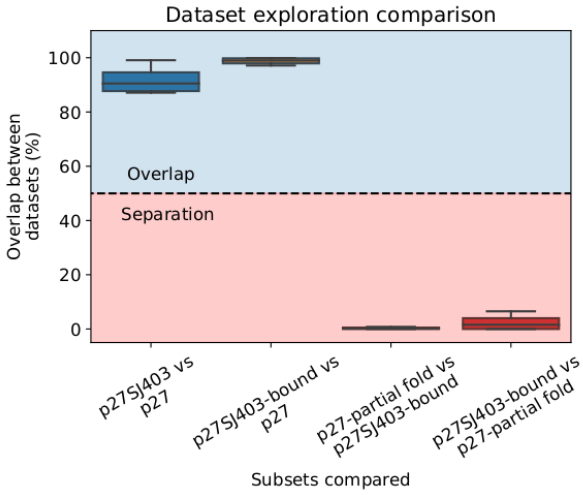

Figure 3: **Count overlap between subsets.** Overlap measures the percentage of counts of a subset (A) sharing common space with another subset (B). Spaces are calculated as described in *Comparing MD runs* section in *Methodology*. An horizontal dashed line at 50% represents an arbitrary boundary between sets occupying similar regions of the space (*overlapping*) or disconnected (*separated*).

Conclusion

The analysis of MD data in the context of IDP ligand modulation was able to provide p27-SJ403 bound conformations in line with experimental information. These protein-ligand complexes retain part of their intrinsic flexibility, while short linear stretches of amino acids become stiffer after binding. On the other hand, rigid protein conformations also appear. However, they do not raise a stable state as they incur in a high entropic cost that is not enthalpically compensated.

In the most populated bound state, residues W60 and W76 are crucial for the complex formation, but, in contrast to folded proteins, the pose is stabilized by multiple varying contacts. This is due to the hydrophobic nature of the interaction. For the most part, the stacking of the small molecule between tryptophan rings maintains the interaction. Typically, IDPs become fully rigid upon interacting with their folded partners by folding after binding. The main interactions driving such processes rely on the burial of hydrophobic residues

in predefined binding interfaces, as demonstrated by protein-IDP complexes included p53-Mdm2[41] or KIX-cMyb.[42] In other cases, such as polyelectrolyte chains, protein complexes are mediated by polar interactions and charge complementarity, and they do not require specific residue-residue contacts or binding sites,[43] thus remaining largely disordered upon interaction. Therefore, p27-SJ403 binding observed here represents an intermediate point between these two behaviors. It configures a scenario where the coupling of both molecules precise residue specific hydrophobic interactions while remaining structurally flexible.

Acknowledgements

The authors thank volunteers at GPUGRID.net for contributing with computational resources and Acellera for funding. G.D.F. acknowledges support from MINECO (Unidad de Excelencia María de Maeztu MDM-2014-0370 and BIO2017-82628-P) and FEDER. This project received funding from the European Union's Horizon 2020 Research and Innovation Programme under Grant Agreement 823712 (CompBioMed2 Project).

Author contributions statement

G.D.F. and P.H.N conceived the experiments, P.H.N and A.P. analysed the results. P.H.N wrote the manuscript. All authors reviewed the manuscript.

References

(1) Shaw, D. E.; Deneroff, M. M.; Dror, R. O.; Kuskin, J. S.; Larson, R. H.; Salmon, J. K.; Young, C.; Batson, B.; Bowers, K. J.; Chao, J. C., et al. Anton, a special-purpose machine for molecular dynamics simulation. *Communications of the ACM* **2008**, *51*, 91–97.

(2) Harvey, M. J.; Giupponi, G.; Fabritiis, G. D. ACEMD: accelerating biomolecular dynamics in the microsecond time scale. *Journal of chemical theory and computation* **2009**, *5*, 1632–1639.

(3) Doerr, S.; De Fabritiis, G. On-the-fly learning and sampling of ligand binding by high-throughput molecular simulations. *Journal of chemical theory and computation* **2014**, *10*, 2064–2069.

(4) Zimmerman, M. I.; Bowman, G. R. FAST conformational searches by balancing exploration/exploitation trade-offs. *Journal of chemical theory and computation* **2015**, *11*, 5747–5757.

(5) Zimmerman, M. I.; Porter, J. R.; Sun, X.; Silva, R. R.; Bowman, G. R. Choice of adaptive sampling strategy impacts state discovery, transition probabilities, and the apparent mechanism of conformational changes. *Journal of chemical theory and computation* **2018**, *14*, 5459–5475.

(6) Martínez-Rosell, G.; Giorgino, T.; Harvey, M. J.; de Fabritiis, G. Drug discovery and molecular dynamics: methods, applications and perspective beyond the second timescale. *Current topics in medicinal chemistry* **2017**, *17*, 2617–2625.

(7) Pérez, A.; Martínez-Rosell, G.; De Fabritiis, G. Simulations meet machine learning in structural biology. *Current opinion in structural biology* **2018**, *49*, 139–144.

(8) Buch, I.; Giorgino, T.; De Fabritiis, G. Complete reconstruction of an enzyme-inhibitor binding process by molecular dynamics simulations. *Proceedings of the National Academy of Sciences* **2011**, *108*, 10184–10189.

(9) Dror, R. O.; Pan, A. C.; Arlow, D. H.; Borhani, D. W.; Maragakis, P.; Shan, Y.; Xu, H.; Shaw, D. E. Pathway and mechanism of drug binding to G-protein-coupled receptors. *Proceedings of the National Academy of Sciences* **2011**, *108*, 13118–13123.

(10) Ferruz, N.; Doerr, S.; Vanase-Frawley, M. A.; Zou, Y.; Chen, X.; Marr, E. S.; Nelson, R. T.; Kormos, B. L.; Wager, T. T.; Hou, X., et al. Dopamine D3 receptor antagonist reveals a cryptic pocket in aminergic GPCRs. *Scientific reports* **2018**, *8*, 897.

(11) Plattner, N.; Doerr, S.; De Fabritiis, G.; Noé, F. Complete protein–protein association kinetics in atomic detail revealed by molecular dynamics simulations and Markov modelling. *Nature chemistry* **2017**, *9*, 1005.

(12) Martínez-Rosell, G.; Harvey, M. J.; De Fabritiis, G. Molecular-simulation-driven fragment screening for the discovery of new CXCL12 inhibitors. *Journal of chemical information and modeling* **2018**, *58*, 683–691.

(13) Oates, M. E.; Romero, P.; Ishida, T.; Ghalwash, M.; Mizianty, M. J.; Xue, B.; Dosztanyi, Z.; Uversky, V. N.; Obradovic, Z.; Kurgan, L., et al. D2P2: database of disordered protein predictions. *Nucleic acids research* **2012**, *41*, D508–D516.

(14) Van Der Lee, R.; Buljan, M.; Lang, B.; Weatheritt, R. J.; Daughdrill, G. W.; Dunker, A. K.; Fuxreiter, M.; Gough, J.; Gsponer, J.; Jones, D. T., et al. Classification of intrinsically disordered regions and proteins. *Chemical reviews* **2014**, *114*, 6589–6631.

(15) Wright, P. E.; Dyson, H. J. Intrinsically disordered proteins in cellular signalling and regulation. *Nature reviews Molecular cell biology* **2015**, *16*, 18–29.

(16) Uversky, V. N.; Oldfield, C. J.; Dunker, A. K. Intrinsically disordered proteins in human diseases: introducing the D2 concept. *Annu. Rev. Biophys.* **2008**, *37*, 215–246.

(17) Hammoudeh, D. I.; Follis, A. V.; Prochownik, E. V.; Metallo, S. J. Multiple independent binding sites for small-molecule inhibitors on the oncoprotein c-Myc. *Journal of the American Chemical Society* **2009**, *131*, 7390–7401.

(18) Erkizan, H. V.; Kong, Y.; Merchant, M.; Schlottmann, S.; Barber-Rotenberg, J. S.; Yuan, L.; Abaan, O. D.; Chou, T.-h.; Dakshanamurthy, S.; Brown, M. L., et al. A small molecule blocking oncogenic protein EWS-FLI1 interaction with RNA helicase A inhibits growth of Ewing's sarcoma. *Nature medicine* **2009**, *15*, 750.

(19) Neira, J. L.; Bintz, J.; Arruebo, M.; Rizzuti, B.; Bonacci, T.; Vega, S.; Lanas, A.; Velázquez-Campoy, A.; Iovanna, J. L.; Abián, O. Identification of a drug targeting an intrinsically disordered protein involved in pancreatic adenocarcinoma. *Scientific reports* **2017**, *7*, 39732.

(20) De Mol, E.; Fenwick, R. B.; Phang, C. T.; Buzón, V.; Szulc, E.; De La Fuente, A.; Escobedo, A.; García, J.; Bertoncini, C. W.; Est'ebanez-Perpiñ'a, E., et al. EPI-001, a compound active against castration-resistant prostate cancer, targets transactivation unit 5 of the androgen receptor. *ACS chemical biology* **2016**, *11*, 2499–2505.

(21) Krishnan, N.; Koveal, D.; Miller, D. H.; Xue, B.; Akshinthala, S. D.; Kragelj, J.; Jensen, M. R.; Gauss, C.-M.; Page, R.; Blackledge, M., et al. Targeting the disordered C terminus of PTP1B with an allosteric inhibitor. *Nature chemical biology* **2014**, *10*, 558.

(22) Jin, F.; Yu, C.; Lai, L.; Liu, Z. Ligand clouds around protein clouds: a scenario of ligand binding with intrinsically disordered proteins. *PLoS computational biology* **2013**, *9*, e1003249.

(23) Iconaru, L. I.; Ban, D.; Bharatham, K.; Ramanathan, A.; Zhang, W.; Shelat, A. A.; Zuo, J.; Kriwacki, R. W. Discovery of Small Molecules that Inhibit the Disordered Protein, p27 Kip1. *Scientific reports* **2015**, *5*, 15686.

(24) Russo, A. A.; Jeffrey, P. D.; Patten, A. K.; Massagué, J.; Pavletich, N. P. Crystal structure of the p27Kip1 cyclin-dependent-kinase inibitor bound to the cyclin A–Cdk2 complex. *Nature* **1996**, *382*, 325.

(25) Ban, D.; Iconaru, L. I.; Ramanathan, A.; Zuo, J.; Kriwacki, R. W. A small molecule causes a population shift in the conformational landscape of an intrinsically disordered protein. *Journal of the American Chemical Society* **2017**, *139*, 13692–13700.

(26) Pérez, A.; Herrera-Nieto, P.; Doerr, S.; Fabritiis, G. D. AdaptiveBandit: A multi-armed bandit framework for adaptive sampling in molecular simulations. 2020.

(27) Chaudhury, S.; Lyskov, S.; Gray, J. J. PyRosetta: a script-based interface for implementing molecular modeling algorithms using Rosetta. *Bioinformatics* **2010**, *26*, 689–691.

(28) Hornak, V.; Abel, R.; Okur, A.; Strockbine, B.; Roitberg, A.; Simmerling, C. Comparison of multiple Amber force fields and development of improved protein backbone parameters. *Proteins: Structure, Function, and Bioinformatics* **2006**, *65*, 712–725.

(29) Eastman, P.; Friedrichs, M. S.; Chodera, J. D.; Radmer, R. J.; Bruns, C. M.; Ku, J. P.; Beauchamp, K. A.; Lane, T. J.; Wang, L.-P.; Shukla, D., et al. OpenMM 4: a reusable, extensible, hardware independent library for high performance molecular simulation. *Journal of chemical theory and computation* **2012**, *9*, 461–469.

(30) Jorgensen, W. L.; Chandrasekhar, J.; Madura, J. D.; Impey, R. W.; Klein, M. L. Comparison of simple potential functions for simulating liquid water. *The Journal of chemical physics* **1983**, *79*, 926–935.

(31) 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.

(32) Schrödinger, N. Y. N., LLC Maestro Release 2017-4. 2017.

(33) Galvelis, R.; Doerr, S.; Damas, J. M.; Harvey, M. J.; De Fabritiis, G. A Scalable Molecular Force Field Parameterization Method Based on Density Functional Theory and Quantum-Level Machine Learning. *Journal of chemical information and modeling* **2019**, *59*, 3485–3493.

(34) Doerr, S.; Harvey, M.; Noé, F.; De Fabritiis, G. HTMD: high-throughput molecular dynamics for molecular discovery. *Journal of chemical theory and computation* **2016**, *12*, 1845–1852.

(35) Buch, I.; Harvey, M. J.; Giorgino, T.; Anderson, D. P.; De Fabritiis, G. High-throughput all-atom molecular dynamics simulations using distributed computing. *Journal of chemical information and modeling* **2010**, *50*, 397–403.

(36) Pérez-Hernández, G.; Paul, F.; Giorgino, T.; De Fabritiis, G.; Noé, F. Identification of slow molecular order parameters for Markov model construction. *The Journal of chemical physics* **2013**, *139*, 07B604_1.

(37) Pedregosa, F.; Varoquaux, G.; Gramfort, A.; Michel, V.; Thirion, B.; Grisel, O.; Blondel, M.; Prettenhofer, P.; Weiss, R.; Dubourg, V., et al. Scikit-learn: Machine learning in Python. *Journal of machine learning research* **2011**, *12*, 2825–2830.

(38) Röblitz, S.; Weber, M. Fuzzy spectral clustering by PCCA+: application to Markov state models and data classification. *Advances in Data Analysis and Classification* **2013**, *7*, 147–179.

(39) Blöchliger, N.; Caflisch, A.; Vitalis, A. Weighted distance functions improve analysis

of high-dimensional data: application to molecular dynamics simulations. *Journal of chemical theory and computation* **2015**, *11*, 5481–5492.

(40) Heller, G. T.; Sormanni, P.; Vendruscolo, M. Targeting disordered proteins with small molecules using entropy. *Trends in biochemical sciences* **2015**, *40*, 491–496.

(41) Kussie, P. H.; Gorina, S.; Marechal, V.; Elenbaas, B.; Moreau, J.; Levine, A. J.; Pavletich, N. P. Structure of the MDM2 oncoprotein bound to the p53 tumor suppressor transactivation domain. *Science* **1996**, *274*, 948–953.

(42) Zor, T.; De Guzman, R. N.; Dyson, H. J.; Wright, P. E. Solution structure of the KIX domain of CBP bound to the transactivation domain of c-Myb. *Journal of molecular biology* **2004**, *337*, 521–534.

(43) Borgia, A.; Borgia, M. B.; Bugge, K.; Kissling, V. M.; Heidarsson, P. O.; Fernandes, C. B.; Sottini, A.; Soranno, A.; Buholzer, K. J.; Nettels, D., et al. Extreme disorder in an ultrahigh-affinity protein complex. *Nature* **2018**, *555*, 61.

Supplementary Information: Small molecule modulation of intrinsically disordered proteins using molecular dynamics simulations

Pablo Herrera-Nieto[a], Adrià Pérez[a], and Gianni De Fabritiis[a,b]

[a]Computational Science Laboratory,Barcelona biomedical research park (PRBB), Universitat Pompeu Fabra, C Dr Aiguader 88, Barcelona 08003, Spain
[b]Institució Catalana de Recerca i Estudis Avançats (ICREA), Passeig Lluis Companys 23, 08010 Barcelona, Spain

April 14, 2020

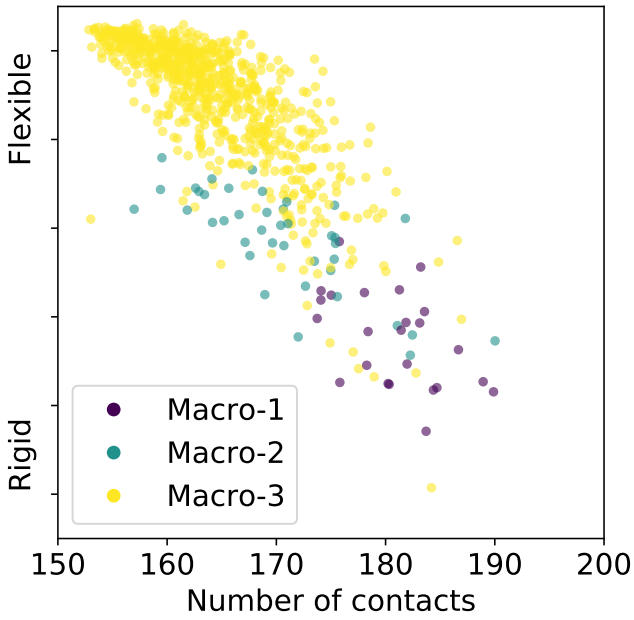

Figure 1: **p27 MSM microstate distribution** according to the number of contacts *x-axis*, flexibility *y-axis* measured as the sum of the deviation of backbone angles, and colored by macrostates assignment.

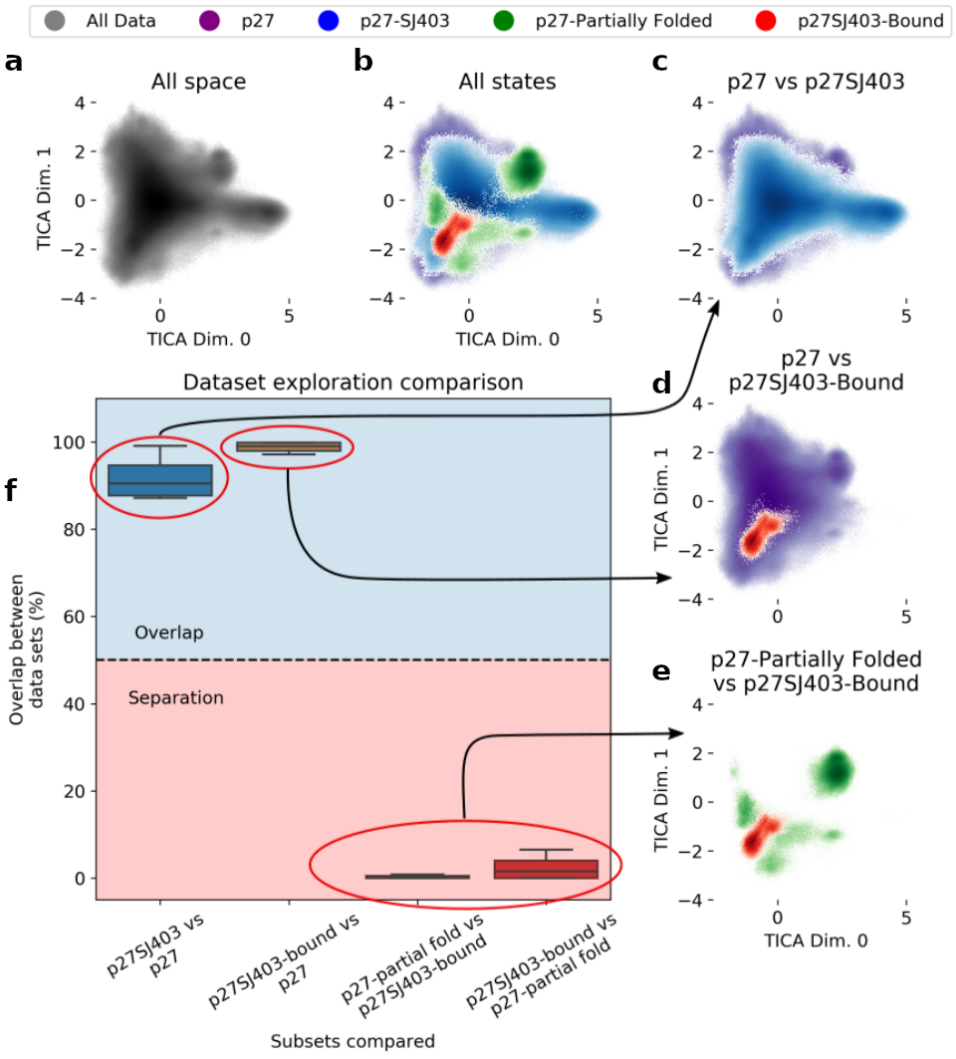

Figure 2: **Overlap calculation**. **a) - e) Counts representation** on the first two TICA dimensions of the dihedral data of the combined data set. They exemplified the overlap or separation between both data sets or part of them. They represent: a) all counts, b) counts colored by macrostates, c) counts of *p27SJ403* and *p27* dataset, d) *p27* data set and protein-ligand bound states, and e) partially folded states and protein-ligand bound states. **f) Overlap measurements** were performed with 4 different metric of three dimensions each (for reference purposes, a-e plots only show the first two dimensions of one of them) and measure the percentage of counts of a given subset found within the boundaries of another one. Arrows relate the overlap percentage to the count plot it represents.

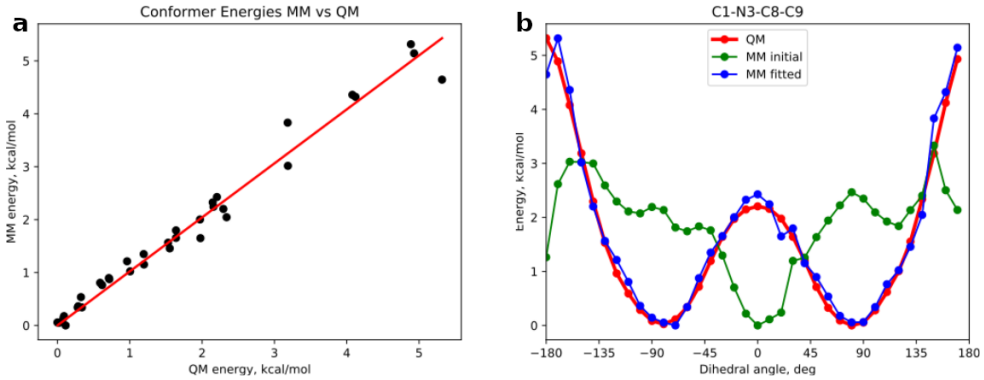

Figure 3: **Molecule parametrization. a) Conformer energies. b) Fit of dihedral angles energy profile.** Comparison of the quantum mechanics (in red) energy profile and the molecular mechanics before (in green) and after (in blue) being fitted.

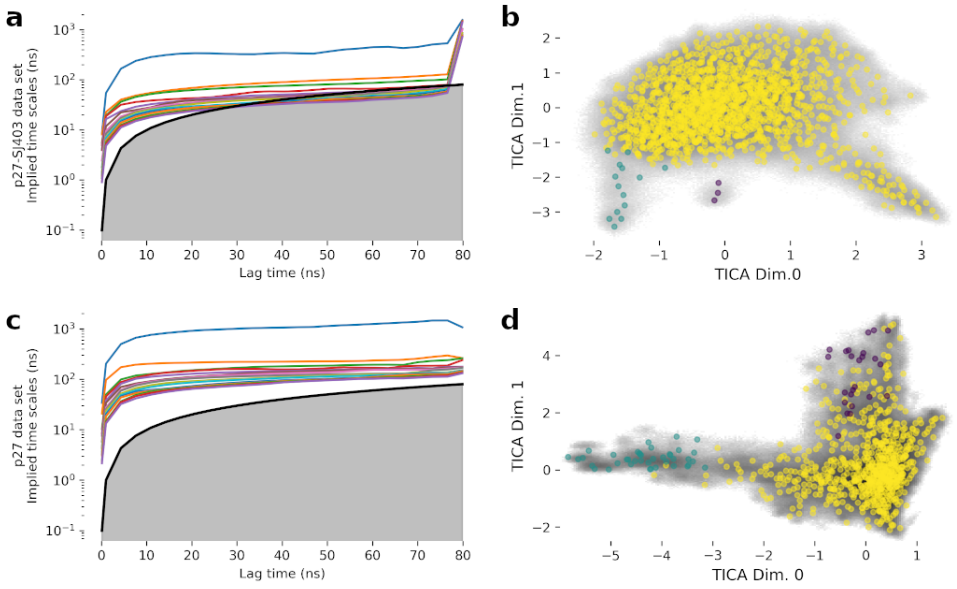

Figure 4: **Models summary a) and c) Implied times scales** for the *p27SJ403* and the *p27* data set, respectively. **b) and d)** Microstates distribution over the first two TICA dimensions.

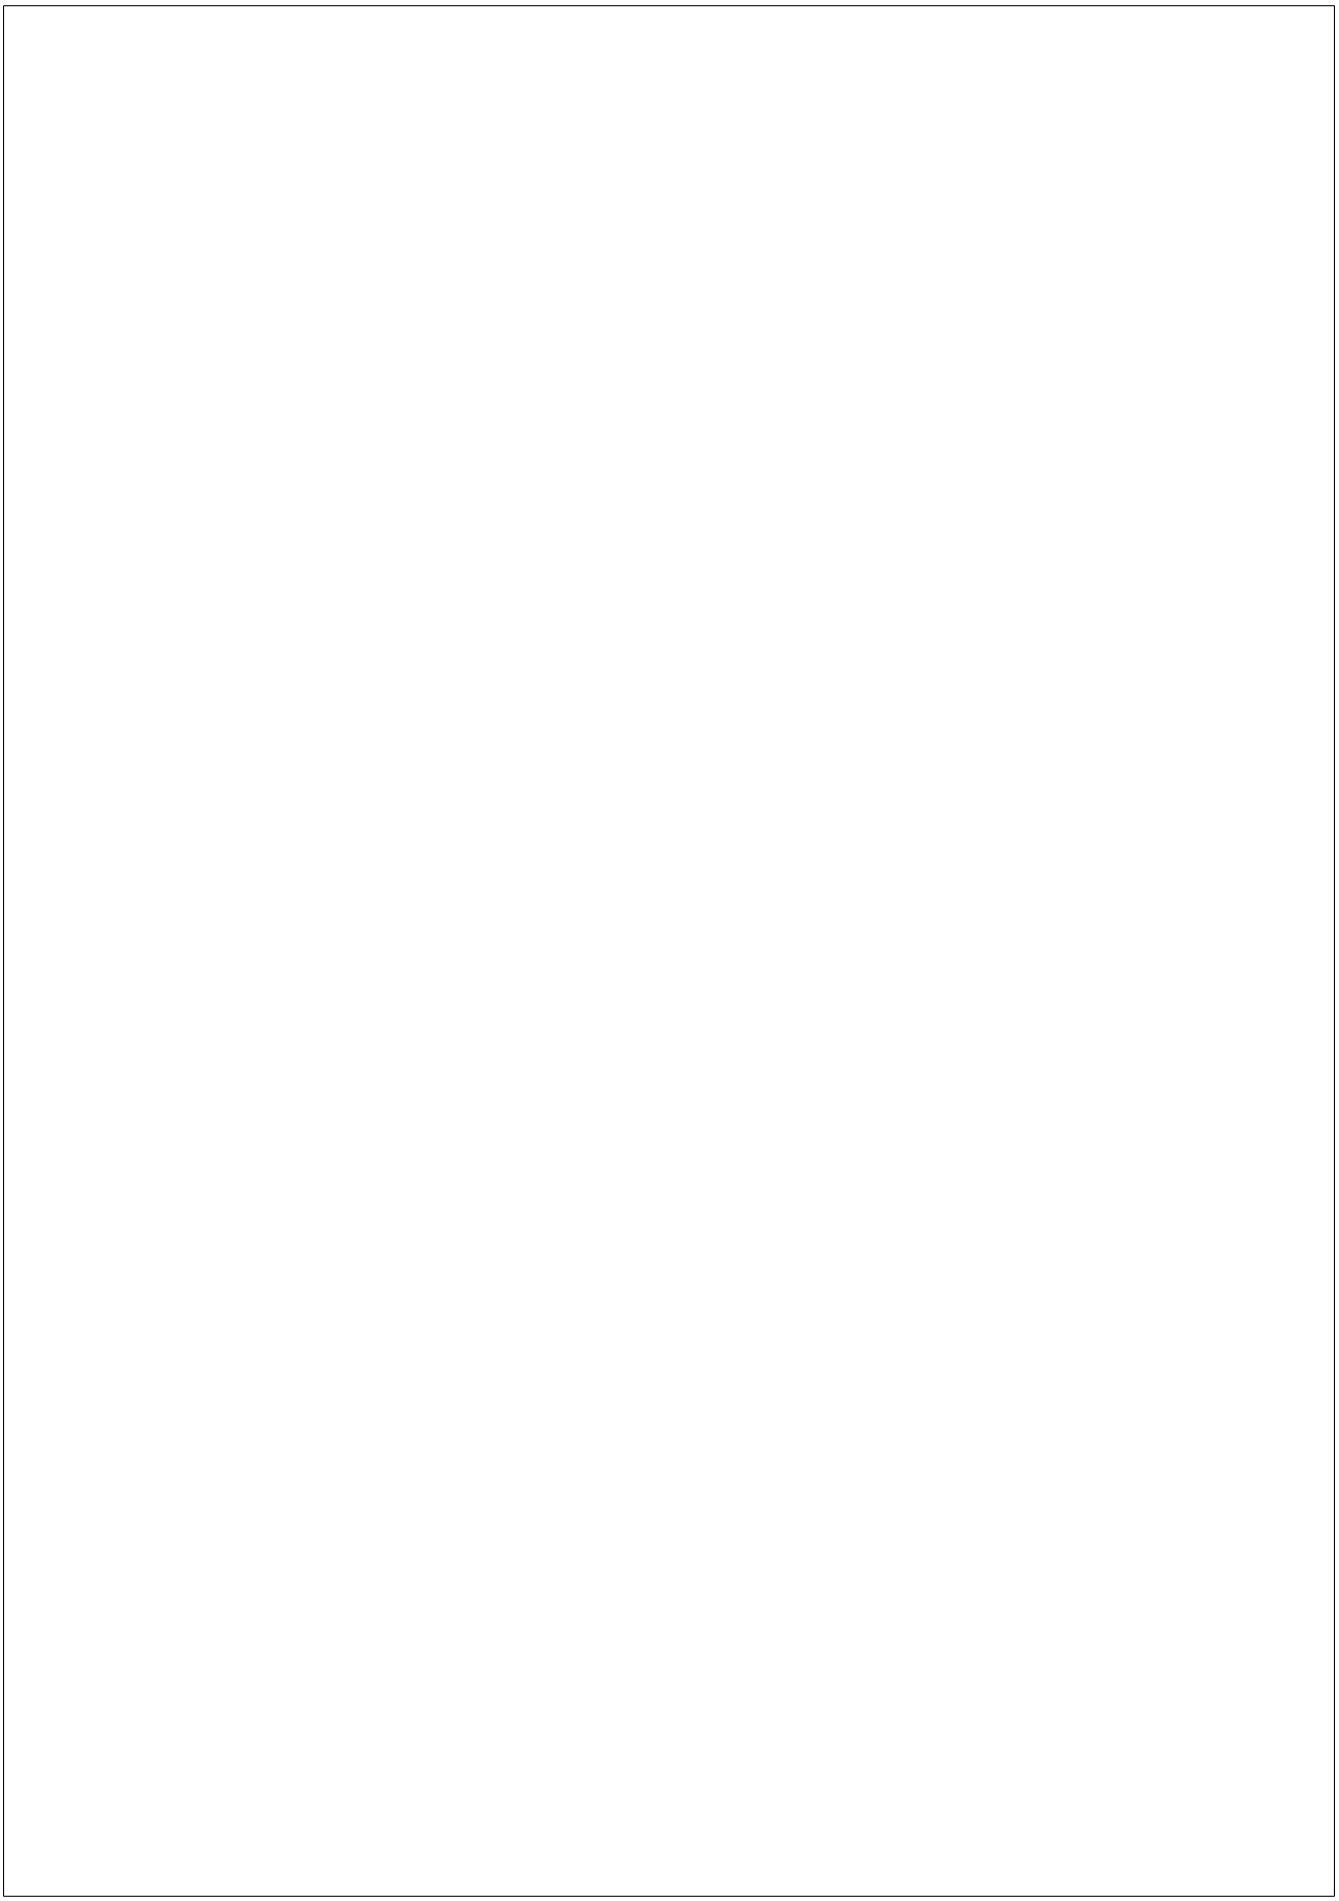

Chapter 4

DISCUSSION

4.1. Partially ordered states within IDPs

The first study presented in the *Results* section focuses on the exploration of the conformational landscape of an IDP in solution. The presence of partially ordered states has been proposed by other studies. However, it has never been exhaustively explored employing MD simulations. Here, by using high throughput MD simulations, it was possible to propose several of such partially ordered states that accounted for a significant proportion of the population once at equilibrium.

Scaling up a similar experiment in a more general project would be needed whether such behavior is other IDPs. One step in that direction is to use more advanced sampling techniques, such as the ones covered in *Section 1.2.3*. The extended study included $\sim$20 IDPs. They were selected based mainly on two criteria. Firstly, their prediction of disorder and, secondly, the existence of a similar sequence (in complex with other proteins) in the PDB.

This approach was also able to identify similar secondary structure enriched segments within several IDPs, as illustrated in Figure 4.1. The N-terminal region of p53 was also included in the selection in order to

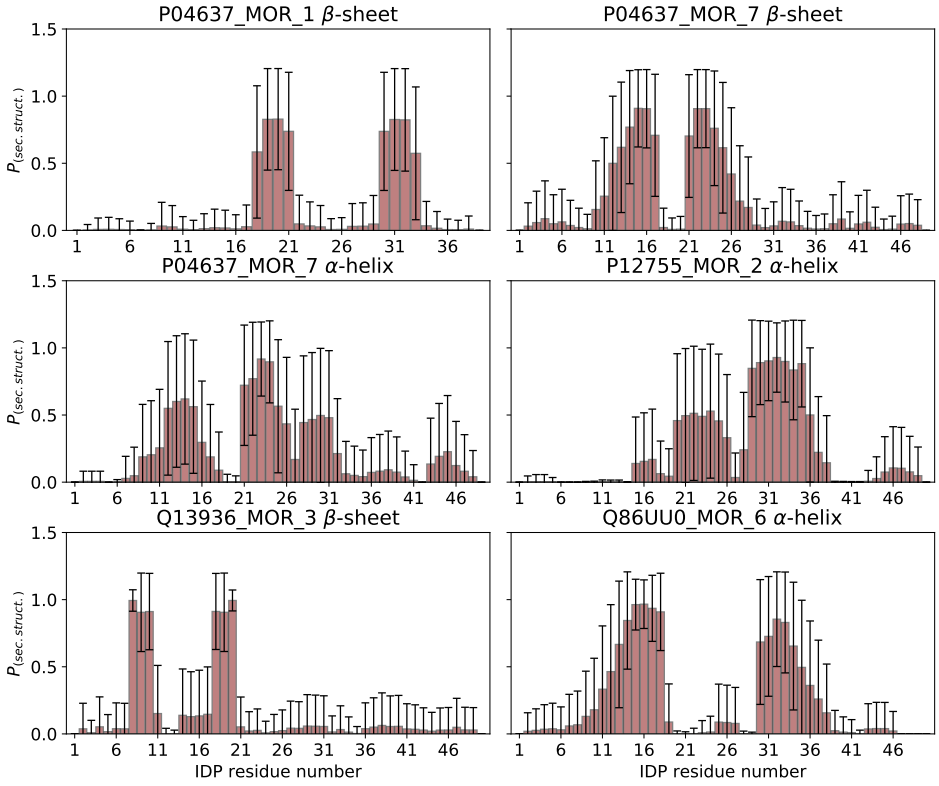

Figure 4.1: **Secondary structure profiles** illustrating the emergence of potential partially ordered states within IDPs using high throughput MD in combination with MSMs. Each plot represents the by-residue secondary structure probability (type in the title).

reproduce previous results. For this particular system, similar β-stranded states were also found (Figure 4.1 top left). Finally, in general, all partially states discovered so far are mediated by short-range contacts (mainly sheets or helices). In contrast, other scenarios, such as the formation of

rigid conformations exclusively by long-range contacts, are not observed.

There are two significant points to be considered when expanding this principle to a bigger scale. The first one is how to ensure the biological relevance of newly discovered partially ordered states rather than artifacts derived from the force fields. Exhaustive validation with experimental information could be a way to address this issue. The second most limiting factor is the computational resources needed to simulate systems with increasingly longer IDPs. Both the original and the expanded study are restricted to work with a maximum protein length of 50 amino acids, but experimental methods can extend to the hundred residues range.

4.2. KIX—c-Myb sampling problem

Reconstruction of binding events is one of the most significant challenges faced by MD in recent years. There have been many success studies focus on protein-ligand (80; 106), protein-protein (84), and protein-IDP binding (93; 95; 107). For the latter, most of the attention has been gathered by the p53-Mdm2 system, which includes a relatively short IDP of 12 amino acids, and a single well-characterized bound conformation. In that context, MD studies have especially focused on the reconstruction of the p53-Mdm2 bound conformation, using that system as a benchmark for novel simulation methodologies (93; 95; 94; 107). Here, by choosing KIX—c-Myb, the focus was to discriminate binding mechanisms, as well as to reconstruct several binding poses using full atom unbiased MD. This has represented a significant challenge throughout the thesis. In essence, out of a total simulation time of $\sim 4\ ms$ invested in several experiments, only one data set of $\sim 500\ \mu s$ was able to reconstruct the binding process partially.

As a summary, the results of three different sampling algorithms (AdaptiveBandit, ϵ-greedy (49), and FAST (83; 108)) are shown in Figure 4.2. AdaptiveBandit is inspired by RL and the multi-armed bandit problem and was covered in *Section 1.1.4*. ϵ-greedy is a heuristic method based on spawning new simulations from less explored states. Finally, FAST make use of prior knowledge in the form of a goal function that ranks all structures.

Data derived from each run were used to build its MSM. The three of them were compared in terms of exploration of the bound region of the conformational space (defined as that with lowest C_α RMSD mean and deviation). After 225.0 μs, only the AdaptiveBandit run was able to explore the region of NMR-like bound conformations. ϵ-greedy performs a reduced exploration of the space, getting stuck on an RMSD around 7 Å.

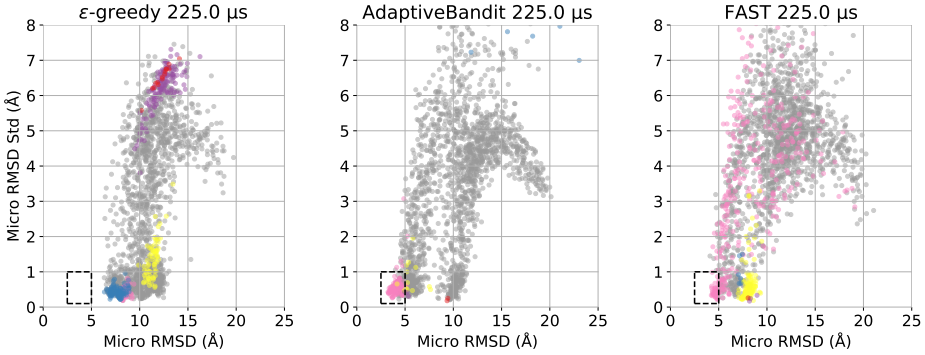

Figure 4.2: **KIX—c-Myb binding models.** Performance on the exploration of the conformational space of KIX—c-Myb biding of **a)** ϵ-greedy, **b)** AdaptiveBandit, and **c)** FAST. Both AdaptiveBandit and ϵ-greedy were run without any prior knowledge, and the goal function used to guide FAST minimize the euclidean distance between the C_α-C_β distance vector between MD and NMR structures. All MSMs were built using the pairwise distances between $C_\alpha + C_\beta$ between KIX and c-Myb and self distances of C_α of c-Myb, 2 000 clusters at 30 ns lag time and 6 macrostates.

In FAST, the exploration of the bound-like region is limited and includes many structurally diverging microstates.

This exemplifies how it is possible to overcome the restrictions on the sampling of complex energy surfaces with the right tools. In particular, the study of the KIX—c-Myb system illustrates how the evolution of sampling approaches affects performance, where more advanced approximations succeed in tasks where the old ones could not.

However, although better than the previous ones, the MSM presented in *Section 3.2* does not yield kinetic parameters in agreement with experimental data. In particular, the most diverging value is found in k_{off}, with

a difference of several orders of magnitude. Such discrepancy points to the lack of stability of the bound conformation obtained, which might be closer to a *pre bound* conformation rather than the NMR fully bound structure. Whether these deviations are due to MSM building issues, sampling problems, force field inaccuracies, or even if they represent more precisely the binding at physiological temperature is still a matter of discussion.

4.3. KIX—c-Myb: MD vs. experimental data

The analysis presented in *Section 3.2* provides a detailed molecular description of binding of c-Myb to the primary interface of KIX. The process is summarized as a two-step process. Initially, the N-terminal region of c-Myb binds with a preferred helical conformation, allowing the formation of native contacts and, in the last step, the C-terminal folds and binds. The study of the fluxes derived from the MSM showed the relevance of residue Leu302 not only in the final bound structure but also as the responsible of establishing the first contacts and serving as an anchor point between c-Myb and KIX. The first identified contact features Leu302 of c-Myb and KIX's residue Leu603 and the next steps comprise the formation of additional native contacts involving Leu298 and Leu302. The role of Leu298 and Leu302 as the driving force for the interaction was originally described in the report on the NMR structure (15). Additionally, mutagenesis studies showed the abolition of binding upon their mutation to alanine (98).

One of the main questions is whether coupled folding and binding takes place via conformational selection or induced-fit. For KIX—c-Myb, there is experimental evidence supporting the latter: the use of α-helix stabilizers did not affect binding rates (103) as well as helicity stabilizers mutations (98; 109). According to our model, conformational selection would only affect residues 298 - 302. Therefore, mutations outside this region would not have a high impact on binding kinetics despite affecting global helicity. This is the case of the flanking proline residues Pro289 and Pro316, which upon mutation to alanine, show a significant increase in helicity, but do not impact affinity (98; 109). On the other hand, disruption of the helicity of residues 298-302 should negatively affect binding. On the available experimental information (98; 109), we can find that L300P and L300G, have a decrease in helicity and 10 times slower k_{on} indeed, but L300A has little impact on these parameters.

A similar model to the one described above was first proposed by *Arai et al.* (100). It also pointed out the influence of c-Myb’s helicity and the combination of both binding mechanisms. Finally, binding before folding is also observed in our MSM. It accounts for a smaller fraction of the total flux, thus establishing an induced-fit only pathway that coexists with conformational selection.

In this way, the MSM allows us to reinterpret experimental data in the light of a atomically detailed model.

4.4. Druggability of IDPs

Modulation of protein by small molecules is ongoing problem faced in structural biology and by pharmaceutical companies. This task is even more problematic when the target does not have a well-defined pocket, as in the case of IDPs. A proof of this is the limited number of examples found in the literature covering this topic (and recently reviewed in (110)). They add up to a few tens compared to the hundreds devoted to rational drug design every year. Thus, spotting a feasible system for free ligand binding MD studies like the one shown in *Publication 3.3* is complicated. In the first place and as discussed above, one of the most significant limitations is protein length. To simulate a system big enough to fit an IDP longer than 50 amino acids would be too computationally expensive to be performed in the micro to millisecond time scale. At the same time, some of the IDPs discussed in the literature have hundreds of residues. Secondly, experimental information, such as the molecular determinants of the interaction (i.e., primary residues involved) or kinetic parameters, is also needed to understand and validate our data. So far, the only system to fulfill such features is p27—SJ403.

On the other hand, performing MD free ligand binding assays with IDPs can also be seen as a natural next step in the application of this technique in unexplored contexts. It initially raised many questions that are not found when working with folded proteins. In particular, how to simultaneously assess and characterize both IDP-ligand interaction and IDP structure. This was done at several levels: residue-molecule and residue-heavy atom contacts, dihedral variation, and intra-protein contacts. Additionally, the other principal question was how to relate simulations of the isolated protein with those performed in the presence of the ligand, a pivotal point to evaluate the impact of small molecule binding. By combining MD data from various data sets, it was possible to approximate this problem and

to draw observations in line with NMR data regarding differences in the behavior of the IDP in the presence/absence of the ligand.

Studies of this kind also arise similar concerns to the ones exposed in *Section 4.1* regarding force field capabilities to accurately capture the behavior of IDPs. For the particular case of the interaction between p27-SJ403, bound conformations also correspond with the most rigid ones. In contrast, NMR data (34) suggests that binding increases protein flexibility, as residues near Y88 free upon protein-ligand interaction. Specially adapted force fields for IDP simulations, such as the one employed throughout all studies, may also over compact proteins. This issue, however, does not rule out our results as there are many points in common with experimental data, too, mainly the identification of residues involved in the interaction. Testing different force fields for IDP simulations is out of the scope of the present thesis. However, advances in this area are also fundamental. In essence, it would be pointless to have state-of-the-art sampling and analysis capabilities paired with suboptimal parameters.

Once again, IDP modulation is still a field to be defined. Here, the main objective was to explore the performance of an existing successful computational methodology and to put it into practice. Overall, results provide a promising starting point. They agree with the experimental data available and can capture the main features retained by an IDP when interacting with a small molecule. However, adapting other pipelines for drug design steps, especially hit discovery and optimization, both computational and experimental, remains as the needed foundations to settle the field of IDP modulation.

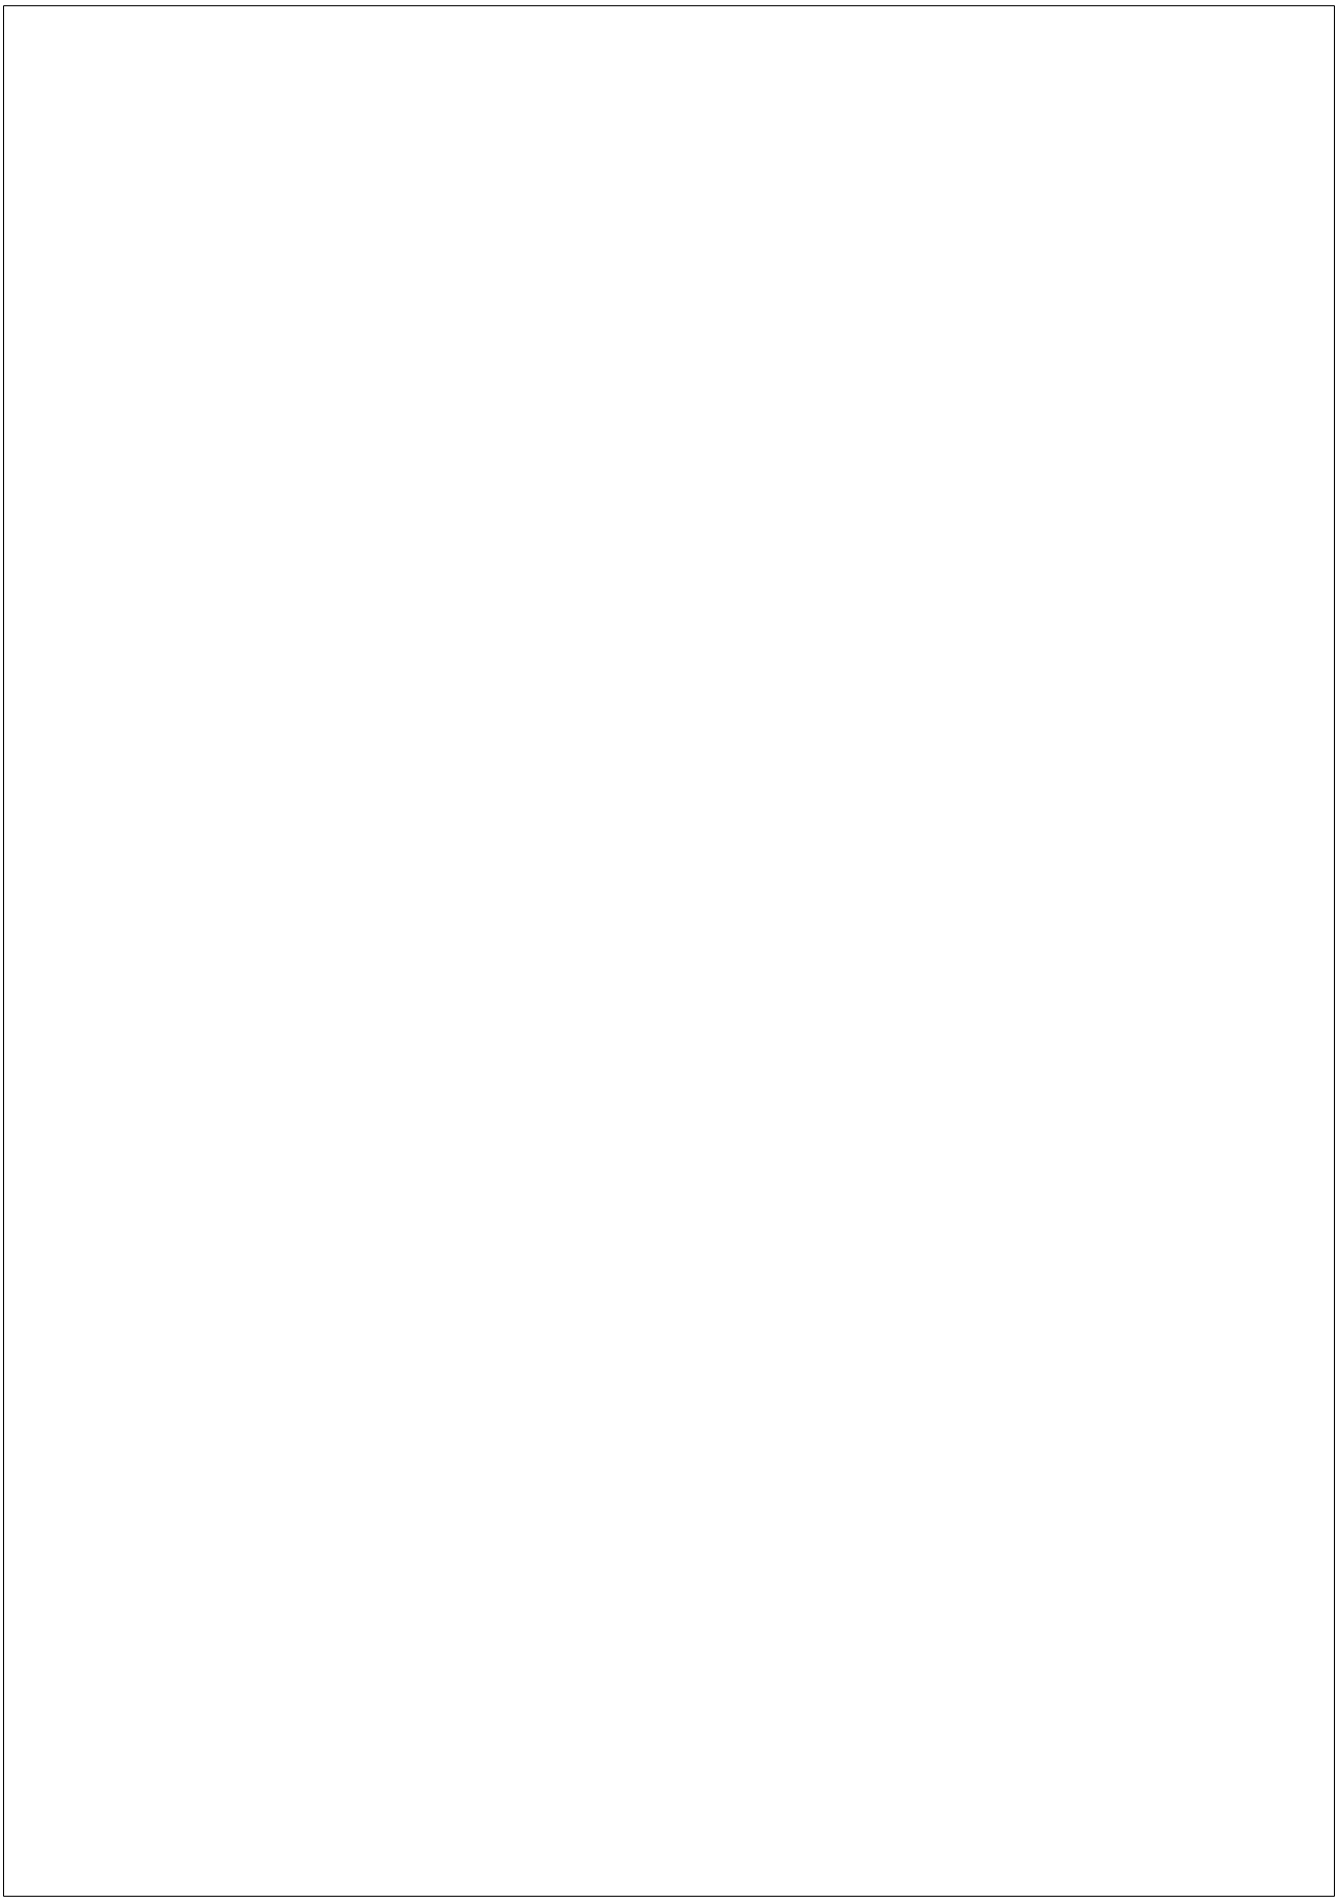

Chapter 5

CONCLUSIONS

1. High-throughput MD simulations, in combination with Markov state models, revealed partially ordered states within the conformational landscape of the disordered N-terminal region of p53.

2. Novel adaptive sampling algorithms can reconstruct the coupled folding and binding between the folded protein KIX and its disordered partner c-Myb.

3. c-Myb binds KIX in a two-steps mechanism. In the first place, the N-terminal of c-Myb binds either as a pre-folded helix via conformational selection or by induced fit. In either case, the central region is the first to established key native contacts. Lastly, binding of the C-terminal tail occurs via induced fit.

4. Free ligand binding studies can be successfully applied in IDP-ligand scenarios. Protein-ligand bound states feature residues specific interactions while retaining part of their intrinsic flexibility. Binding of the small molecule is restricted to a portion of the total protein conformational pool.

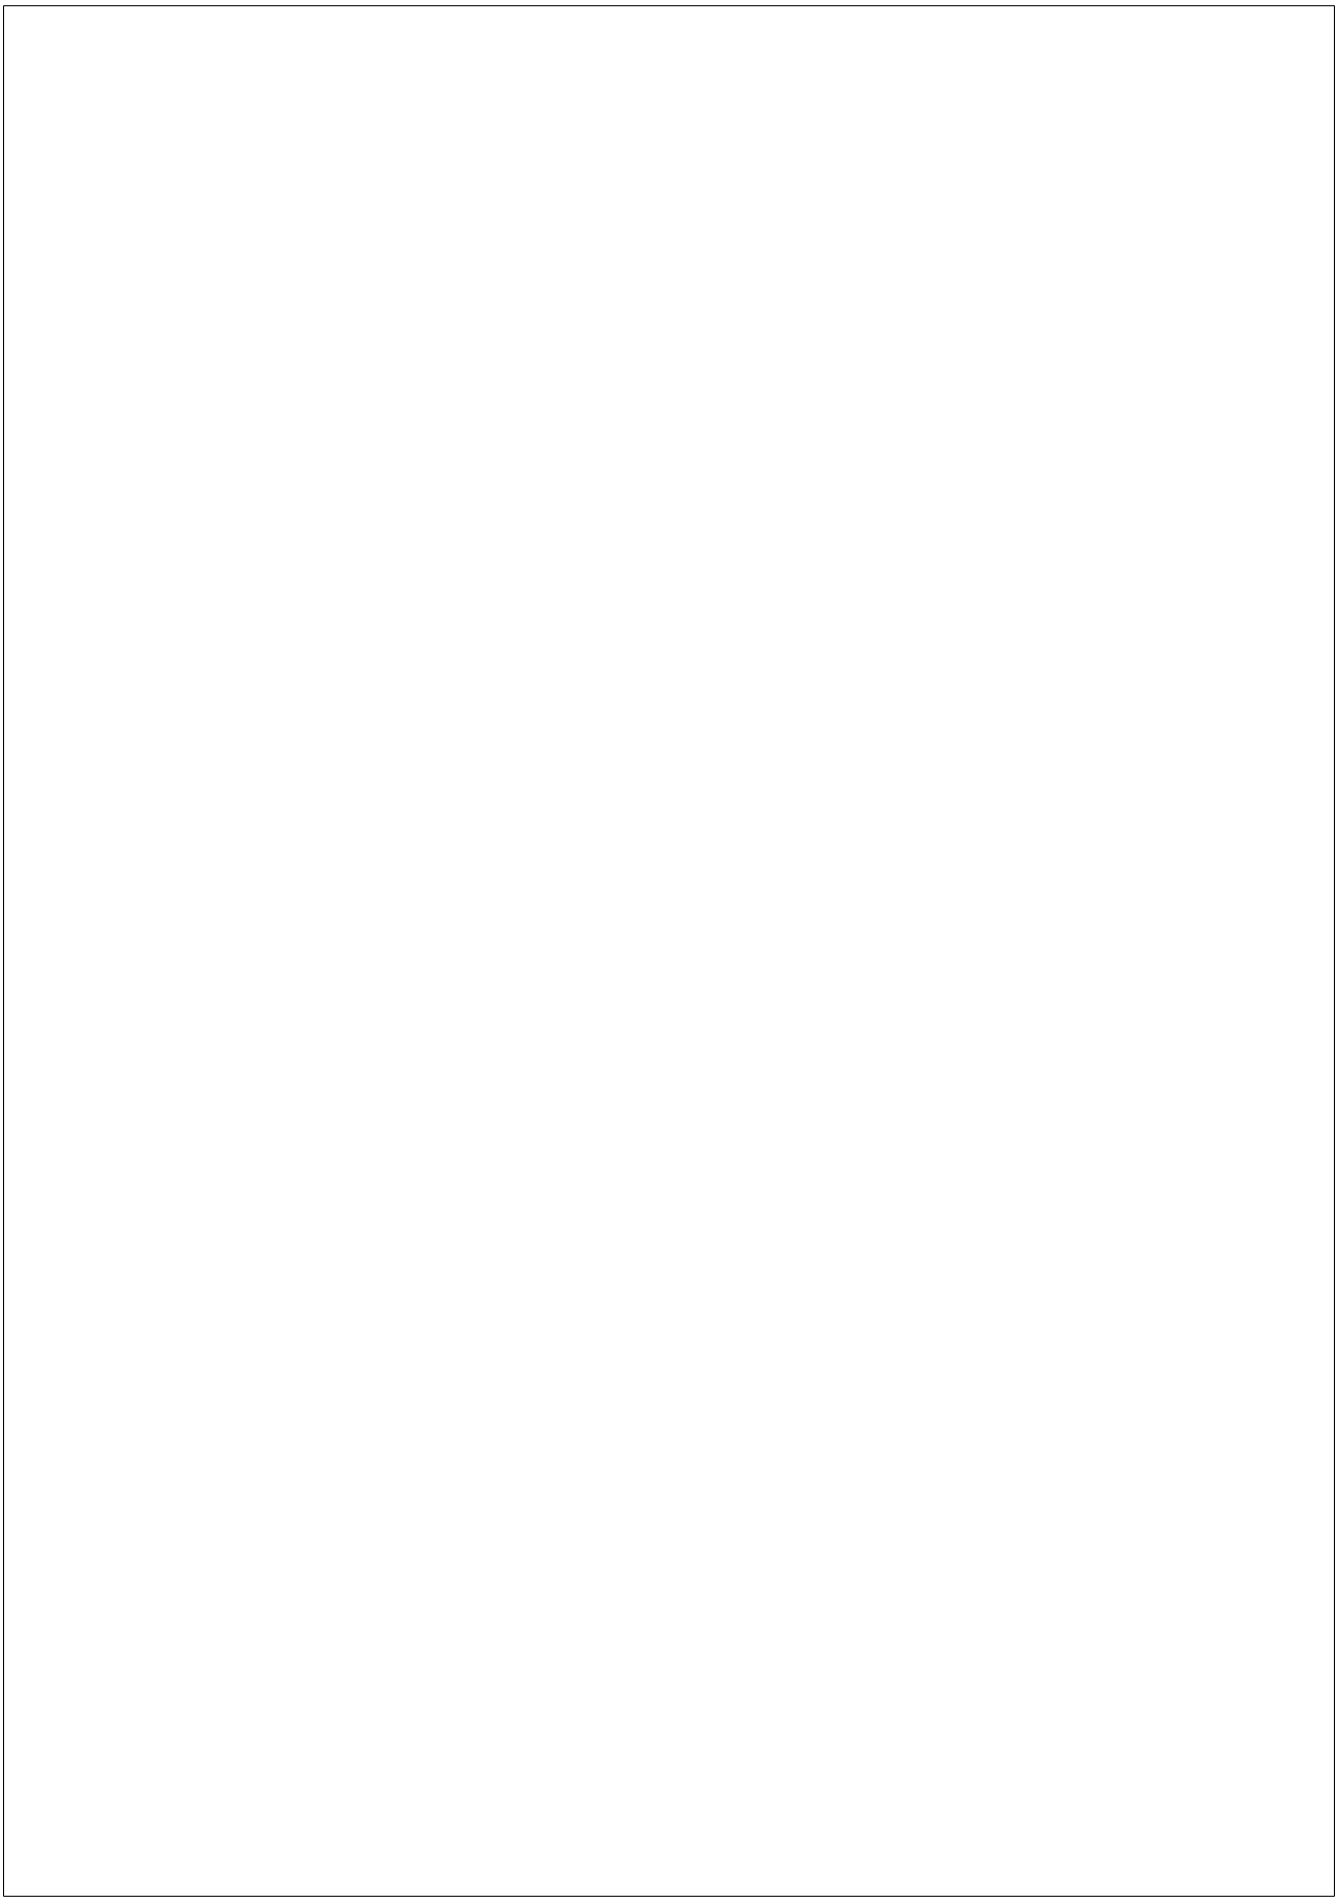

Chapter 6

APPENDIX: OTHER PUBLICATIONS

6.1. AdaptiveBandit: A multi-armed bandit framework for adaptive sampling in molecular simulations

Adrià Pérez, Pablo Herrera-Nieto, Stefan Doerr, Gianni De Fabritiis. Submitted to *Journal of Chemical Theory and Computation*.

Summary

AdaptiveBandit is a novel algorithm devoted to boosting the performance of MD in adaptive sampling scenarios. This new strategy takes its name from the *multi armed bandits* problem, a simplified version of a typical reinforcement learning scenario. In the *multi armed bandits* problem, the objective is to maximize a final reward by choosing from a pool of known states, each of them associated with their reward. Adapting MD sampling to this framework allows introducing already explored solutions

in other contexts (i.e., computer science) to the *explorations vs. exploitation dilemma*. In particular, the solution inspiring AdaptiveBandit is known as *upper confidence bound* (UCB). Here, the theoretical foundations for a solid adaptation of RL principles in MD sampling are introduced and tested in several case scenarios. In the first place, simple 2D potentials are used to show the improved capabilities of AdaptiveBandit to reach and converge in the minima of each potential, regardless of wrong external biases. Finally, AdaptiveBandit is also shown to outperform other methods in a more complex and realistic problem, such as the folding of villin.

AdaptiveBandit: A multi-armed bandit framework for adaptive sampling in molecular simulations

Adrià Pérez ,[†] Pablo Herrera-Nieto,[†] Stefan Doerr,[‡] and Gianni De Fabritiis[*,†,‡,¶]

†*Computational Science Laboratory, Universitat Pompeu Fabra, Barcelona, Spain*

‡*Acellera Ltd., Barcelona, Spain*

¶*Institució Catalana de Recerca i Estudis Avançats, Barcelona, Spain*

E-mail: gianni.defabritiis@upf.edu

Abstract

Sampling from the equilibrium distribution has always been a major problem in molecular simulations due to the very high dimensionality of conformational space. Over several decades, many approaches have been used to overcome the problem. In particular, we focus on unbiased simulation methods such as parallel and adaptive sampling. Here, we recast adaptive sampling schemes on the basis of multi-armed bandits and develop a novel adaptive sampling algorithm under this framework, AdaptiveBandit. We test it on multiple simplified potentials and in a protein folding scenario. We find that this framework performs similarly or better in every type of test potentials compared to previous methods. Furthermore, it provides a novel framework to develop new sampling algorithms with better asymptotic characteristics.

Introduction

In computational biology, macroscopic measurements by computer simulations are obtained by simulating microscopic molecular systems made of the order of a hundred thousand degrees of freedom. Statistical mechanics tells us what is the analytical form of the equilibrium distribution given the macroscopic constraint of the environment, e.g. constant temperature, pressure, and number of atoms. Therefore the problem consists in generating samples from such distribution.

Molecular simulation methods have always been hampered by sampling limitations over the equilibrium distribution due to their computational cost.[1,2] The two main forms to obtain samples are molecular dynamics (MD), a numerical scheme where the propagator of the dynamical system is discretized in time and iterated for billions of steps, and Monte Carlo sampling (MC), where the Monte Carlo rule is used to draw samples from the distribution. These sampling methods are also commonly used in other fields to sample for arbitrary probability distributions, and many of the methods developed for molecular simulations have been exploited in such contexts later, for instance, umbrella sampling,[3] biased Montecarlo methods[4] or biased molecular dynamics like replica-exchange,[5,6] steered MD,[7,8] metadynamics,[9] etc. Progress in molecular simulation sampling has therefore shown its relevance to a broader field of problems. Recently, a new generative method based on normalizing flows[10] has been proposed to sample from the Boltzmann distribution.[11]

Due to the difficulties in determining the bias a priori, practically equivalent to having a good prior, unbiased methods such as adaptive sampling[12–15] have been recently developed and used successfully.[16,17] Equally, due to the difficulty in generating good Montecarlo moves, molecular dynamics is almost always preferred to Montecarlo methods, largely due to the current efficiency of generating trajectories rooted in the capability of modern hardware. Specialized computer chips like Anton[18] made possible to run long simulations of the order of hundreds of microseconds, sampling reversibly fast processes and exploring longer timescales.[19] The advent of GPUs and GPU molecular dynamics software[20–23] was a notable improvement,

greatly increasing the computational efficiency of simulations. This, combined with Markov state models (MSMs)[24,25] allowed to reconstruct a complete statistical description of the full dynamical system from many shorter trajectories, obtaining a description that is equivalent to reversible sampling, once at convergence.

Running not one, but hundreds or thousands of simulation trajectories[26,27] created a new opportunity to decide the starting conditions of these simulations to obtain the best equilibrium characterization at the minimal computational cost, i.e. adaptive sampling. Initially, adaptive sampling algorithms[12,15] were used to reduce statistical uncertainty by choosing conformations that contributed the most to the error in mean first passage time of an MSM,[12] eigenvalues, and eigenvectors,[13] or choosing low state populations.[14,15] Furthermore, similar algorithms appeared recently which introduced prior knowledge to the selection criteria,[28–30] seeking to further speed up sampling towards equilibrium. One notable example is where contact information is used for protein folding[31] or bound state contacts in protein-ligand or protein-protein binding.[17] Other applications have used alternative geometric features, such as RMSD or pocket volume, to improve conformational exploration[32] and to find cryptic pockets.[33] In general, the adaptive sampling policy was always empirical, not based on any mathematical decision process, even though similarities have been recognized with the multi-armed bandit problem[30,34] and reinforcement learning[35] before.

Here we frame adaptive sampling in terms of a multi-armed bandit problem and propose AdaptiveBandit, an algorithm that uses an action-value function and an upper confidence bound[36,37] selection algorithm, improving adaptive sampling's performance and increasing its versatility when faced against different free energy landscapes. Our main goal is to provide strong fundamentals when facing the exploration-exploitation dilemma by redefining it in terms of reinforcement learning, creating a solid framework from where to easily develop novel algorithms. AdaptiveBandit is available in HTMD (`https://github.com/Acellera/htmd`).[38]

Methods

MD Simulations

The configurational space of a molecular system for MD simulations is given by $\chi = \{x = (\mathbf{r}_1, \ldots, \mathbf{r}_N) \in \mathbb{R}^{3N}\}$, where N is the number of atoms of the system. Experimental observables O are measured as equilibrium expectations $< O >= \int O(x)\,\mu(x)\,dx$, where $\mu(x)$ is the equilibrium distribution. The form of this distribution is known, for instance, the Boltzmann distribution in the canonical ensemble at temperature T is

$$\mu(x) = e^{\frac{-U(x)}{k_B T}}, \tag{1}$$

where $U(x)$ is the molecular potential energy and $k_B T$ is the Boltzmann constant multiplied by the temperature. MD numerically solves Newton's equation over the potential $U(x)$ for the variable x, plus a Langevin stochastic term accounting for thermal fluctuations.[39] Now consider the state $x(t) \in \chi$ as a specific conformation inside the configurational space χ at time t, the probability of finding the molecule in configuration $x_{t+\tau}$ at a later time can be expressed by the conditional transition density function p_τ, $x_{t+\tau} \sim p_\tau(x_{t+\tau}|x_t)$ which describes the probability of finding state $x_{t+\tau}$ given state x_t at time t after a time increment τ. When performing an MD simulation, the dynamics of the molecular system propagates the state x_t across time. Therefore, MD samples from the transition density p_τ given discrete time-steps τ to obtain the next state $x_{t+\tau}$. The process is repeated for many steps, generating a trajectory of conformations.

The main goal when performing MD simulations is to obtain a good representation of the system's equilibrium distribution $\mu(x)$ i.e. the probability to find conformation x under equilibrium conditions, in order to measure the average of observable $< O >$. If an MD trajectory τ is long enough, sampling from p_τ is equivalent to sampling from $\mu(x)$

(Equation (1))

$$\lim_{\tau\to\infty} p_\tau(x_{t+\tau}|x_t) = \mu(x). \tag{2}$$

Generating long enough trajectories is computationally expensive, and often practically impossible when trying to sample slow events. However, long trajectories can be substituted by short parallelized trajectories. While in principle one could model directly the conditional probability in Equation (2), in practice this is not possible given the very high dimensional space. Fortunately, it can be shown that the dynamics can be separated into a slow and fast set of variables,[24] and because contributions of fast variables decay exponentially in τ, a reliable MSM can be constructed in terms of the slow variables to compute thermodynamic averages. Usually, time-independent component analysis (tICA)[40] and clustering methods are used to learn this set of variables during sampling, necessary to build the MSM. Once we obtain the MSM, computed by estimating transition probabilities from discrete conformational states, one can derive thermodynamic and kinetic properties, just assuming local, not global, equilibrium (i.e. τ is much shorter than what is necessary to satisfy Equation (2)).

The multi-armed bandit problem

The multi-armed bandit problem is a simplified reinforcement learning setting where one faces the exploration versus exploitation dilemma. The problem is defined as a tuple $\langle \mathcal{A}, \mathcal{R}, \gamma \rangle$, where $\mathcal{A}$ is a set of k actions $\mathcal{A} = \{a_1, a_2, \ldots, a_k\}$ and $\mathcal{R}$ is an unknown probability distribution $\mathcal{R}^a = \mathbb{P}[r|a]$ of rewards given the chosen action. We choose $\gamma = 0$ for totally discounted rewards. At each time-step t, the agent applies a policy $\pi_a = \mathbb{P}[a]$ to select an action $a_t \in \mathcal{A}$, based on previous actions taken and the respectively obtained rewards. Subsequently, the environment returns a reward $r_t \sim R^{a_t}$. Given that we set $\gamma = 0$, we define the value of an action $Q_\pi(a)$ as its instantaneous mean reward

$$Q_\pi(a) = \mathbb{E}_\pi[r|a]. \tag{3}$$

The goal is to find the optimal policy π^* that maximizes the cumulative reward $\sum_{t=1}^{T} r_t$. Policies must take into account the exploration versus exploitation dilemma and combine both explorative actions, to sample their associated unknown reward function to update their value-estimates, and greedy actions, to increase the total cumulative reward by choosing the action with the highest value-estimate. The main advantage of describing adaptive sampling in terms of a multi-armed bandit is that we can benefit from the extensive literature on bandits to find solutions and replace heuristic policies with more mathematically sound ones.

AdaptiveBandit

Standard adaptive sampling algorithms work by performing several rounds or epochs of short parallel simulations. At each round, the algorithm is faced with the decision to select any of the sampled conformations from where to respawn a new round of simulations. The objective of these decisions is to avoid any redundant sampling and optimize our simulations to obtain the desired goal (which can be anything, from a full equilibrium characterization of a molecular system to sampling a specific conformation or dynamic event) at the minimum computational cost.

Here, we recast adaptive sampling in bandit terms, defining its tuple $\langle \mathcal{A}, \mathcal{R}, \gamma \rangle$. We define the action space $\mathcal{A}$ in terms of all possible conformations that are respawnable, i.e. they have been visited at least once,

$$\mathcal{A} = \mathcal{H}_m = \{x_k \in \mathbb{R}^{3N}, k = 1, \ldots, K_m\} \ , \tag{4}$$

where K_m is the number of sampled configurations at epoch m.

There are different possible choices for the a priori unknown reward function $\mathcal{R}$ that the policy will try to maximize, and it will mostly depend on your objective with the simulation experiment.

Because most of our MD experiments are usually aimed at sampling metastable states of

interest, e.g. folded states of proteins or bound states between proteins and ligands, we have defined the reward $\mathcal{R}$ to be proportional to minus equilibrium distribution so that that the optimal policy always picks conformations from the most stable state. Therefore, we define the reward $\mathcal{R}_a$ of action a as the mean of the minus free energies of each configuration x visited in the trajectory started with action a, i.e.

$$\mathcal{R}_a =< k_B T \log(\mu(x)) >_{(a,x_1,\ldots,x_\tau)} , \tag{5}$$

where $\mu(x)$ is the equilibrium distribution and the average is computed over the succeeding frames in the trajectory starting from a.

The action space would be too large to compute meaningful value-estimations for each conformation, and there is no way to know the exact equilibrium distribution. To address this issue, we take advantage of MSM analysis to redefine the tuple $\langle \mathcal{A}, \mathcal{R}, \gamma \rangle$ in a more practical form. We define a reduced and tractable action space by using the MSM's discretized conformational space and use the stationary distribution of each state to obtain an estimate of their free energy to compute the rewards. We count each trajectory frame as an action taken, and use the succeeding frames to assign the reward. Because rewards strongly depend on how accurate the MSM estimation is, we use the latest MSM to recompute all past rewards from all trajectories at each epoch, differently from common Q-learning approaches.[41] Not only it ensures the best free energy estimation possible, but it also addresses the increasing action space problem, due to new conformations being sampled. Every epoch, the discretized conformational space is redefined, all frames are reassigned and rewards are recomputed on the newly defined states.

Solving the multi-armed bandit problem

With the bandit tuple defined, we now need to deal with the exploration-exploitation trade-off and optimally solve it. To do so, AdaptiveBandit relies on the UCB1 algorithm[37] to optimize

the action-picking policy, which defines the upper confidence bound for action values based on the number of times the agent has picked that action and the total number of actions taken. Therefore, actions are selected based on

$$a_t = \underset{a \in \mathcal{A}}{\operatorname{argmax}} \left[Q_t(a) + c\sqrt{\frac{\ln t}{N_t(a)}} \right], \tag{6}$$

where t denotes the total number of actions taken, $Q_t(a)$ is the estimated action-value for action a, $N_t(a)$ is the number of times action a has been selected (prior to time t) and c is a parameter controlling the degree of exploration. UCB1 follows the principle of *"optimism in face of uncertainty"*, prioritizing actions with uncertain value-estimations, even if those values are not the greatest. To select an action, UCB1 not only takes into account the estimated value of that action, but also the amount of uncertainty on such value. By doing so, the algorithm not only promotes action exploration but also prioritizes the exploration of the most promising ones. In the long term, as our knowledge of action-values increases, the exploration term will decrease, and more greedy actions will be selected. UCB1 has a theoretical bound of $O(\sqrt{kTlog(L_t)})$ on its total regret L_t.[37]

AdaptiveBandit with knowledge-based initialization

AdaptiveBandit also has the option to initialize action-value estimates with external knowledge from the system, providing an initial value estimation to new actions, aiding to prioritize the most valuable actions. While in previous methods[17,30] this is done by forcing the algorithm to sample from conformations based on a fixed empirical ranking, here we use the bandit formalism to initialize Q in Equation (6) with an empirical action-value function. This notably allows for the MSM to correct the initial prior suggestion for Q given enough sampling. This is not true in previous schemes, where a partially wrong prior can affect sampling to the point of non-convergence to the intended results due to its degeneracy, i.e. even just some wrong contact information could kinetically bias the simulations far from the folding funnel.

We demonstrate this aspect in the result section. The initial prior $Q_{prior}(a)$ is computed as the average goal score from all frames in a state, and it is recalculated at each epoch, after re-clustering. The states are assigned with an initial pseudo-count $N_0(a)$, representing the statistical certainty of $Q_{prior}(a)$.

Other adaptive sampling algorithms

To evaluate AdaptiveBandit's performance, we have tested it against several different adaptive sampling strategies, mainly the standard low-counts adaptive sampling, FAST[30] and Exploration-Exploitation.

The low-counts adaptive sampling is a simple and intuitive strategy that is optimal in pure exploration scenarios.[38] The method works by selecting conformations from the least populated clusters at each adaptive epoch. The other two methods, FAST and Exploration-Exploitation, are goal-oriented, were external knowledge on the system is used to guide sampling.

FAST is also inspired by the multi-armed bandit problem, but the implementation differs as it uses an acquisition function to rank discrete conformational states rather than a reward function by definition, and actions (and their outcomes) do not influence their value-estimates. The acquisition function contains an exploitation term, defined by the goal scoring function that assigns a fixed value to each state, and an exploration term, based on state counts. The FAST implementation we used works as

$$\rho_i = \alpha\phi_i + (1-\alpha)\psi_i \ , \tag{7}$$

where ρ_i is the score for state i, ϕ is the exploitative value obtained from the goal function for state i, ψ is the exploration value defined by state i counts (as in low-counts adaptive sampling) and α is a parameter regulating the weight of both terms. Both ϕ and ψ terms are scaled to values that range from 0 to 1. The states are defined as the microstates obtained

by the constructed Markov model at each epoch.

Lastly, we have Exploration-Exploitation, a strategy inspired by the popular method for multi-armed bandits ϵ-greedy, implemented in HTMD's AdaptiveGoalEG.[38] Simulations are restarted ϵ times from the top goal ranking states, and $1-\epsilon$ times from the least sampled states (i.e the low-counts strategy).

Langevin dynamics on 2D Potentials

We designed a set of experiments in a simple simulation set-up, performing Langevin dynamics on a single point mass of 1000 amu and a diffusion coefficient of 10Å^2/ns at 300 K on two different potentials: a 2-wells potential (Figure 1a) inspired from Ref. 42, given by

$$\begin{aligned} U(x,y) = &-3e^{-(x-1)^2-y^2} - 3e^{-(x+1)^2-y^2} + 15e^{-0.32(x^2+y^2+20(x+y)^2)} \\ &+ 0.0512(x^4+y^4) + 0.4e^{-2-4y} \end{aligned} \tag{8}$$

and a funnel potential (Figure 1c) given by

$$U(x,y) = 2\cos(2\sqrt{x^2+y^2}) - 8e^{-(x^2+y^2)} + 0.2((x/8)^2+(y/8)^2)^3 \tag{9}$$

A reference baseline for each 2D potential was calculated using an MSM built with 10μs and 500μs of aggregate simulation time for the 2-wells and funnel potential respectively, spawning trajectories from conformations covering the whole surface. Equilibrium probability was determined to be 50% and 85% respectively on each global minima.

A total of 1 μs were simulated for each combination of method and potential, spawning 25 trajectories of 0.1 ns at each epoch for a total of 400 epochs. Performance at each epoch was measured as the mean of the equilibrium probabilities for the macrostate containing the targeted minimum for 10 independent MSMs built with 80% of bootstrapped data. All the MSMs calculations were performed using HTMD.[38]

For the goal methods, we simulated a total of 2 μs for each method, spawning 10 trajectories per epoch with trajectories of 0.05 ns. Values of $\alpha = 0.1$ for FAST and $\epsilon = 0.1$ for Exploration-Exploitation were selected. In AdaptiveBandit the exploration rate was set to $c = 0.01$ and the initial pseudo-counts to $N_0(a) = 50$.

MD simulation set-up

Simulation system for the chicken villin headpiece (PDB:2F4K) was built with HTMD.[38] We solvated villin in a 64Å cubic box with a NaCl concentration of 0.05 M. Starting unfolded conformations for the runs were selected from a villin unfolding trajectory at high temperature (500 K).

In this context, we tested AdaptiveBandit with $c = 0.01$ and $N_0(a) = 100$, against two different FAST setups, $\alpha = 0.5$ and $\alpha = 0.1$. A goal scoring function was used to guide the algorithms, based on the number of native Cα contacts formed. For each setup, we ran parallel simulations of 10 ns, with 5 to 10 simulations per epoch, until we reached a total aggregate time of 4 μs. All simulations were run with ACEMD,[22] using the CHARMM22* force-field[43] on a local GPU cluster. A short HTMD code listing is provided as an example to run AdaptiveBanditfor villin simulations (Listing 1).

Results

Performance testing on 2D Potentials

The initial objective is to compare the performance of a set of adaptive sampling algorithms in a simple environment defined by 2D potentials. For this purpose, we performed Langevin dynamics on two different potentials: the 2-wells potential, composed of two minima separated by a high energetic barrier (Figure 1a), and a funnel potential, comprised of concentric isoenergetic regions with the global minimum located at its center (Figure 1c). The funnel potential is a useful benchmark to test the exploration-exploitation balance, as a purely

```
from htmd.ui import *
from sklearn.cluster import MiniBatchKMeans
from jobqueues.localqueue import LocalGPUQueue
from goals import goalFunction

refmol = Molecule('villin_2f4k.pdb')
md = AdaptiveBandit()
md.app = LocalGPUQueue()
md.generatorspath = './generators'

md.clustmethod = MiniBatchKMeans
md.projection = MetricSelfDistance('protein and name CA')
md.goalfunction = delayed(goalFunction)(refmol)
md.ticadim = 3
md.nmin=5
md.nmax=10
md.nframes = 1000000

md.exploration = 0.01   ## "c" value
md.goal_init = 100      ## prior initialization value

md.run()
```

Listing 1: Example AdaptiveBandit code

exploratory strategy would tend to guide towards the outer circular wells, while the minimum is in the center. The objective of these experiments is to predict the equilibrium population of the targeted minima. The equilibrium populations are computed with MSM analysis to assess how different sampling strategies affect the MSM estimation.

First, AdaptiveBandit was compared with two other common sampling policies, based on simple heuristics: random selection and the low-counts policy. Results for the 2-wells potential (Figure 1b) show a similar performance for the low-counts policy and AdaptiveBandit. Both converge at the baseline population (50%) while random sampling underestimates it. Because the potential just contains two large minima, comprising almost the entire conformational space, a fully explorative heuristic algorithm like the low-counts is optimal, as there is no need to prioritize anything besides exploring the two minima. AdaptiveBandit is able to reach the same optimal performance.

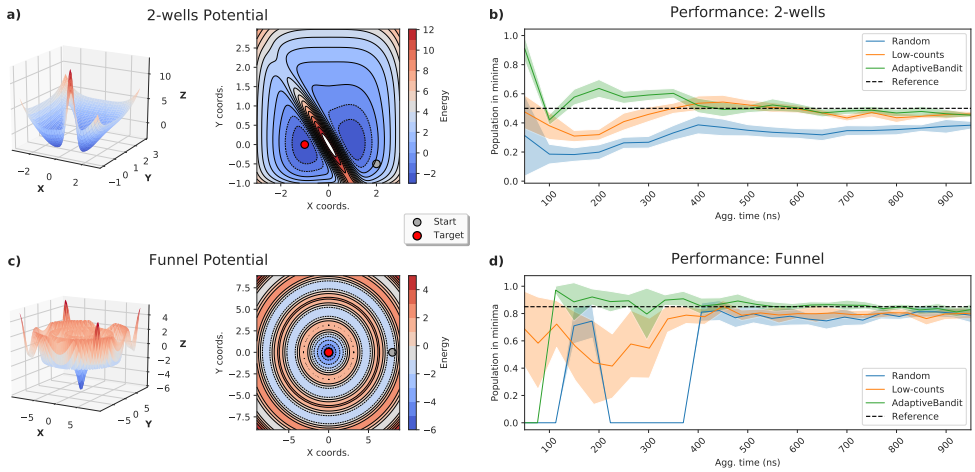

Figure 1: **Performance comparison between random, low-counts and AdaptiveBandit in the experiments with 2D potentials. a), c)** 3D view and top view of the 2-wells and funnel potentials. Global minima are located at (-1, 0) and (0, 0) coordinates, respectively. Blue dot indicates starting points for simulations and red dot indicates the target global minima where population is measured at every epoch. **b), d)** Performance comparison of total aggregate simulation time needed for random, low-counts and AdaptiveBandit sampling methods in 2-wells and funnel potential, respectively, to achieve correct population estimates at their global minimum.

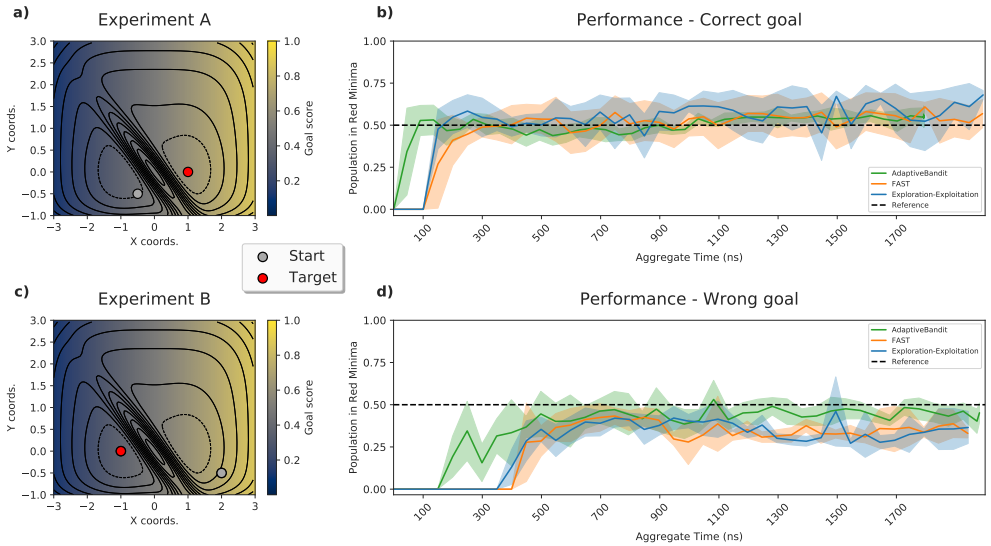

Figure 2: **Performance comparison between goal-oriented algorithms FAST, epsilon-greedy and AdaptiveBandit in the experiments with 2D potentials. a), c)** Top view of 2-wells potential. Goal distribution across the potential is shown. Blue dots indicate the starting conformations for the runs. Red dots indicate the minima where population is measured. **b), d)** Performance comparison of total aggregate simulation time needed for FAST, Exploration-Exploitation and AdaptiveBandit methods to correctly estimate populations at their target minimum.

For the funnel potential (Figure 1d), the relative size of the minima is much smaller compared to the conformational space, hence its detection by random sampling is more inefficient than for the other two algorithms. The low-counts method is able to reach the minima faster, as it is to cover the space quickly. Both these algorithms obtain a slight underestimation of the equilibrium population. On the other hand, AdaptiveBandit achieves a more accurate estimation and reaches convergence with 4 times less aggregate time than the other algorithms, highly reducing the computational resources needed to obtain accurate estimations of the equilibrium distribution.

This first test here showcases how introducing an exploitation term to quantify an action’s value, besides the exploration term, either increases or equals the performance of fully exploratory algorithms on obtaining correct equilibrium estimations in the tested systems.

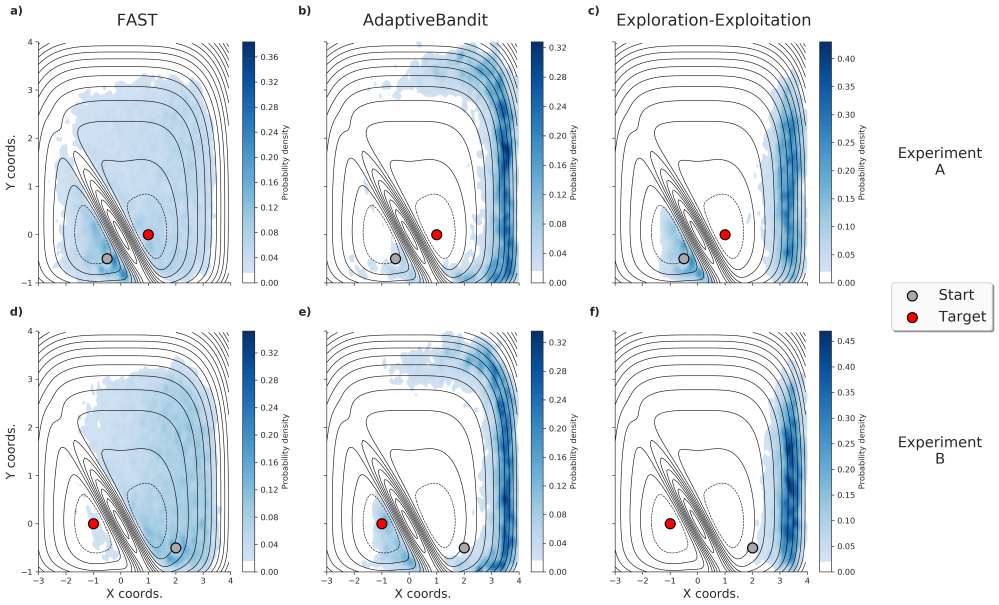

Figure 3: **Simulation re-spawning distribution by algorithm across the 2-wells potential.** Each plot depicts the probability distribution of selected conformations throughout the runs, obtained by kernel density estimate.[44] Starting points for each run are represented with a blue dot and target minimum with a red one. Goal distribution (not shown) is the same as in Figure 2. Subplots **a), b), c)** represent the spawning probability distribution across the potential surface on experiment A, target minima at coordinates (1, 0), for FAST, AdaptiveBandit and Exploration-Exploitation algorithms, and **d), e), f)** for experiment B, target minima at coordinates (0, -0.5).

Value-estimations of each action help on prioritizing sampling on the most relevant areas of the conformational space, rather than just exploring everything and sampling irrelevant conformations. While in the 2-wells potential this does not make a big difference, it does in the funnel potential, where AdaptiveBandit focuses sampling on the minima by identifying its relevance with action value-estimates and does not waste resources on exploring irrelevant conformations.

Using system external knowledge

Next, we want to test how AdaptiveBandit performs in the 2-wells potential against two existing methods that incorporate an exploitative term by employing external knowledge on the system. The pair of tested algorithms, also known as goal-oriented methods, are FAST[30] and Exploration-Exploitation. To make sure AdaptiveBandit is at the same level of system knowledge as the other methods, the information provided by the goal-function was used in AdaptiveBandit through knowledge-based initialization (as explained in *Methods*).

The goal function employed in the experiments with the 2-wells potential increases the score linearly with the x axis (Figure 2a,c), thus creating a gradient of reward pushing to the right boundary of the potential. Two tests were performed in different scenarios. In the first test, the target minimum has a greater score than the starting coordinate (Experiment A, Figure 2a). In the second one, the target minimum has a lower goal than the starting conformations, and therefore requires opposition to the goal's influence to obtain accurate estimations on the target minimum (Experiment B, Figure 2c).

For experiment A, all methods reached the reference population, with AdaptiveBandit needing slightly less simulation time to reach the correct population estimation in the target (2b). Differences in the algorithms can be visualized by a distribution plot of the spawning conformations in Figure 3. During the initial epochs, both FAST and Exploration-Exploitation follow the goal, spawning new simulations pushing against the energy barrier. AdaptiveBandit, on the other hand, quickly discovers the target minima and starts exploring other areas and not only directs sampling on the high score region but also in its surroundings. Even though the performance of all three algorithms is similar, differences in the spawning patterns between the three algorithms can be appreciated throughout the experiment. FAST presents a more explorative behavior and respawns simulations from all along the conformational space (Figure 3a). On the other hand, Exploration-Exploitation presents a highly exploitative behavior, strongly focusing on the highest goal-scoring region once it is discovered (Figure 3c). In between, AdaptiveBandit presents an overall greedy behavior, but with higher

levels of exploration than the Exploration-Exploitation method which translates into a small boost in its performance. It is interesting to point out the few resources invested by AdaptiveBanditin the origin minima, which demonstrates that the algorithm quickly identifies it as a non-interesting area (Figure 3b).

For experiment B, AdaptiveBandit reaches the target minimum faster and equilibrium populations are estimated more accurately (Figure 2d). Both Exploration-Exploitation and FAST require more simulation time to reach the target minima and fail to converge on the correct equilibrium populations. In this scenario, Exploration-Exploitation is greatly focused in the high scoring region (Figure 3f) resulting in a marginal exploration of the target minimum, while FAST and specially AdaptiveBandit do perform a more significant search on it (Figure 3d,e). Comparison between AdaptiveBandit and FAST spawning patterns (Figure 3d,e) reveals the differences in the exploration profile, where again FAST thoroughly spawns conformations from every explored point in the surface, while AdaptiveBandit, following the goal, explores the boundaries of the conformational space. Even if differences in performance are not substantially large, the experiment shows us the inability of FAST and Exploration-Exploitation to update the initial action-value estimates, translating into a lack of adaptation to the system being sampled. In opposition, AdaptiveBandit is able to correct the prior action-value estimates and readjust the sampling policy to a more optimal one, as it uses exploitation intrinsically based on MSM estimations from the available simulation data and external knowledge is introduced as prior information, rather than as the function to optimize. The ability to update the system knowledge at each epoch is crucial in experiments where the goal scoring function used has high levels of degeneracy or is directly wrong. Asymptotically, AdaptiveBandit should always be better as it is logarithmically bound on the number of trials to the total regret[37] (the difference between the maximum possible reward and the current reward), whereas Exploration-Exploitation and FAST are linearly bound.

Testing on protein folding simulations

Besides testing in simple 2D potentials, we explored AdaptiveBandit's performance on a more realistic and challenging scenario. AdaptiveBandit was tested on protein folding simulations, using villin as a benchmark. The chicken villin headpiece consists of a chain of 35 residues that folds into a three α-helical bundle, sharing a common hydrophobic core.[45] It is known to have a fast folding rate of $(0.7\ \mu s)^{-1}$.[45] Our target for this test is to reach the folded state with the minimum amount of aggregate time and compare how AdaptiveBandit and FAST distribute sampling across the conformational space of villin. Because we are testing the algorithm's effect rather than the technical capabilities of reaching villin's folding state with MD, we set up very short simulation times to increase the number of epochs and ensure we are evaluating the algorithm's performance. The goal function used for the algorithms maximizes the number of native Cα contacts formed to guide sampling on to the folded state. 30 μs of villin folding simulations were used to build some reference tICA dimensions to evaluate the sampled conformational space from each method. The first two TICA dimensions reveal three main states (Figure 4a): the unfolded state (random coil), the folded structure, and a misfolded state.

Figure 4b shows the distinctive behavior of AdaptiveBandit and FAST while sampling the folding path. AdaptiveBandit clearly reaches the crystal structure. FAST struggles to do so due to the very short trajectories used, which produces a sampling bias, as indicated in.[46] The results showcase how AdaptiveBandit is able to select the most relevant conformations to reach the folded state, prioritizing the most promising actions from the subset of undersampled actions. On the contrary, FAST, even in its most greedy setting ($\alpha = 0.1$), is not able to correctly prioritize the most relevant states and keeps exploring over random coil states, even in the latest epochs (Figure 4c). The greedy setting also presents a slight misdirection towards the misfolded state, which suggests that the used goal scoring function has degeneracy and it does not differentiate enough between native-like structures and misfolded structures that are very far dynamically. As commented in the previous experiment using external knowledge

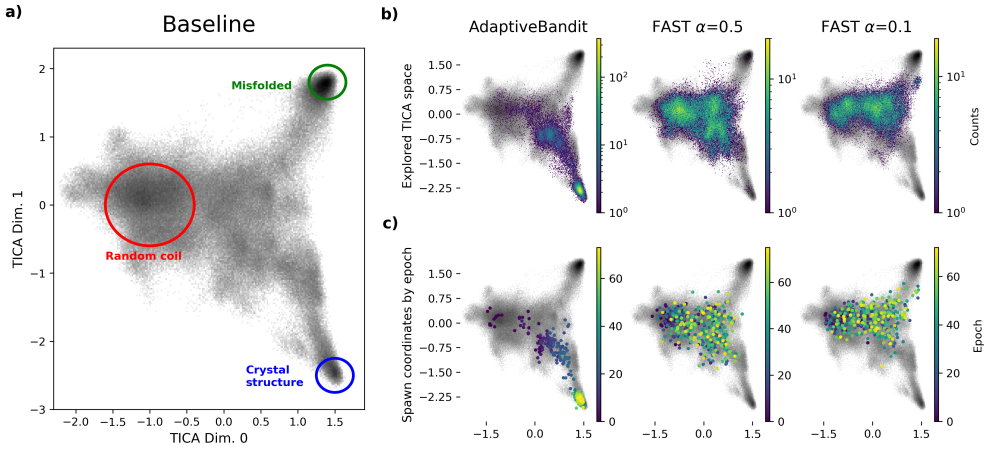

Figure 4: **Villin folding simulations. a)** Conformational space for folding of villin on the baseline data set. The tICA space includes large regions of random coil (initial conformation are located within the red circle), misfolded conformations (green circle), and crystal-like structures (blue circle). **b)** Exploration of the conformational space by sampling algorithms. Each plot includes the baseline exploration depicted on gray and the explored space with a colored heatmap. **c)** Spamming coordinates for new epochs. Scattered points indicate starting conformations for new epochs, colored from first (purple) to last (yellow).

on the 2-wells potential, methods like FAST or Exploration-Exploitation that rely only on external information can be severely hampered when the provided information does not represent the true energetic gradient. AdaptiveBandit prevents that by updating the prior information with rewards coming from interacting with the system and observing its response to our actions.

Conclusion

AdaptiveBandit formally introduces adaptive sampling into reinforcement learning by describing it in terms of multi-armed bandits and builds upon it to deliver a novel algorithm with increased performance and flexibility across different energy landscapes. AdaptiveBandit is able to perform equally or better than previous adaptive sampling algorithms in a diverse set of systems, and it has demonstrated the ability to learn from simulation results. Adaptive-

Bandit works both with and without external knowledge of the system, and it can update prior beliefs in the system based on the results obtained during the experiment.

Goal-oriented adaptive sampling methods as in Ref. 30 also get inspiration from exploration-exploitation strategies, like *epsilon*-greedy. The context, however, is quite different as there is not a definition of a multi-armed bandit framework and a reward per action, rather it is more akin to directly define an acquisition function. Furthermore, the greediness is towards predetermined states given from external knowledge on the system. AdaptiveBandit, as used here, uses exploitation intrinsically without requiring external information. It is, however, a possibility to do so and use experimental data to provide a prior for the sampling.

We have exemplified here cases were AdaptiveBandit works better due to its adaptability and flexibility, but that does not mean that it could underperform in other scenarios. Our implementation of AdaptiveBandit relies on good MSM estimates, and therefore the action-value estimates carry on with errors caused not only by discretization and dimensionality reduction but also by the sampling bias, especially on estimations of equilibrium populations.[46] Additionally, AdaptiveBandit's performance also depends on the c hyperparameter to regulate exploration and it is not very intuitive, as it must be tuned according to the scale of both terms in Equation (6).

The version of AdaptiveBandit presented here defines a reward proportional to the free energy of each state and utilizes the UCB1 algorithm to optimize the action-picking policy. However, this is not the only possible way to apply AdaptiveBandit and the algorithm can be changed to better adapt the experiment and systems. We hope that our work inspires the development of new adaptive sampling algorithms built under theoretical fundamentals instead of using simple heuristic policies.

Acknowledgement

G.D.F. acknowledges support from MINECO (Unidad de Excelencia María de Maeztu MDM-

2014-0370 and BIO2017-82628-P) FEDER and Secretaria d'Universitats i Recerca de la Generalitat de Catalunya. This project received funding from the European Union's Horizon 2020 Research and Innovation Programme under Grant Agreement 675451 (CompBioMed Project).

References

(1) Martinez-Rosell, G.; Giorgino, T.; Harvey, M. J.; de Fabritiis, G. *Current topics in medicinal chemistry* **2017**, *17*, 2617–2625.

(2) Pérez, A.; Martínez-Rosell, G.; De Fabritiis, G. *Current opinion in structural biology* **2018**, *49*, 139–144.

(3) Torrie, G. M.; Valleau, J. P. *Journal of Computational Physics* **1977**, *23*, 187–199.

(4) Frenkel, D., Smit, B., Eds. *Understanding Molecular Simulation: From Algorithms to Applications*, 1st ed.; Academic Press, Inc.: Orlando, FL, USA, 1996.

(5) Sugita, Y.; Okamoto, Y. *Chemical physics letters* **1999**, *314*, 141–151.

(6) Fukunishi, H.; Watanabe, O.; Takada, S. *The Journal of chemical physics* **2002**, *116*, 9058–9067.

(7) Izrailev, S.; Stepaniants, S.; Isralewitz, B.; Kosztin, D.; Lu, H.; Molnar, F.; Wriggers, W.; Schulten, K. *Computational molecular dynamics: challenges, methods, ideas*; Springer, 1999; pp 39–65.

(8) Isralewitz, B.; Gao, M.; Schulten, K. *Current opinion in structural biology* **2001**, *11*, 224–230.

(9) Laio, A.; Parrinello, M. *Proceedings of the National Academy of Sciences* **2002**, *99*, 12562–12566.

(10) Rezende, D. J.; Mohamed, S. *arXiv preprint arXiv:1505.05770* **2015**,

(11) Noé, F.; Olsson, S.; Köhler, J.; Wu, H. *Science* **2019**, *365*, eaaw1147.

(12) Singhal, N.; Pande, V. S. *The Journal of chemical physics* **2005**, *123*, 204909.

(13) Hinrichs, N. S.; Pande, V. S. *The Journal of chemical physics* **2007**, *126*, 244101.

(14) Pronk, S.; Larsson, P.; Pouya, I.; Bowman, G. R.; Haque, I. S.; Beauchamp, K.; Hess, B.; Pande, V. S.; Kasson, P. M.; Lindahl, E. Copernicus: A new paradigm for parallel adaptive molecular dynamics. Proceedings of 2011 International Conference for High Performance Computing, Networking, Storage and Analysis. 2011; p 60.

(15) Doerr, S.; De Fabritiis, G. *Journal of chemical theory and computation* **2014**, *10*, 2064–2069.

(16) Noé, F.; Schütte, C.; Vanden-Eijnden, E.; Reich, L.; Weikl, T. R. *Proceedings of the National Academy of Sciences* **2009**, *106*, 19011–19016.

(17) Plattner, N.; Doerr, S.; De Fabritiis, G.; Noé, F. *Nature chemistry* **2017**, *9*, 1005.

(18) Shaw, D. E.; Deneroff, M. M.; Dror, R. O.; Kuskin, J. S.; Larson, R. H.; Salmon, J. K.; Young, C.; Batson, B.; Bowers, K. J.; Chao, J. C.; et al., *Commun. ACM* **2008**, *51*, 91–97.

(19) Lindorff-Larsen, K.; Piana, S.; Dror, R. O.; Shaw, D. E. *Science* **2011**, *334*, 517–520.

(20) Friedrichs, M. S.; Eastman, P.; Vaidyanathan, V.; Houston, M.; Legrand, S.; Beberg, A. L.; Ensign, D. L.; Bruns, C. M.; Pande, V. S. *Journal of computational chemistry* **2009**, *30*, 864–872.

(21) Harvey, M.; De Fabritiis, G. *Journal of chemical theory and computation* **2009**, *5*, 2371–2377.

(22) Harvey, M. J.; Giupponi, G.; Fabritiis, G. D. *Journal of chemical theory and computation* **2009**, *5*, 1632–1639.

(23) Eastman, P.; Swails, J.; Chodera, J. D.; McGibbon, R. T.; Zhao, Y.; Beauchamp, K. A.; Wang, L. P.; Simmonett, A. C.; Harrigan, M. P.; Stern, C. D.; Wiewiora, R. P.; Brooks, B. R.; Pande, V. S. *PLoS Comput. Biol.* **2017**, *13*.

(24) Prinz, J.-H.; Wu, H.; Sarich, M.; Keller, B.; Senne, M.; Held, M.; Chodera, J. D.; Schütte, C.; Noé, F. *The Journal of chemical physics* **2011**, *134*, 174105.

(25) Bowman, G. R.; Pande, V. S.; Noé, F. *An introduction to Markov state models and their application to long timescale molecular simulation*; Springer Science & Business Media, 2013; Vol. 797.

(26) Buch, I.; Giorgino, T.; De Fabritiis, G. *Proceedings of the National Academy of Sciences* **2011**, *108*, 10184–10189.

(27) Martínez-Rosell, G.; Harvey, M. J.; De Fabritiis, G. *Journal of chemical information and modeling* **2018**, *58*, 683–691.

(28) Sabbadin, D.; Moro, S. *Journal of chemical information and modeling* **2014**, *54*, 372–376.

(29) Perez, A.; MacCallum, J. L.; Dill, K. A. *Proceedings of the National Academy of Sciences* **2015**, *112*, 11846–11851.

(30) Zimmerman, M. I.; Bowman, G. R. *Journal of chemical theory and computation* **2015**, *11*, 5747–5757.

(31) Ovchinnikov, S.; Park, H.; Varghese, N.; Huang, P.-S.; Pavlopoulos, G. A.; Kim, D. E.; Kamisetty, H.; Kyrpides, N. C.; Baker, D. *Science* **2017**, *355*, 294–298.

(32) Zimmerman, M. I.; Hart, K. M.; Sibbald, C. A.; Frederick, T. E.; Jimah, J. R.; Knoverek, C. R.; Tolia, N. H.; Bowman, G. R. *ACS central science* **2017**, *3*, 1311–1321.

(33) Cruz, M. A.; Frederick, T. E.; Singh, S.; Vithani, N.; Zimmerman, M. I.; Porter, J. R.; Moeder, K. E.; Amarasinghe, G. K.; Bowman, G. R. *bioRxiv* **2020**,

(34) Zimmerman, M. I.; Porter, J. R.; Sun, X.; Silva, R. R.; Bowman, G. R. *Journal of chemical theory and computation* **2018**, *14*, 5459–5475.

(35) Shamsi, Z.; Cheng, K. J.; Shukla, D. *The Journal of Physical Chemistry B* **2018**, *122*, 8386–8395.

(36) Lai, T. L.; Robbins, H. *Advances in applied mathematics* **1985**, *6*, 4–22.

(37) Auer, P. *Journal of Machine Learning Research* **2002**, *3*, 397–422.

(38) Doerr, S.; Harvey, M.; Noé, F.; De Fabritiis, G. *Journal of chemical theory and computation* **2016**, *12*, 1845–1852.

(39) Loncharich, R. J.; Brooks, B. R.; Pastor, R. W. *Biopolymers: Original Research on Biomolecules* **1992**, *32*, 523–535.

(40) Pérez-Hernández, G.; Paul, F.; Giorgino, T.; De Fabritiis, G.; Noé, F. *The Journal of chemical physics* **2013**, *139*, 07B604_1.

(41) Sutton, R. S.; Barto, A. G. *Reinforcement learning: An introduction*; 2018.

(42) Pan, A. C.; Roux, B. *The Journal of chemical physics* **2008**, *129*, 064107.

(43) Piana, S.; Lindorff-Larsen, K.; Shaw, D. E. *Biophysical journal* **2011**, *100*, L47–L49.

(44) Waskom, M. seaborn.kdeplot. `https://seaborn.pydata.org/generated/seaborn.kdeplot.html/`, [Online; accessed 11-January-2020].

(45) Kubelka, J.; Chiu, T. K.; Davies, D. R.; Eaton, W. A.; Hofrichter, J. *Journal of molecular biology* **2006**, *359*, 546–553.

(46) Wan, H.; Voelz, V. A. *The Journal of Chemical Physics* **2020**, *152*, 024103.

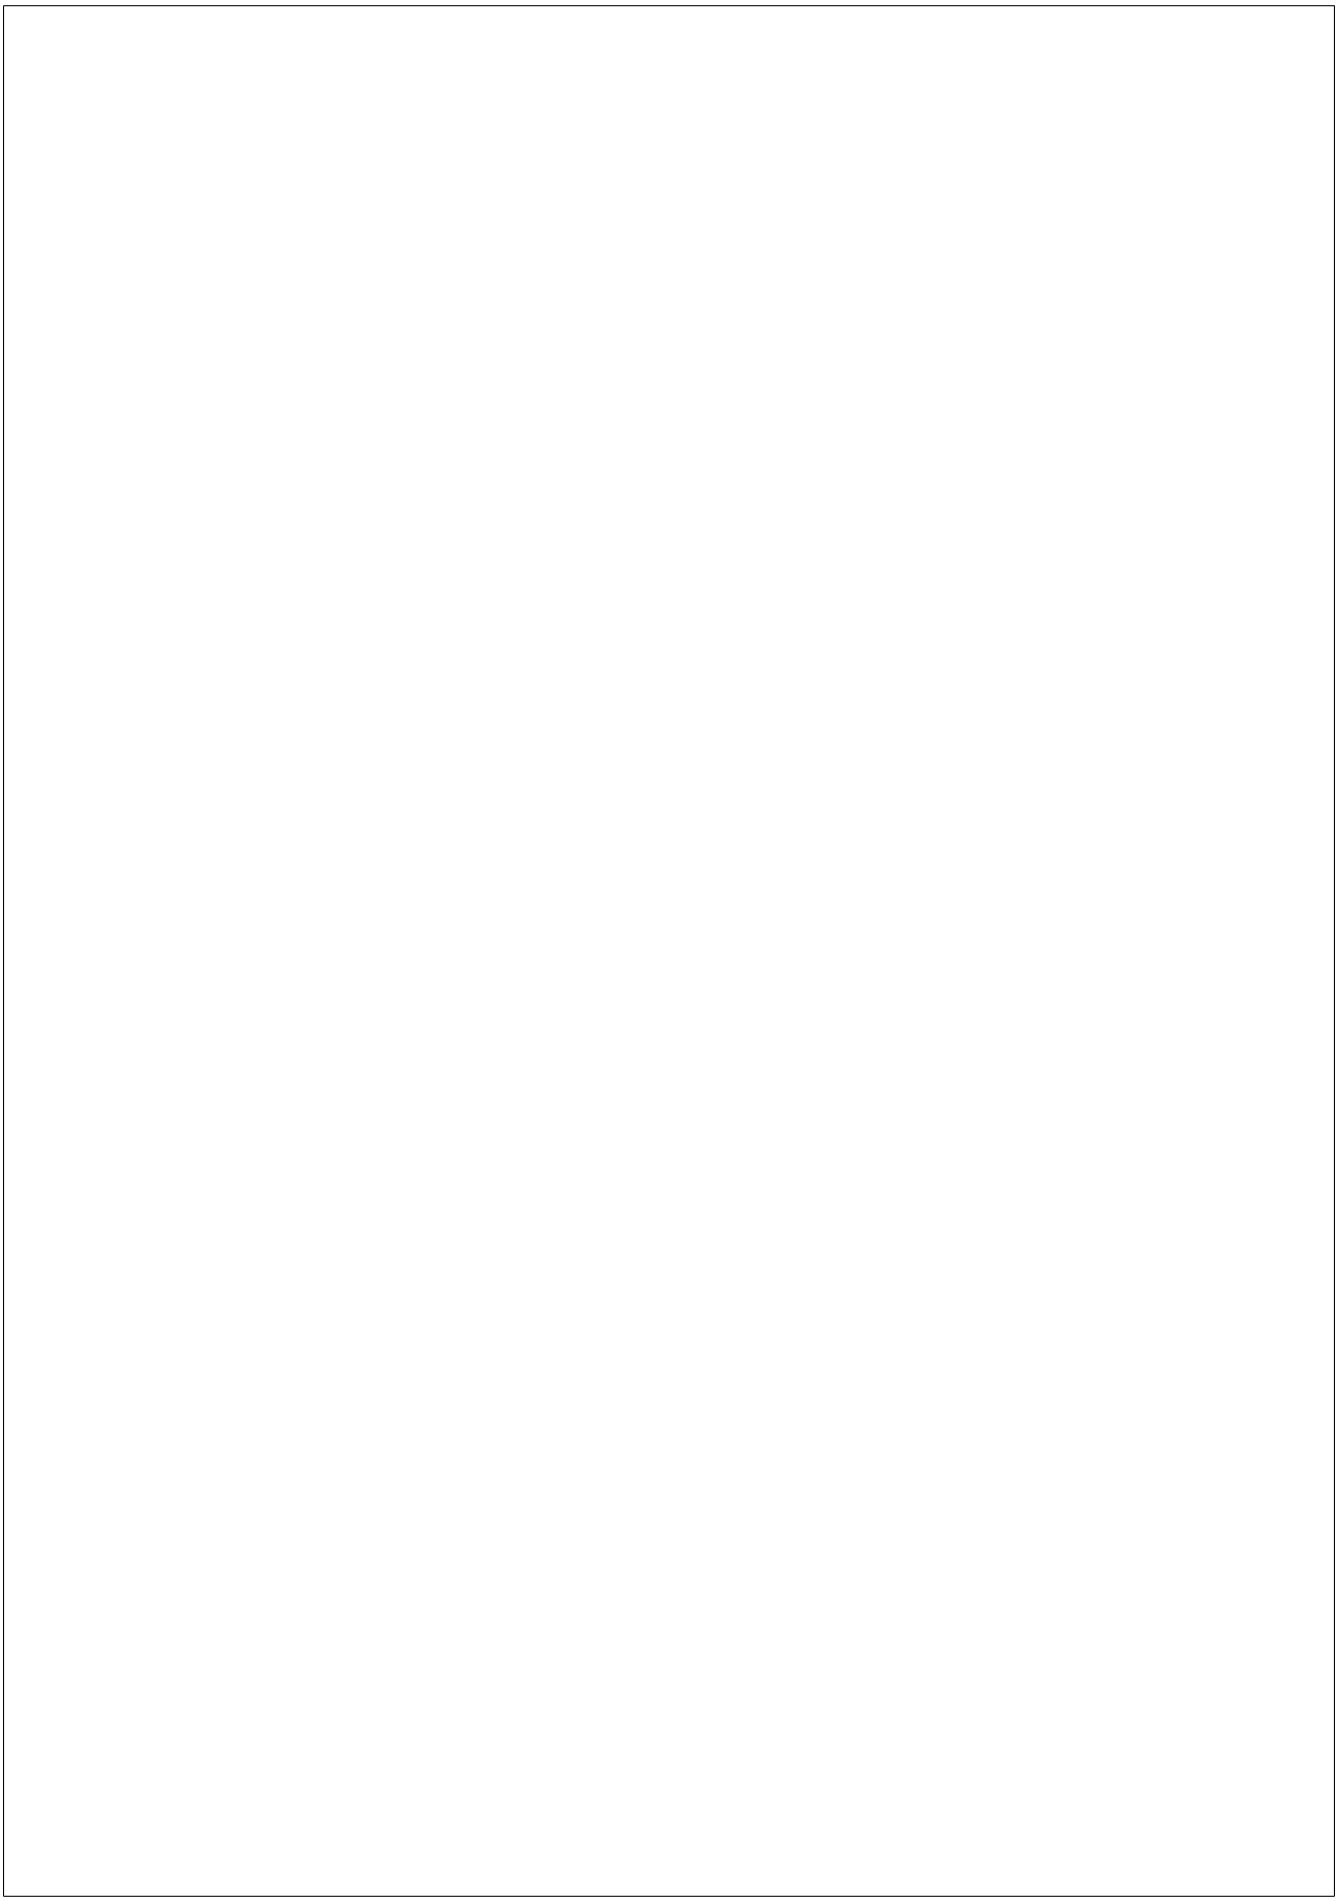

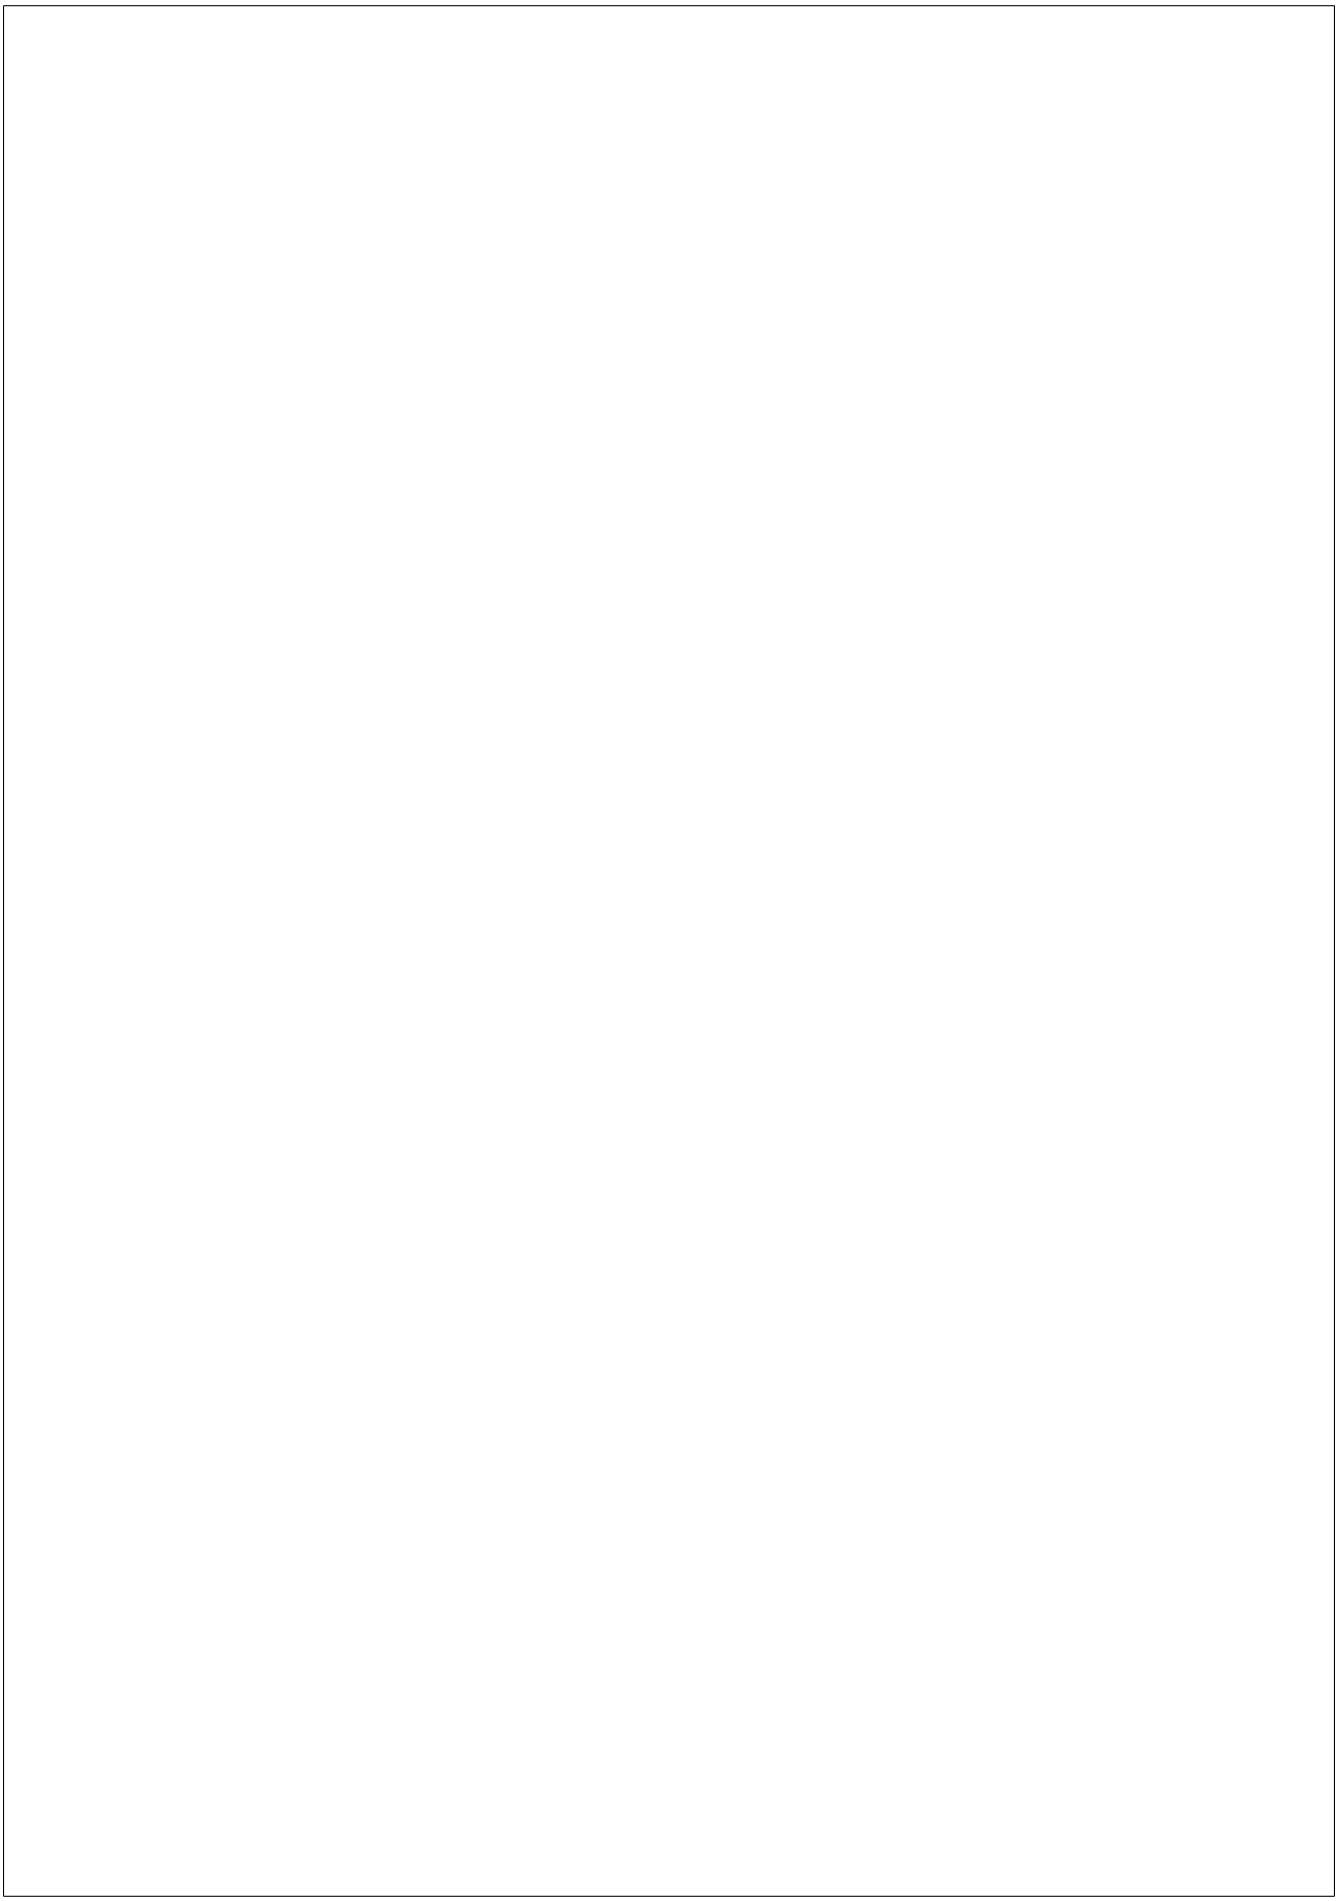

Bibliography

[1] Wright PE, Dyson HJ. Intrinsically unstructured proteins: re-assessing the protein structure-function paradigm. Journal of molecular biology. 1999;293(2):321–331.

[2] van Der Lee R, Buljan M, Lang B, Weatheritt RJ, Daughdrill GW, Dunker AK, et al. Classification of intrinsically disordered regions and proteins. Chemical reviews. 2014;114(13):6589–6631.

[3] Romero P, Obradovic Z, Kissinger C, Villafranca J, Dunker A. Identifying disordered regions in proteins from amino acid sequence. In: Proceedings of International Conference on Neural Networks (ICNN’97). vol. 1. IEEE; 1997. p. 90–95.

[4] Dunker AK, Lawson JD, Brown CJ, Williams RM, Romero P, Oh JS, et al. Intrinsically disordered protein. Journal of molecular graphics and modelling. 2001;19(1):26–59.

[5] Iakoucheva LM, Brown CJ, Lawson JD, Obradović Z, Dunker AK. Intrinsic disorder in cell-signaling and cancer-associated proteins. Journal of molecular biology. 2002;323(3):573–584.

[6] Liu J, Perumal NB, Oldfield CJ, Su EW, Uversky VN, Dunker AK. Intrinsic disorder in transcription factors. Biochemistry. 2006;45(22):6873–6888.

[7] van der Lee R, Lang B, Kruse K, Gsponer J, de Groot NS, Huynen MA, et al. Intrinsically disordered segments affect protein half-life in the cell and during evolution. Cell reports. 2014;8(6):1832–1844.

[8] Buljan M, Chalancon G, Dunker AK, Bateman A, Balaji S, Fuxreiter M, et al. Alternative splicing of intrinsically disordered regions and rewiring of protein interactions. Current opinion in structural biology. 2013;23(3):443–450.

[9] Stanley N, Esteban-Martín S, De Fabritiis G. Kinetic modulation of a disordered protein domain by phosphorylation. Nature communications. 2014;5:5272.

[10] Bah A, Forman-Kay JD. Modulation of intrinsically disordered protein function by post-translational modifications. Journal of Biological Chemistry. 2016;291(13):6696–6705.

[11] Burger V, Gurry T, Stultz C. Intrinsically disordered proteins: where computation meets experiment. Polymers. 2014;6(10):2684–2719.

[12] Kussie PH, Gorina S, Marechal V, Elenbaas B, Moreau J, Levine AJ, et al. Structure of the MDM2 oncoprotein bound to the p53 tumor suppressor transactivation domain. Science. 1996;274(5289):948–953.

[13] Radhakrishnan I, Pérez-Alvarado GC, Parker D, Dyson HJ, Montminy MR, Wright PE. Solution structure of the KIX domain of CBP bound to the transactivation domain of CREB: a model for activator: coactivator interactions. Cell. 1997;91(6):741–752.

[14] De Guzman RN, Martinez-Yamout MA, Dyson HJ, Wright PE. Interaction of the TAZ1 Domain of the CREB-Binding Protein with the Activation Domain of CITED2 Regulation by competition between intrinsically unstructured ligand for non-identical binding sites. Journal of Biological Chemistry. 2004;279(4):3042–3049.

[15] Zor T, De Guzman RN, Dyson HJ, Wright PE. Solution structure of the KIX domain of CBP bound to the transactivation domain of c-Myb. Journal of molecular biology. 2004;337(3):521–534.

[16] De Guzman RN, Goto NK, Dyson HJ, Wright PE. Structural basis for cooperative transcription factor binding to the CBP coactivator. Journal of molecular biology. 2006;355(5):1005–1013.

[17] Bah A, Vernon RM, Siddiqui Z, Krzeminski M, Muhandiram R, Zhao C, et al. Folding of an intrinsically disordered protein by phosphorylation as a regulatory switch. Nature. 2015;519(7541):106.

[18] Krois AS, Ferreon JC, Martinez-Yamout MA, Dyson HJ, Wright PE. Recognition of the disordered p53 transactivation domain by the transcriptional adapter zinc finger domains of CREB-binding protein. Proceedings of the National Academy of Sciences. 2016;113(13):E1853–E1862.

[19] Borgia A, Borgia MB, Bugge K, Kissling VM, Heidarsson PO, Fernandes CB, et al. Extreme disorder in an ultrahigh-affinity protein complex. Nature. 2018;555(7694):61.

[20] Uversky VN, Dunker AK. Understanding protein non-folding. Biochimica et Biophysica Acta (BBA)-Proteins and Proteomics. 2010;1804(6):1231–1264.

[21] Liu Z, Huang Y. Advantages of proteins being disordered. Protein Science. 2014;23(5):539–550.

[22] Davey NE, Van Roey K, Weatheritt RJ, Toedt G, Uyar B, Altenberg B, et al. Attributes of short linear motifs. Molecular BioSystems. 2012;8(1):268–281.

[23] Shammas SL, Crabtree MD, Dahal L, Wicky BI, Clarke J. Insights into coupled folding and binding mechanisms from kinetic studies. Journal of Biological Chemistry. 2016;291(13):6689–6695.

[24] Heller GT, Sormanni P, Vendruscolo M. Targeting disordered proteins with small molecules using entropy. Trends in biochemical sciences. 2015;40(9):491–496.

[25] Russo AA, Jeffrey PD, Patten AK, Massagué J, Pavletich NP. Crystal structure of the p27Kip1 cyclin-dependent-kinase inibitor bound to the cyclin A–Cdk2 complex. Nature. 1996;382(6589):325.

[26] Hyman AA, Weber CA, Jülicher F. Liquid-liquid phase separation in biology. Annual review of cell and developmental biology. 2014;30:39–58.

[27] Owen I, Shewmaker F. The role of post-translational modifications in the phase transitions of intrinsically disordered proteins. International journal of molecular sciences. 2019;20(21):5501.

[28] Nott TJ, Petsalaki E, Farber P, Jervis D, Fussner E, Plochowietz A, et al. Phase transition of a disordered nuage protein generates environmentally responsive membraneless organelles. Molecular cell. 2015;57(5):936–947.

[29] Vassilev LT, Vu BT, Graves B, Carvajal D, Podlaski F, Filipovic Z, et al. In vivo activation of the p53 pathway by small-molecule antagonists of MDM2. Science. 2004;303(5659):844–848.

[30] Andersen RJ, Mawji NR, Wang J, Wang G, Haile S, Myung JK, et al. Regression of castrate-recurrent prostate cancer by a small-molecule inhibitor of the amino-terminus domain of the androgen receptor. Cancer cell. 2010;17(6):535–546.

[31] De Mol E, Fenwick RB, Phang CT, Buzón V, Szulc E, De La Fuente A, et al. EPI-001, a compound active against castration-resistant prostate cancer, targets transactivation unit 5 of the androgen receptor. ACS chemical biology. 2016;11(9):2499–2505.

[32] Neira JL, Bintz J, Arruebo M, Rizzuti B, Bonacci T, Vega S, et al. Identification of a drug targeting an intrinsically disordered protein involved in pancreatic adenocarcinoma. Scientific reports. 2017;7:39732.

[33] Iconaru LI, Ban D, Bharatham K, Ramanathan A, Zhang W, Shelat AA, et al. Discovery of Small Molecules that Inhibit the Disordered Protein, p27 Kip1. Scientific reports. 2015;5:15686.

[34] Ban D, Iconaru LI, Ramanathan A, Zuo J, Kriwacki RW. A small molecule causes a population shift in the conformational landscape of an intrinsically disordered protein. Journal of the American Chemical Society. 2017;139(39):13692–13700.

[35] Hammoudeh DI, Follis AV, Prochownik EV, Metallo SJ. Multiple independent binding sites for small-molecule inhibitors on the oncoprotein c-Myc. Journal of the American Chemical Society. 2009;131(21):7390–7401.

[36] Krishnan N, Koveal D, Miller DH, Xue B, Akshinthala SD, Kragelj J, et al. Targeting the disordered C terminus of PTP1B with an allosteric inhibitor. Nature chemical biology. 2014;10(7):558.

[37] Dror RO, Dirks RM, Grossman J, Xu H, Shaw DE. Biomolecular simulation: a computational microscope for molecular biology. Annual review of biophysics. 2012;41:429–452.

[38] Kikhney AG, Svergun DI. A practical guide to small angle X-ray scattering (SAXS) of flexible and intrinsically disordered proteins. FEBS letters. 2015;589(19PartA):2570–2577.

[39] Gibbs EB, Cook EC, Showalter SA. Application of NMR to studies of intrinsically disordered proteins. Archives of biochemistry and biophysics. 2017;628:57–70.

[40] Dyson HJ, Wright PE. Perspective: the essential role of NMR in the discovery and characterization of intrinsically disordered proteins. Journal of Biomolecular NMR. 2019;73(12):651–659.

[41] Gibbs E, Kriwacki R. Direct detection of carbon and nitrogen nuclei for high-resolution analysis of intrinsically disordered proteins using NMR spectroscopy. Methods. 2018;138:39–46.

[42] Zor T, Mayr BM, Dyson HJ, Montminy MR, Wright PE. Roles of phosphorylation and helix propensity in the binding of the KIX domain of CREB-binding protein by constitutive (c-Myb) and inducible (CREB) activators. Journal of Biological Chemistry. 2002;277(44):42241–42248.

[43] Sugase K, Dyson HJ, Wright PE. Mechanism of coupled folding and binding of an intrinsically disordered protein. Nature. 2007;447(7147):1021.

[44] Kendrew JC, Bodo G, Dintzis HM, Parrish R, Wyckoff H, Phillips DC. A three-dimensional model of the myoglobin molecule obtained by x-ray analysis. Nature. 1958;181(4610):662–666.

[45] McCammon JA, Gelin BR, Karplus M. Dynamics of folded proteins. Nature. 1977;267(5612):585.

[46] Martinez-Rosell G, Giorgino T, Harvey MJ, de Fabritiis G. Drug discovery and molecular dynamics: methods, applications and perspective beyond the second timescale. Current topics in medicinal chemistry. 2017;17(23):2617–2625.

[47] Prinz JH, Wu H, Sarich M, Keller B, Senne M, Held M, et al. Markov models of molecular kinetics: Generation and validation. The Journal of chemical physics. 2011;134(17):174105.

[48] Chodera JD, Noé F. Markov state models of biomolecular conformational dynamics. Current opinion in structural biology. 2014;25:135–144.

[49] Doerr S, De Fabritiis G. On-the-fly learning and sampling of ligand binding by high-throughput molecular simulations. Journal of chemical theory and computation. 2014;10(5):2064–2069.

[50] Chen J, Brooks III CL, Khandogin J. Recent advances in implicit solvent-based methods for biomolecular simulations. Current opinion in structural biology. 2008;18(2):140–148.

[51] Jorgensen WL, Chandrasekhar J, Madura JD, Impey RW, Klein ML. Comparison of simple potential functions for simulating liquid water. The Journal of chemical physics. 1983;79(2):926–935.

[52] MacKerell Jr AD, Bashford D, Bellott M, Dunbrack Jr RL, Evanseck JD, Field MJ, et al. All-atom empirical potential for molecular modeling and dynamics studies of proteins. The journal of physical chemistry B. 1998;102(18):3586–3616.

[53] Cornell WD, Cieplak P, Bayly CI, Gould IR, Merz KM, Ferguson DM, et al. A second generation force field for the simulation of proteins, nucleic acids, and organic molecules. Journal of the American Chemical Society. 1995;117(19):5179–5197.

[54] Duan Y, Wu C, Chowdhury S, Lee MC, Xiong G, Zhang W, et al. A point-charge force field for molecular mechanics simulations of proteins based on condensed-phase quantum mechanical calculations. Journal of computational chemistry. 2003;24(16):1999–2012.

[55] MacKerell Jr AD, Banavali N, Foloppe N. Development and current status of the CHARMM force field for nucleic acids. Biopolymers: Original Research on Biomolecules. 2000;56(4):257–265.

[56] Galindo-Murillo R, Robertson JC, Zgarbova M, Sponer J, Otyepka M, Jurecka P, et al. Assessing the current state of amber force field modifications for DNA. Journal of chemical theory and computation. 2016;12(8):4114–4127.

[57] Notman R, Anwar J. Breaching the skin barrier—Insights from molecular simulation of model membranes. Advanced drug delivery reviews. 2013;65(2):237–250.

[58] Hospital A, Goñi JR, Orozco M, Gelpí JL. Molecular dynamics simulations: advances and applications. Advances and applications in bioinformatics and chemistry: AABC. 2015;8:37.

[59] Monticelli L, Kandasamy SK, Periole X, Larson RG, Tieleman DP, Marrink SJ. The MARTINI coarse-grained force field: extension to proteins. Journal of chemical theory and computation. 2008;4(5):819–834.

[60] Shaw DE, Deneroff MM, Dror RO, Kuskin JS, Larson RH, Salmon JK, et al. Anton, a special-purpose machine for molecular dynamics simulation. Communications of the ACM. 2008;51(7):91–97.

[61] Laio A, Parrinello M. Escaping free-energy minima. Proceedings of the National Academy of Sciences. 2002;99(20):12562–12566.

[62] Dickson BM. Approaching a parameter-free metadynamics. Physical Review E. 2011;84(3):037701.

[63] Hansen HS, Hünenberger PH. Using the local elevation method to construct optimized umbrella sampling potentials: calculation of

the relative free energies and interconversion barriers of glucopyranose ring conformers in water. Journal of computational chemistry. 2010;31(1):1–23.

[64] Harvey MJ, Giupponi G, Fabritiis GD. ACEMD: accelerating biomolecular dynamics in the microsecond time scale. Journal of chemical theory and computation. 2009;5(6):1632–1639.

[65] Buch I, Harvey MJ, Giorgino T, Anderson DP, De Fabritiis G. High-throughput all-atom molecular dynamics simulations using distributed computing. Journal of chemical information and modeling. 2010;50(3):397–403.

[66] Huang J, MacKerell Jr AD. Force field development and simulations of intrinsically disordered proteins. Current opinion in structural biology. 2018;48:40–48.

[67] Best RB, Zheng W, Mittal J. Balanced protein–water interactions improve properties of disordered proteins and non-specific protein association. Journal of chemical theory and computation. 2014;10(11):5113–5124.

[68] Piana S, Donchev AG, Robustelli P, Shaw DE. Water dispersion interactions strongly influence simulated structural properties of disordered protein states. The journal of physical chemistry B. 2015;119(16):5113–5123.

[69] Song D, Luo R, Chen HF. The IDP-specific force field ff14IDPSFF improves the conformer sampling of intrinsically disordered proteins. Journal of chemical information and modeling. 2017;57(5):1166–1178.

[70] Liu H, Song D, Lu H, Luo R, Chen HF. Intrinsically disordered protein-specific force field CHARMM 36 IDPSFF. Chemical biology & drug design. 2018;92(4):1722–1735.

[71] Piana S, Robustelli P, Tan D, Chen S, Shaw DE. Development of a force field for the simulation of single-chain proteins and protein-protein complexes. Journal of Chemical Theory and Computation. 2020;.

[72] Piana S, Lindorff-Larsen K, Shaw DE. How robust are protein folding simulations with respect to force field parameterization? Biophysical journal. 2011;100(9):L47–L49.

[73] Pérez-Hernández G, Paul F, Giorgino T, De Fabritiis G, Noé F. Identification of slow molecular order parameters for Markov model construction. The Journal of chemical physics. 2013;139(1):07B604_1.

[74] Pearson K. LIII. On lines and planes of closest fit to systems of points in space. The London, Edinburgh, and Dublin Philosophical Magazine and Journal of Science. 1901;2(11):559–572.

[75] Scherer MK, Trendelkamp-Schroer B, Paul F, Pérez-Hernández G, Hoffmann M, Plattner N, et al. PyEMMA 2: A software package for estimation, validation, and analysis of Markov models. Journal of chemical theory and computation. 2015;11(11):5525–5542.

[76] Pedregosa F, Varoquaux G, Gramfort A, Michel V, Thirion B, Grisel O, et al. Scikit-learn: Machine learning in Python. Journal of machine learning research. 2011;12(Oct):2825–2830.

[77] Deuflhard P, Weber M. Robust Perron cluster analysis in conformation dynamics. Linear algebra and its applications. 2005;398:161–184.

[78] Noé F, Schütte C, Vanden-Eijnden E, Reich L, Weikl TR. Constructing the equilibrium ensemble of folding pathways from short off-equilibrium simulations. Proceedings of the National Academy of Sciences. 2009;106(45):19011–19016.

[79] Voelz VA, Bowman GR, Beauchamp K, Pande VS. Molecular simulation of ab initio protein folding for a millisecond folder NTL9 (1-39). Journal of the American Chemical Society. 2010;132(5):1526–1528.

[80] Buch I, Giorgino T, De Fabritiis G. Complete reconstruction of an enzyme-inhibitor binding process by molecular dynamics simulations. Proceedings of the National Academy of Sciences. 2011;108(25):10184–10189.

[81] Dror RO, Pan AC, Arlow DH, Borhani DW, Maragakis P, Shan Y, et al. Pathway and mechanism of drug binding to G-protein-coupled receptors. Proceedings of the National Academy of Sciences. 2011;108(32):13118–13123.

[82] Weinan E, Vanden-Eijnden E. Towards a theory of transition paths. Journal of statistical physics. 2006;123(3):503.

[83] Zimmerman MI, Bowman GR. FAST conformational searches by balancing exploration/exploitation trade-offs. Journal of chemical theory and computation. 2015;11(12):5747–5757.

[84] Plattner N, Doerr S, De Fabritiis G, Noé F. Complete protein–protein association kinetics in atomic detail revealed by molecular dynamics simulations and Markov modelling. Nature chemistry. 2017;9(10):1005.

[85] Singhal N, Pande VS. Error analysis and efficient sampling in Markovian state models for molecular dynamics. The Journal of chemical physics. 2005;123(20):204909.

[86] Hinrichs NS, Pande VS. Calculation of the distribution of eigenvalues and eigenvectors in Markovian state models for molecular dynamics. The Journal of chemical physics. 2007;126(24):244101.

[87] Sutton RS, Barto AG. Reinforcement learning: An introduction. MIT press; 2018.

[88] Pérez A, Herrera-Nieto P, Doerr S, Fabritiis GD. AdaptiveBandit: A multi-armed bandit framework for adaptive sampling in molecular simulations; 2020.

[89] Auer P. Using confidence bounds for exploitation-exploration trade-offs. Journal of Machine Learning Research. 2002;3(Nov):397–422.

[90] Doerr S, Harvey M, Noé F, De Fabritiis G. HTMD: high-throughput molecular dynamics for molecular discovery. Journal of chemical theory and computation. 2016;12(4):1845–1852.

[91] Humphrey W, Dalke A, Schulten K. VMD: visual molecular dynamics. Journal of molecular graphics. 1996;14(1):33–38.

[92] Camilloni C, De Simone A, Vranken WF, Vendruscolo M. Determination of secondary structure populations in disordered states of proteins using nuclear magnetic resonance chemical shifts. Biochemistry. 2012;51(11):2224–2231.

[93] Zwier MC, Pratt AJ, Adelman JL, Kaus JW, Zuckerman DM, Chong LT. Efficient atomistic simulation of pathways and calculation of rate constants for a protein–peptide binding process: application to the MDM2 protein and an intrinsically disordered p53 peptide. The journal of physical chemistry letters. 2016;7(17):3440–3445.

[94] Morrone JA, Perez A, MacCallum J, Dill KA. Computed binding of peptides to proteins with MELD-accelerated molecular dynamics. Journal of chemical theory and computation. 2017;13(2):870–876.

[95] Zhou G, Pantelopulos GA, Mukherjee S, Voelz VA. Bridging microscopic and macroscopic mechanisms of p53-MDM2 binding with kinetic network models. Biophysical journal. 2017;113(4):785–793.

[96] Paul F, Noé F, Weikl TR. Identifying conformational-selection and induced-fit aspects in the binding-induced folding of PMI from Markov state modeling of atomistic simulations. The Journal of Physical Chemistry B. 2018;122(21):5649–5656.

[97] Wells M, Tidow H, Rutherford TJ, Markwick P, Jensen MR, Mylonas E, et al. Structure of tumor suppressor p53 and its intrinsically disordered N-terminal transactivation domain. Proceedings of the National academy of Sciences. 2008;105(15):5762–5767.

[98] Giri R, Morrone A, Toto A, Brunori M, Gianni S. Structure of the transition state for the binding of c-Myb and KIX highlights an unexpected order for a disordered system. Proceedings of the National Academy of Sciences. 2013;110(37):14942–14947.

[99] Shammas SL, Travis AJ, Clarke J. Remarkably fast coupled folding and binding of the intrinsically disordered transactivation domain of cMyb to CBP KIX. The journal of physical chemistry B. 2013;117(42):13346–13356.

[100] Arai M, Sugase K, Dyson HJ, Wright PE. Conformational propensities of intrinsically disordered proteins influence the mechanism of binding and folding. Proceedings of the National Academy of Sciences. 2015;112(31):9614–9619.

[101] Dahal L, Kwan TO, Shammas SL, Clarke J. pKID binds to KIX via an unstructured transition state with nonnative interactions. Biophysical journal. 2017;113(12):2713–2722.

[102] Dahal L, Shammas SL, Clarke J. Phosphorylation of the IDP KID modulates affinity for KIX by increasing the lifetime of the complex. Biophysical journal. 2017;113(12):2706–2712.

[103] Gianni S, Morrone A, Giri R, Brunori M. A folding-after-binding mechanism describes the recognition between the transactivation

domain of c-Myb and the KIX domain of the CREB-binding protein. Biochemical and biophysical research communications. 2012;428(2):205–209.

[104] Shammas SL, Travis AJ, Clarke J. Allostery within a transcription coactivator is predominantly mediated through dissociation rate constants. Proceedings of the National Academy of Sciences. 2014;111(33):12055–12060.

[105] Jin F, Yu C, Lai L, Liu Z. Ligand clouds around protein clouds: a scenario of ligand binding with intrinsically disordered proteins. PLoS computational biology. 2013;9(10):e1003249.

[106] Ferruz N, Doerr S, Vanase-Frawley MA, Zou Y, Chen X, Marr ES, et al. Dopamine D3 receptor antagonist reveals a cryptic pocket in aminergic GPCRs. Scientific reports. 2018;8(1):897.

[107] Paul F, Wehmeyer C, Abualrous ET, Wu H, Crabtree MD, Schöneberg J, et al. Protein-peptide association kinetics beyond the seconds timescale from atomistic simulations. Nature communications. 2017;8(1):1095.

[108] Zimmerman MI, Porter JR, Sun X, Silva RR, Bowman GR. Choice of adaptive sampling strategy impacts state discovery, transition probabilities, and the apparent mechanism of conformational changes. Journal of chemical theory and computation. 2018;14(11):5459–5475.

[109] Poosapati A, Gregory E, Borcherds WM, Chemes LB, Daughdrill GW. Uncoupling the folding and binding of an intrinsically disordered protein. Journal of molecular biology. 2018;430(16):2389–2402.

[110] Ruan H, Sun Q, Zhang W, Liu Y, Lai L. Targeting intrinsically disordered proteins at the edge of chaos. Drug discovery today. 2019;24(1):217–227.