Quantitative Models of Molecular Dynamics from Sparse Simulation and Experimental Data

Christopher Kolloff[a] **and Simon Olsson**[a, *]

[a]Chalmers University of Technology, Department of Computer Science and Engineering, Rännvägen 6, 412 58 Gothenburg, Sweden

This manuscript was compiled on May 29, 2023

Studying the long-timescale behavior of proteins is the key to understanding many of the fundamental processes of life. Molecular Dynamics (MD) simulations and biophysical experiments probe the dynamics of such systems. However, while MD aims to simulate the processes detected in experiments, their predictions are often not in quantitative agreement. Reconciling these differences is a significant opportunity to build quantitative mechanistic models of these systems. To this end, here we present dynamic Augmented Markov Models (dynAMMo), a new approach to integrate dynamic experimental observables, such as NMR relaxation dispersion data, with a Markov state model derived from MD simulation statistics. We find that integrating experimental data that are sensitive to dynamic processes allows us to accurately recover the unbiased kinetics from biased MD simulations. Further, we show that dynAMMo can recover exchange processes not observed in MD data and yield a kinetic model reconciling experiment and disconnected simulations, something which has not yet been possible. We demonstrate the effectiveness of dynAMMo using well-controlled model systems and show the broad applicability of the method on a well-studied protein system. Our approach opens up a wealth of new opportunities to quantitatively study protein structure and dynamics from a mechanistic point of view.

Markov state models | protein dynamics | Molecular Dynamics simulations | relaxation dispersion NMR spectroscopy |

Understanding the triad of protein structure–function–dynamics is of paramount importance in many fields, including biochemistry, biophysics, and medicine (1–7). Thanks to extensive studies of the bovine pancreatic trypsin inhibitor (BPTI), for example, we now understand the atomic details of its essential role in inhibiting serine proteases (8). This knowledge has been possible by combining findings from the fields of X-ray crystallography (9) and Nuclear Magnetic Resonance (NMR) (10–13) with Molecular Dynamics (MD) simulations (14–16). However, reconciling experimental and simulation data in a systematic manner often poses problems due to technical and resource limitations. Enabling such a merger would yield a significant opportunity for quantitative structural biology and biophysics.

Typically, we model dynamic experiments, (17) such as NMR relaxation dispersion, single molecular Förster Resonance Energy Transfer (FRET), dynamic neutron scattering, or X-ray Photon Correlation Spectroscopy using simple n-site jump models (18–22). This approach yields forward and reverse exchange rates for the different states as well as site populations. However, modeling dynamics this way beyond a two-state exchange is challenging, due to experimental limitations and poor timescale separation. In effect, we are limited to highly simplified models of the complex underlying dynamics of our data where the structure of states often remain elusive or ambiguous (23–28).

Over the past few decades, molecular dynamics (MD) simulations have become increasingly popular in the field of biophysics, providing atomistic insights into the behavior of biological systems at high temporal and spatial resolutions (29–31). Force field models are steadily improving in quality and their scope is boarding to include disordered proteins and nucleic acids (32–34). Although not broadly available, the development of special-purpose computers, like Anton (35–37), makes it possible to study millisecond timescale molecular processes. GPU-accelerated simulations as well as distributed computing initiatives, such as Folding@home (38) or GPUGrid are more widely available allow us to access processes on the micro- to millisecond timescale, in particular when large ensembles of simulations are analyzed using Markov state models (MSM) (39–43). MSMs represent the *molecular kinetics* fully: the relevant structural states, their thermodynamic weights and their mutual exchange rates. MSMs have enabled studies of many biological processes, such as protein folding, enzymatic activity, or protein–protein interactions (42, 44–50).

Despite these advances in force field accuracy and simulation technologies, we often observe a systematic discrepancy between the experimental values predicted from Molecular Dynamics (MD) trajectories and experimental data (51–54). The origin of these discrepancies is two fold. First, imperfections in the force field models remain which lead to skewed populations and altered dynamics. Second, simulations still do not cover the range of biological timescales of interest. Together these limitations prevent us from directly comparing to many experiments and thus gaining a mechanistic interpretation of our data.

The integration of simulation and experimental data is a big challenge with a long history (55). Previous methods include *post hoc* reweighing or sub-selection of simulations data (56, 57), modeling kinetics with generated ensembles (58), biasing simulations with stationary experimental data (59–68) and dynamic experimental data (69), and building Markov state models using experimental and simulation data (70), Augmented Markov models. Augmented Markov Models (AMMs) combine stationary experimental observables with simulation data to correct for bias in MD data, which improved agreement with complementary stationary and dynamic data (70). However, AMMs cannot take into account dynamic data, and consequently also cannot deal with situations where our simulations have not sampled processes which are important to explain the data.

Here, we present dynamic Augmented Markov Models

C.K. and S.O. conceptualized, designed, and performed research as well as wrote the manuscript; C.K. analyzed and visualized the data and performed statistical analysis; S.O. provided supervision, project administration and funding acquisition.

The authors declare no competing interests.

*To whom correspondence should be addressed. E-mail: simonols@chalmers.se

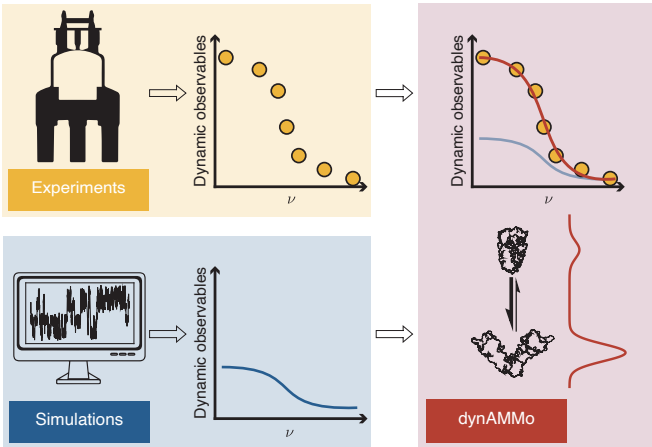

Fig. 1. Schematic of dynamic Augmented Markov Models (dynAMMo). Simulations often fail to reproduce experimentally observed dynamic observables, such as NMR relaxation dispersion data due to force field inaccuracies and finite simulation length. Experiments, on the other hand, are often noisy and often don't reveal the mechanistic details of kinetic exchange. With dynamic Augmented Markov Models, we take into account, both, simulation and experimental data to obtain a single mechanistic kinetic model.

(dynAMMo), a new approach that accounts for stationary and dynamic experimental data, such as $R_{1\rho}$ or Carr-Purcell-Meiboom-Gill (CPMG) relaxation dispersion experiments (20, 71, 72), to estimate a Markov model. By combining constrained optimization with the principle of maximum entropy, we are able to correctly recover experimental timescales from biased simulations and are able to model exchange between states not seen in the simulation data as long as the states themselves are known. The method therefore broadens the scope for future research in understanding the complex dynamics of biomolecular systems in general and brings the field closer to the development of data-driven models that accurately capture the underlying mechanisms-of-action.

Theory

Dynamic Augmented Markov Models. Markov state models are based on the discretization of the state space Ω of a molecule into n states. By following the traversal of a MD trajectory through these states we can estimate the transition probabilities p_{ij} from states i to states j through the analysis of transition counts $c_{ij}(\tau)$ with a lag time τ (39, 40, 43). The resulting transition matrix $\mathbf{T}(\tau) \in \mathbb{R}^{n\times n}$ encodes the *molecular kinetics* of the system, including the populations of the states and the rates of exchange between them. This information is accessible through the spectral components of $T(\tau)$, the eigenvectors $\mathbf{R}$ and eigenvalues $\boldsymbol{\lambda}$, as well as the stationary distribution $\boldsymbol{\pi}$ (see Supporting Information, 'Theory') (43).

Augmented Markov Models (70) aim to estimate a MSM which matches stationary experimental observables, such as NMR ^3J-coupling or Residual Dipolar Coupling (RDC) data that probe the "true" Boltzmann distribution, by reweighing the relative populations of the states through the maximum entropy principle. Even though this approach does not directly take into account information about the kinetic rates between the states, Olsson and Noé observe that the integration of stationary observables has an effect on the prediction of dynamic observables, such as $R_{1\rho}^{\text{ex}}$ relaxation dispersion. However, in general we cannot expect AMMs to match dynamic experimental data, nor can they consider cases in which not all states of the MSM are connected.

Here, we address these limitations with dynamic Augmented Markov models that combine simulation data, in the form of one or more count matrices $\mathcal{C}$ (Supplementary Information, algorithm 1, line 1), and dynamic and stationary experimental observables $\mathbf{o}^{\text{exp}}$ to a single kinetic model. By combining these sources of information, we aim to obtain a more accurate and comprehensive representation of biomolecular dynamics. Using experimental data that report on the conformational exchange kinetics, we can directly estimate the forward and reverse rate constants of switching from one state to another. Unlike Brotzakis et al., no prior knowledge of the kinetic rates are required nor are we limited to a two-state exchange.

Connection between experiments and simulations. For each experimental observable $\mathbf{o}^{\text{exp}}$, we assume that there is a corresponding observable function $f(\cdot)$ (*forward model*) available that maps all configurations, $\mathbf{x} \in \Omega$, to a complex or real vector space, $\mathcal{V}$, however, often just a scalar, e.g., a distance or a chemical shift. For MSMs, we can average these values over the n Markov states yielding $\mathbf{a} \in \mathcal{V}^n$ (70). Here, we focus on dynamic experiments, where we measure time correlations of these observables either directly or through a transformation. From an MSM, $\mathbf{T}(\tau)$, we can compute the time correlation of f:

$$\langle f(\mathbf{x}(0))^\top f(\mathbf{x}(k\tau))\rangle \approx o^{\text{dynamic}}(k) = \mathbf{a}^\top \boldsymbol{\Pi}\mathbf{T}(\tau)^k\mathbf{a}, \quad [1]$$

where $\boldsymbol{\Pi}$ is the diagonal matrix of the stationary distribution $\boldsymbol{\pi}$, and $^\top$ is the transpose or the complex conjugate. We can compare this quantity directly to experimentally measured counterparts and thereby use it to drive the estimation of MSMs. Many dynamic observables, however, are transformations of the time-correlation, rather than the time-correlation itself. This includes, among others, Carl-Purcell-Meiboom-Gill (CPMG) and $R_{1\rho}^{\text{ex}}$ relaxation dispersion, which measure the convolution of the time-correlation function with a spin-lock field. Here we assume fast chemical exchange, and use previously described closed-form expression for Markov models (74, 75) to predict data and drive MSM estimation (see Supporting Information, 'Theory' for a more detailed explanation).

Estimation of dynamic Augmented Markov Models. We estimate dynAMMo models by optimizing a loss function which includes the transition counts from the between states i and j and the sum of the mean-square difference between the predictions and experimental data, $\mathcal{D}$, of the lth observable at the kth lag-time,

$$\arg\min_{\hat{\boldsymbol{\lambda}},\hat{\mathbf{R}},\hat{\boldsymbol{\pi}}} \mathcal{L}(\hat{\mathbf{T}}(\tau) \mid \mathcal{D}, \mathbf{C}(\tau)) = -\sum_{ij} c_{ij}\log p_{ij} + \sum_{l,k}(o_{l,k}^{\text{pred, dyn}} - o_{l,k}^{\text{exp, dyn}})^2. \quad [2]$$

The loss is computed with respect to the spectrum of $\hat{\mathbf{T}}(\tau)$, i.e., the eigenvalues $\hat{\boldsymbol{\lambda}}$, eigenvectors $\hat{\mathbf{R}}$, and the stationary distribution $\hat{\boldsymbol{\pi}}$ and is subject to several constraints. To estimate $\hat{\boldsymbol{\lambda}}$ and $\hat{\mathbf{R}}$, we use gradient descent with additional orthogonality constraints for optimization of the eigenvectors. Rather

than enforcing orthogonality directly with a penalty term in the loss function, dynAMMo optimizes the eigenvectors on the Stiefel manifold through Riemannian optimization (76). Following AMMs, we further have the option to include stationary experimental observables as described previously (70). The estimation procedure as well as the theoretical details are explained in more detail in Supporting Information, 'Theory.'

Results and Discussion

Enforcing dynamic experimental constraints on a kinetic model rescues biased simulation data. To demonstrate the power of dynamic Augmented Markov Models, we will first examine our model by applying it to two one-dimensional energy potentials: the Prinz potential (43) (Fig. 2, brown background) and the three-well potential (Fig. 2, teal background). In both model systems, there is one slow transition as well as one or more faster transition(s) that we aim to model. The Prinz potential has four states with comparable populations with one slow transition between the first and last two states, whereas the three-well potential has two fast-interchanging low-energy states and one state with a high energy barrier. In both model systems, we used the slowest eigenvectors as an observable function as they encode near perfect information about the slowest process.

We will first examine the scenario where we have biased simulations, which we compare with the experimental data derived from a "ground truth" model. In this case, all states were reversibly sampled in the simulations. However, due to, for example, force field inaccuracies and finite sampling, the timescales of exchange between the processes and the thermodynamics of the system do not correspond to the "true" ensemble. By investigating the free energy profiles of the two systems (Fig. 2A/J), we see that the populations of the MSM (blue) match the ground truth well (yellow). Consequently, we would not expect reweighing considering only the stationary observables, as is done in AMMs, will not have a big effect on the prediction of the dynamic observables since the timescales show significant discrepancies (Fig. 2B/K) between the MSM and the ground truth. Using dynAMMo, we can perfectly match the slowest timescales (Fig. 2B/K). We find a similar mismatch between the MSMs estimated on the biased data and the "ground truth data" (Fig. 2C/L). Since our observable function inherently informs about the slowest process we find that, dynAMMo does not substantially modify the timescale of the faster processes (Fig. 2B). This implies that the model does not introduce unnecessary bias into the estimation if it is not reflected in the observable. As opposed to the Prinz potential, we find that the predicted kinetics in the model trained on biased data from the three-well potential (Fig. 2, panels J–L) are accelerated compared to the ground truth (Fig. 2K). This mismatch in timescales manifests itself as poor agreement with the observable time correlation functions (Fig. 2L) between the ground truth (yellow) and the naive MSM (blue). Integrating the correlation function data and the biased simulation data with dynAMMo, we are in agreement with the ground-truth data and match the underlying rates.

Disconnected simulations can be combined to a single Markov model using dynamic constraints. Many systems are characterized by timescales which are impractical to sample with

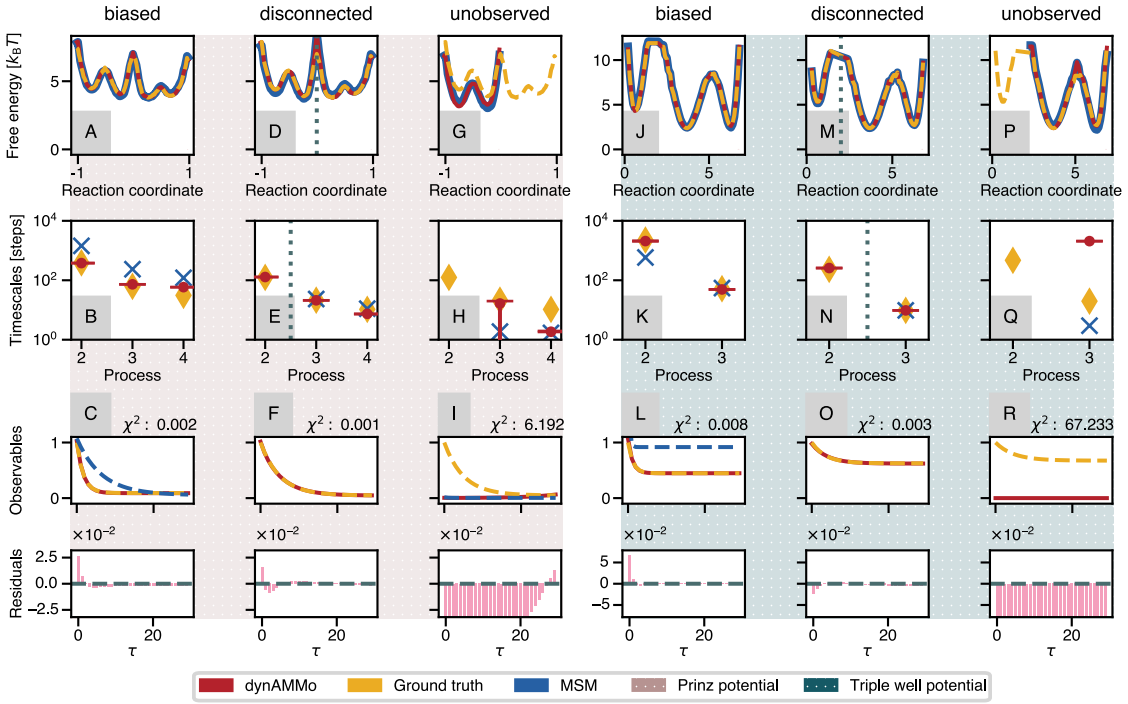

Fig. 2. Overview of model system benchmark results. Three different scenarios using two model systems, the Prinz potential (brown background) and the three-well potential (teal background) are shown. The three different scenarios include biased simulations (A–C, J–L), disconnected trajectories (D–F, M–O), and unobserved or missing states (G–I, P–R). Each of the scenarios show the free energy potential, timescales of exchange, and observable plots. χ^2 values and residuals between the model predictions and the ground truth data are also shown for the observable plots. The model results are shown in red, the 'ground truth' experimental data in yellow, and the naive MSM predictions in blue. ΔG values of the slowest transition for the different systems are given in the Supporting Information (Table S4).

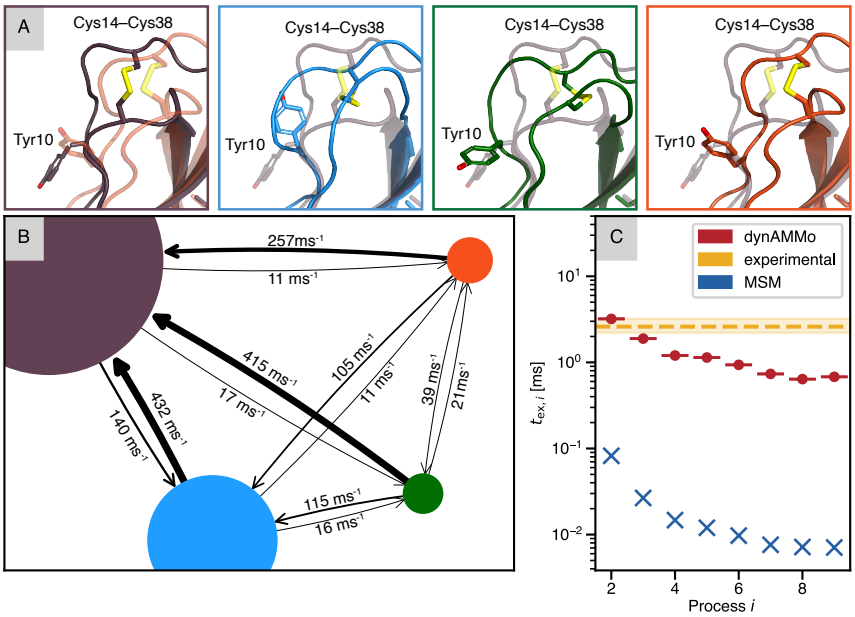

Fig. 3. Integrating BPTI simulations and CPMG NMR data to build a quantitative kinetic model with dynAMMo. (A) Representative structures of the major BPTI states with the structural characteristics highlighted. Aromatics are shown for orientation. (B) Kinetic network between macrostates. The colors of the nodes correspond to the structural representatives shown above. The size of the nodes and the arrow widths are proportional to the size of the populations as well as the reaction rate, respectively. Rates are shown above / below the arrows in ms^{-1}. (C) Timescales of exchange as a function of the slowest processes. dynAMMo is shown in red and the MSM from the simulation data is colored blue. The time constant of the experimentally determined exchange is shown in yellow with the standard deviation shown as shaded area.

statistical confidence. However, we may have access to multiple experimental structures of each of the states in isolation, some of which we can sample transitions between, others which are infeasible to sample. An example is the bovine pancreatic trypsin inihibitor (BPTI), for which numerous studies have reported slow millisecond timescale dynamics (73, 77–79), and a millisecond long MD simulation only sampled the suspected slow transition once (16). In many other cases, sampling such a transition in an unbiased fashion remains impractical.

To test such a scenario, we designed two experiments where we discard the transition counts of the slowest transition (Fig. 2D–E/M–N, gray dotted line) and split the trajectory in two. We build two Markov state models corresponding to the, now, disconnected subregions of the state space (Supporting Information, 'Materials and Methods'). After reweighing the populations using the stationary AMM procedure (70), we perfectly match the model (red) and the ground truth (yellow) populations (Fig. 2D/M). Despite not having prior information on the slowest process, we can correctly identify the missing timescale using dynamic experimental observables (Fig. 2E/N). This discovery is guided by the correlation function (Fig. 2F/O). Since we can match, both, the kinetics and the thermodynamics of the systems, we can also fit the observable prediction (Fig. 2F/O red, solid) to the ground truth (Fig. 2F/O yellow, dashed). Using experimental dynamic observables, we can thereby merge disconnected simulation statistics and estimate a single model which faithfully reproduces all the available data.

dynAMMo does not overfit when relevant states are missing. Next, we consider the case where one or more states that contribute to a measurable experimental signal is missing from the MD simulation data. This situation is common in MD simulation studies as the simulation time is often insufficient to sample all the relevant states, and is an edge case related to the 'disconnected' situation discussed above. However, contrary to the previous case, we do not have all structural information to support a reliable prediction of the observable here. We therefore anticipate that dynAMMo is unable to yield a model perfectly fitting these data. To simulate this scenario, we discard transition counts from one state completely. Concretely, this procedure discards simulation data about states above 0 (Fig. 2G–I) in the Prinz potential and states below 2.5 in the asymmetric triple well potential (Fig. 2P–R). In this case, models built with dynAMMo cannot match (Fig. 2I/R) the data, which translates into missing timescales (Fig. 2H/Q). Mismatching predictions indicate that the model is failing, which suggests that one or more states that give rise to a measurable signal are missing.

A mechanistic model of BPTI disulfide isomerization dynamics with dynAMMo. To test how our model performs in a realistic scenario, we turned to BPTI as a protein system. BPTI is a 58-residue protein whose dynamics has been extensively studied, both experimentally (77, 78, 80, 81) and computationally (16, 82). BPTI is known to have micro- to millisecond conformational exchange (77, 78, 83), centered mainly on different isomerizations of the disulfide bond between Cys14 and Cys38. In addition, there is a 1-ms long MD trajectory available at a temperature of 300 K (16). In this simulation, all known major conformations of BPTI are sampled, and the transitions between them show a distinct separation of timescales. Analysis of the trajectory shows conformational exchange in the fast microsecond regime (16, 70, 84), which is much faster than what the experimentally determined rates are suspected to reflect these processes. The discrepancy between the experiments and the simulation data makes BPTI an ideal test case for demonstrating dynAMMo as an avenue to reconcile the data.

By integrating simulation (16) and experimental CPMG data (73) from NMR spectroscopy with dynAMMo, we build a kinetic model of BPTI (Fig. 3). In line with previous analyses,(70, 85), we used time-lagged Independent Component Analysis (tICA) (86) to define a low dimensional space which we discretized into 384 states and aggregated into 4 metastable structural states. Consistent with previous analyses, the major structural substates display isomerization of the disulfide bridges between residues 14 and 38 (Fig. 3A).

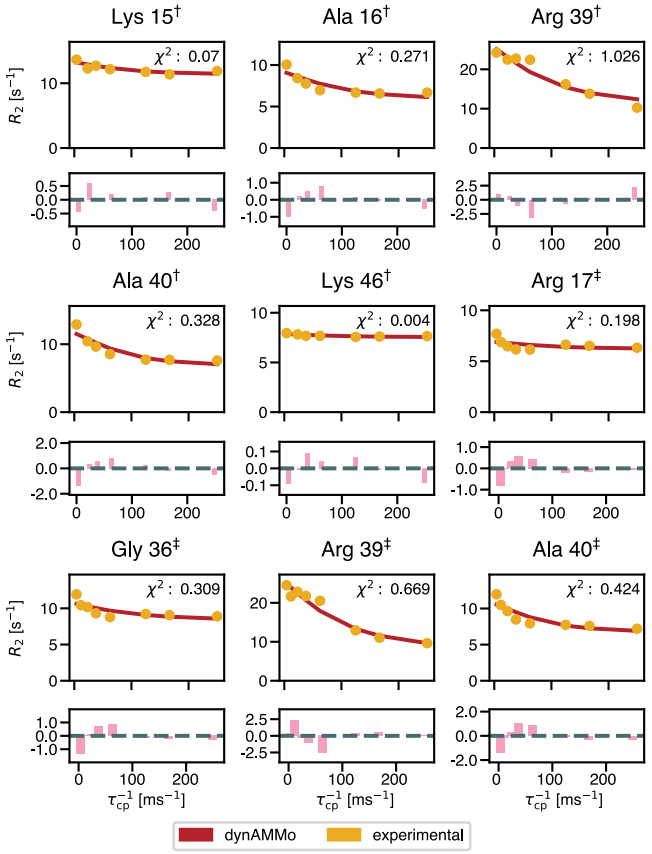

Fig. 4. CPMG relaxation dispersion data of BPTI $^{15}N^{H}$ spins. Nine representative examples of CPMG plots are shown. The fitted model (red, solid) is in high agreement with the experimental data (yellow circles). χ^2 values between the prediction and the experiments are shown for each subplot. The dagger † refers to the dataset that has been recorded at a Larmor frequency of 600 MHz. Conversely, the double dagger ‡ refers to the dataset recorded at 500 MHz (73). Residuals between the experiments and predictions are shown in pink.

We show the most populated states colored purple and light blue (a total population of approximately 90 %), while the remaining population is shared by the green and orange states (Fig. 3A). The state connectivity is dense and rates vary across an order of magnitude (Fig. 3B), which we show with arrows of varying thickness between the states colored by identity and scaled by their relative populations. The slowest rates correspond to the transitions to the two minor states (Supporting Information, Fig. S6a) and we expect to occur in the low millisecond regime. The implied timescales computed from our dynAMMo model are systematically shifted compared to those of the naive Markov state model that only takes simulation statistics into account (Fig. 3C). For the MSM, the slowest processes are barely on the order of hundreds of microseconds (dark blue crosses) (74). On the other hand, the slowest implied timescale estimated by dynAMMo is on the order of magnitude of approximately 3.2 ms. This timescale matches well with those estimated from experimental data at the same temperature using a two-state fit, where Millet et al. determined a chemical exchange of the order of $2 - 3$ ms (73).

In Fig. 4 we show representative examples of some key observables. The plots show CPMG relaxation dispersion curves, which measure the effect of chemical exchange on $^{15}N^{H}$ spins. The chemical shift predictions that were used as observables for the backbone amides were obtained by the PPM algorithm (87). The experimental CPMG data are shown in yellow, whereas the predictions of the model are shown in red and the predictions were scaled according to the values reported in the Supporting Information (Fig. S7). All observables show an excellent overall agreement with the data, suggesting that the underlying model is capable of explaining the data in a meaningful way (see Supporting Information, Fig. S8 and S9). We note that all relevant residues involved in the exchange display a relatively strong dispersion, which we are able to perfectly match using dynAMMo. This observation strengthens the argument that the conformations sampled in the MD simulation constitute the relevant configurations needed to explain the experimental data. We stress that dynAMMo uses all observables to fit one global kinetic model. Therefore, the predictions of the relaxation dispersion curves for the backbone nitrogens differ only by the observable used for each residue. Here, we demonstrate how dynAMMo can be used to combine experimental NMR relaxation data with simulation data from molecular dynamics simulations. Therefore, we obtain a detailed mechanistic explanation of how the different metastable states observed in the MD trajectory contribute to the experimentally probed chemical exchange.

Quantitative molecular kinetics from disconnected simulation statistics with dynAMMo. Above we saw how dynAMMo could recover the correct kinetics on a controlled test system. To evaluate whether dynAMMo generalizes to more complex protein systems, we establish a similar benchmark, systematically removing simulation statistics that connect the major and minor populated states of BPTI. For this case, we similarly find that dynAMMo can quantitatively recover the exchange rates between the disconnected states (Fig. 5A), and recover the implied timescales accurately (Fig. 5B). The MSM in this case is missing the slowest process (Fig. 5A, dashed cross), however, dynAMMo can recover this process and quantitatively predict the timescale. We show a detailed analysis of this scenario in the Supporting Information (Fig. S6b).

Conclusion

Here we have introduced dynamic Augmented Markov Models (dynAMMo), a new method to improve the accuracy of mechanistic biomolecular models by incorporating dynamic experimental measurements to correct for biases in the kinetics and thermodynamics of MD simulations. However, most intriguingly, dynAMMo also allows us to build quantitatively predicted model of molecular kinetics even in the absence of simulation statistics on (slow) conformational transitions. We show the performance of dynAMMo across two well-controlled benchmark systems and later deploy it to two realistic scenarios using data from molecular simulations and NMR spectroscopy on the protein BPTI. As such, dynAMMo opens up the possibility of salvaging sparsely sampled simulation data sampled using biased force fields and repurpose them to build quantitatively predictive models for structural biology.

Materials and Methods

The estimator is implemented in Python and uses PyTorch (88) and deeptime (89) as the main analysis and modeling tools. The

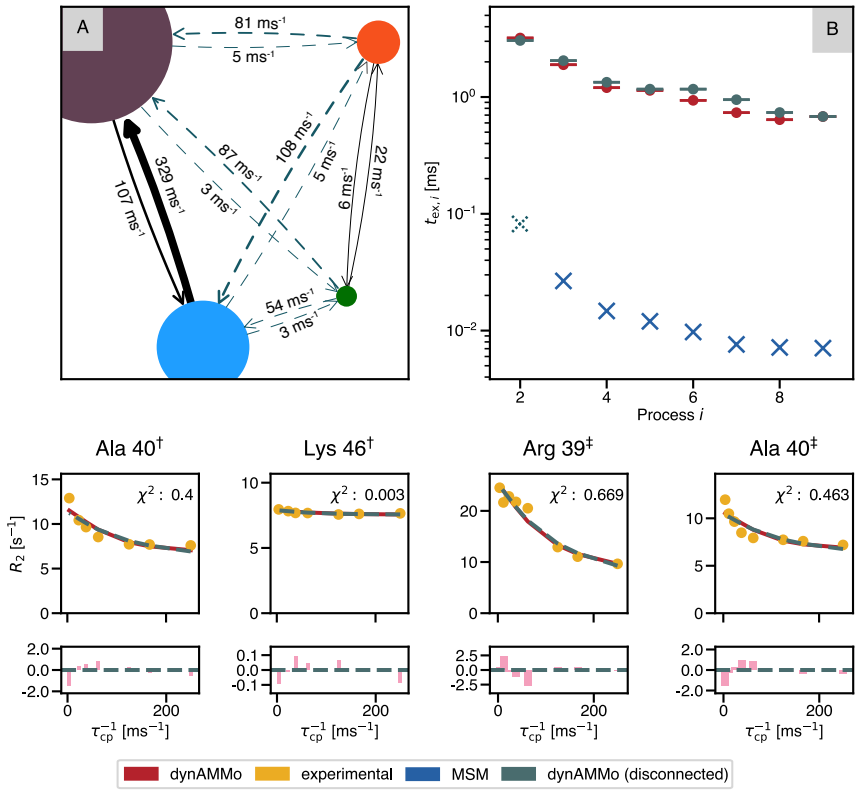

Fig. 5. Overview of simulating disconnected case using BPTI simulations and CPMG data. (A) Kinetic network between the disconnected states. All transitions between the purple/light blue and orange/green clusters were removed (indicated with dashed gray arrows) and two Markov state models were built using only the within-states trajectory data. The colors of the states correspond to the clusters shown in 3A. (B) Implied timescale plot of the disconnected model (gray), the connected model (red), and the MSM (blue) for comparison. The presumed timescale of exchange that was removed in this scenario is indicated as a dashed cross. Lower panel: Four representative examples of CPMG relaxation dispersion predictions of selected backbone nitrogens. The predictions of the disconnected case (dashed gray) are plotted together with the predictions of the connected scenario (red) for comparison. χ^2 values are shown with respect to the predictions of the disconnected case and the experimental data (yellow).

estimation procedure and theory details are provided in Supporting Information, 'Theory'. The code will be made available on https://github.com/olsson-group/dynAMMo. All figures showing molecular structures were made using PyMol (90). All plots were generated using Matplotlib (91). Additional results, such as the slowest estimated eigenvectors, loss function, and stationary distribution are reported in the Supporting Information for all model systems (Fig. S4 and S5) and BPTI (Fig. 6), respectively.

Benchmark model systems. The deeptime implementation of the four-state Prinz potential and the three-well potential datasets was used to simulate the two benchmark systems (89). The parameters used to simulate the trajectories are reported in the Supporting Information (table S1). The estimation of the dynamic Augmented Markov Models were carried out as outlined in *Results and Discussion*. Each scenario uses different parametrizations of the potential and a table with an overview is listed in the Supporting Information (Table S2). Chapman-Kolmogorov tests have been performed on all MSMs used in this study (Supporting Information Fig. S1–S2).

BPTI. The estimation and analysis of the BPTI dynAMMo were conducted as described in *Results and Discussion*. Chapman-Kolmogorov tests were conducted on the MSMs used here to ensure validity of the models Supporting Information (Fig. S3). The estimation parameters of the two scenarios are listed in the Supporting Information (Table S3). Further details are provided in Supporting Information, 'Materials and Methods.'

ACKNOWLEDGMENTS. The authors would like to thank D. E. Shaw Research for sharing the BPTI simulation and Arthur G. Palmer III, for sharing the raw NMR relaxation dispersion data. CK thanks Shanawaz Ahmed for fruitful discussions and for sharing a preliminary implementation of the Cayley transform for the eigenvector estimation. This work was partially supported by the Wallenberg AI, Autonomous Systems and Software Program (WASP) funded by the Knut and Alice Wallenberg Foundation.

1. W Arber, S Linn, DNA modification and restriction. *Annu. review biochemistry* **38**, 467–500 (1969).
2. E Antonini, M Brunori, Hemoglobin. *Annu. Rev. Biochem.* **39**, 977–1042 (1970).
3. L Poretsky, MF Kalin, The gonadotropic function of insulin. *Endocr. reviews* **8**, 132—141 (1987).
4. S Wullschleger, R Loewith, MN Hall, TOR Signaling in Growth and Metabolism. *Cell* **124**, 471–484 (2006).
5. J Monod, J Wyman, JP Changeux, On the nature of allosteric transitions: A plausible model. *J. Mol. Biol.* **12**, 88–118 (1965).
6. DE Koshland, G Némethy, D Filmer, Comparison of Experimental Binding Data. *Theor. models proteins containing subunits* **5**, 365–385 (1966).
7. F Sanger, Chemistry of insulin (1958).
8. P Ascenzi, et al., The bovine basic pancreatic trypsin inhibitor (Kunitz inhibitor): a milestone protein. *Curr. protein & peptide science* **4**, 231–251 (2003).
9. A Wlodawer, J Walter, R Huber, L Sjölin, Structure of bovine pancreatic trypsin inhibitor. Results of joint neutron and X-ray refinement of crystal form II. *J. Mol. Biol.* **180**, 301–329 (1984).
10. G Wagner, A DeMarco, K Wüthrich, Dynamics of the aromatic amino acid residues in the globular conformation of the basic pancreatic trypsin inhibitor (BPTI). *Biophys. structure mechanism* **2**, 139–158 (1976).
11. KD Berndt, P Güntert, LP Orbons, K Wüthrich, Determination of a high-quality nuclear magnetic resonance solution structure of the bovine pancreatic trypsin inhibitor and comparison with three crystal structures. *J. Mol. Biol.* **227**, 757–775 (1992).
12. JW Peng, G Wagner, [20] Investigation of protein motions via relaxation measurements in *Nuclear Magnetic Resonance, Part C*, Methods in Enzymology. (Academic Press) Vol. 239, pp. 563–596 (1994).
13. PE Smith, RC van Schaik, T Szyperski, K Wüthrich, WF van Gunsteren, Internal Mobility of the Basic Pancreatic Trypsin Inhibitor in Solution: A Comparison of NMR Spin Relaxation Measurements and Molecular Dynamics Simulations. *J. Mol. Biol.* **246**, 356–365 (1995).
14. D van der Spoel, AR van Buuren, DP Tieleman, HJC Berendsen, Molecular dynamics simulations of peptides from BPTI: A closer look at amide—aromatic interactions. *J. Biomol. NMR* **8**, 229–238 (1996).
15. V Daggett, M Levitt, A model of the molten globule state from molecular dynamics simulations. *Proc. Natl. Acad. Sci. United States Am.* **89**, 5142–5146 (1992).
16. DE Shaw, et al., Atomic-Level Characterization of the Structural Dynamics of Proteins. *Science* **330**, 341–347 (2010).
17. M Grimaldo, F Roosen-Runge, F Zhang, F Schreiber, T Seydel, Dynamics of proteins in solution. *Q. Rev. Biophys.* **52** (2019).
18. O Trott, AG Palmer, Theoretical study of R1ρ rotating-frame and R2 free-precession relaxation in the presence of n-site chemical exchange. *J. Magn. Reson.* **170**, 104–112 (2004).
19. H Koss, M Rance, AG Palmer, General expressions for R1ρ relaxation for N-site chemical exchange and the special case of linear chains. *J. Magn. Reson.* **274**, 36–45 (2017).
20. AG Palmer, F Massi, Characterization of the dynamics of biomacromolecules using rotating-frame spin relaxation NMR spectroscopy. *Chem. Rev.* **106**, 1700–1719 (2006).
21. B Lindner, Z Yi, JH Prinz, JC Smith, F Noé, Dynamic neutron scattering from conformational dynamics. i. theory and markov models. *The J. Chem. Phys.* **139**, 175101 (2013).

22. J Möller, M Sprung, A Madsen, C Gutt, X-ray photon correlation spectroscopy of protein dynamics at nearly diffraction-limited storage rings. *IUCrJ* **6**, 794–803 (2019).
23. S Hiller, Chaperone-Bound Clients: The Importance of Being Dynamic. *Trends Biochem. Sci.* **44**, 517–527 (2019).
24. B Schiffrin, et al., Skp is a multivalent chaperone of outer-membrane proteins. *Nat. Struct. Mol. Biol.* **23**, 786–793 (2016).
25. BM Burmann, C Wang, S Hiller, Conformation and dynamics of the periplasmic membrane-protein-chaperone complexes OmpX-Skp and tOmpA-Skp. *Nat. Struct. Mol. Biol.* **20**, 1265–1272 (2013).
26. J Thoma, BM Burmann, S Hiller, DJ Müller, Impact of holdase chaperones Skp and SurA on the folding of β-barrel outer-membrane proteins. *Nat. Struct. Mol. Biol.* **22**, 795–802 (2015).
27. DF Gauto, et al., Functional control of a 0.5 MDa TET aminopeptidase by a flexible loop revealed by MAS NMR. *Nat. Commun.* **13** (2022).
28. P Neudecker, et al., Structure of an Intermediate State. *Science* **336**, 362 (2012).
29. AC Kruse, et al., Structure and dynamics of the M3 muscarinic acetylcholine receptor. *Nature* **482**, 552–556 (2012).
30. DM Rosenbaum, et al., Structure and function of an irreversible agonist-β2 adrenoceptor complex. *Nature* **469**, 236–242 (2011).
31. K Lindorff-Larsen, S Piana, RO Dror, DE Shaw, How fast-folding proteins fold. *Science* **334**, 517–520 (2011).
32. J Huang, et al., CHARMM36m: An improved force field for folded and intrinsically disordered proteins. *Nat. Methods* **14**, 71–73 (2016).
33. K Lindorff-Larsen, et al., Improved side-chain torsion potentials for the Amber ff99SB protein force field. *Proteins: Struct. Funct. Bioinforma.* **78**, 1950–1958 (2010).
34. JS Smith, O Isayev, AE Roitberg, {ANI}-1: an extensible neural network potential with {DFT} accuracy at force field computational cost. *Chem. Sci.* **8**, 3192–3203 (2017).
35. DE Shaw, et al., Millisecond-scale molecular dynamics simulations on Anton in *Proceedings of the Conference on High Performance Computing Networking, Storage and Analysis*. pp. 1–11 (2009).
36. DE Shaw, et al., Anton 2: Raising the Bar for Performance and Programmability in a Special-Purpose Molecular Dynamics Supercomputer. *Int. Conf. for High Perform. Comput. Networking, Storage Analysis, SC* **2015-Janua**, 41–53 (2014).
37. DE Shaw, et al., Anton 3: Twenty Microseconds of Molecular Dynamics Simulation before Lunch. *Int. Conf. for High Perform. Comput. Networking, Storage Analysis, SC* **1** (2021).
38. VA Voelz, VS Pande, GR Bowman, Folding@home: Achievements from over 20 years of citizen science herald the exascale era. *Biophys. J.* **122**, 1–12 (2023).
39. GR Bowman, KA Beauchamp, G Boxer, VS Pande, Progress and challenges in the automated construction of Markov state models for full protein systems. *J. Chem. Phys.* **131** (2009).
40. GR Bowman, VA Voelz, VS Pande, Atomistic folding simulations of the five-helix bundle protein λ6-85. *J. Am. Chem. Soc.* **133**, 664–667 (2011).
41. TJ Lane, GR Bowman, K Beauchamp, VA Voelz, VS Pande, Markov State model reveals folding and functional dynamics in ultra-long MD trajectories. *J. Am. Chem. Soc.* **133**, 18413–18419 (2011).
42. VA Voelz, et al., Slow unfolded-state structuring in acyl-CoA binding protein folding revealed by simulation and experiment. *J. Am. Chem. Soc.* **134**, 12565–12577 (2012).
43. JH Prinz, et al., Markov models of molecular kinetics: Generation and validation. *The J. Chem. Phys.* **134**, 174105 (2011).
44. CD Wassman, et al., Computational identification of a transiently open L1/S3 pocket for reactivation of mutant p53. *Nat. Commun.* **4**, 1–9 (2013).
45. 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. United States Am.* **106**, 19011–19016 (2009).
46. Q Qiao, GR Bowman, X Huang, Dynamics of an intrinsically disordered protein reveal metastable conformations that potentially seed aggregation. *J. Am. Chem. Soc.* **135**, 16092–16101 (2013).
47. L Raich, et al., Discovery of a hidden transient state in all bromodomain families. *Proc. Natl. Acad. Sci.* **118**, e2017427118 (2021).
48. KS Chakrabarti, et al., A litmus test for classifying recognition mechanisms of transiently binding proteins. *Nat. Commun.* **13** (2022).
49. K Liebl, M Zacharias, The development of nucleic acids force fields: From an unchallenged past to a competitive future. *Biophys. J.* **122**, 1–11 (2023).
50. D Tan, S Piana, RM Dirks, DE Shaw, RNA force field with accuracy comparable to state-of-the-art protein force fields. *Proc. Natl. Acad. Sci. United States Am.* **115**, E1346–E1355 (2018).
51. K Lindorff-Larsen, et al., Systematic validation of protein force fields against experimental data. *PLoS ONE* **7**, 1–6 (2012).
52. S Piana, JL Klepeis, DE Shaw, Assessing the accuracy of physical models used in protein-folding simulations: Quantitative evidence from long molecular dynamics simulations. *Curr. Opin. Struct. Biol.* **24**, 98–105 (2014).
53. J Henriques, C Cragnell, M Skepö, Molecular Dynamics Simulations of Intrinsically Disordered Proteins: Force Field Evaluation and Comparison with Experiment. *J. Chem. Theory Comput.* **11**, 3420–3431 (2015).
54. P Robustelli, S Piana, DE Shaw, Developing a molecular dynamics force field for both folded and disordered protein states. *Proc. Natl. Acad. Sci. United States Am.* **115**, E4758–E4766 (2018).
55. S Bottaro, K Lindorff-Larsen, Biophysical experiments and biomolecular simulations: A perfect match? *Science* **361**, 355–360 (2018).
56. HTA Leung, et al., A Rigorous and Efficient Method to Reweight Very Large Conformational Ensembles Using Average Experimental Data and to Determine Their Relative Information Content. *J. Chem. Theory Comput.* **12**, 383–394 (2016).
57. R Capelli, G Tiana, C Camilloni, An implementation of the maximum-caliber principle by replica-averaged time-resolved restrained simulations. *J. Chem. Phys.* **148** (2018).
58. CA Smith, et al., Enhancing NMR derived ensembles with kinetics on multiple timescales. *J. Biomol. NMR* **74**, 27–43 (2020).
59. A Cavalli, C Camilloni, M Vendruscolo, Molecular dynamics simulations with replica-averaged structural restraints generate structural ensembles according to the maximum entropy principle. *J. Chem. Phys.* **138** (2013).
60. W Boomsma, J Ferkinghoff-Borg, K Lindorff-Larsen, Combining Experiments and Simulations Using the Maximum Entropy Principle. *PLoS Comput. Biol.* **10**, 1–9 (2014).
61. S Olsson, D Strotz, B Vögeli, R Riek, A Cavalli, The Dynamic Basis for Signal Propagation in Human Pin1-WW. *Structure* **24**, 1464–1475 (2016).
62. KA Beauchamp, VS Pande, R Das, Bayesian energy landscape tilting: Towards concordant models of molecular ensembles. *Biophys. J.* **106**, 1381–1390 (2014).
63. JW Pitera, JD Chodera, On the use of experimental observations to bias simulated ensembles. *J. Chem. Theory Comput.* **8**, 3445–3451 (2012).
64. M Bonomi, C Camilloni, A Cavalli, M Vendruscolo, Metainference: A Bayesian inference method for heterogeneous systems. *Sci. Adv.* **2**, 1–9 (2016).
65. G Hummer, J Köfinger, Bayesian ensemble refinement by replica simulations and reweighting. *J. Chem. Phys.* **143** (2015).
66. S Olsson, J Frellsen, W Boomsma, KV Mardia, T Hamelryck, Inference of Structure Ensembles of Flexible Biomolecules from Sparse , Averaged Data. *PLoS ONE* **8**, 1–7 (2013).
67. K Lindorff-Larsen, RB Best, MA DePristo, CM Dobson, M Vendruscolo, Simultaneous determination of protein structure and dynamics. *Nature* **433**, 128–132 (2005).
68. AD White, GA Voth, Efficient and minimal method to bias molecular simulations with experimental data. *J. Chem. Theory Comput.* **10**, 3023–3030 (2014).
69. Z Faidon Brotzakis, M Vendruscolo, PG Bolhuis, A method of incorporating rate constants as kinetic constraints in molecular dynamics simulations. *Proc. Natl. Acad. Sci. United States Am.* **118** (2021).
70. S Olsson, H Wu, F Paul, C Clementi, F Noé, Combining experimental and simulation data of molecular processes via augmented Markov models. *Proc. Natl. Acad. Sci. United States Am.* **114**, 8265–8270 (2017).
71. S Meiboom, D Gill, Modified Spin-Echo Method for Measuring Nuclear Relaxation Times. *THE REVIEW OF SCIENTIFIC INSTRUMENTS* **28**, 688–690 (1958).
72. Z Luz, S Meiboom, Nuclear magnetic resonance study of the protolysis of trimethylammonium ion in aqueous solution-order of the reaction with respect to solvent. *The J. Chem. Phys.* **39**, 366–370 (1963).
73. O Millet, JP Loria, CD Kroenke, M Pons, AG Palmer, The static magnetic field dependence of chemical exchange linebroadening defines the NMR chemical shift time scale. *J. Am. Chem. Soc.* **122**, 2867–2877 (2000).
74. S Olsson, F Noé, Mechanistic models of chemical exchange induced relaxation in protein NMR. *J. Am. Chem. Soc.* **139**, 200–210 (2017).
75. Y Xue, JM Ward, T Yuwen, IS Podkorytov, NR Skrynnikov, Microsecond time-scale conformational exchange in proteins: Using long molecular dynamics trajectory to simulate NMR relaxation dispersion data. *J. Am. Chem. Soc.* **134**, 2555–2562 (2012).
76. Z Wen, W Yin, A feasible method for optimization with orthogonality constraints. *Math. Program.* **142**, 397–434 (2013).
77. MJ Grey, C Wang, AG Palmer, Disulfide Bond Isomerization in Basic Pancreatic Trypsin Inhibitor: Multisite Chemical Exchange Quantified by CPMG Relaxation Dispersion and Chemical Shift Modeling. *J. Am. Chem. Soc.* **125**, 14324–14335 (2003).
78. F Massi, E Johnson, C Wang, M Rance, AG Palmer, NMR R1ρ Rotating-Frame Relaxation with Weak Radio Frequency Fields. *J. Am. Chem. Soc.* **126**, 2247–2256 (2004).
79. U Weininger, U Brath, K Modig, K Teilum, M Akke, Off-resonance rotating-frame relaxation dispersion experiment for 13C in aromatic side chains using L-optimized TROSY-selection. *J. Biomol. NMR* **59**, 23–29 (2014).
80. VP Denisov, B Halle, Protein hydration dynamics in aqueous solution: A comparison of bovine pancreatic trypsin inhibitor and ubiquitin by oxygen-17 spin relaxation dispersion. *J. Mol. Biol.* **245**, 682–697 (1995).
81. B Brooks, M Karplus, Harmonic dynamics of proteins: Normal modes and fluctuations in bovine pancreatic trypsin inhibitor. *Proc. Natl. Acad. Sci. United States Am.* **80**, 6571–6575 (1983).
82. G Wagner, Characterization of the distribution of internal motions in the basic pancreatic trypsin inhibitor using a large number of internal nmr probes. *Q. Rev. Biophys.* **16**, 1–57 (1983).
83. G Wagner, et al., Exchange of Two-Spin Order in Nuclear Magnetic Resonance: Separation of Exchange and Cross-Relaxation Processes. *J. Am. Chem. Soc.* **107**, 6440–6446 (1985).
84. F Noé, C Clementi, Kinetic Distance and Kinetic Maps from Molecular Dynamics Simulation. *J. Chem. Theory Comput.* **11**, 5002–5011 (2015).
85. MK Scherer, et al., PyEMMA 2: A Software Package for Estimation, Validation, and Analysis of Markov Models. *J. Chem. Theory Comput.* **11**, 5525–5542 (2015).
86. G Pérez-Hernández, F Paul, T Giorgino, G De Fabritiis, F Noé, Identification of slow molecular order parameters for Markov model construction. *J. Chem. Phys.* **139** (2013).
87. DW Li, R Brüschweiler, PPM: A side-chain and backbone chemical shift predictor for the assessment of protein conformational ensembles. *J. Biomol. NMR* **54**, 257–265 (2012).
88. A Paszke, et al., PyTorch: An imperative style, high-performance deep learning library. *Adv. Neural Inf. Process. Syst.* **32** (2019).
89. M Hoffmann, et al., Deeptime: a Python library for machine learning dynamical models from time series data. *Mach. Learn. Sci. Technol.* **3** (2022).
90. L Schrödinger, The PyMOL Molecular Graphics System, Version˜2.0 (2015).
91. JD Hunter, MATPLOTLIB: A 2D GRAPHICS ENVIRONMENT. *Comput. Sci. Eng.* **9**, 90–95 (2007).