

## ABSTRACT

Title of Dissertation: **INVESTIGATING BIOMOLECULAR RARE  
EVENTS WITH ARTIFICIAL INTELLIGENCE  
AUGMENTED MOLECULAR DYNAMICS**

**Yihang Wang**  
Doctor of Philosophy, 2022

Dissertation Directed by: **Professor Pratyush Tiwary**  
**Department of Chemistry and Biochemistry**  
**Institute for Physical Science and Technology**

To decipher the biomolecular interaction mechanism play an important role in understanding the mystery of life. Understanding the mechanism of receptor-ligand interactions is of great importance both in context of fundamental biology and medical applications. Limited by the spatial or temporal resolution, wet-lab experiments may not be able to capture enough details to unravel the complicated interaction mechanisms. Molecular dynamics (MD) simulation, as computational tool to study many-body systems with atomic resolution, has emerged as a powerful tool to investigate the physical or biochemical properties of biomolecules.

Though MD simulation has its advantage in terms of its high spacial and temporal resolution over experimental methods, a big gap still remains between the time scale that can be reached by MD simulations and the time scale of the biological process that we want to study. In this thesis, I explore two frameworks to utilize the power of statistical mechanics, molecular dynamics, and machine learning to rectify this gap, and

thus enable the simulation study of ligand-receptors interaction without overwhelming demand on computational resources.

Firstly, I propose the reweighted autoencoded variational Bayes for enhanced sampling (RAVE) method, a new iterative scheme that uses the deep learning framework of information bottleneck to enhance sampling in molecular simulations. RAVE involves iterations between molecular simulations and deep learning in order to produce an increasingly accurate probability distribution along a low- dimensional latent space that captures the key features of the molecular simulation trajectory. RAVE determines an optimum, yet nonetheless physically interpretable, reaction coordinate and optimum probability distribution. Both then directly serve as the biasing protocol for a new biased simulation, which is once again fed into the deep learning module with appropriate weights accounting for the bias, the procedure continuing until estimates of desirable thermodynamic observables are converged. The usefulness and reliability of RAVE is demonstrated by applying it to two test-pieces, studying processes slower than milliseconds, calculating free energies, kinetics and critical mutations.

I also systematically study the following questions: (a) the choice of a predictive time-delay in RAVE, or how far into the future should the machine learning model try to predict the state of a given system output from MD, and (b) for short time-delays, how much of an error is made in approximating the biased propagator for the dynamics as the unbiased propagator. I demonstrate through a master equation framework as to why the exact choice of time-delay is irrelevant as long as a small non-zero value is adopted. I also derive a correction to the objective function by reweighting the biased propagator, which better approximates the unbiased objective function without incurring



extra computational overhead.

To promote our understanding of RNA-ligand interaction at the molecular level, I use RAVE and collaborate with experimentalists to study the interplay between two ligands and PreQ<sub>1</sub>, which is a widely-studied model for RNA-small molecule recognition. I show that site-specific flexibility profiles from our simulations are in excellent agreement with *in vitro* measurements of flexibility using Selective 2' Hydroxyl Acylation analyzed by Primer Extension and Mutational Profiling (SHAPE-MaP). And with orders of magnitude simulation speedup attained by RAVE, I can directly observe ligand dissociation for cognate and synthetic ligands from a PreQ<sub>1</sub> riboswitch system. The artificial intelligence-argumented simulations reproduce known binding affinity profiles for the cognate and synthetic ligands, and pinpoint how both ligands make use of different aspects of riboswitch flexibility. On the basis of the dissociation trajectories, I also make and validate predictions of pairs of mutations for both the ligand systems that would show differing binding affinities. These mutations are distal to the binding site and could not have been predicted solely on the basis of structure.

Secondly, I develop a framework based on statistical mechanics and generative Artificial Intelligence to use simulations or experiments performed at some set of temperatures to learn about the physics or chemistry at some other arbitrary temperature. Specifically, I use denoising diffusion probabilistic models, and show how these models in combination with replica exchange molecular dynamics achieve superior sampling of the biomolecular energy landscape at temperatures that were never even simulated without assuming any particular slow degrees of freedom. The key idea is to treat the temperature as a fluctuating random variable and not a control parameter as is usually done. This

allows us to directly sample from the joint probability distribution in configuration and temperature space. The results are demonstrated for a chirally symmetric peptide and single-strand ribonucleic acid undergoing conformational transitions in all-atom water. I demonstrate how we can discover transition states and metastable states that were previously unseen at the temperature of interest, and even bypass the need to perform further simulations for wide range of temperatures. At the same time, any unphysical states are easily identifiable through very low Boltzmann weights. The procedure while shown here for a class of molecular simulations should be more generally applicable to mixing information across simulations and experiments with varying control parameters.

INVESTIGATING BIOMOLECULAR RARE EVENTS WITH  
ARTIFICIAL INTELLIGENCE AUGMENTED MOLECULAR  
DYNAMICS

by

Yihang Wang

Dissertation submitted to the Faculty of the Graduate School of the  
University of Maryland, College Park in partial fulfillment  
of the requirements for the degree of  
Doctor of Philosophy  
2022

Advisory Committee:  
Professor Pratyush Tiwary, Chair/Advisor  
Professor John D. Weeks  
Professor Silvina R. Matysiak  
Professor Garegin Papoian  
Professor Jeffery B. Klauda

## Dedication

Dedicated to my dearest mom and dad, who have been standing behind me.

## Acknowledgments

I have never wondered I could have such fortune being loved and accompanied by so many great folks at stage of my life. It is my honor to express my gratitude to them in this thesis.

First, I'd like to thank Prof. Pratyush Tiwary, my dear advisor. Before joining our lab, I had no idea what a MD simulation was. The past few years' experiences with Prof. Tiwary gave me the confidence and knowledge to decide what type of career to pursue. What I learned from him also goes way beyond just how to do science. It is my gratefulness to witness and learn how he set up the whole lab, facing and solving all the great challenges. As a graduate student, I do have the privilege of focusing on and enjoying learning and doing research in an undisturbed environment, which is undoubtedly due to his tireless efforts. I know his office door is always open whenever I have any questions about science or life.

I want to thank Dr. Jay Schneekloth and the members in his lab. In my view, implementing experiments is one of the toughest and most impressive parts of the procedure of science. Being able to collaborate with such talented experimentalists and learn from them is a unique experience for me.

One important component of my graduate studies is my gifted colleagues. Thanks for the intense discussions, neat ideas, funny jokes, proofreading my drafts, and making

our office such a wonderful place to explore science. In particular, I would like to thank Dedi Wang for his brilliant ideas in improving RAVE to a new level. Also, I want to thank Sun-Ting Tsai, Zack Smith, Ziyue Zou, and Shams Mehdi, the people I bounce ideas off.

I would also like to express my gratitude for the help and support from Prof. John Weeks, Prof. Silvina Matysiak, Prof. Garegin Papoian, Prof. Jeffery Klauda, and Prof. Chris Jarzynski. Thanks for being great teachers and serving on my committees, as well as giving me suggestions. All of them have been and will always be the model scientists that I want to be and pursue to be.

I will never forget the fantastic time I spent with my friends at College Park. Special thanks to Zhiyu, Mingshu, Yixu, Yalun, and Xiangyu for being my closest friends. Without these lovely people like mirrors reflecting each beautiful moment of life, I might get lost in the anxiety and pressure of facing an unforeseen future. Also, I would like to thank Zhongyuan, Zhiyuan, Lingqi, Jiameng, Sangyu, and Jiahui for being willing to share their stories and listen to mine. Although we are far from each other in terms of physical distances, their optimism and enterprise keep pushing me forward.

There is a long list of teachers and mentors I would also like to thank, for educating me with the basic tools to understand this world and nourishing my early interest in knowledge. In particular, I want to thank Mr. Shaobo Zheng, Mr. Dong Xiang, and Mr. Wuhui Zhang for their influence on my personality. Thanks to Prof. Jiansheng Wu and Prof. Lang Chen, who gave me the first glimpse of the colorful physics world.

Finally, I owe my deepest thanks to my family. I want to express my gratitude to my parents and grandparents, who never hesitate to support my decisions. Their unconditional love is the main reason I can chase my dream with great courage. As I

always know, there is a place to shelter me from the storm outside. To my dear maternal grandfather, sorry for not being by your side in the last days of your life. I really hope to share my happiness on this fantastic journey with you!

## Table of Contents

Dedication	ii
Acknowledgements	iii
Table of Contents	vi
List of Tables	ix
List of Figures	x
List of Abbreviations	xii
Chapter 1: Introduction	1
1.1 Understanding ligand-receptor interaction mechanisms	1
1.1.1 Biological meaning of ligand-receptor interaction	1
1.1.2 A two-state model of ligand-receptor interaction	3
1.2 Molecular dynamics simulation	5
1.2.1 Introduction to molecular dynamics simulation	5
1.2.2 Simulations of the thermodynamic and kinetic properties of biomolecules	6
1.2.3 Reaction coordinate	8
1.2.4 Challenges in MD simulations	9
1.3 Solving rare event problem with machine learning	9
1.3.1 Rare event problems	9
1.3.2 Enhanced sampling methods	10
1.3.3 Machine learning approaches for enhancing MD simulations	15
1.4 Overview of the thesis	22
Chapter 2: Reweighted autoencoded variational Bayes for enhanced sampling (RAVE)	25
2.1 Introduction	25
2.2 Using latent variables of VAE as RCs	27
2.3 Principle of Past–Future Information Bottleneck	28
2.3.1 Variational inference and neural network architecture	30
2.3.2 Variational inference on unbiased and biased trajectories	34
2.3.3 Patching it all	35
2.4 Application of RAVE to two benchmark systems	38



2.4.1	Conformation transitions in a model peptide . . . . .	38
2.4.2	Benzene dissociation from T4-L99A lysozyme . . . . .	40
2.4.3	Predicting critical residues . . . . .	44
2.5	Discussion . . . . .	49
Chapter 3: Understanding the role of predictive time delay and biased propagator in RAVE . . . . .		51
3.1	Introduction . . . . .	51
3.2	Dependence of predictive information bottleneck on $\Delta t$ . . . . .	52
3.3	Correcting the propagator and the information bottleneck when using biased trajectory . . . . .	56
3.4	Results . . . . .	59
3.4.1	System set-up . . . . .	60
3.4.2	RAVE set-up . . . . .	61
3.4.3	Calculated RC under different time-delays and approximation schemes . . . . .	62
3.5	Conclusion . . . . .	64
Chapter 4: Interrogating RNA-small molecule interactions with RAVE . . . . .		66
4.1	Introduction . . . . .	66
4.2	Ligand binding induces flexibility change in riboswitch aptamer . . . . .	68
4.3	Ligand binding affinity with riboswitch aptamer . . . . .	70
4.4	Predicting critical nucleotides for mutagenesis experiments . . . . .	75
4.5	Validating predicted critical nucleotides . . . . .	78
4.6	Discussion . . . . .	83
Chapter 5: Mixing physics across temperatures with generative AI . . . . .		86
5.1	Introduction . . . . .	86
5.2	Temperature as a random variable . . . . .	88
5.3	Denoising diffusion probabilistic model . . . . .	90
5.4	Results . . . . .	93
5.5	Discussion . . . . .	104
Chapter 6: Summary and future directions . . . . .		107
6.1	AI-augmented MD simulations . . . . .	107
6.2	Study of ligand dissociation from RNA . . . . .	108
6.3	Future work . . . . .	109
Appendix A: . . . . .		112
A.1	MD setups . . . . .	112
A.1.1	Alanine dipeptide . . . . .	112
A.1.2	L99A T4L-benzene . . . . .	113
A.1.3	AIB <sub>9</sub> . . . . .	113
A.1.4	GACC . . . . .	114
A.1.5	Architecture of RAVE . . . . .	115

A.2	Definitions of quantities in information theory and variational lower bound	116
A.2.1	Mutual information . . . . .	117
A.2.2	Shannon entropy . . . . .	117
A.2.3	Cross entropy . . . . .	117
A.2.4	Kullback-Leibler Divergence . . . . .	117
A.2.5	Exact decoder definition . . . . .	118
A.2.6	Gibbs's inequality and variational lower bound . . . . .	118
A.3	PIB objective $\mathcal{L}'$ for unbiased trajectory . . . . .	119
A.4	PIB objective $\mathcal{L}'$ for biased trajectory . . . . .	120
	List of Publications	121
	Bibliography	122

## List of Tables

2.1	Order parameters for T4-L99A lysozyme . . . . .	42
2.2	List of order parameters used to construct RC and their weights in $\chi_1$ and $\chi_2$ . . . . .	48
5.1	Definition of dihedral angles . . . . .	101

## List of Figures

1.1	An illustration of a free energy profile . . . . .	7
1.2	Workflow of methods that use machine learning to analyze and enhance MD simulation . . . . .	16
2.1	Network architecture used for learning predictive information bottleneck .	31
2.2	RAVE one alanine dipeptide . . . . .	39
2.3	Definition of Order parameters for Benzene-Lysozyme dissociation . . . .	41
2.4	Order parameters weights as functions of time delay for Benzene-lysozyme dissociation . . . . .	43
2.5	The 2-component Predictive Information Bottleneck for benzene-lysozyme dissociation . . . . .	44
2.6	Estimating benzene-lysozyme dissociation constant by Poisson fit . . . .	45
2.7	Free energy surfaces for Benzene-Lysozyme dissociation . . . . .	45
2.8	Critical residue analysis for benzene-lysozyme complex . . . . .	47
2.9	Robustness test for RAVE on studying benzene-lysozyme dissociation . .	48
3.1	Free energy profile along each degree of freedom of the model system . .	61
3.2	Reaction coordinates learned by RAVE as functions of predictive time delay	63
4.1	Ligand-PreQ <sub>1</sub> riboswitch interaction . . . . .	68
4.2	SHAPE reactivities for the <i>Tt</i> PreQ <sub>1</sub> riboswitch aptamer . . . . .	71
4.3	Weights for different order parameters for PreQ <sub>1</sub> riboswitch . . . . .	72
4.4	Free energy profiles of riboswitch with different ligands . . . . .	73
4.5	Dissociation pathways of cognate and synthetic ligand of PreQ <sub>1</sub> riboswitch	75
4.6	Critical nucleotides prediction . . . . .	76
4.7	The change of RMSD of each nucleotide during the dissociation process .	79
4.8	Affinity measurement of PreQ <sub>1</sub> cognate synthetic ligands . . . . .	80
4.9	Free energy profile from 2 $\mu$ s unbiased simulation for wild tpye and mutated PreQ <sub>1</sub> riboswitch . . . . .	82
4.10	Hydrogen bond occupancy for wild type and mutated systems . . . . .	83
5.1	DDPM network architecture . . . . .	91
5.2	AIB <sub>9</sub> structure and free energy surface . . . . .	92
5.3	Samples of AIB9 from REMD and DDPM . . . . .	95
5.4	$\Delta G$ between ground state and excited state in AIB9 . . . . .	97
5.5	$\Delta G$ between ground state and transition state in AIB9 . . . . .	97

5.6	Interpolation or extrapolation in DDPM . . . . .	98
5.7	Projection of samples and free energy profile on dihedral angles $\zeta$ and $\gamma$ of A2 in GACC at 325 K . . . . .	99
5.8	DDPM outperforms direct reweighting . . . . .	100
5.9	REMD samples and free energy profile the C2 in GACC . . . . .	102
5.10	REMD samples and free energy profile of the C3 in GACC . . . . .	103
5.11	Free energy profiles of all dihedral angles in GACC . . . . .	104

## List of Abbreviations

AI	Artificial intelligence
ANN	Artificial neural networks
DDPM	Denoising diffusion probabilistic model
FRET	Fluorescence resonance energy transfer
FIA	Fluorescence intensity assay
MC	Monte Carlo
MD	Molecular dynamics
MESA	Molecular enhanced sampling with autoencoders
MST	Microscale thermophoresis
PIB	Predictive information bottleneck
RAVE	Reweighted autoencoded variational Bayes for enhanced sampling
REAP	Reinforcement learning based adaptive sampling
RC	Reaction coordinate
SHAPE-MaP	Seletive 2' Hydroxyl Acylation analyzed by Primer Extension and Mutational Profiling
SRV	State-free reversible VAMPnets
TAE	Time-lagged autoencoders
VAC	variational approach to conformation dynamics
VAMP	Variational approach for Markov processes
VDE	Variational Dynamic encoder
VAE	Variational autoencoder
WT	Wild type

## Chapter 1: Introduction

### 1.1 Understanding ligand-receptor interaction mechanisms

#### 1.1.1 Biological meaning of ligand-receptor interaction

One of the defining characteristics of living systems is their ability to respond to and interact with the environment. For example, cells can sense the change of ion concentration and open a channel through the membrane that allows specific ions to pass through. The molecular biology basis of this ability is often the interaction between biomolecular receptors and their cognate ligands[1, 2]. Usually, the binding of a ligand will change the biochemical properties of receptors such as stabilizing or unstabilizing their structures. These changes can either promote or inhibit the downstream processes and thus regulate the whole biological system. Similarly for nucleotides, recent studies have shown that structural RNA can react to the binding of small molecules to regulate the downstream gene expression processes[3, 4].

Different small molecules can trigger different cell responses by activating different signaling pathways and thus ensure that the whole living system can function properly. Moreover, different ligands can bind to the same receptor but trigger different pathways. Such a mechanism, which is usually referred to as functional selectivity, is both common

and important for cell signaling. For example, receptor tyrosine kinases (RTKs), a family of cross-membrane proteins, can trigger different responses such as phosphorylation of itself or other proteins when binding to different ligands. Such different reactions will further initiate a chain of downstream biomolecular reactions and finally pass the information to gene regulator biomolecules such as transcription factors and thus influence cell fates such as differentiation and proliferation[5]. However, we still lack a more detailed understanding of the underlying mechanism. In many cases, it still remains elusive why the receptor can recognize its cognate ligand and how the binding of specific ligand makes the receptors tend to adopt one structure over others? Another interesting phenomenon related to RTKs is the allosteric effect, in which the binding of ligands happens at the receptor domain lies outside the membrane, while the conformational changes occur at the functional domain lies in the intracellular part of the membrane. Such a property of transferring the effect of ligand binding from one end of a molecule to the other end, though very common in ligand-receptor interactions, is not well-understood and is often referred to as the second secret of life[6, 7].

Besides its importance in answering fundamental biological problems, a better understanding of the ligand-receptor mechanism is also crucial for practical purposes. The disfunction of the signaling process, which is usually associated with the gain- or loss-of-function of the mutated or misfolding of biomolecules, can lead to various diseases, including cancers[8, 9]. Therefore, small molecules that can selectively bind to their target biomolecules and modulate their functionalities will have the potential to be used as drugs to cure diseases. However, the lack of understanding of the ligand-receptor interaction mechanism restricts us to use brute force screening methods, leading



to substantial financial costs in the search for small-molecule drugs[10].

With the emergence of experimental, computational, and theoretical tools, much progress has been made to understand the mystery of ligand-receptor interactions. At the same time, we have to admit that challenges still remain. In the following sections, I will first introduce basic concepts and methods in studying ligand-receptor interaction. Particularly, I will mainly discuss computational methods and related unresolved problems.

### 1.1.2 A two-state model of ligand-receptor interaction

Understanding the behavior of systems consisting of over thousands or millions of atoms jiggling around is very challenging. One way to simplify this problem is to coarse-grain the system into two states: the bound state where the ligand and receptor form a complex and the unbound state where there are little or no interactions between ligand and receptor. Though it is just a simplified picture, it still permits a very profound understanding and has great explanatory power on what we observe in reality. One example is the binding affinity which is defined to characterize how likely a ligand can bind to the receptor. Macroscopically, we can describe the ligand binding and unbinding process by the reaction equation:



Where R, L and RL represent receptor, ligand and ligand-receptor complex separately. The  $k_{\text{on}}$  and  $k_{\text{off}}$  are the binding rate and the dissociation rate. The dissociation constant

$K_D = \frac{k_{\text{off}}}{k_{\text{on}}}$  also determines the ratio of concentration of species by the relationship  $K_D = \frac{[R][L]}{[LR]}$  where the concentration is denoted by the square brackets. Also, according to the law of mass action, we can also write down the dissociation constant as  $K_D = e^{\frac{\Delta G}{k_B T}} = \frac{P_{\text{unbind}}}{P_{\text{bind}}}$  of where  $P$  is the equilibrium distribution of bound and unbound state and  $\Delta G$  is the free energy difference between bound and unbound state.

These quantities play an important role in deciphering the ligand-receptor interaction mechanism and drug development. For example, a low  $K_D$  value indicates that the ligand can bind to the receptor which relatively low concentration and thus associate with high binding affinity. An ideal drug should have low  $K_d$  with its target receptor but high  $K_d$  with other biomolecules to selectively binds to its target receptor. Moreover, a previous study of drug efficacy reveals that the residence time of the drug, which is the inverse of  $k_{\text{off}}$ , is more relevant to the percent survival compared to its binding affinity[11].

All these quantities can be measured by experimental methods. For example, isothermal titration calorimetry (ITC)[12] measures the heat absorbed or released during the titration of the ligands to estimate the dissociation constant  $K_D$ . Surface plasmon resonance (SPR) [13] monitor the refractive index change on a surface caused by the binding events to measure the binding or unbinding rates. Fluorescence polarization (FP) labels the ligand with fluorescent and measures the polarization of the light emitted from the fluorescent[14]. A higher degree of polarization indicates that the ligands are bound to the receptor as it rotates much slower than being unbound. However, limited by the spatial or temporal resolution, wet-lab experiments may not be able to capture enough details to unravel the complicated interaction mechanisms[15]. Sometimes we might want to go

beyond these quantities and look at the system in more detail. For example, we might want to ask the following questions:

1. Which types of interaction make certain ligand binds tightly to a receptor?
2. How do the interactions induce structural and functional changes in the receptor?

To answer these questions, we need methodologies to monitor the ligand-receptor systems at very high spatial and temporal resolutions, which unfortunately might not be achievable with current experimental methods. But with the rapid development of computational power over the decades, molecular dynamics (MD) simulations provide an opportunity to overcome such limitations[16]. In many studies, MD simulations not only directly measure important quantities such as binding affinities and dissociation rates but also provide an atomic-level description of the binding/unbinding processes[17, 18, 19, 20]. In the next section, I will briefly introduce how we can use simulation methods to study biomolecular systems and discuss challenges in this field.

## 1.2 Molecular dynamics simulation

### 1.2.1 Introduction to molecular dynamics simulation

Molecular dynamics (MD) simulation is a computational method to study the time evolution of molecular systems consisting of many atoms. In MD simulations, we consider a generic atomic system comprising  $N$  atoms, where the  $3N$  position coordinates are denoted by  $\mathbf{X}$ , and the interaction between atoms is described by the potential  $U(\mathbf{X})$ . According to Newton's laws of motion, the time evolution of atom coordinates should

follow the differential equation:

$$\frac{d^2\mathbf{X}}{dt^2} = -\frac{1}{m}\nabla U(\mathbf{X}) \quad (1.2)$$

However, a system follows Eq.1.2 will have conserved total energy and could not be a good model to describe what happens in reality. In reality, a molecular system is usually not isolated from the environment. It can interact with the environment by transferring heat, doing work or even exchange particles. In order to mimic such effects, algorithms called barostatting and thermostatting are introduced and are separately used to control the pressure and temperature in an MD simulation[21, 22, 23, 24]. Though differ in their implementation details, different versions of barostatting and thermostatting algorithms all bring stochastic to the simulations and make it possible to simulate different ensembles like canonical (NVT) ensemble and isothermal–isobaric (NPT) ensemble which are widely used in biomolecular simulations.

### 1.2.2 Simulations of the thermodynamic and kinetic properties of biomolecules

When studying biological systems with MD simulations, we are typically interested in cases relevant to biophysical systems characterized by the existence of metastable states. For example, as showed in Fig.1.1, if we project the system along a collective variable (CV), we can see some low free energy basins, which are usually referred to as metastable states. During a simulation, a system usually spends extended amounts of time each metastable states and infrequently moves among those states. Typical objective

in performing MD is to evaluate the following two broad types of quantities for some generic low-dimensional CVs  $s(\mathbf{X})$ :

1. equilibrium properties, such as the free energy  $F(s)$ .
2. dynamic properties, such as the mean first passage time for escaping any metastable state.

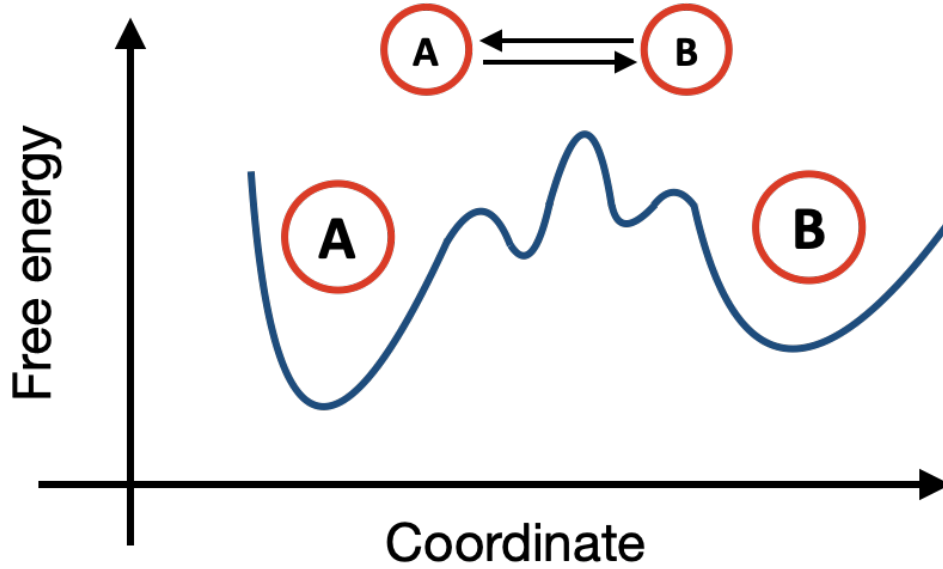


Figure 1.1: An illustration of a biomolecular system's free energy profiles after projecting along one collective variable (CV).

The free energy is related to the probability of the CV by  $P(s) \equiv e^{-\beta F(s)} = \int d\mathbf{X} \delta[s - s(\mathbf{X})] P_0(\mathbf{X})$ , where  $P_0(\mathbf{X}) \propto e^{-\beta U(\mathbf{X})}$  is the equilibrium probability of a microstate  $\mathbf{X}$  and  $\beta = \frac{1}{k_B T}$  is the inverse temperature. From a set of  $N$  samples  $\{\mathbf{x}_0, \mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_N\}$  from an MD simulation, the probability of the CV can be estimated by:

$$P(s) = \langle \delta[s - s(\mathbf{X})] \rangle \approx \frac{\sum_i^N \delta(s - s(\mathbf{x}_i))}{N} \quad (1.3)$$

In the case of ligand-receptor interaction, if we use  $A$  and  $B$  to denote the bound and

unbound states which are two subsets of the state space. We can calculate the dissociation constant discussed in Sec. 1.1.2 by  $K_D = \exp(\beta(F_B - F_A)) = \frac{\int_{s \in \{s(x)|x \in B\}} P(s) ds}{\int_{s \in \{s(x)|x \in A\}} P(s) ds}$ .

To estimate the transition rate between metastable states, we can first define  $\tau_{A \rightarrow B}$  to be the time that it takes for a trajectory starts from  $A$  first reaches  $B$ :

$$\tau_{A \rightarrow B} = \min \{t \geq 0 \mid x(0) \in A, x(t) \in B\} \quad (1.4)$$

The first passage time from  $A$  to  $B$  is the expectation value of  $\tau_{A \rightarrow B}$  over trajectories. Again, if  $A$  is the bound state and  $B$  is the unbound is state,  $k_{off} = 1/\tau_{A \rightarrow B}$  and  $k_{on} = 1/\tau_{B \rightarrow A}$ .

### 1.2.3 Reaction coordinate

One important ingredient of the calculation mentioned above is the CV. In principle, we arbitrarily project the system on some variables and consider them as CVs. However, most of these projections are not very useful. They cannot serve as good variables to describe the relevant reaction processes and thus can not be used to define the metastable states or calculate the thermodynamic and kinetic quantities. Good CVs representing the reaction process are usually called reaction coordinates (RCs). Many studies focused on defining what a good RC for molecular systems is[25]. But even though the definition exists, calculating RCs can still be challenging. For example, we can define isocommittor[26]  $q_{a \rightarrow b}(x) = P(\tau_B > \tau_A \mid x_0 = x)$ , which is the probability that a trajectory initiate at  $x$  will first commit to metastable state  $A$  instead of  $B$ .  $q_{a \rightarrow b}(x)$  can be calculated by initiating simulations starts at each point  $x$  in state space get the

statistics of which state they first commit to. This calculation is usually too expensive to be applied to molecular systems.

#### 1.2.4 Challenges in MD simulations

Besides the challenge in finding good RCs, to accurately simulate a biomolecular system to use computation results to reveal what is happening in reality, we need to reduce both the statistical and systematic errors[27]. The systematic error is caused by the deviation of classical molecular models from the real quantum world. How to adjust the force field to better approximate will not be the focus of this thesis. Instead, I will focus on discussing how to reduce the statistical error in MD simulation, which is caused by the insufficient sampling of the processes of interest. In Sec. 1.3.1, I will discuss the sampling problem in detail and show how it connects to the problem of finding good RCs.

### 1.3 Solving rare event problem with machine learning

#### 1.3.1 Rare event problems

On the application of MD simulations to study ligand-receptor interaction, one of the main challenges is to overcome the gap between the time scale that can be reached by MD simulations and the time scale of the events that we want to study. More specifically, the dissociation of a ligand from its receptor usually happens at a time scale from microseconds to seconds[28]. On the contrary, MD simulation, which has to numerically solve Newton's law of motion at femtosecond ( $10^{-15}$  s), might need years to simulate an unbinding event with very powerful computational resources. Though

compared with less than two decades ago when the longest MD simulation performed was only  $2\ \mu s$ [29, 30], MD simulation has become more efficient with the development of special-purpose supercomputers like ANTON2 ( $100\ \mu s$  per day for systems with a million atoms) [31] and better combination between MD software and hardware like Graphics processing units (GPUs)[32], it is still associated with high energy and time cost and could not support high throughput study of a large amount of systems. Moreover, if you want to simulate the system in a more accurate manner by considering some degrees of freedom of electrons[33, 34, 35, 36] or even fully considering the quantum effect[37], the gap of the time scales will be much bigger. To overcome these limitations and make MD applicable to generic systems with arbitrary complexity and timescales, algorithms need to be developed to enhance sampling when there are limited computational resources. In the following section, I will introduce some widely used enhanced sampling methods and discuss their limitations.

### 1.3.2 Enhanced sampling methods

Different enhanced sampling methods have been proposed to solve the time scale problem. We can group these methods into three broad categories:

1. **Raising the temperature of the system:** By running the system at higher temperatures, we can enhance its fluctuation and thus accelerate the sampling of states. As an example, Replica-exchange molecular dynamics (REMD)[38] runs many replicas of the system at different temperatures. Every given number of simulation steps, configurations in different replicas are exchanged according to



probabilities that satisfy detailed balances:

$$p(\text{accept}) = \min \left( 1, e^{(E_i - E_j) \left( \frac{1}{kT_i} - \frac{1}{kT_j} \right)} \right) \quad (1.5)$$

Here  $E$  is the total energy of the system,  $T$  is the temperature. Systems in the high-temperature replicas are more likely to explore the high free energy regions or cross the energy barriers, while the replica with the lowest temperature can still give the right ensemble of states. One signification drawback of REMD is its poor scaling with the system size. According to Eq.1.5, the acceptance probability becomes very small if the difference is much bigger than the thermal fluctuation. As a consequence, for systems of  $N$  particles, the temperature difference between neighbour replicas should be small enough to ensure that that the energy difference, which is proportional to  $N$ , is comparable to the thermal fluctuation, which is proportional to  $\sqrt{N}$ . Therefore, it becomes impractical to apply REMD on large systems as it might require a huge amount of replicas. To overcome this problem, methods like Replica Exchange with Solute Tempering (REST)[39, 40] has been developed. In REST, instead of directly run many replicas with different temperature, systems with different Hamiltonians scaled by the control parameter  $\beta_m$ :

$$E_m^{\text{REST}}(\mathbf{X}) = \frac{\beta_m}{\beta_0} E_{\text{pp}}(\mathbf{X}) + \sqrt{\frac{\beta_m}{\beta_0}} E_{\text{pw}}(\mathbf{X}) + E_{\text{ww}}(\mathbf{X}) \quad (1.6)$$

Here the subscripts pp, pw and ww denote protein intramolecular energy, the interaction energy between the protein and water, and the self-interaction energy

between water molecules separately. Such a scaling operation can be considered heating only part of the system and almost equivalent to reducing the  $N$ , thus improving the acceptance rate of exchanging configurations between replicas. Other methods try to improve REMD includes introducing nonequilibrium switches between replicas to increase phase space overlap [41].

Besides needs for many replicas, REMD has been rarely applied to studying unbinding problems. Because raising temperature might induce conformational change of the system, leading to big perturbations on the binding or unbinding pathways.

2. **Taking advantage of parallel computing:** Instead of running a very long trajectory, methods based on parallel computing are developed to extract useful information from multiple short trajectories. In order to guarantee an unbiased and efficient sampling of the ensemble, these methods need to design smart strategies about how to initiate, terminate, and analyze the trajectories. As an example, the weighted ensemble (WE) sampling[42, 43] uses the idea of resampling to calculate the free energy and transition rate. In WE, many trajectories are run in parallel. Each trajectory can be split into multiple trajectories or terminated according to the number of trajectories in each pre-defined bin. These trajectories carry different weights to account for the effects of these splitting or terminating operations. By enforcing the number of trajectories in each bin, WE guarantees that the low probability regions can still be well sampled. Another example is milestone sampling[44], which is used to get the transition rate between metastable states.

In the milestone method, a sequence of hypersurfaces that are perpendicular to a RC is defined as milestones. Trajectories are initiated from each hypersurface and the distributions of first-passage-time of arriving at neighbour hypersurfaces are calculated. With these distributions, systems' long time kinetic property can be estimated. One disadvantage of methods of this type is that they are still relatively more expensive in terms of computational resources. Also, the choice of RCs is critical for the success of these methods. In WE, bins are usually defined by a partition of a space spanned by a few RCs; In milestone, the assumption that dynamic along RC is much slower than the dynamic along the directions perpendicular to it should be satisfied[44].

3. **Modify the Hamiltonian of the system:** The transitions between metastable states become exponentially rare as the free energy of transition states increases[45]. By adding external potentials to reduce the free energy differences between transition states and metastable states, these methods are designed to make it easier for the system to across the energy barriers in a simulation. Metadynamics[46, 47] is one example of this type, where during the simulation, biasing potentials of Gaussian form will be added to the system every given number of MD steps. As the biasing potential accumulating in the regions that have been frequently visited, it will be easier for the system to escape from free energy basins to sample more metastable states. Based on this framework, several strategies have been proposed to improve the performance of metadynamics. One widely used variant is well-tempered metadynamics, where instead of adding a Gaussian biasing potential with

a constant height, the new Gaussian biasing potential will be scaled by a factor  $\exp \left[ -\frac{1}{\gamma-1} \beta V_{n-1}(\mathbf{s}_n) \right]$  where  $V_{n-1}(\mathbf{s}_n)$  is the external potential added at previous steps and  $\mathbf{s}_n$  denotes the center of new Gaussian biasing potential. Therefore, the new biasing potential added to the system will become smaller and smaller on the places where large external potential has been added. This strategy brings a better convergence property and allow a simple way to estimate the unbiased average of observables:

$$\langle O(\mathbf{R}) \rangle = \langle O(\mathbf{R}) e^{\beta[V(\mathbf{s}(\mathbf{R}),t)-c(t)]} \rangle_V c(t) = \frac{1}{\beta} \log \frac{\int d\mathbf{s} \exp \left[ \frac{\gamma}{\gamma-1} \beta V(\mathbf{s}, t) \right]}{\int d\mathbf{s} \exp \left[ \frac{1}{\gamma-1} \beta V(\mathbf{s}, t) \right]} \quad (1.7)$$

in addition to this, as shown in reference [48], if the biasing potential is added at relatively low frequency to ensure that the biasing potential added on the transition states is negligible, the transition time between metastable states can also be estimated by:

$$\tau_{\text{transition}} = \sum_i^N dt e^{\beta V(\mathbf{s}(t_i), t_i)} \quad (1.8)$$

where  $N$  is the total number of MD steps.

I would like to point out that for Metadynamics, or similar methods like umbrella sampling[49], the basing potentials are defined as functions of RCs, which is discussed in Sec. 1.2.3. Only when the RCs capture the slow mode of the system will these methods be able to accelerate the sampling of rare events.

As we can see, in most cases, the success of methods in categories (2) and (3) relies on the choice of RCs. However, in many studies, the RCs were selected in an intuitive, *ad*

*hoc* way. This is mainly due to the challenge of learning the RCs for a system with many degrees of freedom as discussed in Sec. 1.2.2. Many other methods were developed to solve problems of this kind. Their central theme is finding a simpler description of the systems to better understand their physical properties, such as free energy difference or characteristic time scale of different modes. However, most of these methods still suffer from the drawback that they need well-sampled unbiased MD trajectories as input.

In summary, to use MD simulations to study long-timescale processes, we need first to know RCs that capture the system's slow modes. However, methods developed to learn the slow modes require trajectories from simulations that have already given good sampling of the process. The main question that I will address to answer in this thesis is how to sample rare events in MD simulations when neither long trajectories nor optimal RC is available.

### 1.3.3 Machine learning approaches for enhancing MD simulations

The rapid development of artificial intelligent (AI) and machine learning (ML) techniques has shed light on solving challenging problems both in daily life and scientific research. Compared with traditional methods, it has unprecedented ability to learn and find relationship from complex dataset. This makes it an ideal tool to study biomolecular systems. By ML here I mean methods incorporating artificial neural networks (ANN).

A large category of ML methods developed for MD simulations follow some variant of the scheme shown in Fig.1.2. That is, ML is used to find a projection from the high dimensional structure space to a low dimensional feature space. In these methods, ML has




Figure 1.2: A schematic illustrating the typical workflow of some of the methods that use machine learning to analyze and enhance MD simulation. High dimensional data which describes the time evolution of the system in configuration space is used as the input of an ANN. The ANN is trained to project the input to a low dimensional space. Depending on the structure of NN and objective function, the low dimensional representation captures different features that are considered to be important, such as slow modes. In some methods, the feature learnt are used to further enhanced the sampling of MD simulation as shown by the arrow in the bottom.

been used in various forms to analyze long MD trajectories and learn relevant slow modes.

Ma and Dinner in 2005 used ANNs for constructing a RC[50], the idea being to sample configurations along transition pathways and tabulate an ensemble of committor values, with ANNs trained to fit CVs to these values. Recent work has revived and somewhat generalized this framework[51].

While a committor might be the most rigorous description of a slow mode, there are other reasonably accurate and computationally cheaper principles defining what constitutes slow modes. One of such a principle is based on the eigendecomposition of the time-propagation of the system. The dynamics in the full high-dimensional  $\mathbf{X}$ -space can be safely assumed to be Markovian, and thus the time-propagation of the system in  $\mathbf{X}$ -space can be described for a given lag time  $\tau$  using a transfer operator  $K$  with eigenvalues  $1 = \lambda_0 > \lambda_1 \geq \dots$  and corresponding eigenvectors  $\psi_0(\mathbf{X}) = 1, \psi_1(\mathbf{X}), \psi_2(\mathbf{x}), \dots$ [52].

The eigenvectors with indices  $i \geq 1$  correspond to the slow modes of the system with corresponding timescales  $t_i = -\frac{\tau}{\log \lambda_i}$ . For a molecular system with very high-dimensional dynamics, there is no simple way to solve the eigenvalues and eigenvectors of their transfer operator.

### 1.3.3.1 Variational approach to conformation dynamics

One formalism to solve the eigendecomposition problem is the variational approach to conformation dynamics (VAC)[53]. VAC calculates the eigenvalues and eigenvectors of the transfer operator starting with  $\psi_1$ , by solving a variational principle which states that any trial eigenvector  $\psi'_1$  will have a time-lagged autocorrelation  $\langle \psi'_1 | K | \psi'_1 \rangle \leq 1$ , as long as it is orthonormal to  $\psi_0$ , with equality holding iff  $\psi'_1 = \psi_1$ . Thus we can approximate the first nontrivial eigenfunction  $\psi_1$  by searching for a  $\psi'_1$  that maximizes its autocorrelation function subject to the orthonormal conditions. Ref.[53, 54] show how to calculate such matrix elements given unbiased MD data. Further modes can be learned in a similar manner but with more orthonormality conditions. Methods such as TICA[55] implement VAC directly by learning slow eigenvectors as linear combinations of pre-selected basis functions, and are at the heart of building MSMs. The Variational approach for Markov processes (VAMP) principle[56] generalizes the mathematical framework in the VAC principle to nonstationary and nonreversible processes, while keeping the same key underlying intent of learning a transformation of  $\mathbf{X}$  in which the dynamics is as Markovian as possible. VAMPnets use ANNs to implement the VAMP principle and automate the various steps involved in construction of a MSM. Some improvements

on VAMPnets includes methods such as state-free reversible VAMPnets (SRV), which combine strengths of VAC and VAMPnets for equilibrium systems, and use ANNs to construct a full hierarchy of slow modes expressible as non-linear functions of input coordinates.

Since the training of ML models involves finding optimal parameters for the model from a space with many local optimas. A smart design of objective function can make a model more likely to converge to an optimal one when there are limited amount of data. For example, slow modes  $s$  as per VAC may also be equivalently defined as a mapping  $s(\mathbf{X})$  that maximizes autocorrelation  $A(s)$  for a lag time  $\tau$  :

$$A(s) = \frac{\mathbf{E} [\tilde{s}(\mathbf{X}_t) \tilde{s}(\mathbf{X}_{t+\tau})]}{\sigma(s(\mathbf{X}_t)) \sigma(s(\mathbf{X}_{t+\tau}))} \quad (1.9)$$

where  $\tilde{s} = s - \mathbf{E}(s)$  is the mean-free latent variable and  $\sigma(s)$  is the standard deviation of  $s$ . Time-lagged autoencoders (TAE) use an encoder-decoder framework to find the slow component so defined. However, instead of calculating the precise expectation value in Eq.1.9, TAE approximates the slow mode by minimizing the reconstruction loss  $\mathcal{L}_R$ :

$$\mathcal{L}_R \equiv \mathcal{L}_{\text{TAE}} = \sum_t \|\mathbf{X}_{t+\tau} - D(E(\mathbf{X}_t))\|^2 \quad (1.10)$$

Here  $E$  is the encoder which maps the input configuration  $\mathbf{X}$  to low-dimensional  $s$  and  $D$  is the decoder mapping it back to coordinate space. Maximizing Eq.1.9 and minimizing Eq.1.10 are identical if the encoder and decoder are linear[57].



### 1.3.3.2 Variational Dynamic encoder

Variational Dynamic encoder (VDE)[58] also uses an encoder–decoder pair with time-lagged data, but instead of having only a term related to reconstruction error  $\mathcal{L}_R$ , VDE includes two additional terms. Firstly, it takes the idea of variational autoencoder (VAE) to enforce a prior distribution of the latent variable  $s_t$ . The use of a Gaussian prior helps the distribution of  $s'_t$  to be smooth, allowing meaningful interpolation between states in latent space. Secondly, an autocorrelation loss  $\mathcal{L}_{AC} = -A(s)$  is introduced. It encourages the learning of modes with high autocorrelation and makes the training process easier to converge.

Work from Olsson and Noé has extended the notion of encoding a global molecular configuration,  $\mathbf{X}$ , into encoding several local configurations  $\mathbf{X}^1, \mathbf{X}^2, \dots, \mathbf{X}^j$ , each representing a partition of the original global molecular structure into a local substructure[59]. To model the time evolution of  $\mathbf{X}$  the propagator or conditional distribution  $p(\mathbf{X}_t|\mathbf{X}_{t-\tau})$  is written in terms of the substructures:

$$p(\mathbf{X}_t|\mathbf{X}_{t-\tau}) \propto e^{\sum_i \mathbf{X}_t^i (\sum_j J_{ij}(\tau) \mathbf{X}_{t-\tau}^j + h_i(\tau))} \quad (1.11)$$

where  $J_{ij}$  is the coupling parameter between the  $i^{th}$  and  $j^{th}$  subsystem and  $h_i$  describes the coupling between the  $i^{th}$  subsystem with an external field. With the choice of this model, the problem of determining the coupling parameters can then be reduced to  $N$  logistic regression problems. A notable feature of this approach is that it seems capable of predicting molecular configurations that have not been incorporated into the training

data.

The above motioned methods focus on learning the slow mode from MD data. Now I will review approaches that use ML to not just analyze existing MD generated structures and trajectories, but also actively enhance the sampling capacity of MD. In other words, these approaches use ML to not just learn from given data, but actually generate statistically accurate information when the underlying processes are so slow that they simply cannot be sampled in unbiased MD even with the best available computing resources.

#### 1.3.3.3 Molecular enhanced sampling with autoencoders

The molecular enhanced sampling with autoencoders (MESA) approach uses an ANN with non-linear encoder and decoder to learn the RC from input data, which itself is generated through umbrella sampling along trial RCs[60]. Every round of ML leads to an improved RC along which new umbrella sampling is performed. The iterations are continued until the free energy from umbrella sampling no longer varies with further iterations. Similar to MESA, nonlinear RCs learned by methods like VDE[61] can also be used to perform enhanced sampling, typically using TICA modes as input variables.

#### 1.3.3.4 Neural networks-based variationally enhanced sampling

The NN based variationally enhanced sampling (VES) method also stands out with respect to many of the other ML methods mentioned here as it does not try to learn slow modes, but instead tries to express the bias as a smoothly differentiable ANN potential as

a function of pre-selected small number of CVs[62]. To learn such a bias, it optimizes the objective function introduced in [63].

### 1.3.3.5 Reinforcement learning based methods

The reinforcement learning based adaptive sampling (REAP) method of Ref.[64] learns relevant CVs on-the-fly as exploration of the landscape is carried out. REAP starts with a dictionary of OPs  $\mathbf{v}$  and associated trial weights. A round of unbiased MD is carried out, clustered into states, after which the weights are adjusted in order to maximize a reward function. In a nutshell, the reward function is designed to favor the least populated clusters and is iteratively adjusted as new clusters are visited.

### 1.3.3.6 Boltzmann Generators

Boltzmann Generator is a very recent deep learning based approach that learns the equilibrium probability  $P_0(\mathbf{X})$  without resorting to running long trajectories[65]. It leverages recent advances in generative modeling in which an invertible coordinate transformation  $F_{xz}(\mathbf{X})$  that maps microstate  $\mathbf{X}$  onto a random variable  $\mathbf{z}$  is learnt[66, 67]. If  $\mathbf{z}$  follow a distribution that is straightforward to sample, we can pass a sample from  $P(\mathbf{z})$  and map it back to a microstate  $\mathbf{X}$  with  $F_{zx}(\mathbf{z}) = F_{xz}^{-1}$ . The probability density of  $\mathbf{z}$  is  $P_{\mathbf{z}}(\mathbf{z}) = P_{\mathbf{x}}(\mathbf{X}) |J_{zx}|$ , where  $J_{zx}$  is the Jacobian of  $F_{zx}$ . The invertible transformation is learned such that  $P_x(\mathbf{X})$  is closed to the Boltzmann distribution,  $P_0(\mathbf{X}) \propto e^{-\beta U(\mathbf{X})}$ . Such invertible transformation replaces the Monte Carlo sampling on  $\mathbf{X}$ -space to Monte Carlo sampling on  $\mathbf{z}$ -space where the Monte Carlo moves can be easily designed. The model

can use this strategy to generate more samples in  $\mathbf{x}$ -space and optimize the invertible transformation by learning from the new samples. Once the algorithm converges,  $P_x(\mathbf{X})$  will give a good approximation of  $P_0(\mathbf{X})$  and be used to generate samples. The author also propose a reweighting scheme to assign weights to samples. So even  $P_x(\mathbf{X})$  might slightly deviate from  $P_0(\mathbf{X})$ , the weighted samples can still be used to give estimation of the free energy profiles of the systems. Later improvement on Boltzmann Generators includes introducing stochastic transformations to support more flexible learning[68] and design a network structure that can adjust the learned distribution to satisfy the distribution at different temperatures[69].

## 1.4 Overview of the thesis

In this thesis, I will show my attempts at solving the problems mentioned in the previous sections. I have explored the possibility of taking help from artificial intelligence (AI) and statistical mechanics to develop frameworks that can accelerate the sampling of the configuration space of complex (bio)molecular systems, making it possible to calculate various thermodynamic and kinetic observables. I also applied these enhanced sampling methods to study the dissociation of ligand from proteins and RNA, and compare the computational results with experimental results to better understand the mechanism of the ligand-receptor interactions. A summary of these efforts is as follows.

1. In Chapter 2, I will discuss the method Reweighted autoencoded variational Bayes for enhanced sampling (RAVE), which can automatically learn a low dimensional description of systems with hundreds of thousands of atoms, and use

this information to sample the events that could not be studied with normal MD. RAVE is based on the framework of information bottleneck[70, 71, 72]. I will use variational inference to reformulate predictive information bottleneck (PIB) as a neural network and take advantage of AI to overcome the challenges faced by other sampling methods. This AI-based sampling method is applicable to generic biomolecular systems and it is not limited by the lack of prior knowledge about the system.

2. In Chapter 3, I will focus on theoretical study of RAVE. I will discuss the choice of the important parameter time delay in the model, and show that it can be chosen from a relatively wide range without significantly change the RCs learned by the model. I will also derive a correction term to address of the problem of learning from biased MD.
3. In Chapter 4, I will apply RAVE to study the thermodynamics, kinetics and mechanisms of unbinding of two different ligands from PreQ<sub>1</sub> riboswitch. In this study, by applying the AI augmented MD and with the help of experimental techniques, I aim to better understand the interplay between RNA structure and its interaction with different ligands. I will show how MD simulation can reproduce quantities measured in experiment. Moreover, I will make predictions about critical mutants based on what we learn from MD simulations and show how it has been validated in experiment.
4. In Chapter 5, I will explore using the generative AI to learn from very high dimensional data and generate samples. I use a generative AI model called

Denoising diffusion probabilistic model (DDPM) to improve the sampling quality of replica exchange sampling. I will show that by treating temperature as a random variable instead of a control parameter, we can learn a joint distribution to discover transition states and metastable states that are previously unseen at the temperature of interest.

5. In Chapter 6, I will give a summary of works presented in this thesis and suggest possible direction for future research.

## Chapter 2: Reweighted autoencoded variational Bayes for enhanced sampling (RAVE)

### 2.1 Introduction

To address this time scale problem discussed in the Sec.1.3.1, I developed an AI-based enhanced sampling method called Reweighted autoencoded variational Bayes for enhanced sampling (RAVE). In the following sections, I will introduce this method and show its application to different systems. I will first demonstrate the method on sampling alanine dipeptide and calculate the full dissociation process of benzene from L99A mutant of the T4 lysozyme protein[73, 74]. In the later system, with use of all-atom MD simulations taking barely a few hundred nanoseconds in total, and with the minimal use of prior human intuition as in other related methods, I used the method to obtain accurate thermodynamic and kinetic information for a process that takes few hundred milliseconds in reality.

This method is based on the ansatz is that efficient sampling of energy landscapes of molecular systems has the same key underlying challenge as one faced by a fly as it goes about surviving[75], or the human brain trying to process how to catch a

moving baseball[70]. Namely, given limited storage and computing resources, which memories to preserve and which ones to ignore in order to be best prepared for various possible future challenges? This can be paraphrased as the ability to rapidly learn a low-dimensional representation of a complex system that carries maximal information about its future state. Since storing and processing large amounts of information can be computationally and thus energetically expensive for the brain, it has been suggested that neurons in the brain separate predictive information from the non-predictive background in a way that by encoding and processing a minimum amount of relevant information, the brain can still be maximally prepared of future outcomes. The past–future (or Predictive) Information Bottleneck framework introduced and developed in many forms [70, 71, 72, 75] involves implementing such neuronal models from an information theoretic basis that can originally be traced back to Shannon’s rate distortion theory.

Reweighted autoencoded variational Bayes for enhanced sampling (RAVE) [76, 77] is an iterative ML-MD method motivated by the observation that many feature learning methods, in addition to classifying features, also provide the probability distribution in feature space [78]. The learnt features and their probability distribution can then respectively be used as RC and its fixed or static bias can then be applied to  $U(\mathbf{X})$  leading to more ergodic exploration. To learn these features and their probabilities, RAVE uses a past-future information bottleneck optimization scheme that outputs a minimally complex yet maximally predictive model. An ANN decoder is trained to predict the future state of the system instead of only trying to recover the input data, and a linear encoder is used to get an interpretable projection from the space of order parameters to the RC. The active enhanced sampling (AES) approach [79] is another approach similar to RAVE, that uses



well-tempered metadynamics [80, 81] and a stochastic kinetic embedding formalism.

## 2.2 Using latent variables of VAE as RCs

Our first attempt on solving the rare event problem start with learning RCs by the unsupervised machine learning method Variational Autoencoder(VAE)[58] and construct a biasing potential as the learned RCs to speed up the simulation[76].

The VAE is comprised of an encoder,  $q(z|\mathbf{X})$ , mapping the original high-dimensional data into its low-dimensional latent space representation and a decoder,  $p(\mathbf{X}|z)$ , mapping such a latent variable representation back into the original dataspace. By enforcing  $z$  to be close to a simple prior distribution such as Gaussian distribution, the  $z$  learnt by the model will be able to capture the data’s main features. In the context of the MD simulations, the latent variable representation will describe a low-dimensional manifold for the molecular simulation trajectories within configuration space and thus could serve as a good approximation of the RCs. So the learnt variables and their probability distribution can then respectively be used as RC and its static bias can then be applied to leading to more ergodic exploration.

The RAVE protocol involves iterating between machine learning and MD, with each iteration learning a more refined low-dimensional RC representation and bias potential. These iterations then continue until these biasing parameters are converged and sufficient to sample the desired ligand dissociation. One benefit of this approach is that a time-independent bias is produced that can be used to launch multiple independent production MD runs, which is helpful in extracting kinetics information with statistical robustness.

In this proof-of-concept study, RAVE shows its usefulness on several model systems, including two analytical potentials as well as a hydrophobic buckyball-substrate system in explicit water. All these systems have extremely high barriers (between 5 kBT and 30 kBT), and using RAVE, we can obtain near-ergodic sampling and converged free energy profiles both accurately and efficiently[76]. However, I would like to note that there are some shortcomings if only use the framework of VAE to learn the RCs. By construction, VAE is designed to reproduce the original sample, so it complete ignores the kinetic of MD trajectories dataset. So, instead of learning the slow mode of a system, it might focus on learning the mode with high variance. Also, the restriction enforce the distribution of the latent variable to be closed to a Gaussian prior makes hard to learn free energy landscape with many basins. To overcome these limitations, I use a more general framework called information bottleneck to replace the VAE in RAVE and thus allow RAVE to find the RCs by learning from the kinetic of the system.

### 2.3 Principle of Past–Future Information Bottleneck

Here I interpret the RC in molecular systems as such a past–future information bottleneck [71]. I develop a sampling method that for small biomolecules, simultaneously and with minimal use of human intuition, learns this bottleneck, its thermodynamics and kinetics. The central idea is that not all features of the past carry predictive value for the future. A complex model can be made to be very predictive, however it will often obscure physical interpretability and also end up capturing noise. In order to address this task, I set up an optimization problem and demonstrate how to solve it through

the principle of variational inference[82] implemented through deep neural networks. This makes it possible to learn a predictive information bottleneck[72], which we can interpret as the RC, that given a molecule’s past trajectory is maximally predictive of its future behavior. The net product is an iterative framework on the lines of Ref.[76] that starts from a short MD simulation, and given this data, makes an estimate of the RC, its Boltzmann probability, and its associated causal Green’s function valid for short times. This information is leveraged to perform systematically biased simulations with enhanced exploration of phase space, which can then be used to re-learn the RC along with its probability and propagator, and iterating between MD and variational inference until optimization is achieved. At this point we have converged estimates of the most informative degrees of freedom, associated metastable states and their equilibrium probabilities. Finally, through the use of a generalized transition state theory based framework on the lines of Ref.[83], we recover the unbiased kinetics for moving between different metastable states.

I formalize this problem in terms of a high-dimensional signal  $\mathbf{X}$  characterizing the state of a  $N$ -particle system under some generic set of thermodynamic conditions. I take  $\mathbf{X}$  as some  $d$  generalized coordinates or basis set elements, where  $1 \ll d \ll N$ . Let the value of this signal measured at time  $t$ , or the past, be denoted by  $\mathbf{X}_t$  and at time  $t + \Delta t$ , or the future by  $\mathbf{X}_{t+\Delta t}$ . We call  $\Delta t$  the prediction time delay. We assume that  $\mathbf{X}_t$  and  $\mathbf{X}_{t+\Delta t}$  are jointly distributed as per some probability distribution  $P(\mathbf{X}_t, \mathbf{X}_{t+\Delta t})$ . The mutual information  $I(\mathbf{X}_t, \mathbf{X}_{t+\Delta t})$  (See appendix for more details) quantifies how much an observation at one instant of time  $t$  can tell us about an observation at another instant of time  $t + \Delta t$ . Furthermore, here we can restrict our attention to stationary

systems, hence omit the choice of time origin and write down  $\mathbf{X}_t$  as  $\mathbf{X}$  and  $\mathbf{X}_{t+\Delta t}$  as  $\mathbf{X}_{\Delta t}$ . The principle of Predictive information bottleneck (PIB)[71, 72] postulates a bottleneck variable  $s$  which is related to  $\mathbf{X}$  by an encoder function  $P(s|\mathbf{X})$ . Given the bottleneck variable  $s$ , predictions of the future  $\mathbf{X}_{\Delta t}$  can be made with a decoder  $P(\mathbf{X}_{\Delta t}|s)$ . PIB says that the optimal bottleneck variable is one which is as simple as possible in terms of the past it needs to know, yet being as powerful as possible in terms of the future it can predict correctly. This intuitive principle can be formally stated through the optimization of an objective function  $\mathcal{L}$  which is a difference of two mutual informations:

$$\mathcal{L} \equiv I(s, \mathbf{X}_{\Delta t}) - \gamma I(\mathbf{X}, s) \quad (2.1)$$

The above objective function quantifies the trade-off between complexity and prediction through a parameter  $\gamma \in [0, \infty)$ .

### 2.3.1 Variational inference and neural network architecture

Typically, both the encoder  $P(s|\mathbf{X})$  and the decoder  $P(\mathbf{X}_{\Delta t}|s)$  can be implemented by fitting deep neural networks [84] to data in form of time-series of  $\mathbf{X}$ . This work stands out in three fundamental ways to typical implementations of the information bottleneck principle and in general of AI methods to sampling biomolecules[85, 86, 87, 88]. Firstly, I use a stochastic deep neural network to implement the decoder  $P(\mathbf{X}_{\Delta t}|s)$ , but use a simple deterministic linear encoder  $P(s|\mathbf{X})$  (see Fig. 2.1). The simple encoder ensures that the information bottleneck or RC we learn is actually physically interpretable, which




Figure 2.1: Network architecture used for learning predictive information bottleneck  $s$ . The decoder  $Q(\mathbf{X}_{\Delta t}|\mathbf{s})$  is a stochastic deep neural network, while the encoder  $P(\mathbf{s}|\mathbf{X})$  is of a simple deterministic and thus directly interpretable linear form.

is notably hard to achieve in machine learning. On the other hand by introducing noise in the decoder, I can control the capacity of the model to ensure that the neural network can delineate useful feature from useless information instead of just memorizing the whole dataset. Secondly, now that our encoder is a simple linear model, I completely drop the complexity term in Equation (2.1) and set  $\gamma = 0$ . Due to a reduced number of variables, this leads to a simpler and more stable optimization problem. Finally, the rare event nature of processes in biomolecules makes it less straightforward to use of information bottleneck/AI methods for enhanced sampling. Here I develop a framework on the lines of [76] that makes it possible to maximize the objective function in Equation (2.1) through the use of simulations that are progressively biased using importance sampling as an increasingly accurate information bottleneck variable is learnt.

The typical starting point is an unbiased MD trajectory  $\mathbf{X} = \{\mathbf{X}^1, \dots, \mathbf{X}^M\}$  with  $M$

data points. We want to develop a low-dimensional mapping  $\mathbf{s}$  of this high-dimensional space, that maximizes the objective function  $\mathcal{L} = I(\mathbf{s}(\mathbf{X}), \mathbf{X}_{\Delta t})$ . At the heart of this mutual information lies the calculation of the decoder  $P(\mathbf{X}_{\Delta t}|\mathbf{s})$ , which can in principle be done exactly using Bayes' theorem (See Appendix ). However this becomes impractical as soon as the dimensionality of  $\mathbf{X}$  increases, due to a fundamental problem in statistical mechanics and machine learning: intractability of the partition function in high dimensions [89]. The principle of variational inference is one such elegant and powerful approach to surmount this problem [90].

Let's consider a generic encoder given by some conditional probability  $P_{\theta}(\mathbf{s}|\mathbf{X})$  where  $\theta$  is a set of parameters. Our objective then is to find the optimal RC or equivalently, the encoder  $\theta$  which optimizes the PIB objective:

$$\theta^* = \underset{\theta}{\operatorname{argmax}} \mathcal{L}(\theta) \quad (2.2)$$

As mentioned above, this optimization problem is intractable for almost all cases of practical interest. However, it is possible to perform an approximate inference problem by assuming an approximate decoder  $Q_{\phi}(\mathbf{X}_{\Delta t}|\mathbf{s})$  parametrized by the vector  $\phi$ . For any choice of  $\phi$ , I make a straightforward use of Gibbs' inequality[82] to write down (See Appendix for derivation ):

$$\begin{aligned} I(\mathbf{s}, \mathbf{X}_{\Delta t}) &= H(P_{\theta}(\mathbf{X}_{\Delta t})) - H(P_{\theta}(\mathbf{X}_{\Delta t}|\mathbf{s})) \\ &\geq H(P_{\theta}(\mathbf{X}_{\Delta t})) - C(P_{\theta}(\mathbf{X}_{\Delta t}|\mathbf{s})||Q_{\phi}(\mathbf{X}_{\Delta t}|\mathbf{s})) \end{aligned} \quad (2.3)$$

Here  $H$  and  $C$  denote Shannon entropy and cross entropy respectively. Take note that the first term in Equation (2.3) is independent of our model parameters and hence can be completely ignored from the optimization. Focusing on the second term in Equation (2.3), we thus obtain a variational lower bound on the predictive information bottleneck objective function:

$$\mathcal{L} \geq \mathcal{L}' = -C(P_{\theta}(\mathbf{X}_{\Delta t}|\mathbf{s})||Q_{\phi}(\mathbf{X}_{\Delta t}|\mathbf{s})) \quad (2.4)$$

Thus  $\mathcal{L}'$  is a tractable lower bound bound to the true Predictive Information Bottleneck objective function  $\mathcal{L}$ , that involves a variational approximation through the trial decoder parametrized by  $\phi$ . It has a simple physical interpretation. We are attempting to learn a decoder probability function  $Q$  that mirrors the actual Bayesian inverse probability function  $P$  in terms of predicting the future state  $\mathbf{X}_{\Delta t}$  of the system, given knowledge of the RC  $\mathbf{s}$ . The difference between the two probability distributions is calculated as a cross-entropy. By maximizing the right hand side of Equation (2.4) simultaneously with respect to the decoder and encoder parameters  $\phi$  and  $\theta$  respectively, we can then solve the actual optimization problem posited in Equation (2.4) rigorously and identify the optimal RC.

It is clear that a model of a dynamical system  $\mathbf{X}$  that attempts to capture just its stationary probability  $P(\mathbf{X})$  will be less informative and useful than one that captures the joint past-future probability distribution  $P(\mathbf{X}, \mathbf{X}_{\Delta t})$ . This is simply because the stationary probability can always be calculated by integrating  $P(\mathbf{X}, \mathbf{X}_{\Delta t})$  over future outcomes  $\mathbf{X}_{\Delta t}$ . What is however less clear is the choice of the time-delay  $\Delta t$  [87]. In

biomolecular systems, it is likely that there will be a hierarchy of time-scales and thus time-delays relevant to different types of structural and functional details. In principle, our formulation allows us to probe these various time-delays in a systematic manner. Here, for the purpose of enhanced sampling, I propose an approach for selecting  $\Delta t$  that is rooted in the reactive flux formalism of chemical kinetics[91, 92, 93]. This formalism applies to any system with stochastic transitions on a network of microstates with arbitrary, complex connectivity. Summarily, it states that the correlation function for a trajectory’s population in any given state can be partitioned into 3 parts: (a) an initial inertial part, (b) an exponential decay, and (c) an intermediate plateau region between (a) and (b). A key insight from this formalism is that capturing (c), i.e. the plateau part of a system’s state to state dynamics accurately is necessary and sufficient to capture the temporal evolution at any timescale. By paraphrasing this argument in the context of the present work, we propose to learn our PIB model for gradually increasing values of the predictive time-delay  $\Delta t$ , and stop when the calculated bottleneck variable converges.

### 2.3.2 Variational inference on unbiased and biased trajectories

I now show how to calculate  $\mathcal{L}'$  in practice. For a given unbiased trajectory  $\{\mathbf{X}^1, \dots, \mathbf{X}^{M+k}\}$  with large enough  $M$ , we can easily show (See Appendix for derivation):

$$\mathcal{L}' = \frac{1}{M} \sum_{n=1}^M \log Q(\mathbf{X}^{n+k} | \mathbf{s}^n) \quad (2.5)$$

where  $\mathbf{s}^n$  is sampled from  $P(\mathbf{s} | \mathbf{X}^n)$  and the time interval between  $\mathbf{X}^n$  and  $\mathbf{X}^{n+k}$  is  $\Delta t$ . For practical rare event systems however, a typical MD trajectory will be trapped



in the state where it was started. Here we use our current best estimate of the PIB to perform importance sampling of the landscape, so that the system is more likely to sample different regions in configuration space, and use this enhanced sampling to iteratively improve the quality of the RC. However, the data so generated is biased per definition, and we need to reweight out the effect of the bias. We suppose that along with the time series  $\{\mathbf{X}^1, \dots, \mathbf{X}^{M+k}\}$  we also have been provided the corresponding time-series for the bias  $V$  applied to the system  $\{V^1, \dots, V^{M+k}\}$ . We can then use the principle of importance sampling[94] to write our PIB objective function  $\mathcal{L}'$  as follows (See Appendix for derivation):

$$\mathcal{L}' = \left\{ \sum_{n=1}^M e^{\beta V^n} \right\}^{-1} \sum_{n=1}^M e^{\beta V^n} \log Q(\mathbf{X}^{n+k} | \mathbf{s}^n) \quad (2.6)$$

where  $\beta$  is inverse temperature. The above equation is however approximate, as it assumes  $P_{\text{biased}}(\mathbf{X}^{n+k} | \mathbf{s}^n) \approx P_{\text{unbiased}}(\mathbf{X}^{n+k} | \mathbf{s}^n)$ . This is exact as  $\Delta t \mapsto 0$ , and can be expected to be reasonably valid for small  $\Delta t$ , where we expect that the system on average would not have diffused too far from its starting position at the beginning of that interval. If the bias varies smoothly enough that its natural variation length scale is smaller or comparable to this diffusion distance, then for small enough  $\Delta t$  we can indeed make the aforementioned approximation. This means that we select the smallest possible  $\Delta t$  at which the RC estimate plateaus.

### 2.3.3 Patching it all

I will discuss our complete sampling algorithm, that accomplishes in a seamless manner, the identification of the RC together with the sampling of its thermodynamics

and kinetics. The first step is to perform an initial round of unbiased MD. This trajectory, expressed in terms of  $d$  order parameters  $\{\chi_1, \dots, \chi_d\}$  (where  $1 \ll d \leq N$ ), is fed to a deep learning module (Fig. 2.1). The deep learning module implements the optimization of  $\mathcal{L}'$  in Equation (2.6) through the use of multi-layer feed-forward neural network for the stochastic decoder  $Q$ , and a physically interpretable linear map for the deterministic encoder  $P$  (Fig. 2.1). Unlike the decoder, the encoder has no noise term and always maps  $\{\chi_1, \chi_2, \dots, \chi_d\}$  to  $\sum_i c_i \chi_i$ , where  $\{c_i\}$  denote the weights of different order parameters. I perform this optimization for gradually increasing values of the predictive time-delay  $\Delta t$ , and estimate RC  $\mathbf{s}$  (given by the values of the weights  $c_i$ ) as seen by the first plateau in terms of when it ceases to depend on choice of  $\Delta t$ . This value of  $\Delta t$  is then kept constant for different rounds of our protocol. At this point we have an initial estimate of  $\mathbf{s}$  and also its unbiased probability distribution  $P^u(\mathbf{s})$ . These are both used to construct a bias potential  $V_{\text{bias}}(\mathbf{s})$  for the next iteration of MD:

$$V_{\text{bias}}(\mathbf{s}) = k_B T \log P^u(\mathbf{s}) \quad (2.7)$$

With this bias potential added to the original Hamiltonian of the system, we run a biased MD simulation. This explores an increased amount of configuration space since we have applied a bias along our estimated slow degree of freedom, viz. the PIB or the RC. This next round of MD trajectory is again fed to the deep learning module, but this time each data point carries a weight  $w = e^{\beta V_{\text{bias}}}$  to compensate for the applied bias. This now identifies an improved RC  $\mathbf{s}$  and its unbiased probability through the use of importance

sampling:

$$P^u(\mathbf{s}) \propto \frac{\langle w \delta(\mathbf{s} - \mathbf{s}(t)) \rangle_b}{\langle w \rangle_b} \equiv e^{-\beta F(\mathbf{s})} \quad (2.8)$$

where the subscript  $b$  denotes sampling under a biased ensemble with weight  $w = e^{\beta V_{\text{bias}}}$  and  $F(\mathbf{s})$  is the free energy along  $\mathbf{s}$ . From here, using the bias as  $-F(\mathbf{s})$  our algorithm can now enter into further iterations of MD–deep learning–MD–... This looping continues until both the RC  $\mathbf{s}$  and the free energy estimate  $F(\mathbf{s})$  along the RC have converged. We have thus obtained an optimized reaction coordinate and its Boltzmann probability density, or equivalently the free energy. Through these we can directly demarcate the relevant metastable states and quantify their relative propensities. Furthermore, we can also calculate the transition rates for moving between these metastable states. The central idea is to keep all transition states between the different metastable states, as identified through the RC, devoid of any bias. As we show in examples, this can be easily achieved when implementing equation (2.8), by ensuring that any barriers in the unbiased probability distribution of the estimated RC are completely bias-free. Once we have done this, we take into account that by virtue of it being the PIB, the RC already encapsulates any relevant, predictive modes in the system. Thus the hidden barriers which have invariably been corrupted through the addition of such a bias do not have any predictive power for the dynamics of the system, and are thus not relevant to the process at hand. This then implies that (i) the biased dynamics preserves the state-to-state sequence one would have seen with unbiased dynamics, and (ii) through the use of a simple time rescaling calculation[83, 95] discussed in Sec.1.3.2, we can calculate the acceleration of rates achieved through biased simulations. We now demonstrate the use of the PIB

framework with two biomolecular case-studies, in both of which we simultaneously learn the RC, free energy and kinetic rate constants. In each case the RC  $s$  is constructed as a linear combination  $s = \sum_i c_i \chi_i$ , where  $\{c_i\}$  denote the weights of different pre-selected order parameters  $\{\chi_1, \chi_2, \dots, \chi_d\}$ .

## 2.4 Application of RAVE to two benchmark systems

### 2.4.1 Conformation transitions in a model peptide

First we consider the well-studied alanine dipeptide system (Fig. 2.2(a)). This system, as characterized by its Ramachandran dihedral angles, can exist in different metastable states with varying stabilities and hard-to-cross intermediate barriers. However, due to its small size it serves as a reliable benchmark where we can perform longer than microsecond unbiased MD simulations to benchmark our PIB calculations.

Here we choose  $\{\cos \phi, \sin \phi, \cos \psi, \sin \psi\}$  as our order parameters, where  $\phi$  and  $\psi$  are the backbone dihedral angles. To avoid problems related to periodic boundary conditions, I take trigonometric functions of dihedral angles, The PIB protocol used here is shown in Fig.2.2. In the initial round, I perform a short unbiased MD simulation (see Fig.2.2(a) for trajectory for technical details). As discussed in Sec. 2.3, the RC is then determined as the linear encoder of the trained neural network with smallest loss function. In Fig.2.2(b), we show how the weights rapidly converge as functions of predictive time delay and reach a plateau in less than 2 ps which we set as  $\Delta t$ . Fig. 2.2(c) shows the RC  $s$  as well as the bias  $V(s)$  learnt along it to be used in the next round of MD. With this biasing potential, I perform biased MD simulation as shown in Fig. 2.2(d). This trajectory

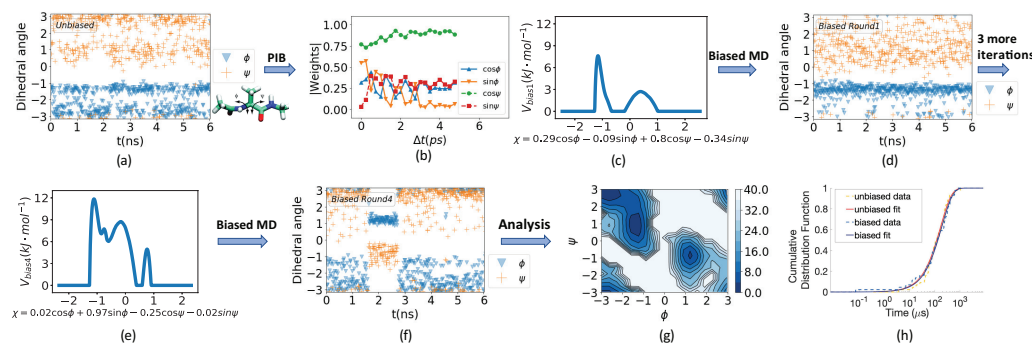


Figure 2.2: Past-future information bottleneck framework results on alanine dipeptide. (a) Unbiased simulation trajectory for dihedral angles  $\Phi$  and  $\Psi$ , in blue triangles and orange pluses respectively. The alanine dipeptide molecule is shown in inset. (b) Absolute weights for different order parameters in the first training round as a function of the predictive time delay  $\Delta t$ . Blue triangles, orange triangles, green circles and red squares correspond to  $\cos\phi$ ,  $\sin\phi$ ,  $\cos\psi$  and  $\sin\psi$  respectively. (c-f) Free energy along the adaptively learnt RC along with the corresponding biased trajectories for different training rounds. (g) Free energy along  $\phi, \psi$  after the RC has converged with energy contours every  $4 \text{ kJ} \cdot \text{mol}^{-1}$ . (h) Kinetics from the post-training biased runs as well as reference unbiased runs. The two are essentially indistinguishable. Orange and blue dashed lines denote unbiased data respectively, while red and black solid lines show corresponding best fits.

through the use of Equation (2.8) leads to a more complicated bias structure as shown in Fig. 2.2(e) along with the improved RC  $\chi$ . Biased simulation with this new RC and bias as shown in Fig. 2.2(f) finally leads to escape from the starting metastable state. The final obtained RC is:  $s = 0.02 \cos \phi + 0.97 \sin \phi - 0.25 \cos \psi - 0.02 \sin \psi$ . It is known for alanine dipeptide that  $\phi$  is more relevant than  $\psi$  for capturing the conformational transitions, and the PIB based RC estimate agrees with that.

Now that we have achieved back-and-forth motion in terms of the rare event we intended to study, I use this final RC and bias to perform multiple sets of longer simulations with no further refinement of the RC. This yields the free energy surface (defined as  $-k_B T \log P(\phi, \psi)$  where  $P$  is the unbiased Boltzmann probability) as shown in Fig. 2.2(g). This is in excellent agreement with previously published benchmarks for

this system[94]. At the same time, I use the acceleration factor to rescale the biased time back to the unbiased time. In Fig. 2.2(h) we show the cumulative distribution functions of the first passage time from the deeper basin as obtained through this approach, and through much longer unbiased MD runs which are feasible given the small size of this system. The distribution functions and their best-fit Poisson curves are nearly indistinguishable, and lead to excellent agreement in values of the escape rate constant, given by  $k = 5.2 \pm 0.8\mu s^{-1}$  and  $5.8 \pm 0.9\mu s^{-1}$  respectively for biased and unbiased simulation.

#### 2.4.2 Benzene dissociation from T4-L99A lysozyme

Now I will apply our framework to a very challenging and important test case, namely the pathway and kinetic rate constant of benzene dissociation from the protein T4-L99A lysozyme in all-atom resolution[73, 74]. I will also demonstrate how the RC calculated through our approach can be directly used to perform a sensitivity analysis of the protein, and predict the most important residues whose mutations could have a significant affect on the stability of the protein-ligand complex. Such an analysis has direct relevance to predicting, for instance, the mutations in a protein which could lead to a pharmacological drug losing its efficacy.

For this problem we choose 11 fairly arbitrary order parameters denoted  $\{\chi_1, \dots, \chi_{11}\}$ . As shown in Fig.2.3 and Table 2.1, eight of these are ligand-protein distances while three are intra-protein distances. The RC is learnt as a linear combination of these order parameters, namely  $s = \sum_i c_i \chi_i$ .




Figure 2.3: Definition of Order parameters for Benzene-Lysozyme dissociation: Our 11 order parameters comprise 8 protein-ligand contacts and 3 protein-protein contacts. The protein-ligand contacts are implemented through the distances between the centres-of-mass of ligand to different protein atoms labeled in the plot. The protein-protein contacts are implemented through the distance between the centres-of-mass of ligand to centres-of-mass of helices 1 labeled in the plot. Further details are provided in Table 2.1

For this problem as well I start with a short unbiased MD simulation. As shown in Fig. 2.4, the weights of different order parameters in the RC learnt from this trajectory change as a function of the predictive time-delay  $\Delta t$ , but converge quickly. On the basis of this plot, I set  $\Delta t = 2ps$  for all further calculations. I then iterate – using the same neural network architecture as for alanine dipeptide (Fig. 2.1) – between rounds of learning an iteratively improved RC  $s_1$  together with its probability distribution, and

running biased MD using the iteration's RC and probability distribution as bias  $V_1(s_1)$  (Equation (2.7)). After 9 rounds, the bias saturates as a function of training rounds. That is, no further enhancement in egodicity is achieved by performing additional rounds of the aforementioned iteration. This corresponds to the system reaching configuration where the previous PIB ceases to be effective. To learn a new PIB, I use the washing out trick from [96] to learn a second RC  $s_2$  conditioned on our knowledge of the first RC  $s_1$ . In the next few rounds of learning-MD iterations, I (a) keep  $s_1$  and  $V_1(s_1)$  fixed, and (b) do not account for  $V_1(\chi_1)$  when using Equation (2.8). Through this I construct a bias  $V(s_1, s_2) = V_1(s_1) + V_2(s_2)$ . In principle we can lift this assumption and learn more complicated non-separable  $V(s_1, s_2)$ . In a few rounds of training  $s_2$ , I observed

OP $\chi$	Type	Definition	Weight $c_i$ in optimized $s_1 = \sum c_i^2 \chi_i$	Weight in optimized $\chi_2 = \sum c_i^2 \chi_i$
$d_1$	Protein-ligand	Y88CA–ligand	0.4930	0.6524
$d_2$	Protein-ligand	A99CA–ligand	0.5059	0.3984
$d_3$	Protein-ligand	L133CA–ligand	-0.4002	-0.4853
$d_4$	Protein-ligand	L118CA–ligand	0.2335	-0.0792
$d_5$	Protein-ligand	V111CA–ligand	0.0300	0.0399
$d_6$	Protein-ligand	A130CA–ligand	-0.4009	-0.3037
$d_7$	Protein-ligand	N140CA–ligand	0.1229	-0.1544
$d_8$	Protein-ligand	A146CA–ligand	0.2313	-0.2240
$hd_{12}$	Protein-protein	Helix 1 (A82-S90) – Helix 2 (T115-123Q)	-0.2081	-0.0378
$hd_{23}$	Protein-protein	Helix 2 (T115-123Q) – Helix 3 (W126-A134)	-0.1114	-0.0073
$hd_{34}$	Protein-protein	Helix 3 (W126-A134) – Helix 4 (K147-T155)	-0.0204	0.0670

Table 2.1: List of order parameters used to construct RC and their weights in  $\chi_1$  and  $\chi_2$ .



spontaneous dissociation of the ligand from the protein. I then use the RC  $(s_1, s_2)$  (shown in Fig. 2.5) and its bias  $V(s_1, s_2)$  learnt to directly study the pathway and kinetics of ligand dissociation.




Figure 2.4: Order parameters weights as functions of time delay for Benzene-lysozyme dissociation. The scheme shows the absolute value of weights for 11 order parameters in the first round as a function of predictive time delay.

For this I launch 20 independent biased simulations using  $(s_1, s_2)$  as RC and  $V_1(s_1) + V_2(s_2)$  as bias. I calculate the acceleration factor to recover the original timescale of the first passage time. As shown in Fig. 2.6, I fit the cumulative distribution function to a Poisson process and get an escape rate constant of  $3.3 \pm 0.8 s^{-1}$ , which are in good agreement with other methods [20, 97, 98, 98]. I also obtain a range of free energies viewed as functions of different order parameters. These are in excellent agreement with previously published results [20, 96, 97, 98].

A comment I would like to make here concerns the magnitude of  $\Delta t$  especially




Figure 2.5: The 2-component Predictive Information Bottleneck for benzene-lysozyme dissociation, where colors red and blue correspond to  $s_1$  and  $s_2$  respectively. The optimized weights for different order parameters are illustrated after scaling all weights to keep  $c_1 = 1$  in  $s_1$ .

in connection to the decorrelation time of the MD thermostat. Here I implemented a canonical ensemble using the velocity rescaling thermostat [22] with a time-constant of 0.1 ps. The predictive time-delay  $\Delta t$  is thus at least 20 times longer than the times for which the history of the thermostat would persist. Interestingly, as can be seen in Fig. 2.4 the estimate of the RC converges with longer time-delays which would be even more accurate from the perspective of not having thermostat-induced noise, but would be less accurate due to biasing related errors as explained earlier.

### 2.4.3 Predicting critical residues

On the basis of the predictive information bottleneck that I have now calculated, we can directly predict which protein residues have the most critical effect on the system. To do so, the guiding principle is that the residues which carry higher mutual information with the PIB are more likely to have an impact on the stability of the system, for instance,




Figure 2.6: Poisson fit to the 15 independent transition times obtained for the benzene-lysozyme dissociation constant with net maximum bias  $V_1^{max} + V_2^{max} = 52 \text{ kJ} \cdot \text{mol}^{-1}$ . The  $k_{off}$  reported in the main text is calculated as the inverse of the fitted time constant here,  $303 \pm 76 \text{ ms}$ .




Figure 2.7: Representative selected free energy surfaces for Benzene-Lysozyme dissociation obtained from our approach: (a) Free energy along  $d_1$  and  $d_2$ . (b) Free energy along  $d_2$  and  $d_3$ . (c) Free energy along  $d_1$  and  $d_8$ . All energies are in units of  $\text{kJ} \cdot \text{mol}^{-1}$  with contours every  $5 \text{ kJ} \cdot \text{mol}^{-1}$ . These surfaces are in excellent agreement with previous benchmarks on this system [96].

if these residues were to be mutated. I use the backbone dihedral angles to describe the motion of each residue. Here I denote them as  $\phi_i, \psi_i$ , where  $i$  is the index of a particular

residue. We use  $I(\theta, \chi)$  to quantify the correlation between dihedral angle  $\theta$  (where  $\theta$  could be  $\phi$  or  $\psi$ ) and the unbinding process, where  $I(\theta, \chi)$  is mutual information between a dihedral angle  $\theta$  and the RC  $\chi$ . For each residue, we have two dihedral angles. At the same time, for one dihedral angle, I calculate its mutual information with  $\chi_1$  or  $\chi_2$ . So we have four quantities for each residue(  $I(\phi, \chi_1), I(\psi, \chi_1), I(\phi, \chi_2)$  and  $I(\psi, \chi_2)$  ). I rank the importance of each residue by the maximum of the four quantities. Here I only consider the parts of trajectory that is bias-free( $\chi_1 > 0.4$  and  $\chi_2 > 0.3$ ). Because we require that the energy barriers between metastable states should be bias-free to ensure we can get the correct reweighted dynamics of the unbinding process. However, we CAN only guarantee that the main energy barrier between bound and unbound state has zero bias when we perform biased MD simulation while bias can still be added on barriers between other metastable states. Looking at the unbiased region allows us to reduce the influence of the biasing potential on the dynamics of the system and focus more on the transition states. By performing such a scan of the mutual information between the predictive information bottleneck and the backbone dihedral angles of different residues, we can rank them as being most critical to least critical. As shown in Fig. 2.8, some of the important residues are (in order of decreasing relevance): Ser136, Lys135, Asn132, Leu133, Ala134, Phe114, Val57, Asp20, Leu118 and Val131. These residues can be classified in two broad groups. Firstly, we have group (a) comprising residues 114, 118, and 131–136 - together these contribute to breathing movement between the two helices through which the ligand leaves. Secondly we have group (b) comprising residues 20 and 57, which lie in different disordered regions of the protein, and have no obvious interpretation. The roles of groups (a) have been hinted at in previous works[20, 99, 100]




Figure 2.8: Critical residue analysis for benzene-lysozyme complex. The plot on left shows for every residue the maximal mutual information between the PIB and either of the Ramachandran angles  $\phi, \psi$  of that residue. The top 10 residues are highlighted through markers and in the right plot, illustrated relative to the ligand in a typical intermediate pose.

and are thus yet another validation of our approach. The role of group (b) during the unbinding process remains to be seen.

In order to demonstrate the robustness of this calculation, I performed the full MD–deep learning–MD–deep learning iterative protocol with a new set of order parameters (defined in Table.2.2) that considered new protein-ligand distances and completely excluded any protein-protein distances, as the selection of latter require a more significant role of human intuition in anticipating protein breathing for instance. In this new set of calculation, as shown in Fig.2.9, I again obtained same critical residues, including a residue from the same disordered region as in group (b) above. Whether the disordered

Order parameter	Type	Definition	Weight $c_i$ in optimized $s_1 = \sum c_i^1 \chi_i$	Weight in optimized $s_2 = \sum c_i^2 \chi_i$
$d_1$	Protein-ligand	V87CA–ligand	0.0297	0.8604
$d_2$	Protein-ligand	N101CA–ligand	0.1448	0.2590
$d_3$	Protein-ligand	G110CA–ligand	-0.0583	-0.0813
$d_4$	Protein-ligand	A119CA–ligand	-0.0261	0.2256
$d_5$	Protein-ligand	A129CA–ligand	-0.4076	0.0282
$d_6$	Protein-ligand	Y139CA–ligand	0.8846	-0.3548
$d_7$	Protein-ligand	V149CA–ligand	0.1593	-0.0921

Table 2.2: List of order parameters used to construct RC and their weights in  $\chi_1$  and  $\chi_2$ .

regions are biophysically relevant to the unbinding of the ligand with the existence of a long-range allosteric communication pathway, or if these residues are picked due to just noise from calculations, needs more detailed mutagenesis study in the future.

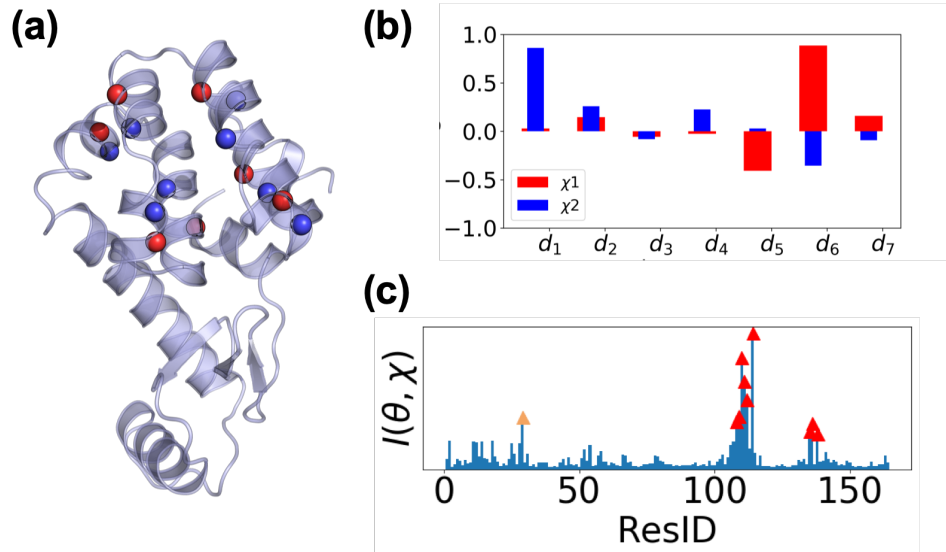


Figure 2.9: (a) Alpha carbons used to define two sets of basis functions. Carbon atoms used in Table 2.1 are colored in blue and the carbon atoms used in Table 2.2 are colored in red. (b) The 2-component Predictive Information Bottleneck for benzene-lysozyme dissociation. Order parameters are defined by Table 2.2. (c) Critical residue analysis from new set of order parameters.

## 2.5 Discussion

In this chapter, I introduced a new framework for the simultaneous sampling of the reaction coordinate, free energy and rate constants in biomolecules with rare events. This work is grounded in the Predictive Information Bottleneck framework, which is an information theoretic approach for building minimally complex yet maximally predictive models from data. Such a framework has previously been found useful for modeling fruit fly movement and human vision. Here I exploit the commonality between these diverse problems and that of sampling complex biomolecular systems, namely the need to quickly predict the future state of a system given noisy and high-dimensional information. This method implements this framework through the use of a unique linear encoder–stochastic decoder model, where the latter is a deep neural network with inbuilt noise. Here I demonstrated the applicability of the method by studying conformational transitions in a model peptide in vacuum and ligand dissociation from a protein in explicit water, with both systems in all-atom resolution. Through extremely short and computationally cheap simulations, I obtained thermodynamic and kinetic observables for slow biomolecular processes in excellent agreement with other methods, experiments and long unbiased MD. Last but not the least, by virtue of having captured the most predictive degrees of freedom in the system, I could also make direct predictions of how protein sequence can impact dissociation dynamics - namely, which mutations in the protein would be most deleterious to the dissociation process.

I would also like to discuss here some limitations of this approach that I will address in the following chapters. As discussed in Sec. [2.3.2](#), since I didn't consider

the correction for propagator, Eq. 3.8 is a good approximation only when  $\Delta t$  is small enough that the system on average would not have diffused too far from its starting position. But it is also required that  $\Delta t$  should be taken from the range where RC reaches a plateau. Though I showed that in the two benchmark systems, we can choose a relatively small  $\Delta t$  as the RCs quickly converge to a constant as  $\Delta t$  increase, we can still ask the question of whether this is generally true for other systems. In chapter 3, I will discuss the  $\Delta t$  dependence of the objective function and drive a new formulation which also considers that correction to the propagator. In addition to figuring out how to better approximate the objective function from biased simulation, one other issue that we will need to address is the choice of input for RAVE. Here I use of a simple linear encoder which preserves the interpretability of the RC. However this comes at the cost of using smartly designed non-linear basis functions, or order parameters, which can often be domain-dependent. For example, different classes of basis functions were needed here for alanine dipeptide conformation change (namely, torsions) and ligand dissociation (namely, distances). Using a linear encoder on distances for alanine dipeptide, or torsions for ligand dissociation, leads to no discernible enhancement in sampling as we iterate through rounds of deep learning and MD. Thus, bad choices of basis functions for the protocol can be ruled out at least in a heuristic manner by quantifying whether these led to more ergodic sampling or not. In Chapter 4, I will show that when applying RAVE to more challenging systems, we can use method AMINO[101] developed in our group to automatically select the basis functions.



## Chapter 3: Understanding the role of predictive time delay and biased propagator in RAVE

### 3.1 Introduction

In this chapter, I will take a closer look at the underlying formulation of RAVE with respect to crucial aspects which were introduced in our past work using intuitive arguments. Here I am going to put these aspects on a rigorous footing, identifying their shortcomings and when applicable introducing simple corrections. These are as follows. RAVE[77] involves learning the PIB through optimization of an objective function that minimizes model complexity while still maximizing predictive power. Firstly, a key parameter in constructing the RAVE objective function is the so-called predictive time-delay  $\Delta t$ , or how far into the future should the algorithm try to predict. In the previous chapter, I had shown through numerical benchmarks that the precise choice of  $\Delta t$  did not matter, as long as  $\Delta t > 0$  and was kept small. That is, the predicted PIB changed but rapidly plateaued with increase in  $\Delta t$ . Here, by using a master Equation based formulation of the system's dynamics, I will demonstrate rigorously why such a dependence on  $\Delta t$  is a direct and simple consequence of Markovianity along the learnt low-dimensional representation. In fact, the absence of a plateau in how the PIB depends

on  $\Delta t$  can even be considered a tell-tale sign of incorrect results from RAVE.

Secondly, at the initial stage of RAVE, when the input trajectory is unbiased, the RAVE objective function can be computed exactly by sampling over the input trajectory. However, after the first iteration once the PIB is learnt and used to perform a round of biased MD, it becomes tricky to calculate the objective function exactly through sampling. In the application mentioned in Sec.2.4 and, I introduced an intuitive approximation, namely that at short timescales the propagator of the system along the putative PIB is invariant between biased and unbiased dynamics. Naturally this statement is exactly true as  $\Delta t \mapsto 0$ , but its fate for  $\Delta t > 0$  is not at all obvious. Here, again by taking recourse in a master equation formalism, I will derive a framework for using biased sampling to estimate a reweighted, unbiased propagator needed in RAVE that is valid for small  $\Delta t > 0$ . I will show how it reduces to the original expression in limiting cases, and how to compute it. I also demonstrate through numerical examples when such a correction might matter and when it might not.

### 3.2 Dependence of predictive information bottleneck on $\Delta t$

Maximizing the mutual information  $I(s, \mathbf{X}_{\Delta t})$  in Eq. 2.1 forces the learnt linear PIB or RC  $s$  to capture features that can be used to build a predictive model. In addition, introducing a slight time delay in RAVE i.e. setting  $\Delta t > 0$  emphasizes the contribution of not just the features that are related to static distribution of the data, but specifically of the features that persist with time. As demonstrated for related work on time-lagged autoencoders[102], this amounts to learning slowly decorrelating aspects of the feature

space  $\mathbf{X}$ .

Two questions immediately arise after introducing such a non-zero time delay  $\Delta t$ . First, what precise value of  $\Delta t$  should we adopt when constructing the objective function? Second, how to approximate unbiased estimates of  $P(\mathbf{X}_{\Delta t}|\mathbf{X})$  from a biased MD trajectory which provides us  $P_{\text{biased}}(\mathbf{X}_{\Delta t}|\mathbf{X})$ ? The second question arises because access to unbiased estimates of  $P(\mathbf{X}_{\Delta t}|\mathbf{X})$  is critical to calculating the objective function  $\mathcal{L}$  in Eq. 2.1.

In this work we answer, with relative rigor, both the questions above. For the first question, we demonstrate that given markovianity in the higher dimensional space  $\mathbf{X}$ , the markovianity of the PIB follows quite naturally. In other words, there exists a range of non-zero time-delay  $\Delta t$  values for which the PIB is independent of the exact choice of  $\Delta t$ . Thus any small enough  $\Delta t$  can be used to construct the PIB. For the second question, we propose a new objective function which corrects for the influence of bias on the short-time propagator  $P(\mathbf{X}_{\Delta t}|\mathbf{X})$ , and test it on model systems. Reassuringly we find that the effect of the correction derived here is minimal, and our intuitive approximation of  $P_{\text{biased}}(\mathbf{X}_{\Delta t}|\mathbf{s}) \approx P_{\text{unbiased}}(\mathbf{X}_{\Delta t}|\mathbf{s})$  made in Sec. 2.4 was not a bad one.

As discussed in Sec. 2.3, estimating the PIB, or minimizing the objective function  $\mathcal{L}$  in Eq. 2.1 is the same as maximizing the cross entropy  $\mathcal{L}'$  between  $P_{\theta}(\mathbf{X}_{\Delta t}, \mathbf{s})$  and  $P_{\theta}(\mathbf{X}_{\Delta t}|\mathbf{s})$ :

$$\mathcal{L}' = - \int P_{\theta}(\mathbf{X}_{\Delta t}, \mathbf{s}) \ln P_{\theta}(\mathbf{X}_{\Delta t}|\mathbf{s}) d\mathbf{X}_{\Delta t} d\mathbf{s} \quad (3.1)$$

where  $\theta$  indicates the parameters of neural network. To understand how this objective function  $\mathcal{L}'$  depends on the predictive time delay  $\Delta t$ , we need to analyze the  $\Delta t$ -

dependence of  $P_\theta(\mathbf{X}_{\Delta t}|\mathbf{s})$  and  $P_\theta(\mathbf{X}_{\Delta t}, \mathbf{s})$ . To do so we start with a master equation framework for the propagator in the high-dimensional space  $\mathbf{X}$ , where we assume markovianity holds true:

$$\begin{aligned} P(\mathbf{X}_2, t + \Delta t | \mathbf{X}_1, t) = & \delta(\mathbf{X}_2 - \mathbf{X}_1)[1 - a(\mathbf{X}_1, t)\Delta t] \\ & + W_t(\mathbf{X}_2 | \mathbf{X}_1)\Delta t + O[(\Delta t)^2] \end{aligned} \quad (3.2)$$

Here  $W_t(\mathbf{X}_2 | \mathbf{X}_1)$  is the transition probability per unit time from  $\mathbf{X}_1$  to  $\mathbf{X}_2$  at time  $t$ , and  $a(\mathbf{X}_1, t) = \int W_t(\mathbf{X}_2 | \mathbf{X}_1) d\mathbf{X}_2$ .

At this point we can make two assumptions: (a) the distribution is stationary so  $P(\mathbf{X}_2, t + \Delta t | \mathbf{X}_1, t)$  can be denoted as  $P(\mathbf{X}_{\Delta t} | \mathbf{X})$ , and (b)  $\Delta t$  is small enough so higher order terms in Eq. 3.2 can be ignored giving:

$$P(\mathbf{X}_{\Delta t} | \mathbf{X}) = \delta(\mathbf{X}_{\Delta t} - \mathbf{X})[1 - a(\mathbf{X})\Delta t] + W(\mathbf{X}_{\Delta t} | \mathbf{X})\Delta t \quad (3.3)$$

In addition, as a direct consequence of the markovianity assumption in  $\mathbf{X}$  space, the following property must hold true:

$$P(\mathbf{X}_{\Delta t} | \mathbf{s}, \mathbf{X}) = P(\mathbf{X}_{\Delta t} | \mathbf{X}) \quad (3.4)$$

This property reflects that  $\mathbf{X}$  contains all the information needed to predict  $\mathbf{X}_{\Delta t}$ , and in addition using knowledge of  $\mathbf{s}$  can not improve the quality of prediction. With this

additional assumption and the use of Eq. 3.3, we have:

$$\begin{aligned}
P(\mathbf{X}_{\Delta t}, \mathbf{s}) &= \int P(\mathbf{s}, \mathbf{X}) P(\mathbf{X}_{\Delta t} | \mathbf{s}, \mathbf{X}) d\mathbf{X} \\
&= P(\mathbf{s}, \mathbf{X}_{\Delta t}) [1 - a(\mathbf{X}_{\Delta t}) \Delta t] \\
&\quad + \Delta t \int P(\mathbf{s}, \mathbf{X}) W(\mathbf{X}_{\Delta t} | \mathbf{X}) d\mathbf{X}
\end{aligned} \tag{3.5}$$

By rearranging Eq. 3.5,  $P(\mathbf{X}_{\Delta t}, \mathbf{s})$  can be expressed in terms of  $W(\mathbf{X}_{\Delta t} | \mathbf{X})$  and  $P(\mathbf{s}, \mathbf{X})$ :

$$\begin{aligned}
P(\mathbf{X}_{\Delta t}, \mathbf{s}) &= \frac{\int P(\mathbf{s}, \mathbf{X}) W(\mathbf{X}_{\Delta t} | \mathbf{X}) d\mathbf{X}}{a(\mathbf{X}_{\Delta t})} \\
&= \frac{\int P(\mathbf{s}, \mathbf{X}) W(\mathbf{X}_{\Delta t} | \mathbf{X}) d\mathbf{X}}{\int W(\mathbf{X}_{\Delta t'} | \mathbf{X}_{\Delta t}) d\mathbf{X}_{\Delta t'}}
\end{aligned} \tag{3.6}$$

In Eq. 3.6 both the numerator and denominator have any dependence on the predictive time-delay  $\Delta t$  only through the transition probability per unit time  $W$ , which in turn per construction does not depend on  $\Delta t$ . As such,  $P(\mathbf{X}_{\Delta t}, \mathbf{s})$  does not depend on  $\Delta t$ . As a direct consequence of this, the PIB estimated by maximizing  $\mathcal{L}'$  in Eq. 3.1 as well should not depend on the choice of  $\Delta t$ , as long as  $\Delta t$  is small enough that the master equation formalism of Eq. 3.3 is valid.

### 3.3 Correcting the propagator and the information bottleneck when using biased trajectory

So far the formalism assumes we have access to various unbiased estimates needed in order to maximize  $\mathcal{L}'$  in Eq. 3.1. However, the whole framework in RAVE is based on iterations between machine learning and progressively biased MD, we need to be able to construct PIB from biased trajectories as well. I am going to address the question of how to do so, and develop a corrected form of Eq. 3.1. In Sec.2.4 we worked with biased trajectories by making the assumption:

$$P_{\text{biased}}(\mathbf{X}^{n+k}|\mathbf{s}^n) \approx P_{\text{unbiased}}(\mathbf{X}^{n+k}|\mathbf{s}^n) \quad (3.7)$$

which led to the following  $\mathcal{L}'$  for a trajectory with length  $N + k$ :

$$\mathcal{L}' = \left\{ \sum_{n=1}^N e^{\beta V^n} \right\}^{-1} \sum_{n=1}^M e^{\beta V^n} \ln Q(\mathbf{X}^{n+k}|\mathbf{s}^n) \quad (3.8)$$

Here  $\Delta t$  equals the time elapsed in  $k$  MD steps and  $Q(\mathbf{X}^{n+k}|\mathbf{s}^n)$  is the probabilistic mapping given by the approximate decoder to approximate the real posterior probability  $P(\mathbf{X}^{n+k}|\mathbf{s}^n)$ .  $e^{\beta V^n}$  is a reweighting factor to correct for the bias in stationary probability density estimate. The need for an approximate decoder arises from the principle of variational inference wherein  $\mathcal{L}$  is bounded from below by  $\mathcal{L}'$ , and  $\mathcal{L}'$  will reach its maximum if and only if  $Q(\mathbf{X}^{n+k}|\mathbf{s}^n) = P(\mathbf{X}^{n+k}|\mathbf{s}^n)$ . This guarantees that maximizing  $\mathcal{L}'$  will also force the approximate decoder be as close to the real posterior probability

as possible. Thus in the limit that the neural network decoder is flexible enough to approximate the real decoder exactly, optimizing  $\mathcal{L}$  and  $\mathcal{L}'$  will give the same encoder. As such we can assume the decoder is exact when deriving the correction for the propagator.

Eq. 3.8 corrects for only one of the two biased aspects of the trajectory: (i) it reweights out the effect of the bias on the sampling probability at any given moment of time, (ii) it however ignores the effect of bias on the short-term dynamical evolution of the system as it assumes Eq. 3.7. Note that in the limit  $\Delta t \rightarrow 0$ , Eq. 3.7 becomes exact. But as  $\Delta t$  increases,  $P_{\text{biased}}(\mathbf{X}|\mathbf{s})$  gradually deviates from  $P_{\text{unbiased}}(\mathbf{X}|\mathbf{s})$  due to the existence of bias.

To correct for the bias in  $P_{\text{biased}}(\mathbf{X}|\mathbf{s})$ , we can first consider the relationship of transition probability between biased and unbiased MD. By discretizing the Smoluchowski equation along  $\mathbf{X}$ , the transition probability can be written as [103, 104]:

$$W(\mathbf{X}_{\Delta t}|\mathbf{X}) = \frac{D(\mathbf{X}_{\Delta t}) + D(\mathbf{X})}{2\|\mathbf{X}_{\Delta t} - \mathbf{X}\|^2} \exp\left\{-\frac{\beta[U(\mathbf{X}_{\Delta t}) - U(\mathbf{X})]}{2}\right\} \quad (3.9)$$

where  $U(\mathbf{X})$  is the equilibrium, unbiased free energy and  $D$  is the diffusion coefficient. In biased MD with bias  $V(\mathbf{s}(\mathbf{X}))$  applied as a function of a putative RC  $\mathbf{s}$ , the underlying free energy gets replaced by  $U(\mathbf{X}) + V(\mathbf{s}(\mathbf{X}))$ . Note that the bias is function of RC  $\mathbf{s}$  which itself is a well-defined function of basis functions or order parameters  $\mathbf{X}$  that depend on atomic coordinates, so the biasing force added on atoms is continuously differentiable. Assuming that the diffusivity itself stays unchanged due to the addition of bias (a common assumption in dynamical reweighting algorithms such as Ref. [105]), we can then write down a relation between the biased and unbiased transition probabilities

in terms of the bias:

$$W_b(\mathbf{X}_{\Delta t}|\mathbf{X}) = \sqrt{\frac{e^{\beta V(s(\mathbf{X}))}}{e^{\beta V(s(\mathbf{X}_{\Delta t}))}}} W_u(\mathbf{X}_{\Delta t}|\mathbf{X}) \quad (3.10)$$

where the subscripts  $u$  and  $b$  denote unbiased and biased measurements respectively. Eqs.

3.3 and 3.10 together provide the relationship between  $P_b(\mathbf{X}_{\Delta t}|\mathbf{s})$  and  $P_u(\mathbf{X}_{\Delta t}|\mathbf{s})$ :

$$P_b(\mathbf{X}_{\Delta t}|\mathbf{X}) = \sqrt{\frac{e^{\beta V(s(\mathbf{X}))}}{e^{\beta V(s(\mathbf{X}_{\Delta t}))}}} P_u(\mathbf{X}_{\Delta t}|\mathbf{X}) \quad \text{if } \mathbf{X}_{\Delta t} \neq \mathbf{X} \quad (3.11)$$

$$P_b(\mathbf{X}_{\Delta t}|\mathbf{X}) = 1 - \int_{\mathbf{X}'_{\Delta t} \neq \mathbf{X}_{\Delta t}} P_b(\mathbf{X}'_{\Delta t}|\mathbf{X}) d\mathbf{X}'_{\Delta t} \quad \text{if } \mathbf{X}_{\Delta t} = \mathbf{X} \quad (3.12)$$

Finally, we can write down the bias-corrected expression for the objective function

$\mathcal{L}'$  :

$$\begin{aligned} \mathcal{L}' = & \frac{c}{N} \sum_{n=1}^N e^{\beta V^n} \ln Q(\mathbf{X}^n|\mathbf{s}^n) \\ & - \frac{c}{N} \sum_{n=1}^N e^{\beta \frac{V^{n+k} + V^n}{2}} \ln Q(\mathbf{X}^n|\mathbf{s}^n) \\ & + \frac{c}{N} \sum_{n=1}^N e^{\beta \frac{V^{n+k} + V^n}{2}} \ln Q(\mathbf{X}^{n+k}|\mathbf{s}^n) \end{aligned} \quad (3.13)$$

where  $c$  is an irrelevant constant independent of the neural network parameters.

The last term in Eq.3.13 is similar to Eq.3.8, which only considers the correction for sampling points and not their dynamical propagation. The summation of the first two terms can be interpreted as the correction for the propagator. To gain some intuition into these, let's consider the effect of the biasing potential on the dynamics. It pushes the system from high bias region to low bias region, thereby spuriously increasing the



conditional probability for being found in low bias region, given that earlier the system was in high bias region. We thus want to reweight out this effect and only learn from the dynamics induced by the original potential, which is what Eq.3.13 achieves. If the trajectory goes from a state  $\mathbf{X}^n$  with higher bias  $V^n$  to a state  $\mathbf{X}^{n+k}$  with lower bias  $V^{n+k}$ , the summation of the first two terms in Eq.3.13 will be positive. Note that in Eq.3.13 the first two terms together in this regime encourage the decoder to generate features close to  $\mathbf{X}^n$  while the third term favors the decoder to generate features closer to  $\mathbf{X}^{n+k}$ . This will negate the spurious enhancement of probability of  $\mathbf{X}^{n+k}$  due to the biasing profile. Similar argument applies for the case  $V^{n+k} > V^n$ . I will conclude this section by highlighting that is a rough approximation valid for short  $\Delta t$ , and it can not completely correct the effect of bias on dynamics. In the following examples I will discuss the advantages and limitations of this new formula, and also analyze how valid was the intuitive approximation Eq.3.7 made in Sec.2.4.

### 3.4 Results

I am now going to demonstrate the nature and effect of the corrections derived in the previous section through a simple illustrative numerical example with exact results pertaining to the true reaction coordinate and its free energy. More specifically, I generated an unbiased trajectory on a model system and apply RAVE iterations to it. This system is simple enough that we could generate a long enough unbiased trajectory with sufficient sampling of different metastable states. RAVE calculations on this perfect unbiased trajectory serve as our benchmark. In parallel, I add a biasing potential along the

relevant slow mode of the system to generate a biased trajectory. Both of these trajectories are then subjected to different treatments as proposed in Sec. 3.3, and as explained in the remainder of this section, in order to judge how much of a difference our corrections really make.

### 3.4.1 System set-up

The model system in this work has three degrees of freedom and a governing potential energy  $U$  given by:

$$U(x, y, z) = 6(x^2 - 1) + 0.0375[(y - z)^2 - 8]^2 + 45(y + z)^2 \quad (3.14)$$

Fig.3.1 shows the free energy profile along each degree of freedom and the corresponding trajectory from evolving Langevin dynamics[106] with an integration time step of 0.01 units at temperature  $k_B T = 1$ . In agreement with the significantly higher energy barrier along  $x$ , the transitions along  $x$  happen much less frequently and it represents accurately the dominant slow mode in this system.

In this work I will focus on the difference between optimizing Eq.3.8 and Eq. 3.13. Therefore, I perform the biased MD by using the slow mode  $x$  as the reaction coordinate and constructing the bias potential on the basis of its analytical free energy. The bias potential is constructed with a maximum bias of  $3kT$ . The sum of the potential energy of the system and the bias potential is shown as a dashed line in Fig. 3.1 (a) and the biased MD trajectory can be found in supplementary information.

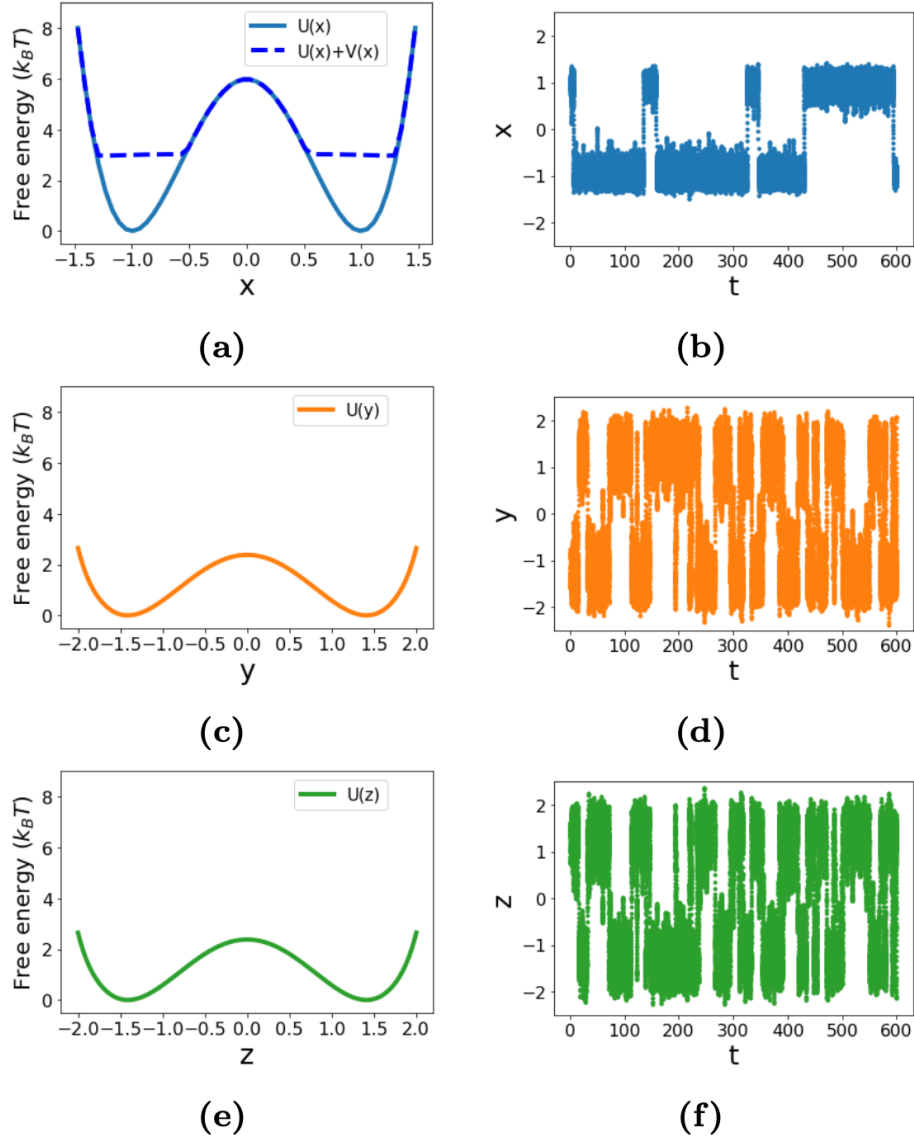


Figure 3.1: Free energy profile along each degree of freedom of the model system ((a), (c) and (e) respectively) and the corresponding trajectories from Langevin dynamics ((b), (d) and (f) respectively). In (a) the sum of underlying and bias potentials are indicated with a dashed line.

### 3.4.2 RAVE set-up

The input to RAVE comprised the three-dimensional time-series of  $\{x, y, z\}$  as obtained from Langevin dynamics. Whitening procedure from Ref. [107] was used to

reduce artifacts from high variance. The RC or PIB was learnt as a linear combination of  $\{x, y, z\}$ , thus characterized by corresponding three weights. Due to the possible non-convex and inherent stochastic nature of the optimization, there is no guarantee that a single training run gives us the most optimal RC. Indeed, here we trained 16 neural networks with randomly initialized weights, and found significant run to run variation in the output weights for  $\{x, y, z\}$ . Out of the 16 choices we selected the run that gave the smallest loss function. In practice, such a high number of trial runs is not needed, and a smaller number might very well be sufficient.

### 3.4.3 Calculated RC under different time-delays and approximation schemes

We use the unbiased and the biased trajectories as input in RAVE with different values of the predictive time-delay  $\Delta t$ . Fig. 3.2 shows the weights of  $x, y$  and  $z$  in the RC as obtained for these different set-ups. For the long well-sampled unbiased trajectory, I use Eq.3.8 setting  $V^n = 0$  (Fig. 3.2 (a)). This serves as a benchmark. For the biased trajectory as well, we first use Eq. 3.8 setting  $V^n = 0$  (Fig. 3.2 (b)). This illustrates the effect of biasing on the different modes. Next, for the biased trajectory, I use Eq. 3.8 taking the bias into account for reweighting the stationary probability, but not the propagator (Fig. 3.2 (c)). Finally, I use the biased trajectory correcting for the effect of bias on the stationary density as well as the propagator by making use of Eq.3.13 (Fig. 3.2 (d)).

The RAVE solution for unbiased trajectory shown in Fig. 3.2 (a) assigns the highest

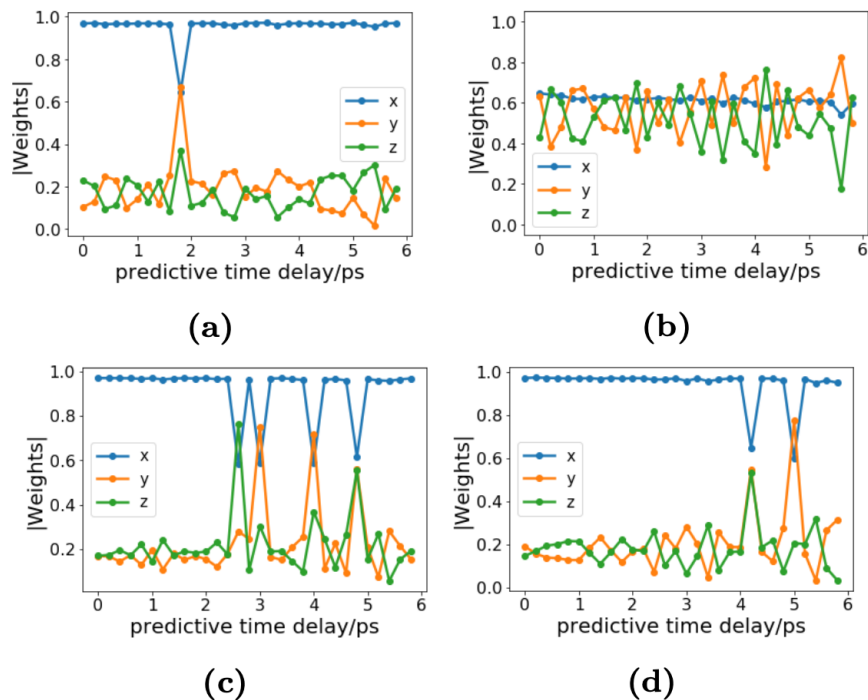


Figure 3.2: Reaction coordinates learned by RAVE as functions of predictive time delay. (a) RC learned from an unbiased trajectory using Eq. 3.8 setting  $V^n = 0$ . (b) RC learned from a biased trajectory using Eq. 3.8 with  $V^n = 0$ . (c) RC learned from a biased trajectory with Eq. 3.8 which takes the bias  $V^n =$  into account for reweighting the stationary probability, but not the propagator. (d) RC learned from a biased trajectory correcting for the effect of bias on the stationary density as well as the propagator by making use of Eq. 3.13.

weight to  $x$ , which is in agreement with our expectation that  $x$  is the slowest mode. The weights for  $y$  and  $z$  always have different signs as they are anticorrelated according to potential energy  $U$  in Eq. 3.14. Here I just want to compare their relative importance so only the absolute values of the weights are shown all throughout Fig. 3.2. Fig. 3.2 (b) shows results after treating the biased trajectory without taking the bias into account i.e. through the use of Eq. 3.8 with  $V^n = 0$ . Due to biasing along  $x$  and subsequent enhancement in fluctuations in this direction, now all 3 modes  $x$ ,  $y$  and  $z$  are comparable in their timescales, and without any reweighting RAVE assigns equal weights

to the three degrees of freedom. This illustrates why reweighting out effects of bias is crucial in RAVE, and arguably in general when using inputs from biased simulations as training data in machine learning. In Figs.3.2(c) and (d) I use Eq. 3.8 and Eq. 3.13 respectively to correct for the effect of bias on the stationary probability and both stationary probability/propagator. Both Eqs. 3.8 and Eq. 3.13 in Figs. 3.2 (c) and (d) respectively give a remarkably similar profile to that obtained from the long benchmark unbiased simulation used as input in RAVE. Thus, at least for this model system, it appears that our intuitive approximation of Eq. 3.7 made in Ref.[77] was reasonable.

### 3.5 Conclusion

In this chapter, I have revisited the parameter time delay in RAVE. Specifically, I first discuss the role of predictive time delay in RAVE demonstrating why its specific value is not relevant as long as a small non-zero value is taken. Secondly, I introduced a correction for the objective function in RAVE that corrects the effect of biasing potential on the dynamical propagator of the system. Our work is grounded in the master equation framework for the dynamics of the system in the true high-dimensional space. I prove that the RC learned from RAVE should not depend on the choice of time delay, as long as time delay is small enough that the master equation formalism of Eq.3.3 is valid. This explains why in 2.4, the RC converged quickly as a function of predictive time delay. Also, by introducing the correction for the transition probability in biased MD, I derive a new objective function, which not only reweights the static distribution as in Sec. 2.4, but also gives a better estimation of the transition probabilities. I find that apart from reducing

the number of outlier solutions, the correction does not significantly improve upon the intuitive approximation introduced in Sec.2.4. It remains to be seen if this holds true in more complex systems, but given that the correction Eq. 3.13 involves no computational overhead relative to Eq.2.6, I will use it as a default in the following chapters.

## Chapter 4: Interrogating RNA-small molecule interactions with RAVE

### 4.1 Introduction

RNA regulates diverse cellular processes, far beyond coding for protein sequence or facilitating protein biogenesis[3]. In this role, RNA can interact with protein[108] or small molecule factors [109, 110] to directly or indirectly impact gene expression and cellular homeostasis[4]. For example, riboswitches are sequences found primarily in bacterial mRNAs that contain aptamers for small molecules and regulate transcription or translation[111]. In humans, mutations in noncoding RNAs can directly impact neurodegenerative diseases and cancer[112, 113]. The regulatory role of bacterial riboswitches, coupled with the broader link between RNA and human disease has also led to interest in RNA as a therapeutic target for small molecules[114, 115, 116]. However, in comparison to proteins, methods to computationally and experimentally interrogate RNA structure are relatively less mature[117, 118]. Thus, more robust tools are needed to better understand RNA structure, dynamics, and recognition of small molecules. In spite of its clear relevance to fundamental science and drug discovery efforts, the highly dynamic nature of RNA makes it difficult to study due to the diverse conformational ensembles it adopts.[119, 120, 121] Luckily, riboswitches represent a valuable model for RNA-small molecule recognition. When ligands bind to riboswitch aptamers, the RNA undergoes



a conformational change that modulates gene expression[122, 123]. The cognate ligand for a riboswitch is typically associated with its upstream gene, often part of the ligand’s biosynthetic pathway.[123, 124] This feedback mechanism, paired with well-established structural analyses, make riboswitches excellent models for RNA-ligand binding since the target ligand is known and often highly specific[123].

Although we can use riboswitches to model RNA-small molecule interactions, visualizing riboswitch-ligand dissociation pathways with high spatiotemporal resolution remains a challenge for *in vitro* and *in silico* experiments. Such a femtosecond and all-atom resolution is hard to directly achieve in *in vitro* experiments, while the associated timescales are several orders of magnitude too slow for molecular dynamics (MD) simulations performed even on the most powerful supercomputers. Previously, Ref. [125] used  $G\bar{o}$ -model simulations of PreQ<sub>1</sub> riboswitches, specifically transcriptional *Bacillus subtilis* (*Bs*) and translational *Thermoanaerobacter tengcongensis* (*Tt*) aptamers. Their work suggests that in respective cases the ligand binds at late and early stages of riboswitch folding, indicating that perhaps these riboswitches respectively fold *via* mechanisms of conformational selection and induced fit. Since Ref. [125], RNA force-fields have become more detailed and accurate, but the timescales involved in RNA-small molecular dissociation continue to stay beyond reach[126].

Encouraged by the success of applying RAVE on the two benchmark systems, I explore the possibility of using RAVE to better understand the ligand-riboswitch dissociation process. Particularly, in my collaboration with experimentalists, we combine (i) enhanced sampling methods combining statistical physics with artificial intelligence (AI)[47, 77, 101, 127] and (ii) experimental techniques to measure RNA flexibility




Figure 4.1: Conformational change of the PreQ<sub>1</sub> riboswitch and chemical structures of PreQ<sub>1</sub> and synthetic dibenzofuran ligands.

and ligand binding[128, 129, 130]. The computational sampling techniques allow us to study the dissociation process in an accelerated but controlled manner, while the Seletive 2' Hydroxyl Acylation analyzed by Primer Extension and Mutational Profiling (SHAPE-MaP), microscale thermophoresis (MST) and fluorescence intensity assay (FIA) provide a rigorous validation of the computational findings. Specifically, I study the *Tt*-PreQ<sub>1</sub> riboswitch aptamer interacting with its cognate ligand PreQ<sub>1</sub> (7-aminomethyl-7-deazaguanine) and a synthetic ligand (2-[(dibenzo[b,d]furan-2-yl)oxy]-N,N-dimethylethan-1-amine)[131]. See Fig. 4.1 for an illustration of the riboswitch and chemical structures of the two ligands.

## 4.2 Ligand binding induces flexibility change in riboswitch aptamer

The first test is to compare the nucleotide-specific flexibility from all-atom simulations to those from SHAPE-MaP measurements. I perform 2  $\mu$ s unbiased MD simulations (4 independent sets of 500 ns each for both ligands) each for PreQ<sub>1</sub> riboswitch bound with the cognate and synthetic ligand respectively. See Appendix for details of simulation set-up. While these simulations are not long enough to

observe ligand dissociation, they can still give a robust measure of how the flexibility varies between the two different ligand-bound systems and provide nucleotide-level insight into how specific ligand binding induces different conformation changes in the riboswitch. Specifically, I compare the fluctuations of the distances between consecutive C2 atoms obtained from these simulations with SHAPE-MaP measurements. Our experimental collaborators performed SHAPE-MaP analyses using a PreQ<sub>1</sub> riboswitch construct consisting of the entire riboswitch (both the aptamer domain and expression platform). Probing was performed using a recently reported SHAPE reagent, 2A3, that provides improved mutagenesis rates during reverse transcription[132]. I analyze data in the absence of any ligand, and in the presence of either cognate ligand (PreQ<sub>1</sub>) or synthetic ligand. The results is in good agreement with previous studies reporting SHAPE-MaP on PreQ<sub>1</sub> riboswitches using different SHAPE reagents and a slightly shorter construct[133]. As discussed previously[134], base pairing interactions are the most important determinant of RNA structural constraints, and thus base-base distance fluctuations are a crucial metric of RNA backbone flexibility. Indeed, as shown in Fig. 4.2, we obtain good agreement between C2-C2 distances from unbiased MD and experimentally measured SHAPE flexibilities for all 3 systems: ligand-free PreQ<sub>1</sub> aptamer, cognate-ligand-bound PreQ<sub>1</sub> aptamer and synthetic-ligand-bound aptamer. I would like to note that though unbiased MD simulations capture most of the fluctuations measured by SHAPE, some of the peaks showed by SHAPE are not observed in MD, which might be due to the time scale limitations of unbiased simulations or the kinetic regime measured by SHAPE reactivity. In addition to the agreement between MD and SHAPE, Fig. 4.2 also shows that, for all three systems, nucleotides in the ranges A13-

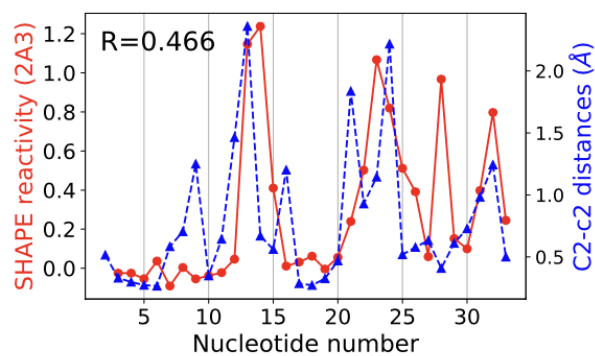
C15, U21-A24, and A32-G33 are more flexible than other regions of the riboswitch[133]. The agreement between MD and SHAPE approaches provides solid validation that our simulations using classical force-fields[135] accurately model the behavior of complex RNAs.

### 4.3 Ligand binding affinity with riboswitch aptamer


Next I aim to calculate the free energy profiles and reaction coordinates for ligand dissociation from PreQ<sub>1</sub>. As these timescales are far beyond unbiased MD, here I use RAVE to learn the dominant slow degrees of freedom and simulate the whole dissociation event.

I use AMINO[101] as a pre-processing method to reduce the number of order parameters fed to RAVE. AMINO achieves the goal of reducing the redundancy of the OP set by using the idea of clustering. AMINO uses a mutual information-base distance metric to quantify the distances between OPs, i.e., those OPs that share similar information will be closer to each other. With this distance metric, k-medoids clustering is used to group each OP into different clusters and the OP from the center of each group is selected as output to represent other OPs in the same group. The number of groups is determined by measuring the rate-distortion functions.


To describe the relative position of ligands to the riboswitch, I choose the pairwise distance between each heavy atom from ligand and the center of mass of each ribose of individual nucleotide. Four 500 ns unbiased trajectory is used as the input to learn the reduced set of OPs. The reduced set of OPs for riboswitch bound with cognate ligand and



(a)



(b)



(c)

Figure 4.2: SHAPE reactivities for the *Tt* PreQ<sub>1</sub> riboswitch aptamer obtained using the 2A3 reagent (red circles joined by solid lines) are compared with the fluctuations of C2 – C2 distances (blue triangles joined by dashed lines) for the PreQ<sub>1</sub> riboswitch (a) in the absence of ligand, (b) bound to cognate ligand, and (c) bound to synthetic ligand. Pearson correlation coefficients  $R$  are shown in upper left corners.

riboswitch bound with synthetic ligand are indicated in Fig. 4.3

After four rounds of such iterations, I obtain converged RCs for both systems,






Figure 4.4: Free energy profiles of riboswitch with different ligands. Coordination number 0 (see Eq. 4.1) indicates ligand fully dissociated from riboswitch. The error bars are shown by the highlighted colors. The approximate range of coordination number corresponding to the unbound pose is highlighted in light green color ( $C \lesssim 5.6$ ).

optimized RC and ran 10 independent biased MD simulations each using well-tempered metadynamics[47] to sample the free energy profiles of the systems. The probability distribution was estimated by using the reweighting technique introduced in Ref.[136]. In Fig. 4.4 I show the free energy profiles as functions of the ligand-riboswitch coordination number for both systems. Note that for both systems there were 2 predominant pathways as indicated through arrows in Fig. 4.5. A key difference between the two paths is that in the first path, the ligand unbinds through the space in between stem 2 and the backbone connecting stem 1 and loop 1 (Fig. 4.5, blue arrows). In the second path, the ligand unbinds through the space in between loop 2 and stem 2 (Fig. 4.5, red arrows). In order to compare the binding free energies of the two systems, I combine the information

from both the pathways by plotting the free energy as a function of the ligand-riboswitch shared coordination number  $C$ . During the dissociation process, the coordination number gradually decreases to 0 irrespective of which path is adopted during dissociation. We can define the coordination number  $C$  as:

$$C = \frac{1}{N} \sum_{i=1}^N \sum_{j=1}^M \frac{1}{1 + \left(\frac{r_{ij}}{r_0}\right)^4} \quad (4.1)$$

where  $N$  and  $M$  are the number of ligand and riboswitch atoms respectively.  $r_{ij}$  is the distance between ligand atom  $i$  and riboswitch atom  $j$  and  $r_0$  is set to be 3 Å. It can be seen from Fig. 4.4 that when bound to the riboswitch, the cognate ligand has higher  $C$  relative to when the synthetic ligand is bound to it, illustrating that the synthetic ligand makes overall fewer contacts per ligand atom with the riboswitch, as also observed in crystal structures (PDB 6E1W and 6E1U). Though I was not able to simulate the rebinding of ligands to get a converged free energy profile, the probability distributions used to construct these free energies are obtained by averaging over multiple independent dissociation events. This leads to the relatively small error bars shown in Fig. 4.4, which indicate that we can use the free energy profile to give a rough estimation of the strength of binding affinity. By considering  $C \lesssim 5.6$  as the unbound states, we can calculate that the binding affinity difference between bound and unbound state is 2~4 kcal/mol higher for the cognate ligand than that for the synthetic ligand. This is in quantitative agreement with relative affinity measurements ( $3.0 \pm 1.4$  kcal/mol) through fluorescence titrations as reported in Ref.[131].






Figure 4.5: Dissociation pathways for (a)-(b) cognate ligand and (c)-(d) synthetic ligand. Stem 1 (S1), stem 2 (S2), loop 1 (L1), loop2 (L2), and loop (L3) are colored green, cyan, magenta, gray, and orange respectively. During the dissociation processes, the ligands at different time frames are colored from white to dark blue and white to red, for the top pathway and bottom pathway respectively. The dissociation pathways are indicated with arrows. The predicted mutation sites U22 and A32 are highlighted in red and blue respectively.

#### 4.4 Predicting critical nucleotides for mutagenesis experiments

Mutations in noncoding RNA can have disastrous consequences on the human body, leading to diseases such as cancer, dementia, and Alzheimer's[137]. An explosion of

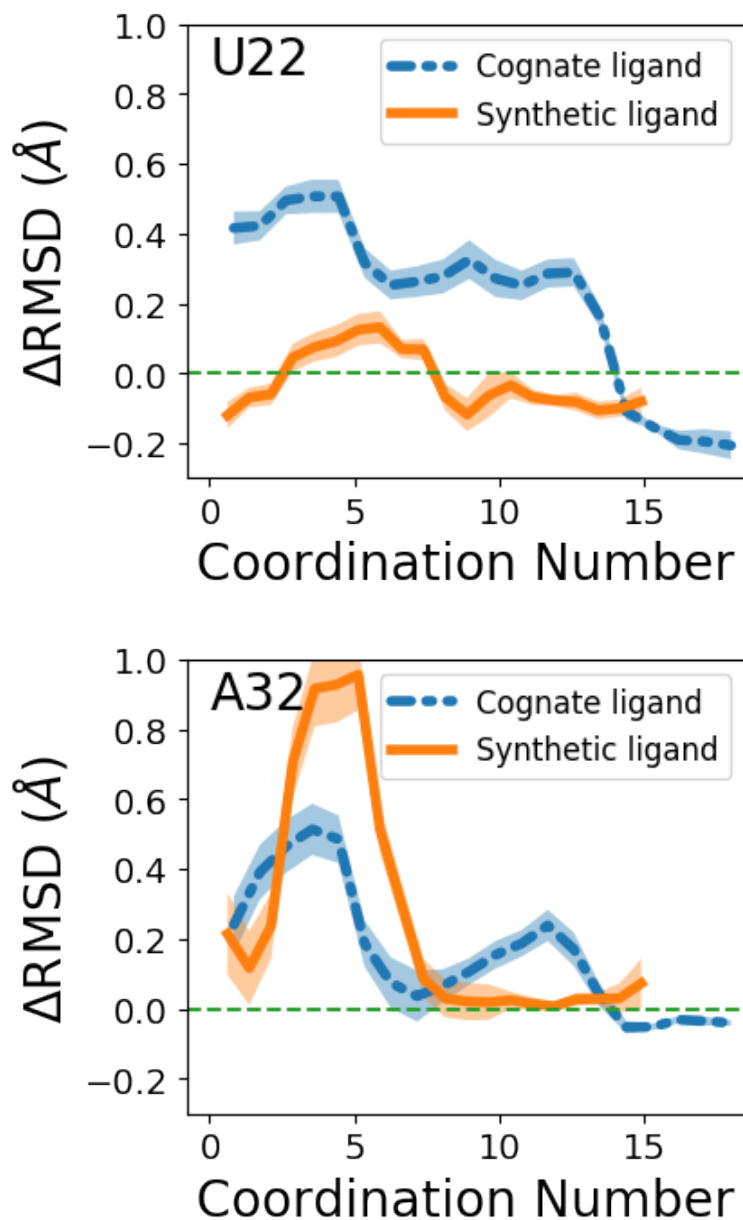


Figure 4.6: Critical nucleotides prediction. Change in RMSD for nucleotides U22 and A32 during the dissociation process, relative to RMSD variation in the apo state. RMSD profiles as function of coordination number are shown by blue dashed line and orange solid line for systems with cognate ligand and synthetic ligand separately. The change in RMSDs is calculated by subtracting the mean RMSDs of the corresponding nucleotide in apo system from the RMSDs measured in the biased simulations.

research exploring the functionality, sequencing, and regulation of noncoding RNA has been published in recent years[138, 139]; however, uncovering the specific nucleotides

involved in target-ligand binding remains a mystery. In addition to its fundamental biochemical relevance, the discovery of nucleotide mutations can also shift the focus of drug design from small molecule therapeutics to highly specific, target-based therapies with lower side effect profiles and higher efficacy.[140] We are interested in not just nucleotides we could have predicted from looking at the static structures of ligand-bound complexes, but also in nucleotides distal from the binding site. Those nucleotides might play an essential role in determining the structural ensemble of a riboswitch. Thus, changing the properties of those nucleotides can alter the interaction between ligand and other nucleotides. Such critical nucleotides are hard to predict from a purely structural perspective, but our all-atom resolution dissociation trajectories are ideally poised to make such predictions. I analyze 10 independent dissociation trajectories each obtained for the synthetic ligand bound and cognate ligand bound systems to predict critical nucleotides in the riboswitch that play a role in the dissociation process. Predictions made in this section are then validated through mutation experiments reported in Sec. 4.5.

From the collection of dissociation trajectories, we can observe that dissociation mainly happens through two pathways (Fig. 4.5) and that for the cognate ligand, 7 out of 10 simulations went through the top/blue pathway, while for the synthetic ligand, 6 out of 10 simulations went through the bottom/red pathway. To more rigorously quantify the role played by different nucleotides, I monitor the change in the RMSD of every nucleotide during the dissociation process, relative to how the respective nucleotide is in the bound complex. My motivation here is that the nucleotides showing greatest relative movement during the dissociation process are the ones most likely to be impacted once mutated. As shown in Fig.4.7, every nucleotide's RMSD changes with the coordination number  $C$

which captures the extent of dissociation. These changes indicate different conformations the riboswitch adopts during ligand dissociation. Through this procedure I identify U22 and A32, highlighted in Fig. 4.5, to be two such critical nucleotides for the cognate and synthetic ligands respectively. Fig. 4.6 shows the changes in the RMSDs of these two selected nucleotides during the dissociation process. For U22, in RAVE simulations the RMSD increases much more for the cognate ligand-bound system, than for the synthetic ligand-bound system. However, the reverse is found to be true for the nucleotide A32, which displays more enhanced movement during the dissociation of the synthetic ligand bound complex relative to the cognate ligand bound complex.

On the basis of the above observations, we can predict that for the cognate ligand, mutating U22 should display more tangible effects on the strength of the complex than mutating A32. For the synthetic ligand, we predict an opposite trend – mutating A32 should have a more pronounced effect on the interaction relative to mutating U22.

## 4.5 Validating predicted critical nucleotides through mutagenesis experiments

In Sec. 4.4 I predicted on the behalf of the dissociation trajectories that mutating nucleotides U22 and A32 will have differing effects on the cognate ligand-bound and synthetic ligand-bound systems. Also, we would like to note that, as we can see from Fig. 4.6, the RMSD changes appear in the range of coordination number that corresponds to the bound state. These changes imply the role of these nucleotides in determining the bound state interactions and suggest that the effects of mutations can be reflected by

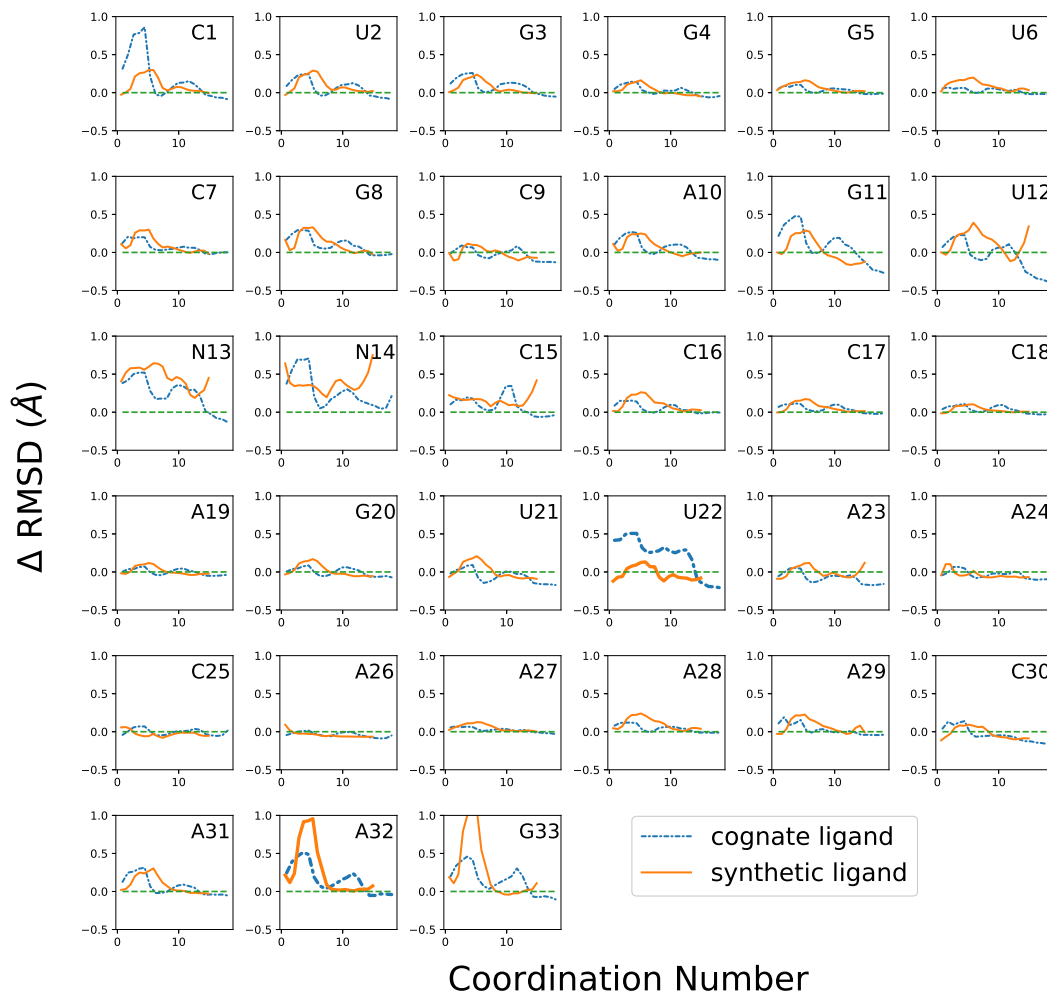


Figure 4.7: The change of RMSD of each nucleotide during the dissociation process. RMSD as a functions of coordination number are shown by blue dashed line and orange solid line for systems with cognate ligand and synthetic ligand separately.

changes in binding affinity. Here I report *in vitro* tests of our predictions. As shown in Fig.4.8, our experimental collaborators measure the equilibrium dissociation constant  $K_D$  for six different complexes, namely the two wild-type complexes, U22A-cognate ligand, U22A-synthetic ligand, A32U-cognate ligand, and A32U-synthetic ligand.

To experimentally test whether the A32U and U22A mutations impacted  $K_D$  values, we used microscale thermophoresis (MST).  $K_D$  values were measured using fluorescently labeled wild type, A32U, and U22A PreQ<sub>1</sub> riboswitch aptamer constructs

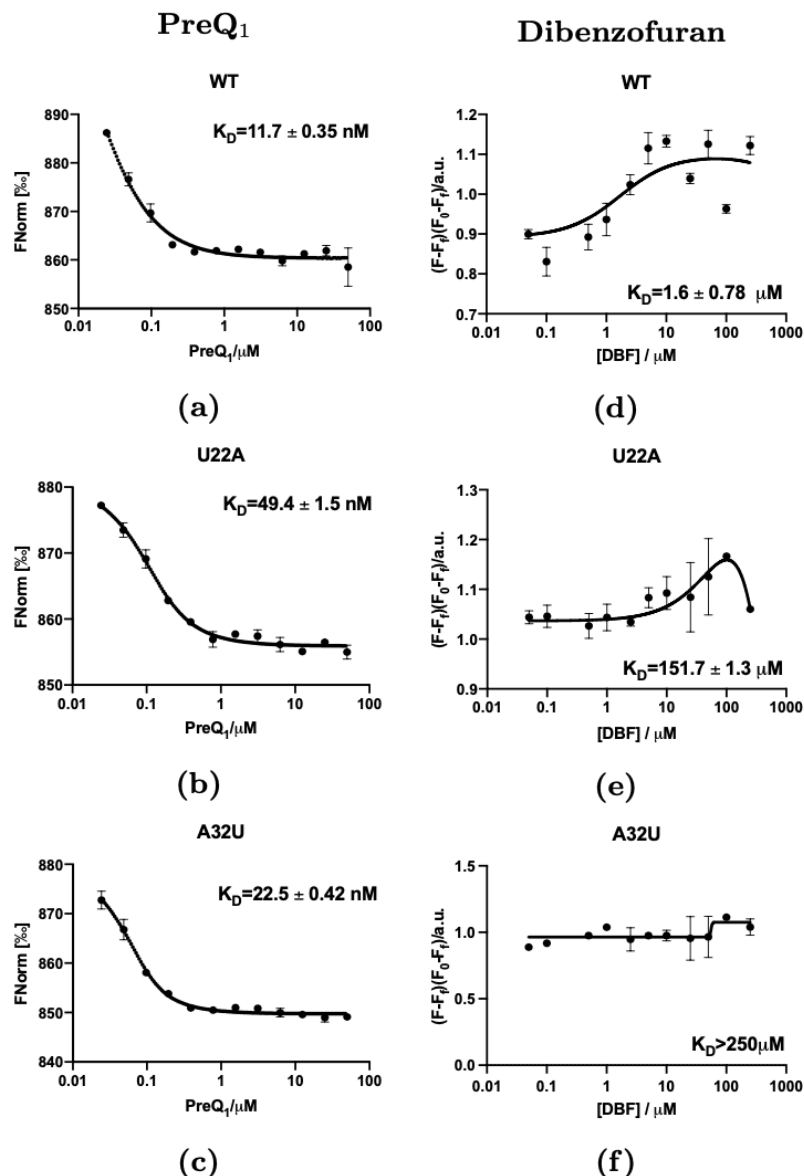


Figure 4.8: Affinity measurement of PreQ<sub>1</sub> (cognate) and dibenzofuran (synthetic) ligands binding to WT *Tte* riboswitch and U22A and A32U mutants. MST of 5'-cy5-labelled (a) WT *Tte*, (b) U22A, and (c) A32U mutants in the presence of PreQ<sub>1</sub> ligand. FIA of 5'-cy5-labelled (d) WT *Tte*, (e) U22A, and (f) A32U mutants in the presence of dibenzofuran. Error bars represent the standard deviation of three replicate experiments.

for the cognate ligand. For PreQ<sub>1</sub>, we observed a  $K_D$  of  $11.7 \pm 0.35$  nM for the WT construct. In contrast, the A32U mutant had a  $K_D$  of  $22.5 \pm 0.42$  nM while the U22A mutant had a  $K_D$  of  $49.4 \pm 1.5$  nM. For the synthetic ligand, changes were observed in

fluorescence of the labeled RNA upon incubation with the ligand, and thus a fluorescence intensity assay (FIA) was used to measure  $K_D$  values rather than MST. Using FIA we observed a  $K_D$  of  $1.6 \pm 0.7 \mu\text{M}$  for the wild type aptamer. In contrast, both the U22A and A32U mutant constructs displayed significantly weaker  $K_D$  values ( $150 \pm 1.3 \mu\text{M}$ , and  $\geq 100 \mu\text{M}$  respectively). These  $K_D$  measurements directly validate the findings from Sec. 4.4. There we had predicted that mutating U22 will effect the cognate ligand-bound system more than mutating A32, and indeed we find that  $K_D$  for U22A is 5 times weaker than WT, while only 2 times weaker for A32U. At the same time, we had predicted that mutating A32 will effect the synthetic ligand-bound system more than mutating U22, and indeed we find that while the  $K_D$  for U22A is 75 times weaker than WT, there is effectively no tangible association at all for A32U with the synthetic ligand.

The *in silico* experiments also support that the mutations alter the bound state ensemble and give more insights into the interaction change in the bound state after the mutations. I performed MD simulations with the mutated systems to validate that the mutations changes the bound state ensemble of both systems. I first used PyMOL[141] to generate the mutants U22A and A32U for riboswitch with cognate or synthetic ligands. I then performed four independent unbiased simulations for each system, with each 500ns long. The structures were saved every 10 ps. In the 500 ns simulations, the ligands stay in the bound pose. The statics of samples from these simulations can be used to illustrate the influence of mutations on the bound state ensemble. Fig.4.9 shows the free energy profiles of these systems and clearly indicates the changes in the bound state ensemble brought by the mutations U22A and A32U. I also study the hydrogen bond occupancy between ligand and nucleotides from riboswitch in  $2\mu\text{s}$  long unbiased MD simulation. This captures the

bound state ensemble better than just the crystal structure and shows how the mutations, even though distal, disrupt the hydrogen binding patterns. The number of hydrogen bonds is calculated by VMD[142] with an angle cutoff of 30 degrees and a distance cutoff 3.5Å. The hydrogen bond occupancy is defined as  $\bar{\rho} = \sum_{i=1}^n N_i/n$  where  $n$  is the number of total frame and  $N_i$  is the number of hydrogen bonds in frame  $i$ . Fig.4.10 shows that different hydrogen bonds between ligand-nucleotide pairs are formed in the bound state ensemble before and after the mutations. For example, the hydrogen bonds between cognate ligand and C16 disappeared after the mutation and the hydrogen bonds between cognate ligand and A28 mainly showed up in the A32U mutant. Also, for synthetic ligand, the hydrogen occupancy with U6 largely increases after the mutations while that with G4 decreases. So even though nucleotides 22 and 32 do not directly form hydrogen bonds with both ligands in the bound pose as seen in the crystal structure, interactions between ligands and other nucleotides are still altered when they are mutated

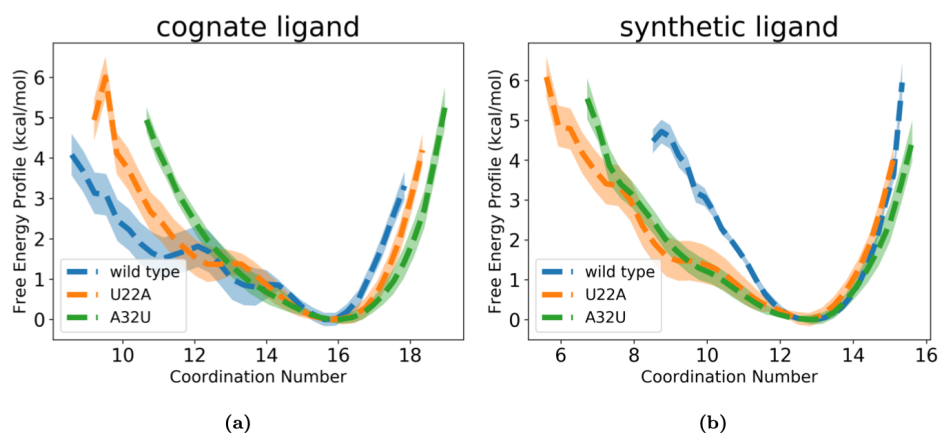


Figure 4.9: Free energy profile from 2 $\mu$ s unbiased simulation for wild type and mutated systems with (a) cognate ligand and (b) synthetic ligand. The error bars are indicated by the shaded region around each curve.



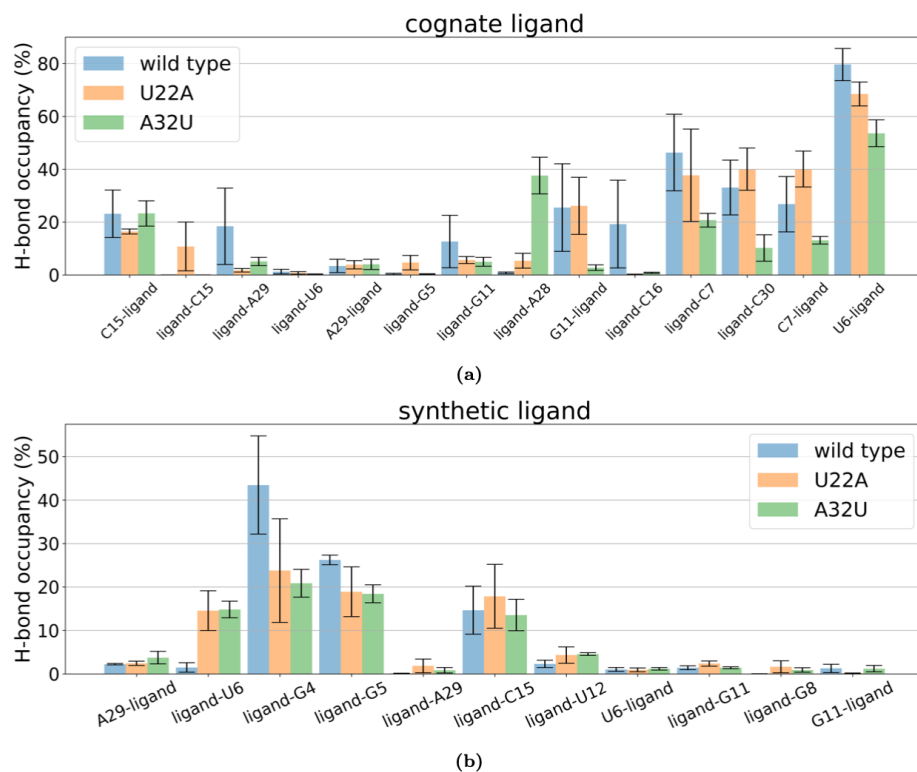


Figure 4.10: Hydrogen bond occupancy for wild type and mutated systems calculated through  $2\mu s$  long unbiased MD simulation with (a) cognate ligand and (b) synthetic ligand. The labels on the x-axis define the hydrogen bond pairs with the format “donor-acceptor”.

## 4.6 Discussion

In this chapter, I showed how one can obtain detailed and robust insights into RNA structural dynamics through AI-augmented molecular simulations that maintain all-atom resolution for RNA, ions and all water molecules. All my predictions are validated through a gamut of complementary and experimental techniques that measure RNA flexibility and ligand binding. I specifically focused on the *Tt*-PreQ<sub>1</sub> riboswitch aptamer system in its apo state and bound to two different ligands, where I first showed that  $2\mu s$  long unbiased MD simulations with classical all-atom force-fields for RNA and water

can provide the same flexibility profile as measured in SHAPE experiments. Since the timescales for ligand dissociation are far slower than MD can typically access, I used the RAVE simulation method[77] that enhances sampling along a low-dimensional reaction coordinate learned on-the-fly through the past-future information bottleneck framework. Through RAVE we obtain multiple independent dissociation trajectories for both cognate and synthetic liganded systems, demonstrating how the cognate and synthetic ligands invoke different aspects of the riboswitch's flexibility in order to dissociate. On this basis, I were able to predict pairs of mutations which would be expected to show contrasting behaviors for the cognate and synthetic ligands. The ability to make such predictions accurately and efficiently is of vital importance, as mutations that can disrupt RNA structural dynamics and ligand recognition processes have been linked to numerous human diseases[143, 144]. My predictions for the *Tt*-PreQ<sub>1</sub> aptamer were validated by designing mutant aptamers and performing affinity measurements.

This work thus demonstrates a pathway to gain predictive insights, as opposed to retrospective validation, from enhanced molecular dynamics RAVE simulations that maintain femtosecond and all-atom resolution while directly observing processes as slow as ligand dissociation. While the focus in this work was on predicting the thermodynamically averaged aspects of riboswitch flexibility, in future work it should also be possible to directly calculate rate constants using complementary approaches and for more complicated RNA systems. More broadly, this work indicates that long time scale MD simulations can accurately model ligand dissociation from RNA, and provide validatable hypotheses about structural features that impact small molecule recognition of RNA. Thus, this strategy will be widely applicable not only to study natural systems

such as riboswitches, but to more completely understand the structural features and dynamics that govern how synthetic small molecules interact with RNA.

## Chapter 5: Mixing physics across temperatures with generative AI

### 5.1 Introduction

Using simulations or experiments performed at some set of temperatures to learn about the physics or chemistry at some other arbitrary temperature is a problem of immense practical and theoretical relevance. Here I develop a framework based on statistical mechanics and generative AI that allows solving this problem.

Specifically, this framework uses generative AI that learns to efficiently sample such high-dimensional, complex probability distributions for  $N$ -body molecular systems valid across temperatures. There are two central ideas guiding this framework. Firstly, I do not treat the temperature  $T$  as just a control parameter. By noting that the temperature is a measure of the average kinetic energy, I instead work directly with the fluctuating kinetic energy, using equipartition theorem to define an instantaneous, effective temperature which we still call  $T$ . For finite  $N$  not yet in the thermodynamic limit,  $T$  so defined will display significant fluctuations proportional to  $1/\sqrt{N}$ . Secondly, if we have experiments or simulations performed at temperatures  $T_1, T_2, \dots, T_K$ , we can view them all together as being sampled from the same, but unknown, joint probability  $p(\mathbf{x}, T)$  for  $N$ -body configurations  $\mathbf{x}$ . I then use a generative AI method, specifically denoising diffusion probabilistic models (DDPM)[145, 146] to generate many more

samples from  $p(\mathbf{x}, T)$ , given the sparse, high-dimensional dataset.

By learning  $p(\mathbf{x}, T)$  I am referring to the ability to generate many more samples from  $p(\mathbf{x}, T)$  given the data set we have from simulations. DDPM has been shown to possess the ability to infer and learn the underlying relationship from complicated data i.e. high-dimensional and with complex correlations, but also noisy and sparse[145, 146]. DDPMs learn two diffusion processes expressed through stochastic deep neural networks, which are called noising and denoising. The forward diffusion or noising converts the samples from high-dimensional, structured and unknown probability distribution into simpler and analytically tractable white noise. The backward diffusion or denoising learns the mapping back from noise to meaningful data. The central idea is that it is easy to generate numerous samples from the noisy distribution, which can then be denoised back to structured data. The process has been demonstrated to be comparable or even outperform other generative AI models for generating high quality samples, and in its ability to model in an unsupervised way the underlying semantics behind meaningful data[145, 146].

Here I focus on REMD[38, 147] introduced in Sec.1.3.2. REMD uses information generated at a ladder of temperatures to swap configurations across temperatures. REMD has been extremely powerful over the decades for the study of molecular systems with rough energy landscapes for fundamental science and practical applications[148, 149]. Numerous advances have been introduced over this basic idea in order to make it more efficient computationally[39, 40, 150, 151, 152, 153, 154, 155] and it continues to be an area of very active research.

I will demonstrate how DDPM applied to all-atom, femtosecond resolution REMD

simulations of a small peptide and a RNA nucleotide very significantly improve the quality of data generated across temperatures, including temperatures at which the simulation was never performed and even at temperatures outside the ladder of temperatures, significantly improves the estimates of free energies made through REMD and providing accurate sampling in parts of configuration that were not previously visited in the lowest temperature replica during REMD. These could be metastable states or transition states. I also show how this can be used to generate samples at temperatures not included in the ladder of replicas. The samples the model generate have thermodynamic relevance and correspond to correct Boltzmann weights as opposed to being just wild hallucinations[146, 156].

## 5.2 Temperature as a random variable

Here let's consider the equivalent  $p(\mathbf{x}, \beta)$  where inverse temperature  $\beta = \frac{1}{k_B T}$  and  $k_B$  is Boltzmann's constant. MD or Monte Carlo (MC) methods allow sampling configurations  $\mathbf{x}$  as per the equilibrium probability  $p_\beta(\mathbf{x}) \equiv e^{-\beta U(\mathbf{x})}/Z$ , where  $U(\mathbf{x})$  is the potential energy of the  $N$ -body system and  $Z = \int d\mathbf{x} e^{-\beta U(\mathbf{x})}$  is the partition function. For systems of practical interest in biology, chemistry and materials science, if  $T$  is not large enough, it becomes nearly impossible to sample reliably from  $p_\beta(\mathbf{x})$  as many regions of interest in configuration space will have  $e^{-\beta U(\mathbf{x})} \approx 0$ . To deal with this problem, in REMD one simulates  $K+1$  replicas of the system at temperatures  $\beta = \beta_0 > \beta_1 > \dots > \beta_K$ . For low enough  $\beta_K$  or equivalently high enough temperature  $T_K$ , the sampling from  $p_{\beta_K}(\mathbf{x}) \equiv e^{-\beta_K U(\mathbf{x})}/Z_K$  is expected to be more ergodic. One

then periodically exchanges conformations between consecutive pairs of replica with a Metropolis-type acceptance probability that depends on the potential energies of the two replicas and their temperatures. This way even the low-temperature  $\beta_0$  replica can explore configurations that it would not have otherwise visited.

I would argue that by treating the temperature just as a control parameter REMD is not making full use of the information gathered across the ladder of temperatures. In most current incarnations of REMD (excluding exceptions such as Ref.[157]), all that the higher temperatures do is to help the lower temperature replicas discontinuously appear in different locations of the configuration space. Here I take a slightly different view of the temperature in REMD. I would like to point out that this prospective can also be easily generalizable to other thermodynamic variables which show fluctuations for finite system size. By working with an approximate, effective temperature, we can treat it as a random variable instead of a control parameter. To do so we can work with an effective, instantaneous temperature of the system associated with its kinetic energy instead of the temperature of the heat bath that we expect the thermostat to enforce on average. More rigorously thus, we are sampling the joint  $p(\mathbf{x}, \kappa)$  where the per-particle kinetic energy  $\kappa$  is related through its ensemble average to the temperature, i.e.  $\langle \kappa \rangle = \frac{3}{2\beta}$ . I would like to emphasize that the effective temperature is not the true thermodynamic temperature - however for the sake of simplicity I still use the symbol  $T$  or  $\beta$  to denote this effective temperature or its inverse respectively.

The motivation in doing so is that all replicas across different temperatures can be viewed as being sampled from the same joint probability  $p(\mathbf{x}, \beta)$  – as opposed to different replicas sampling from respective  $p_{\beta_K}(\mathbf{x})$ . This change of perspective allows us

to combine together the data collected from different temperatures as having arisen from the same, although intractable, probability distribution  $p(\mathbf{x}, \beta)$ .

### 5.3 Denoising diffusion probabilistic model

Our task at hand now is to learn the joint probability  $p(\mathbf{x}, \beta)$  given sampling that has already been performed in REMD across  $\mathbf{x}$ -space and temperatures  $\beta = \beta_0 > \beta_1 > \dots > \beta_K$ . There are two main challenges in this. The first challenge comes from the curse of dimensionality: the memory or computational resources needed to track a very high-dimensional distribution function increase exponentially as the number of degrees of freedom increases. For instance, for REMD of a small 9-residue peptide in explicit water, which we study here, we exchange all 4749 atomic coordinates between the replicas, but for the purpose of analysis we set  $\mathbf{x}$  as 18 Ramachandran dihedral angles. This means we already have a 19-dimensional space where binning procedures are out of the question. The second challenge comes from the sparsity of the data. Most of our samples come from high probability regions  $p(\mathbf{x}, \beta)$  and we have very few samples for low-probability states due to inefficient exchange between replica. In summary we have sparse sampling of data points in very high-dimensional  $(\mathbf{x}, \beta)$ -space and wish to construct  $p(\mathbf{x}, \beta)$  from this information so that one can create many more samples at any temperature of interest.

DDPM can generate many more samples from  $p(\mathbf{x}, \beta)$ . The main idea behind the use of DDPM here is to learn a simple and easy-to-sample-from distribution  $P_{\text{simple}}(\mathbf{x}, \beta)$  that approximates the true  $p(\mathbf{x}, \beta)$ . For notation simplicity, we denote  $\mathbf{s} = \{\mathbf{x}, \beta\}$  and refer to  $P_{\text{simple}}(\mathbf{s})$  henceforth. DDPM does this by learning to reverse a gradual, multi-



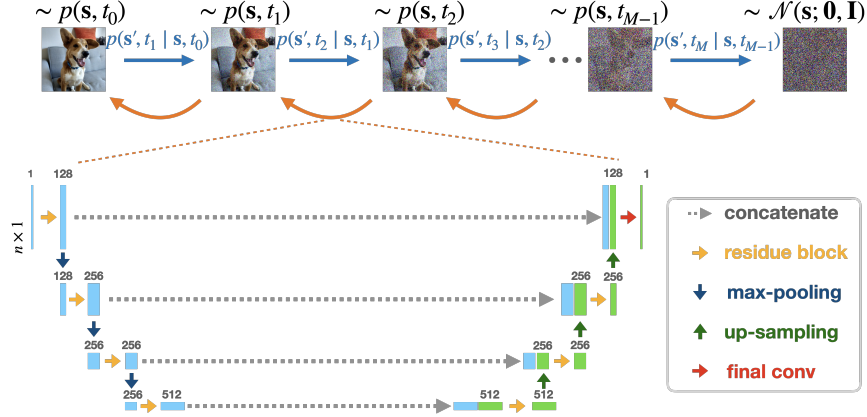


Figure 5.1: Lower panel shows the neural network architecture used in this work. It has the basic structure of the U-Net model[158]. During the diffusion process, which is indicated by the direction of blue arrows in the upper panel, noise is gradually added to the sampled data, in this case the picture of a very good boy, through a diffusion process labeled with diffusion step  $t_i$ . This changes the sampled distribution (for example, more pictures of dogs) to a simpler isotropic Gaussian distribution from which one can easily generate more samples. An AI model is then trained to reverse such a diffusion process and starting from sampled noise, learn to generate images similar to the input image by following the direction indicated by the orange arrows. It takes an 1-d array  $s'$ , which is the noisy sample at diffusion step  $t_i$ , as input and outputs the parameterization of a Gaussian distribution to get the reversed transition kernel  $q(s, t_{i-1} | s, t_i)$ . Each residue block consists of three components: two 1-d convolution operations with kernel size 3 and a group normalization[159] between them. The diffusion step  $t_i$  is added to each convolutional block after being transformed by the sinusoidal position embedding[160]. The final conv label denotes the convolution operation with kernel size equal to 1. Max-pooling reduces size of the features by half, while up-sampling uses the transposed convolution to expand the size of features.

step noising process that starts with the relatively limited number of samples generated from the distribution  $p(s)$  and diffuses to the simpler distribution  $P_{\text{simple}}(s)$  that is easy-to-sample. For instance,  $P_{\text{simple}}(s)$  could be an isotropic Gaussian. In addition to learning this noising process, DDPMs also learn the reverse denoising process which allows us to go back from samples generated using  $P_{\text{simple}}(s)$  to samples that would have been generated from the underlying  $P(s)$ . Both the noising and denoising processes are modeled using diffusion processes that convert probability distributions to one another,

and are implemented using the architecture shown in Fig. 5.1, which is based on the standard architecture for DDPMs as described in Ref.[146].




Figure 5.2: (a) AIB<sub>9</sub> in left-handed (L) and right-handed (R) conformations in explicit TIP3P water. (b) Free energy profile of residues (with residue 5 as an example) with two stable states labeled as L, R and two excited states labeled as La, Ra. When all residues are in L states, the peptide chain is in left-handed state and similar for right-handed state. Free energies in (b) here as well as all through the manuscript and Supporting Information are provided in units of  $k_B T$ .

The noising diffusion process carried out in the space  $\mathbf{s} = \{\mathbf{x}, \beta\}$  that converts  $p(\mathbf{x}, \beta)$  to the simpler  $P_{\text{simple}}(\mathbf{x}, \beta)$  can be decomposed into  $M$  discrete steps denoted by corresponding transition probabilities  $p(\mathbf{s}', t_{i+1} | \mathbf{s}, t_i)$  where  $i \in [0, M]$ ,  $P(\mathbf{s}, t_0) \equiv p(\mathbf{x}, \beta)$ , and  $P(\mathbf{s}, t_M) \equiv P_{\text{simple}}(\mathbf{x}, \beta)$ . In DDPM, this noising diffusion process that converts sampled data to essentially noise, is set to be an Ornstein–Uhlenbeck (OU) process in which the transition probability follows a simple Gaussian form. One can then easily generate samples from  $P(\mathbf{s}, t_0) \equiv p(\mathbf{x}, \beta)$ . The tricky bit now is to convert these samples back to the original distribution. In Ref.[145], it was shown that this transition reversed diffusion kernel  $P(\mathbf{s}, t_i | \mathbf{s}', t_{i+1})$  can also be written in a Gaussian form. A deep neural network (Fig.5.1) is then trained using variational inference to learn the approximate reversed transition kernel  $Q(\mathbf{s}, t_i | \mathbf{s}', t_{i+1}) \approx P(\mathbf{s}, t_i | \mathbf{s}', t_{i+1})$ . Thus by generating samples from a normal Gaussian distribution, which we can easily do in

large numbers, and then passing these through the reversed transition kernel, we can generate samples that follow the target distribution  $p(\mathbf{x}, \beta)$  as desired. Note that instead of learning the joint probabilities  $P(\mathbf{s}_1, \mathbf{s}_2)$ , it can be advantageous to learn the conditional probability  $P(\mathbf{s}_1|\mathbf{s}_2)$ . This can be done through the protocol in Ref.[145] by adding a delta function to allow only a subset of  $\mathbf{s}$  to change during the nosing and denoising process, i.e.  $\delta(\mathbf{s}_2 - \mathbf{s}'_2)P(\mathbf{s}_1, \mathbf{s}_2, t_i|\mathbf{s}'_1, \mathbf{s}'_2, t_{i+1})$ . This is very useful when for instance we are interested in generating samples only at a certain temperature or only in certain regions of the configuration space. In the most general form of diffusion probabilistic models (DPM)[145], the reversed transition kernel  $Q(\mathbf{s}, t_i|\mathbf{s}', t_{i+1})$  is considered as  $Q(\mathbf{s}, t - 1 | \mathbf{s}', t) = \mathcal{N}(\mathbf{s}; \tilde{\boldsymbol{\mu}}(\mathbf{s}', t), \tilde{\boldsymbol{\sigma}}_t(\mathbf{s}', t))$  and the neural network is trained to learn the mean  $\tilde{\boldsymbol{\mu}}(\mathbf{s}', t)$  and variance  $\tilde{\boldsymbol{\sigma}}(\mathbf{s}', t)$ . In practice, however, there are many different ways to choose the Gaussian distribution parameterization. In Ref.[146], DDPM was introduced with a new parameterization approach to reduce the complexity of the training task. DDPM got its name because such a design makes the learning task resemble a denoising score matching procedure[161]. In Ref.[146], it was shown that with such a design, DDPM can generate samples of a quality that are comparable or even better than other generative models.

## 5.4 Results

I going to demonstrate how the above protocol can be applied to mix data collected from different temperatures and configurations in REMD and significantly improve the quality of sampling. Specifically, I consider the following two challenging tasks:

1. Can we improve the sampling quality for the lowest-temperature replica with more

accurate probability estimates than directly seen after REMD? This includes being able to generate samples in low probability regions such as transition and metastable states, and reliably estimating their free energies.

2. Can we generate samples at temperatures that are not even included in the replica ladder? This would include temperatures within the range of the replica ladder and also extrapolation to temperatures outside the range.

#### 5.4.0.1 Peptide conformational transitions

To demonstrate the performance of DDPM, we first study a small peptide chain Aib<sub>9</sub> in explicit TIP3P water[162] using CHARMM36m force-field[163]. This 9 residue system (Fig.5.2) (a) displays rich and complex conformational dynamics[164] including the transitions between fully left-handed and right-handed helices. However, even in 4  $\mu$ s unbiased MD at 400K, one can see only around 2–3 transitions between these two dominant equiprobable conformations, and even fewer, if any, transitions to the higher energy metastable states. To improve the sampling of this system, I perform REMD with 10 replicas at geometrically spaced temperatures ranging from 400K to 518K, with temperature increased by 3% for each replica. The attempt of exchanging configurations was made every 20 ps, with acceptance rate around 1%  $\sim$  2% between neighboring replicas, which is intentionally kept lower than what one usually has in REMD. This is because I want to show that even in the extreme cases where the number of atoms  $N$  is so large that replicas do not have enough overlap, or if one wants to reduce the number of replicas to save computational resources, our DDPM can still do a decent

job of complementing REMD and reconstructing the true probability distribution at any temperature of interest. To benchmark our results, I ran unbiased MD at 400 K for  $4 \mu s$  and at 500 K for  $0.6 \mu s$ .




Figure 5.3: Samples (dots) from REMD (left) and DDPM (right) at 400K. The Boltzmann weights for different samples are indicated through their free energy (contour lines, separated every  $0.74 k_b T$ ). (a) Samples generated from DDPM and free energy profile projected on dihedral angles of residue 5. (b) Samples generated from DDPM and free energy profile projected on dihedral angles of residue 8. DDPM was able to generate samples in states that are not present in the training dataset for both residues, indicated with thick black arrows in the right panels.

To quantify the quality of sampling, we can focus on the 18 dihedral angles corresponding to all 9 residues  $(\Phi_1, \Psi_1, \Phi_2, \Psi_2, \dots, \Phi_9, \Psi_9)$ . As shown in the

Ramachandran plot in Fig. 5.2 (b), the free energy surface along any pair of dihedrals is mainly characterized by four metastable states: two equiprobable low-energy ground states (L, R) and two excited states (La, Ra). The Ramachandran plots for all nine residues in the system look qualitatively similar but the middle residues of the peptide are known to be less flexible with higher energy barriers.[165, 166] I thus focus on the sampling for residues 5 and 8 as shown in Fig.5.3. I train the DDPM on a REMD trajectory. At the lowest temperature of interest (400K), this trajectory has not yet achieved sufficient sampling. As shown in Fig.5.3, DDPM successfully mixes information from all temperatures and configurations, and generates samples in states that are not present in the training dataset for both residues, indicated with thick black arrows in the right panels of Fig.5.3. Specifically, for residue 5 we can see that the transition states between state R and state La, which were not being sampled in the 400K replica, are populated in samples from DDPM. For residue 8 the improvement is even more striking as the state Ra which was simply not sampled in REMD now gets populated after DDPM. To further quantify the improvement gained due to DDPM, I compare the free energy differences between different configurations from both REMD and DDPM against much longer reference unbiased MD at 400K. As shown in Fig.5.5, the free energy difference calculated from samples generated by DDPM are much closer to the that from the reference MD. Thus, DDPMs are able to accomplish the first task highlighted above. I also want to highlight that while some spurious samples are generated, seen through dots outside the free energy contours in Fig. 5.3, their Boltzmann weights are very low. Thus, our model “dreams” new configurations without hallucinating spurious configurations.

I am now moving to the second task from the list above, and test DDPM’s




Figure 5.4: Comparing  $\Delta G$  calculated from samples of REMD, DDPM and long unbiased MD. See the text for precise definitions of various states. The free energy difference between the ground state L (orange box) and excited state Ra (green box). Green empty crosses are reference values from 4  $\mu s$  long unbiased MD at 400 K. Green filled plus signs show values after REMD and DDPM, while orange circles show values after REMD without DDPM.




Figure 5.5: Comparing  $\Delta G$  calculated from samples of REMD, DDPM and long unbiased MD. See main text for precise definitions of various states. The free energy difference between the ground state L (orange box) and the transition state (purple box) at 2 different temperatures, 400K and 500K

ability to generate samples by interpolating or even extrapolating across temperatures not considered in the ladder of replicas. In the first example, I use it to generate samples at 500 K, as in the training set with 10 replicas at geometrically spaced temperatures between 400 and 518K, there is no replica with temperature 500 K. As shown in Fig. 5.5, the  $\Delta G$  calculated from samples of DDPM is in good agreement with that from the reference MD at 500 K. In the second example, I completely removed the samples from 400 K replica in the training set and use it train a new model, which is then used to generate samples at 400 K. Even though in the training set, the lowest temperature is 412K, the model can

make good prediction of free energy difference between states as shown in Fig. 5.6. This finding is particularly encouraging as in general, sampling at lower temperatures tends to be harder than at higher temperatures.




Figure 5.6: Free energy differences between different metastable states  $\Delta G$  calculated from samples from DDPM by extrapolating to a temperature at 400 K, which is lower than any temperature included in the dataset. The DDPM model was trained with REMD trajectories with 9 replicas starting at the lowest temperature 412K. Green empty crosses and purple empty triangles are reference free energy differences  $\Delta G$  from 4  $\mu s$  long unbiased MD at 400 K. Green filled plus signs and purple diamonds show the energy differences from DDPM.  $\Delta G_{L-Ra}$  is the free energy difference between ground state L ( $-2.2 < \Phi < -0.1$ ,  $-1.6 < \Psi < 0.9$ ) and excited state Ra ( $-1.6 < \Phi < -0.2$ ,  $1.4 < \Psi < 3.14$ ); and  $\Delta G_{L-TS}$  is the free energy difference between the ground state L ( $-2.2 < \Phi < -0.1$ ,  $-1.6 < \Psi < 0.9$ ) and a transition state ( $-1.8 < \Phi < -0.8$ ,  $0.8 < \Psi < 1.5$ )

### RNA conformational transitions.

As a second example to illustrate the general applicability of our DDPM+REMD approach, I test this framework on the sampling of RNA conformational ensemble. Rare RNA structures have been previously shown to be biologically relevant [167, 168, 169], but estimating the conformational ensemble still remains computationally intractable using traditional sampling techniques [170]. RNA dynamics may occupy a wide range of timescales - from several hours for conformational changes that require breaking base pairs, to picoseconds for more continuous deformations [169]. As a consequence, identifying rare transient structures and estimating their contribution to the RNA ensemble






Figure 5.7: Projection of samples and free energy profile on dihedral angles  $\zeta$  and  $\gamma$  of A2 in GACC at 325 K. The Boltzmann weights for different samples are indicated through their free energy (contour lines, separated every  $0.74 k_bT$ ). (a) The structure of GACC. (b) Samples from REMD. (c) Samples from the benchmark  $5\mu s$  long unbiased MD. (d) Samples from DDPM. The metastable states that are not present in the training dataset are indicated with thick black arrows.

has proved to be difficult. In this example, I consider a GACC tetranucleotide as our model system (Fig. 5.7(a)). As a single-stranded RNA consisting of four nucleotides labeled G1, A2, C3 and C4, GACC has been previously used as a challenging test system for REMD based sampling methods for its conformational flexibility [171] Despite the fact that GACC has been widely studied, it still remains challenging to effectively sample possible configurations and is a good system to test new methods[170]. For example, in

a previous study, in order to get the converged structural populations, a multidimensional replica exchange molecular dynamics (M-REMD) simulation was performed with 192 replicas with around 1  $\mu s$  of simulation time per replica, thus totalling almost 192 microseconds of all-atom simulations.[171] Here we show that with DDPM, we can better estimate the free energy landscape using fewer computational resources, totalling only 12 microseconds of all-atom simulations adding all replicas. I trained our DDPM

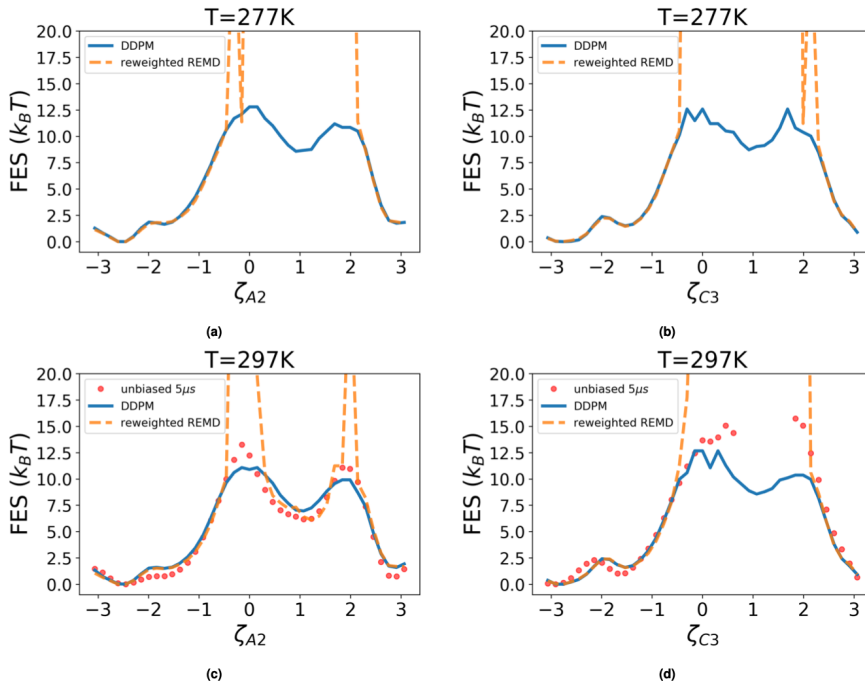


Figure 5.8: DDPM outperforms direct reweighting (Eq. 5.1). Free energy surface (FES) in  $k_B T$  along dihedral angles  $\zeta$  of (a) A2 and (b) C3 from GACC calculated from REMD trajectories after direct reweighting (orange dash line) and DDPM samples (blue solid line) at 277 K. Similarly, (c) and (d) show respective plots at 297K for A2 and C3 respectively. (c) and (d) also shows benchmark results from 5  $\mu s$  unbiased MD (red dots) at 297 K. Unbiased MD at 277K was far too slow to provide any information.

model on 250  $ns$  REMD trajectories from 48 replicas with temperatures ranging from 277 K to 408 K (see Appendix A.1.4 for MD simulations setup for details). In each frame of the trajectory, the structure of GACC is characterized by 6 dihedral angles for

dihedral	definition
$\alpha$	P (H5T) - O5' - C5' - C4'
$\beta$	C4' - C3' - O3' - P (H3T)
$\gamma$	O4' - C1' - C2' - C3'
$\delta$	C1' - C2' - O2' - HO'2
$\epsilon$	C5' - C4' - C3' - O3'
$\zeta$	O5' - C5' - O4' - C3'

Table 5.1: Definition of dihedral angles

each nucleotide (defined in Table 5.1). In these REMD simulations, the sampling is not sufficient, especially for replicas at very low temperatures. Fig. 5.9 and Fig. 5.10 shows the free energy profiles of GACC projected on the  $\gamma$  and  $\zeta$  angles of A2 and C3 at different temperatures. We can see that expected high energy metastable state indicated by black arrows in Fig. 5.9 and Fig. 5.10 were not sampled in REMD. In contrast, as shown in Fig. 5.7 and Fig. 5.11 for different target temperatures, DDPM successfully generates the ensemble of such high free energy states that were never even visited at low temperatures. Here as well, any spurious configurations can be seen through dots outside the free energy contours in Fig. 5.7 with negligible Boltzmann weights.

It is also natural to ask if information at different temperatures could be combined through a Boltzmann-type reweighting where samples from any temperature  $\beta$  and potential energy  $U(\mathbf{x})$  can be treated using weight to estimate the free energy at temperature  $\beta_0$ :

$$w_{\beta_0} = e^{(\beta - \beta_0)U(\mathbf{x})} \quad (5.1)$$

The reference point of potential energy  $U(\mathbf{x})$  is determined such that the minimum of potential energy in each replica is equal to 0. This way we can estimate the free energy by simply mixing samples from different temperatures without the use of any generative

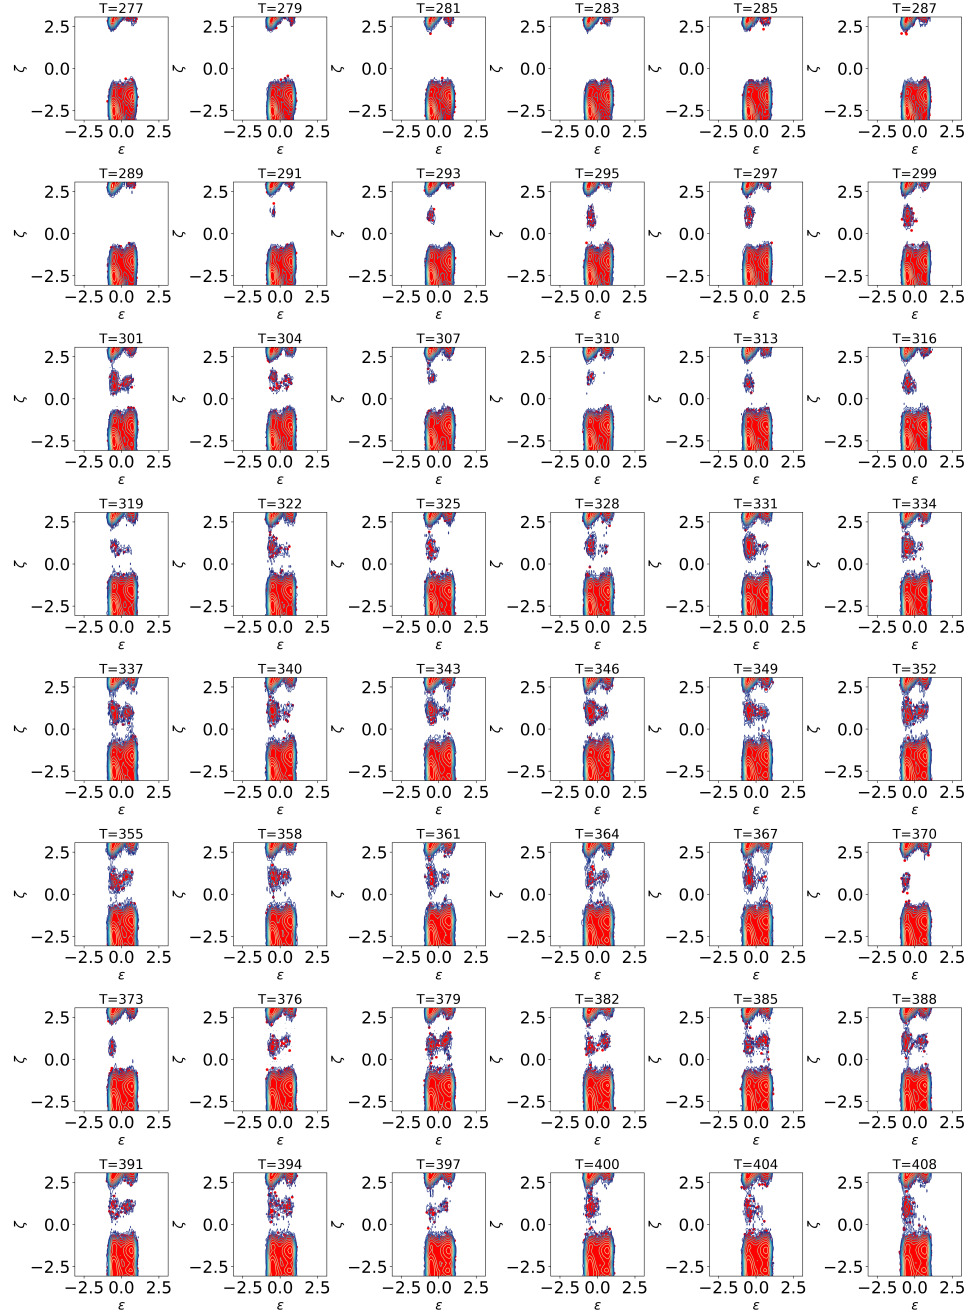


Figure 5.9: Projection of REMD samples and free energy profile on dihedral angles  $\epsilon$  and  $\zeta$  of the A2 in GACC at each replica.

AI. Fig 5.8 shows the results for such an operation and compares them with results from DDPM and long unbiased MD when possible. For  $\zeta_{A2}$  at 297 k, the direct reweighting method gives a good estimation of the free energy of metastable state around  $\zeta_{A2} = 1$  but

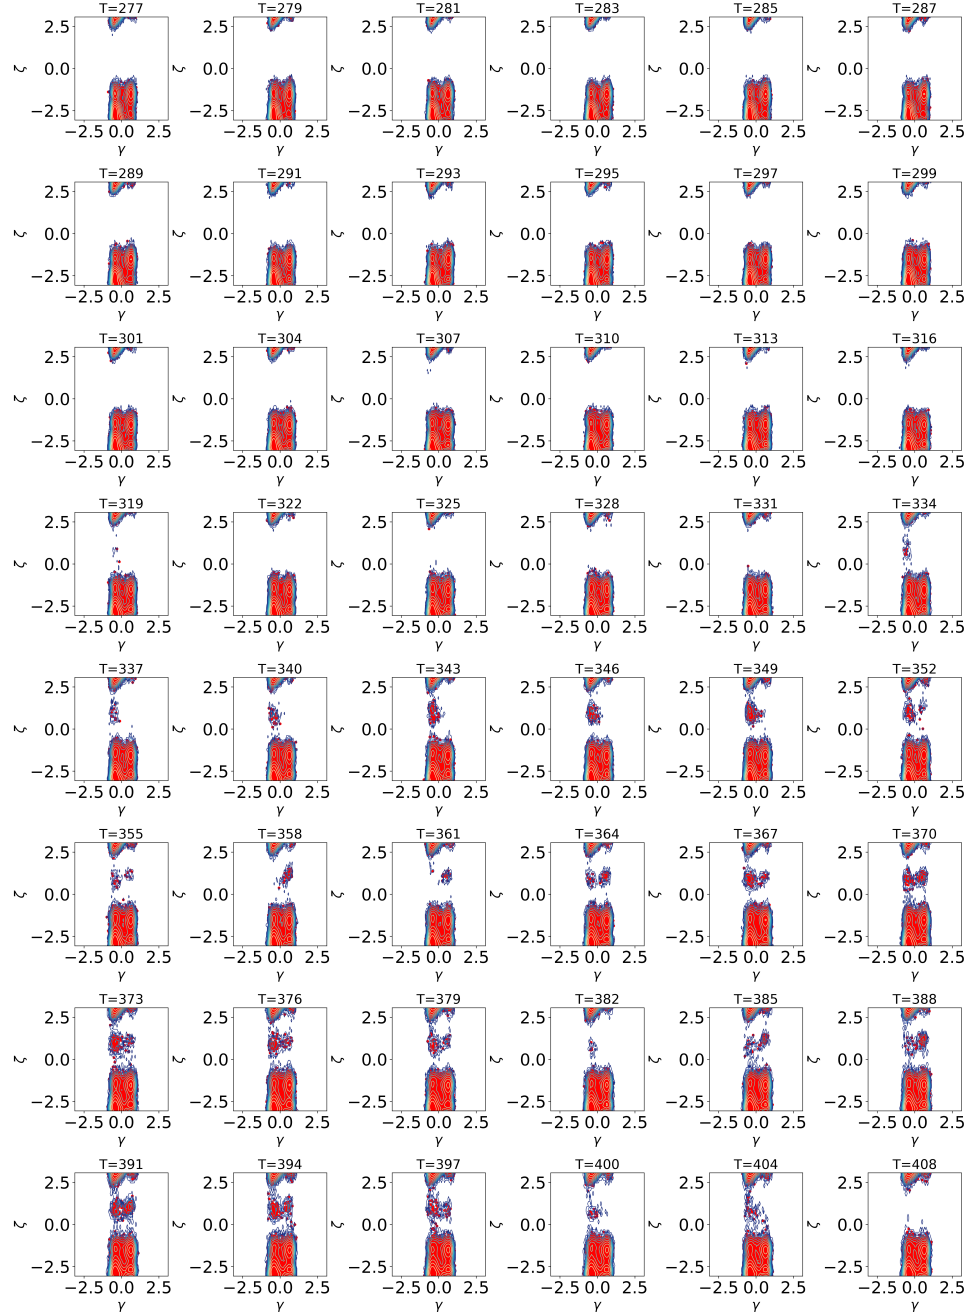


Figure 5.10: Projection of REMD samples and free energy profile on dihedral angles  $\epsilon$  and  $\zeta$  the C3 in GACC at each replica.

almost entirely misses out on any sampling of the barriers, thereby heavily overestimating their heights. For other dihedral angles, even the free energy of the metastable states is overestimated from the direct approach. In conclusion, such a reweighting operation

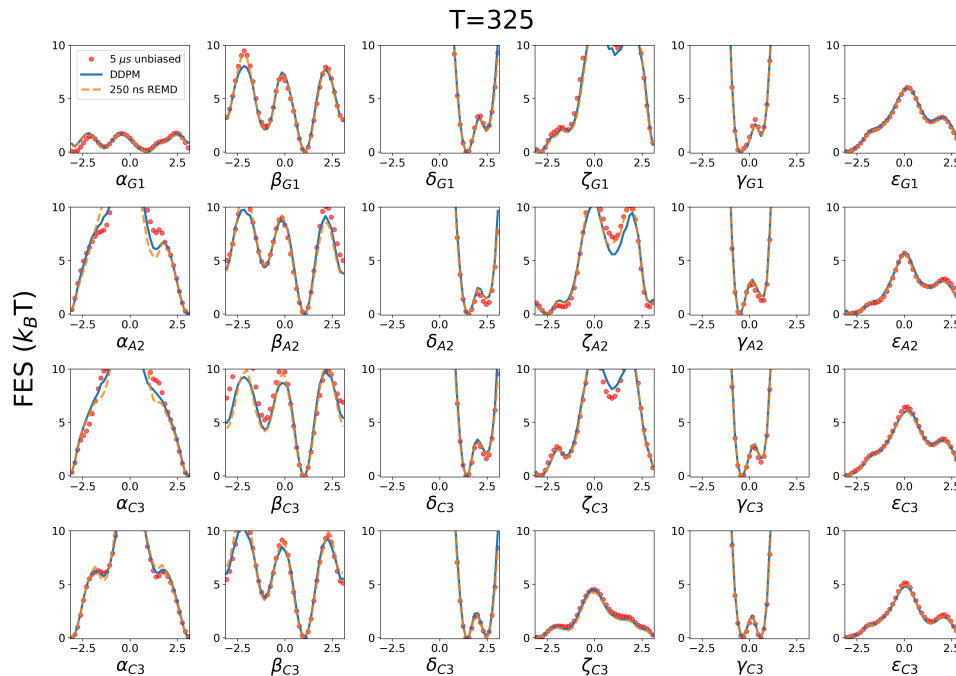


Figure 5.11: Free energy profiles of all dihedral angles in GACC calculated from  $4\mu s$  unbiased MD (in blue dash line),  $250ns$  REMD (in green solid line) and samples generated from DDPM (in orange solid line)

very significantly overestimates free energy difference between metastable states, thus unequivocally demonstrating the usefulness of the DDPM approach. At the temperature where long unbiased MD is possible, the use of direct reweighting (Eq. 5.1) gives an energy scale a least  $10 k_B T$  higher than DDPM or unbiased MD, while there is near perfect agreement between the latter two.

## 5.5 Discussion

I have presented a generative AI-based approach that combines physics from simulations performed at different temperatures to generate reliable new molecular configurations and accurate thermodynamic estimates at any arbitrary temperature even if no actual simulation was performed there. The central idea is to not work with the

thermodynamic temperature  $\beta$  of the system as a parameter set by the heat bath, but instead work with an effective temperature, calculated from the instantaneous kinetic energy of the system. This effective temperature shows non-trivial fluctuations for a finite size system, and on an average equals the thermodynamic temperature. Given sparse sampling from the high-dimensional space comprising configurational space coordinates and the effective temperature, we train a generative AI model that generates countless more samples of configurations at any temperature of interest. While the idea is generally applicable, here I demonstrate its usefulness in the context of the widely used replica exchange molecular dynamics (REMD) framework to improve the sampling of REMD through a post-processing framework. I show how this significantly improves the quality of sampling at low temperatures and even generate samples in states where have not even been visited in the replicas, and at temperatures not considered in the ladder of replicas.

I also show that this approach is far superior to direct reweighting scheme as we do not compute weights involving exponentials of the full system’s kinetic or potential energy, thereby avoiding well-known issues with the convergence of exponentially averaged quantities [172].

The generative AI framework of DDPM used here belongs to the broad class of flow-type methods, which have been shown to have the ability to generate samples from high dimensional space with many interdependent degrees of freedom. Compared with other flow type models such as normalizing flow[65, 173] that use deterministic functions to map from an easy-to-sample distribution to target distribution, the stochastic nature of DDPM avoids the restriction of preserving the topology of configuration space and thus allows the learning of significantly more complicated distributions. At the same time the

design of the transition kernels in DDPM reduces the learning task to just learning means of Gaussian kernels. This makes the training easier compared to other methods[68] while at the same time being able to learn more complicated transition kernels.

I also believe that the generative AI model, while “dreaming” thermodynamically relevant structures at different temperatures, avoids the so-called hallucinations suffered by other generative AI models, i.e. the model do not generate meaningless, unphysical structures with significant thermodynamic weights [156]. I believe this is through the use of relatively simple transition kernels, which avoids overparameterization of the model, and through utilizing molecular basis functions instead of all-atom coordinates, which reduces the space that needs to be sampled. The issue of generating out-of-distribution samples that has been problematic in other methods attempting to generate molecular structure, such as the Boltzmann generator, is usually avoided by discarding the translational and rotational degrees of freedom and reweighting the samples[65, 68]. However, calculating the weight of each sample requires knowledge of all the coordinates of a system which may also become an issue when the system contains explicit water molecules; in such a configuration space the samples will again become sparse. Finally, I would point out that this work shows the possibility of learning generative models in the space of generic thermodynamic ensembles, by following the simple recipe that control parameters can also be viewed as fluctuating variables. As long as one is not in the thermodynamic limit – something we do not have to worry about in molecular simulations – this should be thus a practical and useful procedure for problems far beyond replica exchange molecular dynamics.



## Chapter 6: Summary and future directions

### 6.1 AI-augmented MD simulations

I presented in this thesis two AI-based computational methods to solve the time scale problem in MD simulation study of biomolecular systems.

In the first method RAVE, I showed that based on the framework of information bottleneck, we could use a neural network to learn the system's RCs. By enhancing the fluctuation of RCs, we could accelerate the transition between meta-stable states and thus sample the system better with limited computational power. One significant advantage of this method is that it doesn't require much prior knowledge. By combining OP preselection methods like AMINO[101], RAVE only requires an initial structure as input and automatically learns the optimal RCs by iterating between machine learning and enhanced MD.

In the second method, I used a generative model to address the problem of using high temperature simulations to infer observations about low temperatures. I demonstrated how using generative artificial intelligence to mix information from simulations conducted at a set of temperatures and generate molecular configurations at any temperature of interest, including temperatures at which simulations were never performed. Particularly, the numerical results on a small peptide chain and a nucleotide

chain show that if we treat the instantaneous temperature  $T$  as the random variable and learn its joint distribution with CVs, we can use the learned generate configurations with correct Boltzmann weights at the target temperatures. Since the model can infer the underlying relationship between samples from different temperatures, they can infer the distribution of data in poorly sampled regions, thus significantly enhancing the sampling of REMD. As a method that doesn't require much prior knowledge, REMD can be easily applied to a wide range of systems. The combination of DDPM and REMD can further reduce the computational cost and make it possible to study systems with a bigger size. And it is very suitable to be applied to biomolecular systems where the dynamics are complicated, and simple RCs could not be found. Such as intrinsically disordered proteins. It also has potential in studying conformational ensemble or structure optimization. I also believe the framework is extensible to generic simulations and experiments for mixing control parameters other than temperature.

## 6.2 Study of ligand dissociation from RNA

With the AI-augmented sampling method RAVE, I investigated the dissociation of two ligands from PreQ<sub>1</sub> riboswitch. Unlike ligand dissociation from protein which is widely studied with MD simulation, there are few computational studies of dissociation of ligand from nucleotide due to the lack of solved structure and relatively poor force field parameterization. In this collaborative study with experimentalists, I used MD simulation to provide a high spatiotemporal resolution description interaction between ligands and PreQ<sub>1</sub> riboswitch both in its bound pose and during the dissociation pathways.

The MD simulation successfully compares the stabilization of different regions of PreQ<sub>1</sub> riboswitch, which is also in good agreement with experimental SHAPE measurements. From the dissociation trajectories simulated by RAVE, I found that the existence of two unbinding pathways also gets affinity estimation that is in good agreement with the experiment. Analyzing the trajectories provided insights on how different nucleotides interact with the ligands, and I made prediction mutations that could change the binding affinity based on these observations. Both *in vitro* and *in silico* mutagenesis experiments validated the prediction and revealed how the mutations change the binding ensemble of ligand-RNA complexes. This study indicates the potential of using MD simulations to model and study the interplay between structural features and small molecule recognition of RNA. Such a protocol is also transferable to studying features and dynamics that govern how small synthetic molecules interact with RNA.

### 6.3 Future work

Based on the studies mentioned in this thesis, there are many directions worth exploring in the future. These include improving methods themselves and applying them to more challenging problems, and I will summarize them as follows.

Firstly, there is ample need and space for improvement if we want RAVE to be a method that the broad scientific community can apply to generic systems with arbitrary complexity. One potential improvement might be designing new structures for the encoder to enable the learning of more complicated features. A linear encoder has the advantage that we can directly use the linear combination coefficients to interpret

how different OPs contribute to the RC. Also, the success of applying RAVE to different systems indicates that these linear RCs can give a relatively good approximation of the underlying reaction process. However, this might not be true for more complicated cases. We need multiple linear RCs or a nonlinear RC to capture more complex features in such cases. But I would like to note that introducing a more complicated encoder might also come at the cost of losing both the interpretability and stability of the solutions. One possible solution to this dilemma requires us to develop a strategy to address carefully when one or even two reaction coordinates are insufficient to describe the process of interest. This will allow us to gradually add complexity to the encoder to capture only the essential physics.

Secondly, with the abundant data generated from these enhanced sampling methods, new analysis tools are required for us to interpret the mechanism hiding behind the high dimensional data. The strength and mechanism of the ligand-receptor interaction are the consequence of many features of these molecules. Only a small subset of those features is studied in my study of the dissociation events. Many other features such as the geometry properties of the ligands and receptors, the presence of stacking interactions, and solvent effects are also crucial in determining how ligand and receptor react with each other. Being able to dig information from a long list of such features will help us to understand their different contribution better or even learn new mechanisms. Machine learning is an ideal tool for such an analysis and has already been used to guide human intuition.[\[174, 175, 176\]](#) At the same time, the involvement of enhanced sampling method further enriches this analysis by allowing us to look at the distribution of different features or adding the dimension of time. Moreover, MD simulations can also serve as a flexible

but powerful experimental platform to validate the relationship learned by AI model.

Finally, the power of DPPM in learning complicated distribution and making physical interpolation can also be applied quite generally to data coming from simulations or experiments at different temperatures. It should also be extensible to mixing data from control parameters other than temperatures - possibly including concentration, pressure and volume. Also, it can be extended to learn the system's kinetic from biased simulation. For example, we can apply DDPM on learning a distribution  $P(\mathbf{X}_{t+1}|\mathbf{X}_t, \beta)$  where the  $\beta$  is the control parameters such as the temperature or the magnitude of external potential. Such a learned model can then generate a new trajectory under desired conditions. Moreover, in an ongoing project, I use DDPM to build a model which can convert the time series of one observable to the time series of another observable. More specifically, DDPM is train to learn the conditional distribution  $P(s_N^2 | s_N^1, s_{N-1}^1 \cdots s_0^1)$  where  $s^1$  and  $s^2$  are two variables and the subscripts denote the time steps. This method is grounded on the Takens delay embedding theorem[177] and its probabilistic counterpart[178]. It will be particularly useful in experiment as we can use this model to predict the hard-to-measure variables by measuring the time series of easy-to-measure variables. For example, we can use the fluorescence resonance energy transfer (FRET) signal[179], which reflects the distance between two ends of a molecule, to estimate other quantities that are more of interest but can not be directly measured in experiment.

## Appendix A:

### A.1 MD setups

#### A.1.1 Alanine dipeptide

I follow [63] to set up our simulation for alanine dipeptide in vacuum. The simulations are performed with the software GROMACS 2016/GROMACS 5.0[180, 181], patched with PLUMED 2.4[182]. I constrain bonds involving hydrogens using the LINCS algorithm[183] and employ an integration time-step of 0.002 ps. The temperature is kept constant at 300K using the velocity rescaling thermostat (relaxation time of 0.1 fs)[106]. I employ no periodic boundary conditions and no cut-offs for the electrostatic and non-bonded Van der Waals interactions. The integration time step was 2 fs and order parameters were saved every time step.

In each round, four independent simulations with different initial randomized velocities as per Maxwell-Boltzmann distribution were performed to improve the quality of the free energy sampling.

### A.1.2 L99A T4L-benzene

I follow [96] to set up our simulation for L99A T4L-benzene. The simulations are performed with the software GROMACS 2016/GROMACS 5.0[180, 181] patched with PLUMED 2.4[182]. The simulations are done with the constant number, pressure, temperature (NPT) ensemble with temperature 298 K and pressure 1.0 bar. Constant pressure is maintained using Parrinello-Rahaman barostat while the constant temperature is maintained using the v-rescale thermostat (modified Berendsen thermostat)[22]. The simulation box with periodic boundary condition is filled with TIP3P water. The side lengths of the box are 10 Å and there are around 10,000 water molecules. The interaction is described by the force field CHARMM22\*. The integration time step here as well was taken to be 2 fs.

Similar to the the study of alanine dipeptide, in each round, four independent simulations with different initial randomized velocities as per Maxwell-Boltzmann distribution were performed to improve the quality of the free energy sampling.

### A.1.3 AIB<sub>9</sub>

The simulation of AIB<sub>9</sub> was setup by following a previous study[166]. The PDB file was taken from the authors with permission, and the simulations are done with the CHARMM36m all atom force field [184] using TIP3P water molecules[162], a Parrinello-Rahman barostat[185], and a Nose-Hoover thermostat[21, 186] under the NPT ensemble. Simulations were performed using GROMACS 2016.[180] The structures of AIB<sub>9</sub> were saved every 0.2 ps and the dihedral angles were calculated using PLUMED

2.4[187].

#### A.1.4 GACC

The simulations of GACC were done with the AMBER ff12 all atom force field, using TIP3P water molecules[162], a Parrinello-Rahman barostat[185], and a Bussi-Parrinello velocity rescaling thermostat[106] under the NPT ensemble. The simulations were performed using GROMACS 2016[180]. The AMBER force field was chosen because it exhibits more conformational variability compared with CHARMM force field in a previous study[188]. The structures of GACC were also saved every 0.2 ps and the dihedral angles defined in Table S1 were calculated using PLUMED 2.4[187].

The PDB file of the GACC structure that served as the starting point for our simulation was generated using PyMOL. The initial structure was assumed to be at a temperature of 10K, and the system's energy was minimized with positional restraints of  $25 \text{ kcal mol}^{-1} \text{ \AA}^{-2}$  in a two step process - first the steepest descent algorithm was applied for 1000 steps followed by the conjugate gradient algorithm for another 1000 steps. Next, the system was equilibrated to 150K over 100ps under the NVT ensemble with positional restraints of  $25 \text{ kcal mol}^{-1} \text{ \AA}^{-2}$ . Then, the GACC was equilibrated from 150K to 277K at 1 atm under the NPT ensemble for 100 ps with positional restraints of  $5 \text{ kcal mol}^{-1} \text{ \AA}^{-2}$ . Finally, a long 5 ns equilibration was performed over 5ns at 298K under the NPT ensemble with positional restraints of  $0.5 \text{ kcal mol}^{-1} \text{ \AA}^{-2}$ , after which the system was copied to 48 replicas, and each replica was equilibrated to its target temperature under the NPT ensemble with positional restraints of  $0.5 \text{ kcal mol}^{-1} \text{ \AA}^{-2}$ .



The replica temperatures were geometrically spaced temperatures ranging from 277K to 408K, with temperature increased by 1% for each replica. The attempt of exchanging configurations was made every 10 ps, which is determined by checking the time correlation function of the potential energy.

### A.1.5 Architecture of RAVE

A densely-connected layer without activation function was used to linearly encode the order parameters  $\mathbf{X}$  to reaction coordinates  $\mathbf{s}$ . Gaussian noise was added to  $\mathbf{s}$  before being passed to the decoder. Decoder consisted of 2 hidden layers and an output layer which were all densely-connected. ELU was used as the activation function for hidden layers. We assume  $Q_\phi(\mathbf{X}_{\Delta t}|\mathbf{s}) = \mathcal{N}(\mathbf{X}_{\Delta t}; f_\phi(\mathbf{s}), \sigma^2)$ .  $f_\phi(\mathbf{s})$  corresponds to the decoder part of the neural network, which maps states on the reaction coordinate to states in order parameter space. With this assumption, maximizing the objective function is equivalent to minimizing the mean square error between  $\mathbf{X}_{\Delta t}$  and network prediction  $f_\phi(\mathbf{s})$ .

Hyper-parameters in RAVE included the variance of Gaussian noise, the number of neurons in hidden layers, initializer of weights of each layer, and the learning rate for the RMSprop algorithm [90]. In all three case studies, all these hyper-parameters are set to be the same. The variance of Gaussians was kept 0.005. Each hidden layer had 128 neurons. The learning rate was set to be 0.003. Initial weights of each layer were randomly picked from a uniform within range  $[-0.005, 0.005]$ . The transferability of hyper-parameters between different systems without much tuning reflects that this method is not very sensitive to the choice of neural network hyper-parameters. From our experience, we

suggest that the choice of the variance of Gaussian should not be too big as it will wash out meaningful features. For more complicated systems, a deeper (more layers) or wider (more neurons in each layer) decoder might be needed. The tuning of learning rate was done by looking at how order parameter weights change during the training process. If the learning rate was too high, order parameter weights did not converge.

Independent simulations often explored different configurations due to the finite and small simulation time. We considered the trajectory with the highest variance to have maximum ergodic exploration and used it to train the reaction coordinate for the next round. Similar to other non-convex optimization problems, the results could have converged to a local minimum or even a saddle point. To safeguard against such spurious solutions learnt by the neural network, we performed independent training runs with random initial weights in layers. The RC was then determined as the linear encoder of the trained neural network with smallest loss function. We will be refining this procedure in future work as strictly speaking a lower loss function is no guarantee of reaching a better solution.

## A.2 Definitions of quantities in information theory and variational lower bound

In this section I will define various terms introduced this thesis. In all of these I use  $P(\mathbf{X})$ ,  $P(\mathbf{X}, \mathbf{Y})$  respectively to denote the probability distribution of a random variable  $X$  and the joint probability distribution of two random variables  $\mathbf{X}$  and  $\mathbf{Y}$ . Other variables are the same as defined in the main text.

### A.2.1 Mutual information

This is a commonly used information theoretic measure to describe how much information is shared between two random variables  $X$  and  $Y$ . It is defined as:

$$I(\mathbf{X}, \mathbf{Y}) = \int P(\mathbf{X}, \mathbf{Y}) \ln \frac{P(\mathbf{X}, \mathbf{Y})}{P(\mathbf{X})P(\mathbf{Y})} d\mathbf{X}d\mathbf{Y} \quad (\text{A.1})$$

### A.2.2 Shannon entropy

The Shannon entropy for a random variable  $X$  is defined as:

$$H(\mathbf{X}) = - \int P(\mathbf{X}) \ln P(\mathbf{X}) d\mathbf{X} \quad (\text{A.2})$$

### A.2.3 Cross entropy

The cross entropy between two probability distributions  $P(\mathbf{X})$  and  $Q(\mathbf{X})$  is given by:

$$C(\mathbf{X}, \mathbf{Y}) = - \int P(\mathbf{X}) \ln Q(\mathbf{X}) d\mathbf{X} \quad (\text{A.3})$$

### A.2.4 Kullback-Leibler Divergence

The Kullback-Leibler (KL)  $D_{KL}(P||Q)$  divergence between two probability distributions  $P(\mathbf{X})$  and  $Q(\mathbf{X})$  is given by:

$$D_{KL}(P||Q) = \int P(\mathbf{X}) \ln \frac{P(\mathbf{X})}{Q(\mathbf{X})} d\mathbf{X} \quad (\text{A.4})$$

### A.2.5 Exact decoder definition

The exact decoder can be defined as the conditional probability distribution of  $\mathbf{X}_{\Delta t}$  given a RC  $\mathbf{s}$ . It can be calculated by using Bayes theorem:

$$P(\mathbf{X}_{\Delta t}|\mathbf{s}) = \frac{P(\mathbf{s}, \mathbf{X}_{\Delta t})}{P(\mathbf{s})} = \frac{\int P(\mathbf{s}|\mathbf{X})P(\mathbf{X}, \mathbf{X}_{\Delta t})d\mathbf{X}}{\int \int P(\mathbf{s}|\mathbf{X})P(\mathbf{X}, \mathbf{X}'_{\Delta t})d\mathbf{X}d\mathbf{X}'_{\Delta t}} \quad (\text{A.5})$$

### A.2.6 Gibbs's inequality and variational lower bound

Here I will provide the missing details of various derivations given in the Sec. 2.3.

The bottleneck function  $\mathcal{L}$  defined in the main text for our neural network architecture is given by a difference of two Shannon entropies[82]:

$$\mathcal{L} = I(\mathbf{s}, \mathbf{X}_{\Delta t}) = H(P(\mathbf{X}_{\Delta t})) - H(P_{\theta}(\mathbf{X}_{\Delta t}|\mathbf{s})) \quad (\text{A.6})$$

Gibbs's inequality[82] guarantees that the KL-divergence between two probability distributions is always larger than 0. We thus have:

$$\begin{aligned} D_{KL}(P_{\theta}(\mathbf{X}_{\Delta t}|\mathbf{s})||Q_{\phi}(\mathbf{X}_{\Delta t}|\mathbf{s})) &\equiv \int \int P_{\theta}(\mathbf{X}_{\Delta t}, \mathbf{s}) \ln \frac{P_{\theta}(\mathbf{X}_{\Delta t}|\mathbf{s})}{Q_{\phi}(\mathbf{X}_{\Delta t}|\mathbf{s})} d\mathbf{X}_{\Delta t} d\mathbf{s} \\ &= \int \int P_{\theta}(\mathbf{X}_{\Delta t}, \mathbf{s}) \ln P_{\theta}(\mathbf{X}_{\Delta t}|\mathbf{s}) d\mathbf{X}_{\Delta t} d\mathbf{s} - \int \int P_{\theta}(\mathbf{X}_{\Delta t}, \mathbf{s}) \ln Q_{\phi}(\mathbf{X}_{\Delta t}|\mathbf{s}) d\mathbf{X}_{\Delta t} d\mathbf{s} \\ &= -H(P_{\theta}(\mathbf{X}_{\Delta t}|\mathbf{s})) + C(P_{\theta}(\mathbf{X}_{\Delta t}|\mathbf{s}), Q_{\phi}(\mathbf{X}_{\Delta t}|\mathbf{s})) \geq 0 \end{aligned} \quad (\text{A.7})$$

Only in the limit that our approximate decoder is exactly the same as the exact inverse-Bayes decoder,  $D_{KL}(P_{\theta}||Q_{\phi})$  equals 0. By combining Eq.A.6 and Eq.A.7 , we get the

relationship:

$$\begin{aligned}
\mathcal{L} &\geq \mathcal{L} - D_{KL}(P_{\theta}(\mathbf{X}_{\Delta t}|\mathbf{s})||Q_{\phi}(\mathbf{X}_{\Delta t}|\mathbf{s})) \\
&= H(P(\mathbf{X}_{\Delta t})) - C(P_{\theta}(\mathbf{X}_{\Delta t}|\mathbf{s}), Q_{\phi}(\mathbf{X}_{\Delta t}|\mathbf{s})) \\
&\equiv H(P(\mathbf{X}_{\Delta t})) + \mathcal{L}'
\end{aligned} \tag{A.8}$$

Here,  $H(P(\mathbf{X}_{\Delta t}))$  only depends on the data set and is independent of the parametrization of the encoder and the decoder. Thus the term  $H(P(\mathbf{X}_{\Delta t}))$  can be completely ignored while optimizing the parameters  $\theta$  and  $\phi$ . Maximizing the objective function  $\mathcal{L}' = -C(P_{\theta}(\mathbf{X}_{\Delta t}|\mathbf{s}))$  is then equivalent to maximizing the lower bound of  $\mathcal{L}$ , which is the expression stated in Eq.4 in the main text.

### A.3 PIB objective $\mathcal{L}'$ for unbiased trajectory

For a unbiased trajectory  $\{\mathbf{X}^1, \mathbf{X}^2, \dots, \mathbf{X}^{M+k}\}$ , for given  $\theta$ , we can get corresponding  $\{\mathbf{s}^1, \mathbf{s}^2, \dots, \mathbf{s}^M\}$  using  $\mathbf{s}_i = \sum_i c_i \mathbf{s}_i$ . We also have a sampling of the states of the system after corresponding  $\Delta t$  intervals:  $\{\mathbf{X}^{1+k}, \mathbf{X}^{2+k}, \dots, \mathbf{X}^{M+k}\}$ . Together the pair  $(\mathbf{s}_i, \mathbf{X}_{i+k})$  sampled from the dataset follows the distribution  $P_{\theta}(\mathbf{X}_{\Delta t}|\mathbf{s})$ . With this, we have:

$$\begin{aligned}
\mathcal{L}' &= -C(P_{\theta}(\mathbf{X}_{\Delta t}|\mathbf{s}), Q_{\phi}(\mathbf{X}_{\Delta t}|\mathbf{s})) \\
&= \int \int P_{\theta}(\mathbf{X}_{\Delta t}, \mathbf{s}) \ln Q_{\phi}(\mathbf{X}_{\Delta t}|\mathbf{s}) d\mathbf{X}_{\Delta t} d\mathbf{s} \\
&= \frac{1}{M} \sum_{n=1}^M \log Q(\mathbf{X}^{n+1}|\mathbf{s}^n)
\end{aligned} \tag{A.9}$$

#### A.4 PIB objective $\mathcal{L}'$ for biased trajectory

For a biased trajectory  $\{\mathbf{X}^1, \mathbf{X}^2, \dots, \mathbf{X}^M\}$  with corresponding biasing potential values  $\{V^1, V^2, \dots, V^M\}$ , the unbiased probability distribution of  $\mathbf{X}$  can be calculated by:

$$P(\mathbf{X}) = \frac{\sum_i \delta(\mathbf{X}^i - \mathbf{X}) e^{\beta V^i}}{\sum_i e^{\beta V^i}} \quad (\text{A.10})$$

The encoder  $P(\mathbf{s}|\mathbf{X})$  and decoder  $P(\mathbf{X}_{\Delta t}|\mathbf{s})$  are taken to be independent of the bias. The first assumption is strictly true, while the second is valid for small enough  $\Delta t$  as explained in the main text. Therefore

$$\begin{aligned} \mathcal{L}' &= \int \int P_{\theta}(\mathbf{X}_{\Delta t}|\chi) P(\chi) \ln Q_{\phi}(\mathbf{X}_{\Delta t}|\chi) d\mathbf{X}_{\Delta t} d\chi \\ &= \left\{ \sum_{n=1}^M e^{\beta V^n} \right\}^{-1} \sum_{n=1}^M e^{\beta V^n} \log Q(\mathbf{X}^{n+1}|\chi^n) \end{aligned} \quad (\text{A.11})$$

## Publications from Doctoral Research

1. Ribeiro, João Marcelo Lamim, Bravo, Pablo, **Wang, Yihang**, and Tiwary, Pratyush. "Reweighted autoencoded variational Bayes for enhanced sampling (RAVE)." *The Journal of chemical physics* 149, no. 7 (2018): 072301.
2. Ribeiro, Joao Marcelo Lamim, Tsai, Sun-Ting, Pramanik, Debabrata, **Wang, Yihang**, and Tiwary, Pratyush. "Kinetics of ligand-protein dissociation from all-atom simulations: Are we there yet?." *Biochemistry* 58, no. 3 (2018): 156-165.
3. **Wang, Yihang**, Ribeiro, Joao Marcelo Lamim, and Tiwary, Pratyush. "Past-future information bottleneck for sampling molecular reaction coordinate simultaneously with thermodynamics and kinetics." *Nature communications* 10, no. 1 (2019): 1-8.
4. **Wang, Yihang**, Ribeiro, Joao Marcelo Lamim, and Tiwary, Pratyush. "Machine learning approaches for analyzing and enhancing molecular dynamics simulations." *Current Opinion in Structural Biology* 61 (2020): 139-145.
5. **Wang, Yihang**, and Tiwary, Pratyush. "Understanding the role of predictive time delay and biased propagator in RAVE." *The Journal of Chemical Physics* 152, no. 14 (2020): 144102.
6. Smith, Zachary, **Wang, Yihang**, Ravindra, Pavan, Cooley, Rory, and Tiwary, Pratyush. "Discovering protein conformational flexibility through artificial-intelligence-aided molecular dynamics." *The Journal of Physical Chemistry B* 124, no. 38 (2020): 8221-8229.
7. Pant, Shashank, **Wang, Yihang**, Smith, Zachary, Tajkhorshid, Emad, and Tiwary, Pratyush. "Confronting pitfalls of AI-augmented molecular dynamics using statistical physics." *The Journal of Chemical Physics* 153, no. 23 (2020): 234118.
8. **Wang, Yihang**, Parmar, Shaifaly, Schneekloth, John S., and Tiwary, Pratyush. "Interrogating RNA-small molecule interactions with structure probing and AI augmented-molecular simulations." *bioRxiv* (2021).
9. **Wang, Yihang**, and Tiwary, Pratyush. "From data to noise to data: mixing physics across temperatures with generative artificial intelligence." *arXiv preprint arXiv:2107.07369* (2021).

## Bibliography

- [1] B. Alberts, J.H. Wilson, A. Johnson, T. Hunt, J. Lewis, K. Roberts, M. Raff, and P. Walter. *Molecular Biology of the Cell*. Garland Science, 2008.
- [2] Emilio Gallicchio and Ronald M. Levy. *Recent theoretical and computational advances for modeling protein-ligand binding affinities*, volume 85. 2011.
- [3] A focus on regulatory rnas [editorial]. *Nature Cell Biology*, 21(5):535–535, 2019.
- [4] Anthony M. Mustoe, Charles L. Brooks, and Hashim M. Al-Hashimi. Hierarchy of rna functional dynamics. *Annual Review of Biochemistry*, 83(1):441–466, 2014. PMID: 24606137.
- [5] Xing Du, Yi Li, Yuan-Ling Xia, Shi-Meng Ai, Jing Liang, Peng Sang, Xing-Lai Ji, and Shu-Qun Liu. Insights into protein–ligand interactions: mechanisms, models, and methods. *International journal of molecular sciences*, 17(2):144, 2016.
- [6] Aron W Fenton. Allostery: an illustrated definition for the ‘second secret of life’. *Trends in biochemical sciences*, 33(9):420–425, 2008.
- [7] Hesam N Motlagh, James O Wrabl, Jing Li, and Vincent J Hilser. The ensemble nature of allostery. *Nature*, 508(7496):331–339, 2014.
- [8] F Ulrich Hartl. Protein misfolding diseases. *Annual review of biochemistry*, 86:21–26, 2017.
- [9] Piya Lahiry, Ali Torkamani, Nicholas J Schork, and Robert A Hegele. Kinase mutations in human disease: interpreting genotype–phenotype relationships. *Nature Reviews Genetics*, 11(1):60–74, 2010.
- [10] Jack W Scannell, Alex Blanckley, Helen Boldon, and Brian Warrington. Diagnosing the decline in pharmaceutical r&d efficiency. *Nature reviews Drug discovery*, 11(3):191–200, 2012.
- [11] Robert A Copeland. The drug–target residence time model: a 10-year retrospective. *Nature Reviews Drug Discovery*, 15(2):87–95, 2016.



- [12] Michael M Pierce, CS Raman, and Barry T Nall. Isothermal titration calorimetry of protein–protein interactions. *Methods*, 19(2):213–221, 1999.
- [13] Priyabrata Pattnaik. Surface plasmon resonance. *Applied biochemistry and biotechnology*, 126(2):79–92, 2005.
- [14] William J Checovich, Randall E Bolger, and Thomas Burke. Fluorescence polarization—a new tool for cell and molecular biology. *Nature*, 375(6528):254–256, 1995.
- [15] Xing Du, Yi Li, Yuan-Ling Xia, Shi-Meng Ai, Jing Liang, Peng Sang, Xing-Lai Ji, and Shu-Qun Liu. Insights into protein–ligand interactions: mechanisms, models, and methods. *International journal of molecular sciences*, 17(2):144, 2016.
- [16] Ron O. Dror, Robert M. Dirks, J. P. Grossman, Huafeng Xu, and David E. Shaw. Biomolecular simulation: A computational microscope for molecular biology. *Annual Review of Biophysics*, 41(1):429–452, 2012.
- [17] Ron O. Dror, Albert C. Pan, Daniel H. Arlow, David W. Borhani, Paul Maragakis, Yibing Shan, Huafeng Xu, and David E. Shaw. Pathway and mechanism of drug binding to g-protein-coupled receptors. *Proc. Natl. Acad. Sci.*, 108(32):13118–13123, 2011.
- [18] Yibing Shan, Eric T Kim, Michael P Eastwood, Ron O Dror, Markus A Seeliger, and David E Shaw. How does a drug molecule find its target binding site? *J. Amer. Chem. Soc.*, 133(24):9181–9183, 2011.
- [19] Dow P Hurst, Alan Grossfield, Diane L Lynch, Scott Feller, Tod D Romo, Klaus Gawrisch, Michael C Pitman, and Patricia H Reggio. A lipid pathway for ligand binding is necessary for a cannabinoid g protein-coupled receptor. *J. Biol. Chem.*, pages jbc–M109, 2010.
- [20] Jagannath Mondal, Navjeet Ahalawat, Subhendu Pandit, Lewis E Kay, and Pramodh Vallurupalli. Atomic resolution mechanism of ligand binding to a solvent inaccessible cavity in t4 lysozyme. *PLoS computational biology*, 14(5):e1006180, 2018.
- [21] Denis J Evans and Brad Lee Holian. The nose–hoover thermostat. *The Journal of chemical physics*, 83(8):4069–4074, 1985.
- [22] Giovanni Bussi, Davide Donadio, and Michele Parrinello. Canonical sampling through velocity rescaling. *J Chem Phys*, 126(1):014101–101107, 2007.
- [23] H. J. C. Berendsen, J. P. M. Postma, W. F. van Gunsteren, A. DiNola, and J. R. Haak. Molecular dynamics with coupling to an external bath. *The Journal of Chemical Physics*, 81(8):3684–3690, 1984.
- [24] D Quigley and MIJ Probert. Langevin dynamics in constant pressure extended systems. *The Journal of chemical physics*, 120(24):11432–11441, 2004.

- [25] Baron Peters. Reaction Coordinates and Mechanistic Hypothesis Tests. *Annual Review of Physical Chemistry*, 67(1):669–690, 2016.
- [26] Peter G. Bolhuis, David Chandler, Christoph Dellago, and Phillip L. Geissler. TRANSITION PATH SAMPLING : Throwing Ropes Over Rough Mountain Passes, in the Dark . *Annual Review of Physical Chemistry*, 53(1):291–318, 2002.
- [27] Martin Karplus and Gregory A Petsko. Molecular dynamics simulations in biology. *Nature*, 347(6294):631–639, 1990.
- [28] Maya Shamir, Yinon Bar-On, Rob Phillips, and Ron Milo. Snapshot: timescales in cell biology. *Cell*, 164(6):1302–1302, 2016.
- [29] Karina Martínez-Mayorga, Michael C Pitman, Alan Grossfield, Scott E Feller, and Michael F Brown. Retinal counterion switch mechanism in vision evaluated by molecular simulations. *Journal of the American Chemical Society*, 128(51):16502–16503, 2006.
- [30] Ron O Dror, Robert M Dirks, JP Grossman, Huafeng Xu, and David E Shaw. Biomolecular simulation: a computational microscope for molecular biology. *Annual review of biophysics*, 41:429–452, 2012.
- [31] David E Shaw, Peter J Adams, Asaph Azaria, Joseph A Bank, Brannon Batson, Alistair Bell, Michael Bergdorf, Jhanvi Bhatt, J Adam Butts, Timothy Correia, et al. Anton 3: twenty microseconds of molecular dynamics simulation before lunch. In *Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis*, pages 1–11, 2021.
- [32] Matt J Harvey, Giovanni Giupponi, and G De Fabritiis. Acemd: accelerating biomolecular dynamics in the microsecond time scale. *Journal of chemical theory and computation*, 5(6):1632–1639, 2009.
- [33] Thomas A Halgren and Wolfgang Damm. Polarizable force fields. *Current Opinion in Structural Biology*, 11(2):236 – 242, 2001.
- [34] Justin A Lemkul, Jing Huang, Benoît Roux, and Alexander D MacKerell Jr. An empirical polarizable force field based on the classical drude oscillator model: development history and recent applications. *Chemical reviews*, 116(9):4983–5013, 2016.
- [35] Jay W Ponder, Chuanjie Wu, Pengyu Ren, Vijay S Pande, John D Chodera, Michael J Schnieders, Imran Haque, David L Mobley, Daniel S Lambrecht, Robert A DiStasio Jr, et al. Current status of the amoeba polarizable force field. *The journal of physical chemistry B*, 114(8):2549–2564, 2010.
- [36] Oleg Borodin. Polarizable force field development and molecular dynamics simulations of ionic liquids. *The Journal of Physical Chemistry B*, 113(33):11463–11478, 2009.

- [37] Mark E Tuckerman. Ab initio molecular dynamics: basic concepts, current trends and novel applications. *Journal of Physics: Condensed Matter*, 14(50):R1297, 2002.
- [38] Yuji Sugita and Yuko Okamoto. Replica-exchange molecular dynamics method for protein folding. *Chemical Physics Letters*, 314(1):141 – 151, 1999.
- [39] Pu Liu, Byungchan Kim, Richard A Friesner, and BJ Berne. Replica exchange with solute tempering: A method for sampling biological systems in explicit water. *Proceedings of the National Academy of Sciences*, 102(39):13749–13754, 2005.
- [40] Lingle Wang, Richard A Friesner, and BJ Berne. Replica exchange with solute scaling: a more efficient version of replica exchange with solute tempering (rest2). *The Journal of Physical Chemistry B*, 115(30):9431–9438, 2011.
- [41] Andrew J Ballard and Christopher Jarzynski. Replica exchange with nonequilibrium switches. *Proceedings of the National Academy of Sciences*, 106(30):12224–12229, 2009.
- [42] G.A. Huber and S. Kim. Weighted-ensemble brownian dynamics simulations for protein association reactions. *Biophysical Journal*, 70(1):97 – 110, 1996.
- [43] Daniel M. Zuckerman and Lillian T. Chong. Weighted ensemble simulation: Review of methodology, applications, and software. *Annual Review of Biophysics*, 46(1):43–57, 2017. PMID: 28301772.
- [44] Anton K. Faradjian and Ron Elber. Computing time scales from reaction coordinates by milestoning. *The Journal of Chemical Physics*, 120(23):10880–10889, 2004.
- [45] E.V. Anslyn, D.A. Dougherty, E.V. Dougherty, and University Science Books. *Modern Physical Organic Chemistry*. University Science Books, 2006.
- [46] Alessandro Laio and Michele Parrinello. Escaping free-energy minima. *Proceedings of the National Academy of Sciences*, 99(20):12562–12566, 2002.
- [47] Omar Valsson, Pratyush Tiwary, and Michele Parrinello. Enhancing important fluctuations: Rare events and metadynamics from a conceptual viewpoint. *Annual Review of Physical Chemistry*, 67(1):159–184, 2016. PMID: 26980304.
- [48] Pratyush Tiwary and Michele Parrinello. From metadynamics to dynamics. *Phys. Rev. Lett.*, 111:230602, Dec 2013.
- [49] G.M. Torrie and J.P. Valleau. Nonphysical sampling distributions in monte carlo free-energy estimation: Umbrella sampling. *Journal of Computational Physics*, 23(2):187 – 199, 1977.
- [50] Ao Ma and Aaron R. Dinner. Automatic method for identifying reaction coordinates in complex systems. *The Journal of Physical Chemistry B*, 109(14):6769–6779, 2005. PMID: 16851762.

- [51] Hendrik Jung, Roberto Covino, and Gerhard Hummer. Artificial intelligence assists discovery of reaction coordinates and mechanisms from molecular dynamics simulations. *arXiv preprint arXiv:1901.04595*, 2019.
- [52] Stefan Klus, Feliks Nüske, Péter Koltai, Hao Wu, Ioannis Kevrekidis, Christof Schütte, and Frank Noé. Data-driven model reduction and transfer operator approximation. *Journal of Nonlinear Science*, 28(3):985–1010, 2018.
- [53] Frank Noé and Feliks Nuske. A variational approach to modeling slow processes in stochastic dynamical systems. *Multiscale Modeling & Simulation*, 11(2):635–655, 2013.
- [54] Feliks Nüske, Bettina G Keller, Guillermo Pérez-Hernández, Antonia SJS Mey, and Frank Noé. Variational approach to molecular kinetics. *Journal of chemical theory and computation*, 10(4):1739–1752, 2014.
- [55] Guillermo Pérez-Hernández, Fabian Paul, Toni Giorgino, Gianni De Fabritiis, and Frank Noé. Identification of slow molecular order parameters for markov model construction. *The Journal of chemical physics*, 139(1):07B604.1, 2013.
- [56] Andreas Mardt, Luca Pasquali, Hao Wu, and Frank Noé. Vampnets for deep learning of molecular kinetics. *Nature Communications*, 9(1):5, 2018.
- [57] Wei Chen, Hythem Sidky, and Andrew L. Ferguson. Capabilities and limitations of time-lagged autoencoders for slow mode discovery in dynamical systems. *The Journal of Chemical Physics*, 151(6):064123, Aug 2019.
- [58] Carlos X. Hernández, Hannah K. Wayment-Steele, Mohammad M. Sultan, Brooke E. Husic, and Vijay S. Pande. Variational encoding of complex dynamics. *Phys. Rev. E*, 97:062412, Jun 2018.
- [59] Simon Olsson and Frank Noé. Dynamic graphical models of molecular kinetics. *Proceedings of the National Academy of Sciences*, 116(30):15001–15006, 2019.
- [60] Wei Chen and Andrew L. Ferguson. Molecular enhanced sampling with autoencoders: On-the-fly collective variable discovery and accelerated free energy landscape exploration. *Journal of Computational Chemistry*, 39(25):2079–2102, 2018.
- [61] Carlos X. Hernández, Hannah K. Wayment-Steele, Mohammad M. Sultan, Brooke E. Husic, and Vijay S. Pande. Variational encoding of complex dynamics. *Phys. Rev. E*, 97:062412, Jun 2018.
- [62] Luigi Bonati, Yue-Yu Zhang, and Michele Parrinello. Neural networks-based variationally enhanced sampling. *Proceedings of the National Academy of Sciences*, 116(36):17641–17647, 2019.
- [63] Omar Valsson and Michele Parrinello. Variational approach to enhanced sampling and free energy calculations. *Physical review letters*, 113(9):090601, 2014.

- [64] Zahra Shamsi, Kevin J Cheng, and Diwakar Shukla. Reinforcement learning based adaptive sampling: Reaping rewards by exploring protein conformational landscapes. *The Journal of Physical Chemistry B*, 122(35):8386–8395, 2018.
- [65] Frank Noé, Simon Olsson, Jonas Köhler, and Hao Wu. Boltzmann generators: Sampling equilibrium states of many-body systems with deep learning. *Science*, 365(6457):eaaw1147, 2019.
- [66] Laurent Dinh, Jascha Sohl-Dickstein, and Samy Bengio. Density estimation using real nvp. *arXiv preprint arXiv:1605.08803*, 2016.
- [67] Durk P Kingma and Prafulla Dhariwal. Glow: Generative flow with invertible 1x1 convolutions. In *Advances in Neural Information Processing Systems*, pages 10215–10224, 2018.
- [68] Hao Wu, Jonas Köhler, and Frank Noé. Stochastic normalizing flows. *arXiv preprint arXiv:2002.06707*, 2020.
- [69] Manuel Dibak, Leon Klein, and Frank Noé. Temperature steerable flows and boltzmann generators. *arXiv preprint arXiv:2108.01590*, 2021.
- [70] Stephanie E Palmer, Olivier Marre, Michael J Berry, and William Bialek. Predictive information in a sensory population. *Proc. Natl. Acad. Sci.*, 112(22):6908–6913, 2015.
- [71] Naftali Tishby, Fernando C Pereira, and William Bialek. The information bottleneck method. *Preprint at arXiv:physics/0004057*, 2000.
- [72] Susanne Still. Information bottleneck approach to predictive inference. *Entropy*, 16(2):968–989, 2014.
- [73] A Elisabeth Eriksson, Walter A Baase, Xue-Jun Zhang, Dirk W Heinz, MPBE Blaber, Enoch P Baldwin, and Brian W Matthews. Response of a protein structure to cavity-creating mutations and its relation to the hydrophobic effect. *Science*, 255(5041):178–183, 1992.
- [74] Victoria A Feher, Enoch P Baldwin, and Frederick W Dahlquist. Access of ligands to cavities within the core of a protein is rapid. *Nature Structural and Molecular Biology*, 3(6):516, 1996.
- [75] Gordon J Berman, William Bialek, and Joshua W Shaevitz. Predictability and hierarchy in drosophila behavior. *Proceedings of the National Academy of Sciences*, 113(42):11943–11948, 2016.
- [76] João Marcelo Lamim Ribeiro, Pablo Bravo, Yihang Wang, and Pratyush Tiwary. Reweighted autoencoded variational bayes for enhanced sampling. *J. Chem. Phys.*, 149(7):072301–072309, 2018.

- [77] Yihang Wang, João Marcelo Lamim Ribeiro, and Pratyush Tiwary. Past–future information bottleneck for sampling molecular reaction coordinate simultaneously with thermodynamics and kinetics. *Nature Communications*, 10(1):3573, 2019.
- [78] Sebastian J Wetzel. Unsupervised learning of phase transitions: From principal component analysis to variational autoencoders. *Physical Review E*, 96(2):022140, 2017.
- [79] Jing Zhang and Ming Chen. Unfolding hidden barriers by active enhanced sampling. *Physical review letters*, 121(1):010601, 2018.
- [80] Omar Valsson, Pratyush Tiwary, and Michele Parrinello. Enhancing important fluctuations: Rare events and metadynamics from a conceptual viewpoint. *Annual Review of Physical Chemistry*, 67(1):159–184, 2016. PMID: 26980304.
- [81] Pratyush Tiwary and Michele Parrinello. A time-independent free energy estimator for metadynamics. *The Journal of Physical Chemistry B*, 119(3):736–742, 01 2015.
- [82] David JC MacKay and David JC Mac Kay. *Information theory, inference and learning algorithms*. Cambridge university press, 2003.
- [83] Pratyush Tiwary and Michele Parrinello. From metadynamics to dynamics. *Phys. Rev. Lett.*, 111(23):230602–230606, 2013.
- [84] Alexander A Alemi, Ian Fischer, Joshua V Dillon, and Kevin Murphy. Deep variational information bottleneck. *Preprint at arXiv:1612.00410*, 2016.
- [85] Andreas Mardt, Luca Pasquali, Hao Wu, and Frank Noé. Vampnets for deep learning of molecular kinetics. *Nat. Comm.*, 9(1):5–12, 2018.
- [86] Wei Chen and Andrew L Ferguson. Molecular enhanced sampling with autoencoders: On-the-fly collective variable discovery and accelerated free energy landscape exploration. *Journal of computational chemistry*, 39(25):2079–2102, 2018.
- [87] Christoph Wehmeyer and Frank Noé. Time-lagged autoencoders: Deep learning of slow collective variables for molecular kinetics. *The Journal of Chemical Physics*, 148(24):241703, 2018.
- [88] Mohammad M. Sultan, Hannah K. Wayment-Steele, and Vijay S. Pande. Transferable neural networks for enhanced sampling of protein dynamics. *Journal of Chemical Theory and Computation*, 14(4):1887–1894, 2018. PMID: 29529369.
- [89] Nigel Goldenfeld. *Lectures on phase transitions and the renormalization group*. CRC Press, 2018.
- [90] Ian Goodfellow, Yoshua Bengio, Aaron Courville, and Yoshua Bengio. *Deep learning*, volume 1. MIT press Cambridge, 2016.

- [91] Bruce J Berne, Michal Borkovec, and John E Straub. Classical and modern methods in reaction rate theory. *J. Phys. Chem.*, 92(13):3711–3725, 1988.
- [92] John A Montgomery Jr, David Chandler, and Bruce J Berne. Trajectory analysis of a kinetic theory for isomerization dynamics in condensed phases. *The Journal of Chemical Physics*, 70(9):4056–4066, 1979.
- [93] Christoph Dellago, Peter G Bolhuis, and David Chandler. On the calculation of reaction rate constants in the transition path ensemble. *The Journal of chemical physics*, 110(14):6617–6625, 1999.
- [94] Omar Valsson, Pratyush Tiwary, and Michele Parrinello. Enhancing important fluctuations: Rare events and metadynamics from a conceptual viewpoint. *Annual Review of Physical Chemistry*, 67(1):159–184, 2016. PMID: 26980304.
- [95] Giovanni Bussi, Alessandro Laio, and Pratyush Tiwary. *Metadynamics: A Unified Framework for Accelerating Rare Events and Sampling Thermodynamics and Kinetics*, pages 1–31. Springer International Publishing, Cham, 2018.
- [96] Joao Marcelo Lamim Ribeiro and Pratyush Tiwary. Towards achieving efficient and accurate ligand-protein unbinding with deep learning and molecular dynamics through rave. *Journal of chemical theory and computation*, 2019.
- [97] Yong Wang, Joao Miguel Martins, and Kresten Lindorff-Larsen. Biomolecular conformational changes and ligand binding: from kinetics to thermodynamics. *Chemical science*, 8(9):6466–6473, 2017.
- [98] Zachary Smith, Debabrata Pramanik, Sun-Ting Tsai, and Pratyush Tiwary. Multi-dimensional spectral gap optimization of order parameters (sgoop) through conditional probability factorization. *J. Chem. Phys.*, 2018.
- [99] Guillaume Bouvignies, Pramodh Vallurupalli, D Flemming Hansen, Bruno E Correia, Oliver Lange, Alaji Bah, Robert M Vernon, Frederick W Dahlquist, David Baker, and Lewis E Kay. Solution structure of a minor and transiently formed state of a t4 lysozyme mutant. *Nature*, 477(7362):111, 2011.
- [100] Marcus D Collins, Gerhard Hummer, Michael L Quillin, Brian W Matthews, and Sol M Gruner. Cooperative water filling of a nonpolar protein cavity observed by high-pressure crystallography and simulation. *Proceedings of the National Academy of Sciences*, 102(46):16668–16671, 2005.
- [101] Pavan Ravindra, Zachary Smith, and Pratyush Tiwary. Automatic mutual information noise omission (amino): generating order parameters for molecular systems. *Mol. Syst. Des. Eng.*, pages –, 2020.
- [102] Christoph Wehmeyer and Frank Noé. Time-lagged autoencoders: Deep learning of slow collective variables for molecular kinetics. *J. Chem. Phys.*, 148(24):241703, 2018.

- [103] D. J. Bicout and Attila Szabo. Electron transfer reaction dynamics in non-debye solvents. *J. Chem. Phys.*, 109(6):2325–2338, 1998.
- [104] Pratyush Tiwary and B. J. Berne. Predicting reaction coordinates in energy landscapes with diffusion anisotropy. *J. Chem. Phys.*, 147(15):152701, 2017.
- [105] Edina Rosta and Gerhard Hummer. Free energies from dynamic weighted histogram analysis using unbiased markov state model. *Journal of chemical theory and computation*, 11(1):276–285, 2015.
- [106] Giovanni Bussi and Michele Parrinello. Accurate sampling using langevin dynamics. *Phys. Rev. E*, 75:056707, May 2007.
- [107] Wei Chen, Hythem Sidky, and Andrew L. Ferguson. Capabilities and limitations of time-lagged autoencoders for slow mode discovery in dynamical systems. *J. Chem. Phys.*, 151(6):064123, 2019.
- [108] Natalia Sánchez De Groot, Alexandros Armaos, Ricardo Graña-Montes, Marion Alriquet, Giulia Calloni, R Martin Vabulas, and Gian Gaetano Tartaglia. Rna structure drives interaction with proteins. *Nature communications*, 10(1):1–13, 2019.
- [109] Samantha M Meyer, Christopher C Williams, Yoshihiro Akahori, Toru Tanaka, Haruo Aikawa, Yuquan Tong, Jessica L Childs-Disney, and Matthew D Disney. Small molecule recognition of disease-relevant rna structures. *Chemical Society Reviews*, 49(19):7167–7199, 2020.
- [110] James P Falese, Anita Donlic, and Amanda E Hargrove. Targeting rna with small molecules: from fundamental principles towards the clinic. *Chemical Society Reviews*, 50(4):2224–2243, 2021.
- [111] Brian J Tucker and Ronald R Breaker. Riboswitches as versatile gene control elements. *Current opinion in structural biology*, 15(3):342–348, 2005.
- [112] Evangelia Lekka and Jonathan Hall. Noncoding rnas in disease. *FEBS letters*, 592(17):2884–2900, 2018.
- [113] Manel Esteller. Non-coding rnas in human disease. *Nature reviews genetics*, 12(12):861–874, 2011.
- [114] Kenneth F Blount and Ronald R Breaker. Riboswitches as antibacterial drug targets. *Nature biotechnology*, 24(12):1558–1564, 2006.
- [115] Colleen M Connelly, Michelle H Moon, and John S Schneekloth Jr. The emerging role of rna as a therapeutic target for small molecules. *Cell chemical biology*, 23(9):1077–1090, 2016.
- [116] Aline Umuhire Juru and Amanda E Hargrove. Frameworks for targeting rna with small molecules. *Journal of Biological Chemistry*, page 100191, 2021.



- [117] Jiri Sponer, Giovanni Bussi, Miroslav Krepl, Pavel Banáš, Sandro Bottaro, Richard A Cunha, Alejandro Gil-Ley, Giovanni Pinamonti, Simón Poblete, Petr Jurečka, et al. Rna structural dynamics as captured by molecular simulations: a comprehensive overview. *Chemical reviews*, 118(8):4177–4338, 2018.
- [118] Sandro Bottaro and Kresten Lindorff-Larsen. Biophysical experiments and biomolecular simulations: A perfect match? *Science*, 361(6400):355–360, 2018.
- [119] Nils G. Walter. The blessing and curse of rna dynamics: past, present, and future. *Methods*, 49(2):85–86, 2009. RNA Dynamics.
- [120] Kevin M Weeks. Advances in rna structure analysis by chemical probing. *Current opinion in structural biology*, 20(3):295–304, 2010.
- [121] Hashim M Al-Hashimi and Nils G Walter. Rna dynamics: it is about time. *Current opinion in structural biology*, 18(3):321–329, 2008.
- [122] Alexander Serganov and Evgeny Nudler. A decade of riboswitches. *Cell*, 152(1-2):17–24, 2013.
- [123] Colleen M. Connelly, Tomoyuki Numata, Robert E. Boer, Michelle H. Moon, Ranu S. Sinniah, Joseph J. Barchi, Adrian R. Ferré-D’Amaré, and John S. Schneekloth. Synthetic ligands for preq1 riboswitches provide structural and mechanistic insights into targeting rna tertiary structure. *Nature Communications*, 10(1):1501, 2019.
- [124] Adam Roth and Ronald R. Breaker. The structural and functional diversity of metabolite-binding riboswitches. *Annual review of biochemistry*, 78:305–334, 2009.
- [125] Krishna C Suddala, Arlie J Rinaldi, Jun Feng, Anthony M Mustoe, Catherine D Eichhorn, Joseph A Liberman, Joseph E Wedekind, Hashim M Al-Hashimi, Charles L Brooks III, and Nils G Walter. Single transcriptional and translational preq1 riboswitches adopt similar pre-folded ensembles that follow distinct folding pathways into the same ligand-bound structure. *Nucleic acids research*, 41(22):10462–10475, 2013.
- [126] Griffin M Schroeder, Debapratim Dutta, Chapin E Cavender, Jermaine L Jenkins, Elizabeth M Pritchett, Cameron D Baker, John M Ashton, David H Mathews, and Joseph E Wedekind. Analysis of a preQ1-I riboswitch in effector-free and bound states reveals a metabolite-programmed nucleobase-stacking spine that controls gene regulation. *Nucleic Acids Research*, 48(14):8146–8164, 06 2020.
- [127] Alessandro Barducci, Giovanni Bussi, and Michele Parrinello. Well-tempered metadynamics: A smoothly converging and tunable free-energy method. *Phys. Rev. Lett.*, 100:020603, Jan 2008.

- [128] Nathan A Siegfried, Steven Busan, Gregory M Rice, Julie A E Nelson, and Kevin M Weeks. Rna motif discovery by shape and mutational profiling (shape-map). *Nature Methods*, 11(9):959–965, 2014.
- [129] Matthew J Smola, Gregory M Rice, Steven Busan, Nathan A Siegfried, and Kevin M Weeks. Selective 2'-hydroxyl acylation analyzed by primer extension and mutational profiling (shape-map) for direct, versatile and accurate rna structure analysis. *Nature Protocols*, 10(11):1643–1669, 2015.
- [130] Michelle H Moon, Thomas A Hilimire, Allix M Sanders, and John S Schneekloth Jr. Measuring rna–ligand interactions with microscale thermophoresis. *Biochemistry*, 57(31):4638–4643, 2018.
- [131] Colleen M. Connelly, Tomoyuki Numata, Robert E. Boer, Michelle H. Moon, Ranu S. Sinniah, Joseph J. Barchi, Adrian R. Ferré-D'Amaré, and John S. Schneekloth. Synthetic ligands for preq1 riboswitches provide structural and mechanistic insights into targeting rna tertiary structure. *Nature Communications*, 10(1):1501, 2019.
- [132] Tycho Marinus, Adam B Fessler, Craig A Ogle, and Danny Incarnato. A novel shape reagent enables the analysis of rna structure in living cells with unprecedented accuracy. *Nucleic acids research*, 49(6):e34–e34, 2021.
- [133] Griffin M Schroeder, Debapratim Dutta, Chapin E Cavender, Jermaine L Jenkins, Elizabeth M Pritchett, Cameron D Baker, John M Ashton, David H Mathews, and Joseph E Wedekind. Analysis of a preq1-i riboswitch in effector-free and bound states reveals a metabolite-programmed nucleobase-stacking spine that controls gene regulation. *Nucleic acids research*, 48(14):8146–8164, 2020.
- [134] Giovanni Pinamonti, Sandro Bottaro, Cristian Micheletti, and Giovanni Bussi. Elastic network models for rna: a comparative assessment with molecular dynamics and shape experiments. *Nucleic acids research*, 43(15):7260–7269, 2015.
- [135] Dazhi Tan, Stefano Piana, Robert M. Dirks, and David E. Shaw. Rna force field with accuracy comparable to state-of-the-art protein force fields. *Proceedings of the National Academy of Sciences*, 115(7):E1346–E1355, 2018.
- [136] Pratyush Tiwary and Michele Parrinello. A time-independent free energy estimator for metadynamics. *The Journal of Physical Chemistry B*, 119(3):736–742, 2015.
- [137] Manel Esteller. Non-coding rnas in human disease. *Nature Reviews Genetics*, 12(12):861—874, 2011.
- [138] Qi Yan, Chengming Zhu, Shouhong Guang, and Xuezhu Feng. The functions of non-coding rnas in rrna regulation. *Frontiers in Genetics*, 10:290, 2019.

- [139] Xuanyu Liu, Yi Ma, Kunlun Yin, Wenke Li, Wen Chen, Yujing Zhang, Changsheng Zhu, Tianjiao Li, Bianmei Han, Xuewen Liu, Shuiyun Wang, and Zhou Zhou. Long non-coding and coding rna profiling using strand-specific rna-seq in human hypertrophic cardiomyopathy. *Scientific Data*, 6(1):90, 2019.
- [140] Colleen M. Connelly, Michelle H. Moon, and John S. Schneekloth. The emerging role of rna as a therapeutic target for small molecules. *Cell Chemical Biology*, 23(9):1077–1090, 2016.
- [141] Warren L DeLano and Sarina Bromberg. Pymol user’s guide. *DeLano Scientific LLC*, 629, 2004.
- [142] William Humphrey, Andrew Dalke, and Klaus Schulten. Vmd: visual molecular dynamics. *Journal of molecular graphics*, 14(1):33–38, 1996.
- [143] Laura R Ganser, Megan L Kelly, Daniel Herschlag, and Hashim M Al-Hashimi. The roles of structural dynamics in the cellular functions of rnas. *Nature reviews Molecular cell biology*, 20(8):474–489, 2019.
- [144] Matthew Halvorsen, Joshua S Martin, Sam Broadaway, and Alain Laederach. Disease-associated mutations that alter the rna structural ensemble. *PLoS genetics*, 6(8):e1001074, 2010.
- [145] Jascha Sohl-Dickstein, Eric Weiss, Niru Maheswaranathan, and Surya Ganguli. Deep unsupervised learning using nonequilibrium thermodynamics. In *International Conference on Machine Learning*, pages 2256–2265. PMLR, 2015.
- [146] Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. *arXiv preprint arXiv:2006.11239*, 2020.
- [147] Ulrich HE Hansmann. Parallel tempering algorithm for conformational studies of biological molecules. *Chemical Physics Letters*, 281(1-3):140–150, 1997.
- [148] Cameron Abrams and Giovanni Bussi. Enhanced sampling in molecular dynamics using metadynamics, replica-exchange, and temperature-acceleration. *Entropy*, 16(1):163–199, 2014.
- [149] Robert Abel, Lingle Wang, Edward D Harder, BJ Berne, and Richard A Friesner. Advancing drug discovery through enhanced free energy calculations. *Accounts of chemical research*, 50(7):1625–1632, 2017.
- [150] Andrew J Ballard and Christopher Jarzynski. Replica exchange with nonequilibrium switches. *Proceedings of the National Academy of Sciences*, 106(30):12224–12229, 2009.
- [151] Simon Trebst, Matthias Troyer, and Ulrich HE Hansmann. Optimized parallel tempering simulations of proteins. *The Journal of chemical physics*, 124(17):174903, 2006.

- [152] Walter Nadler and Ulrich HE Hansmann. Optimizing replica exchange moves for molecular dynamics. *Physical Review E*, 76(5):057102, 2007.
- [153] Jaegil Kim, Thomas Keyes, and John E Straub. Generalized replica exchange method. *The Journal of chemical physics*, 132(22):224107, 2010.
- [154] John D Chodera and Michael R Shirts. Replica exchange and expanded ensemble simulations as gibbs sampling: Simple improvements for enhanced mixing. *The Journal of chemical physics*, 135(19):194110, 2011.
- [155] Alejandro Gil-Ley and Giovanni Bussi. Enhanced conformational sampling using replica exchange with collective-variable tempering. *Journal of chemical theory and computation*, 11(3):1077–1085, 2015.
- [156] Ian Goodfellow, Yoshua Bengio, and Aaron Courville. *Deep Learning*. MIT Press, 2016. <http://www.deeplearningbook.org>.
- [157] Emilio Gallicchio, Michael Andrec, Anthony K Felts, and Ronald M Levy. Temperature weighted histogram analysis method, replica exchange, and transition paths. *The Journal of Physical Chemistry B*, 109(14):6722–6731, 2005.
- [158] Olaf Ronneberger, Philipp Fischer, and Thomas Brox. U-net: Convolutional networks for biomedical image segmentation. In *International Conference on Medical image computing and computer-assisted intervention*, pages 234–241. Springer, 2015.
- [159] Yuxin Wu and Kaiming He. Group normalization. In *Proceedings of the European Conference on Computer Vision (ECCV)*, September 2018.
- [160] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. In *Advances in neural information processing systems*, pages 5998–6008.
- [161] Pascal Vincent. A connection between score matching and denoising autoencoders. *Neural computation*, 23(7):1661–1674, 2011.
- [162] William L Jorgensen and Julian Tirado-Rives. Potential energy functions for atomic-level simulations of water and organic and biomolecular systems. *Proceedings of the National Academy of Sciences*, 102(19):6665–6670, 2005.
- [163] Jing Huang, Sarah Rauscher, Grzegorz Nawrocki, Ting Ran, Michael Feig, Bert L De Groot, Helmut Grubmüller, and Alexander D MacKerell. Charmm36m: an improved force field for folded and intrinsically disordered proteins. *Nature methods*, 14(1):71–73, 2017.
- [164] Virgiliu Botan, Ellen HG Backus, Rolf Pfister, Alessandro Moretto, Marco Crisma, Claudio Toniolo, Phuong H Nguyen, Gerhard Stock, and Peter Hamm. Energy transport in peptide helices. *Proceedings of the National Academy of Sciences*, 104(31):12749–12754, 2007.

- [165] Mithun Biswas, Benjamin Lickert, and Gerhard Stock. Metadynamics enhanced markov modeling of protein dynamics. *The Journal of Physical Chemistry B*, 122(21):5508–5514, 2018.
- [166] Shams Mehdi, Dedi Wang, Shashank Pant, and Pratyush Tiwary. Accelerating all-atom simulations and gaining mechanistic understanding of biophysical systems through state predictive information bottleneck, 2021.
- [167] Sylvie Bannwarth and Anne Gatignol. HIV-1 TAR RNA: The target of molecular interactions between the virus and its host. *Current HIV Research*, 3(1):61–71, January 2005.
- [168] Ursula Schulze-Gahmen and James H. Hurley. Structural mechanism for HIV-1 TAR loop recognition by tat and the super elongation complex. *Proceedings of the National Academy of Sciences*, 115(51):12973–12978, December 2018.
- [169] Laura R. Ganser, Megan L. Kelly, Daniel Herschlag, and Hashim M. Al-Hashimi. The roles of structural dynamics in the cellular functions of RNAs. *Nature Reviews Molecular Cell Biology*, 20(8):474–489, June 2019.
- [170] Jiří Šponer, Giovanni Bussi, Miroslav Krepl, Pavel Banáš, Sandro Bottaro, Richard A. Cunha, Alejandro Gil-Ley, Giovanni Pinamonti, Simón Poblete, Petr Jurečka, Nils G. Walter, and Michal Otyepka. RNA structural dynamics as captured by molecular simulations: A comprehensive overview. *Chemical Reviews*, 118(8):4177–4338, January 2018.
- [171] Christina Bergonzo, Niel M Henriksen, Daniel R Roe, and 3rd Cheatham, Thomas E. Highly sampled tetranucleotide and tetraloop motifs enable evaluation of common rna force fields. *RNA (New York, N.Y.)*, 21(9):1578–1590, 09 2015.
- [172] Christopher Jarzynski. Rare events and the convergence of exponentially averaged work values. *Physical Review E*, 73(4):046105, 2006.
- [173] Danilo Rezende and Shakir Mohamed. Variational inference with normalizing flows. In *International conference on machine learning*, pages 1530–1538. PMLR, 2015.
- [174] Giorgio Visani, Enrico Bagli, Federico Chesani, Alessandro Poluzzi, and Davide Capuzzo. Statistical stability indices for lime: Obtaining reliable explanations for machine learning models. *Journal of the Operational Research Society*, pages 1–11, 2020.
- [175] Filip Karlo Došilović, Mario Brčić, and Nikica Hlupić. Explainable artificial intelligence: A survey. In *2018 41st International convention on information and communication technology, electronics and microelectronics (MIPRO)*, pages 0210–0215. IEEE, 2018.

- [176] Alex Davies, Petar Veličković, Lars Buesing, Sam Blackwell, Daniel Zheng, Nenad Tomašev, Richard Tanburn, Peter Battaglia, Charles Blundell, András Juhász, et al. Advancing mathematics by guiding human intuition with ai. *Nature*, 600(7887):70–74, 2021.
- [177] Floris Takens. Detecting strange attractors in turbulence. In *Dynamical systems and turbulence, Warwick 1980*, pages 366–381. Springer, 1981.
- [178] Krzysztof Barański, Yonatan Gutman, and Adam Śpiewak. A probabilistic takens theorem. *Nonlinearity*, 33(9):4940, 2020.
- [179] Robert M Clegg. Fluorescence resonance energy transfer. *Current opinion in biotechnology*, 6(1):103–110, 1995.
- [180] M Abraham, B Hess, D van der Spoel, and E Lindahl. Gromacs reference manual, version 2016.4. *Department of Biophysical Chemistry, University of Groningen*, 2016.
- [181] Mark James Abraham, Teemu Murtola, Roland Schulz, Szilárd Páll, Jeremy C. Smith, Berk Hess, and Erik Lindahl. Gromacs: High performance molecular simulations through multi-level parallelism from laptops to supercomputers. *SoftwareX*, 1-2:19 – 25, 2015.
- [182] Gareth A. Tribello, Massimiliano Bonomi, Davide Branduardi, Carlo Camilloni, and Giovanni Bussi. Plumed 2: New feathers for an old bird. *Computer Physics Communications*, 185(2):604 – 613, 2014.
- [183] Berk Hess, Henk Bekker, Herman J. C. Berendsen, and Johannes G. E. M. Fraaije. Lincs: A linear constraint solver for molecular simulations. *Journal of Computational Chemistry*, 18(12):1463–1472, 1997.
- [184] Jing Huang, Sarah Rauscher, Grzegorz Nawrocki, Ting Ran, Michael Feig, Bert L De Groot, Helmut Grubmüller, and Alexander D MacKerell. Charmm36m: an improved force field for folded and intrinsically disordered proteins. *Nature methods*, 14(1):71–73, 2017.
- [185] M. Parrinello and A. Rahman. Crystal structure and pair potentials: A molecular-dynamics study. *Phys. Rev. Lett.*, 45:1196–1199, Oct 1980.
- [186] William G Hoover. Canonical dynamics: Equilibrium phase-space distributions. *Physical review A*, 31(3):1695, 1985.
- [187] Massimiliano Bonomi. Promoting transparency and reproducibility in enhanced molecular simulations. *Nature methods*, 16(8):670–673, 2019.
- [188] Christina Bergonzo, Niel M. Henriksen, Daniel R. Roe, and Thomas E. Cheatham. Highly sampled tetranucleotide and tetraloop motifs enable evaluation of common RNA force fields. *RNA*, 21(9):1578–1590, June 2015.