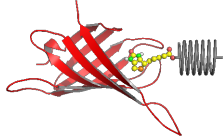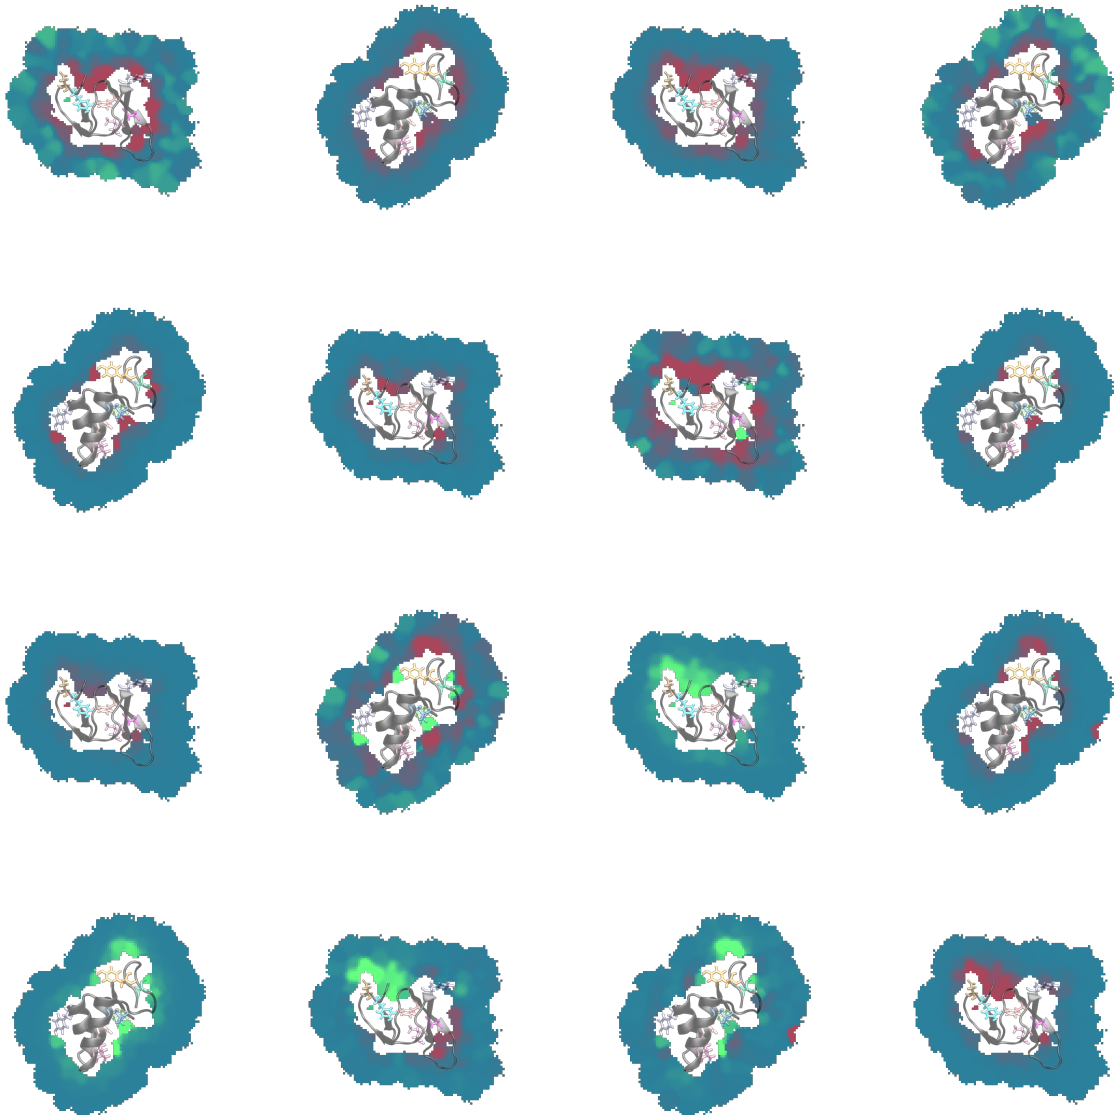# Spatially Resolved Hydration Statistical Mechanics at Biomolecular Surfaces from Atomistic Simulations

Leonard P. Heinz

Spatially Resolved Hydration Statistical Mechanics at Biomolecular Surfaces from Atomistic Simulations

Dissertation

for the award of the degree

Doctor rerum naturalium

of the Georg-August-Universität Göttingen

within the doctoral program

Physics of Biological and Complex Systems

of the Georg-August University School of Science (GAUSS)

submitted by

Leonard P. Heinz

from Karlsruhe (Germany)

Göttingen, 2020

Thesis Committee

1. Prof. Dr. Helmut Grubmüller
   Max Planck Institute for Biophysical Chemistry
   Department of Theoretical and Computational Biophysics

2. Prof. Dr. Marcus Müller
   Georg-August-Universität Göttingen
   Institute for Theoretical Physics
   Statistical Physics of Soft Matter

3. Prof. Dr. Jörg Enderlein
   Georg-August-Universität Göttingen
   Third Institute of Physics (Biophysics)

Members of the Examination Board

Prof. Dr. Helmut Grubmüller (reviewer)
Max Planck Institute for Biophysical Chemistry
Department of Theoretical and Computational Biophysics

Prof. Dr. Marcus Müller (reviewer)
Georg-August-Universität Göttingen
Institute for Theoretical Physics
Statistical Physics of Soft Matter

Prof. Dr. Jörg Enderlein
Georg-August-Universität Göttingen
Third Institute of Physics (Biophysics)

Prof. Dr. Claudia Steinem
Georg-August-Universität Göttingen
Institute of Organic and Biomolecular Chemistry
Biomolekulare Chemie

Prof. Dr. Tim Salditt
Georg-August-Universität Göttingen
Institute for X-Ray Physics

Prof. Dr. Holger Stark
Max Planck Institute for Biophysical Chemistry
Department of Structural Dynamics

Date of the oral examination: December 15, 2020

Affidavit

I confirm that the work of the thesis was solely undertaken by myself and that no help was provided from other sources as those allowed. I confirm that all sections of the thesis that use quotes or describe an argument or concept developed by another author are referenced.

Göttingen, November 3, 2020

_____

Leonard P. Heinz

Abstract

The statistical mechanics of the first few hydration layers is vital for many biophysical processes, such as protein folding and unfolding, protein function, lipid bilayer self-assembly, and ligand binding. These processes are governed by a fine-tuned free energy balance of competing enthalpy and entropy contributions. Despite extensive experimental and theoretical efforts, the molecular mechanisms of solvent-related free energy contributions are often elusive, especially at topologically and chemically heterogeneous surfaces like proteins or lipid bilayers. To better understand, e.g., the effects of individual amino acids on the solvent-related free energy contributions, a spatial resolution of both solvent enthalpy and entropy is necessary.

Whereas the enthalpy component can be readily calculated from a molecular dynamics force field, sampling the entropy contribution presents a significant challenge. In the first part of this thesis, I therefore develop and present a new method, Per|Mut, to calculate spatially resolved solvent entropies from atomistic simulations. The method uses a permutation reduction to increase sampling by the Gibbs factor $N!$ without changing the physics. In addition, Per|Mut employs a third-order mutual information expansion to decompose the solvent entropy into physically interpretable contributions from individual molecules as well as from two- and three-body correlations. The method yielded accurate entropies for test systems with argon and solvated alkanes. When applied to the solvation statistical mechanics of hydrated octanol and the protein crambin, the method revealed the local effects of individual chemical groups or side chains on the solvent entropy. Comparing native-fold crambin and a molten-globule-like conformation, I identified a strongly stabilizing solvent entropy contribution of almost $500 \, \text{kJ} \cdot \text{mol}^{-1}$ to the total free energy difference of $53 \, \text{kJ} \cdot \text{mol}^{-1}$. Remarkably, more than half of the solvent entropy contribution arose from induced water correlations. In addition to understanding protein stability, Per|Mut could prove useful to understand and predict the solvent contributions to ligand binding, which is especially relevant in the realm of computational drug design.

The energetics of lipid headgroup dehydration is furthermore expected to play a major role in the free energy landscape of membrane fusion, a process vital for, e.g., exocytosis during synaptic transmission or fusion of enveloped viruses into the host plasma membrane. As revealed by collaborators, pre-fusion lipid membranes at $< 1 \, \text{nm}$ distances experience a divalent cation-independent, metastable, protein-free fusion intermediate, characterized by local membrane thickening. In the second part of this thesis, I carry out molecular dynamics simulations of double-membrane systems to identify the molecular causes of the structural changes. Through non-equilibrium simulations and response-time analyses, I demonstrate that the lipid bilayer thickening results from an electrostatically favorable lateral area shrinkage, attributed to dehydration-driven lipid headgroup tilting.

List of Abbreviations

2PT	2-Phase Thermodynamics
DoS	density of states
GIST	Grid Inhomogeneous Solvation Theory
IST	Inhomogeneous Solvation Theory
kNN	k-nearest neighbor (in the sense of kNN estimator/method)
MD	molecular dynamics
MI	mutual information
MIE	mutual information expansion
NMSLib	Non-Metric Space Library
NPT	constant number of particles, constant pressure, constant temperature
NVT	constant number of particles, constant volume, constant temperature
Per|Mut	Permutation Reduction and Mutual Information Expansion
QH	quasiharmonic
TI	Thermodynamic Integration

List of Recurring Mathematical Symbols and Constants

k_B	Boltzmann constant[a] $1.380649 \cdot 10^{-23}$ J·K^{-1}
β	$1/(k_B T)$ with temperature T
h	Planck constant[b] $6.62607015 \cdot 10^{-34}$ J·s
$\hbar$	reduced Planck constant $\hbar = h/(2\pi)$

---

[a]https://physics.nist.gov/cgi-bin/cuu/Value?k
[b]https://physics.nist.gov/cgi-bin/cuu/Value?h

Contents

1 | Introduction

Water is the universal solvent in all living things, and without it, life as we know it would not be possible [1–4]. Virtually all cellular processes occur in an aqueous environment, and essential building blocks of life like deoxyribonucleic acid (DNA), ribonucleic acid (RNA), phospholipids, and proteins are typically dissolved in water. Proteins require water to perform their functions, and if the water is replaced by another solvent, most proteins unfold [2, 5]. As such, even the most essential processes like the transcription of genetic information from DNA to RNA and further translation to a protein sequence cannot take place without water, as proteins are required in each step. Water is also vital for the self-assembly of lipid bilayers [6] and, therefore, required for the formation of cellular membranes.

In spite of its importance, the central role of water is not always appreciated in research [2]. In the field of MD (molecular dynamics) simulations, water is often seen as an uninteresting additional burden to simulate, and many times, the information about the water molecules is discarded before analysis. However, the impact of water in MD simulations is underlined by the fact that *in vacuo* simulations fail to accurately reproduce ensemble properties, such as the radius of gyration, the root mean square deviation, and the root mean square fluctuation of protein motion. Similarly, it is very hard to construct accurate implicit solvent models that take all water-effects into account [7, 8], and also slight changes in the force field of explicit water models can result in inaccurate ensemble properties, particularly for intrinsically disordered proteins [9]. Although much work has been done already [10–18], the thermodynamics of the hydration shell, i.e., the closest few water layers around a protein, lipid bilayer, or any other solvent, still deserves a closer look.

Particularly, the effect of structurally and chemically heterogeneous surfaces, such as proteins, on the solvent thermodynamics is often enigmatic. In the first part of this thesis, I will therefore present a new method to calculate the spatially resolved free energy contributions of a solvation shell from atomistic simulations, decomposed into individual enthalpy and entropy terms. The spatial resolution and the decomposition allow the characterization of the solvent contribution to free energy differences of processes like protein unfolding or ligand binding on the molecular level. In the second albeit much smaller part of the thesis, the dehydration-related thickening of close-contact pre-fusion lipid membranes will be examined using MD simulations.

1 Solvent Thermodynamics

The folding thermodynamics of proteins is governed by its free energy landscape. From calorimetric experiments, the free energy difference between folded and unfolded protein states has been determined to be in the order of a few kJ·mol^{-1} to a few 10 kJ·mol^{-1}, mainly depending on the protein size [19–22]. This result is astonishing, as the free energy difference is comparable to the interaction energies of just a handful of hydrogen bonds [23], of which typically hundreds exist within a protein (many of which are broken upon unfolding) and even more between the protein and the solvent water. The stark discrepancy between the expected interaction energy differences and the measured free energy differences showcases the delicate balance between large but almost completely compensating enthalpic (ΔH) and entropic ($-T\Delta S$) contributions to the overall folding free energy difference of a protein ($G^{(\text{unfolded})} - G^{(\text{folded})} = \Delta G = \Delta H - T\Delta S$) [20]. Small perturbations like the mutation of an amino acid or a change of temperature can shift this balance, causing the protein to unfold. One particularly familiar example of such an effect is heat denaturation, where a protein unfolds at elevated temperatures, e.g., during "cooking". Here, the balance is shifted towards the more disordered unfolded state, as the entropy contribution $-T\Delta S$ increases in magnitude for large temperatures T and eventually overpowers the competing enthalpy term. Since heat denaturation is entropy-driven, it is easy and intuitive to argue that cold temperatures should stabilize the more ordered native fold due to a smaller entropic contribution.

Contrary to this argument, proteins also unfold at cold temperatures, in an effect known as cold denaturation[1] [20, 24, 25]. The contradiction arises, because the above argument only considers the thermodynamics of the protein, whereas the solvent that makes up a substantial part of the overall thermodynamic system has been ignored. To explain cold denaturation, both free energy contributions of the entire system (including solvent) need to be considered. The enthalpy and entropy differences can be expressed as

$$\Delta H(T) = \Delta H(T_h) + \Delta C_P(T - T_h)$$

and

$$\Delta S(T) = \frac{\Delta H(T_h)}{T_h} + \Delta C_P \log\left(\frac{T}{T_c}\right),$$

respectively [20]. Here, T_h is the heat denaturation temperature and $\Delta H(T_h)$ is the enthalpy difference at T_h. The heat-capacity difference $\Delta C_P = C_P^{(\text{unfolded})} - C_P^{(\text{folded})}$ has been assumed to be approximately constant within the relevant temperature range [20]. The resulting free energy difference (Fig. 1.1) shows two denaturation temperatures for protein-typical parameters, indicated by the two roots of the function $\Delta G(T)$. As discussed above, heat

---

[1]For most proteins, the cold-denaturation temperature is well below the freezing point of water, making the effect a far less common experience compared to the well-known heat denaturation.
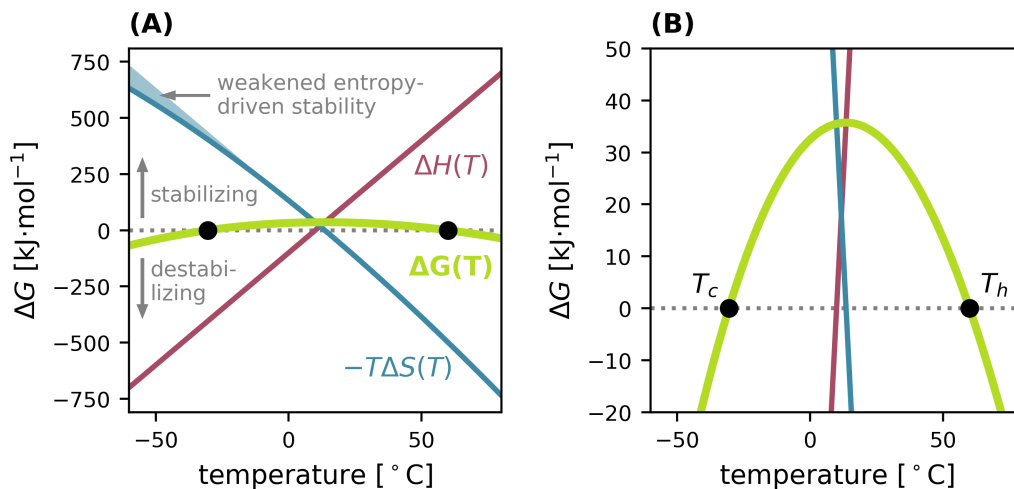
Figure 1.1: Folding free energy difference $\Delta G(T)$, its enthalpy ($\Delta H(T)$) and entropy ($-T\Delta S(T)$) contributions for typical protein values of $T_h = 60\,°\text{C}$, $\Delta H(T_h) = 500\,\text{kJ·mol}^{-1}$, and $\Delta C_P = 10\,\text{kJ·mol}^{-1}\text{·K}^{-1}$, as determined from calorimetry experiments [20, 26]. The shaded area in (A) denotes the weakened entropy-driven stability relative to the opposing $\Delta H(T)$ term. (B) is a zoomed-in version of (A).

denaturation at T_h is entropy-driven. Contrarily, cold denaturation (here at $T_c \approx -30\,°\text{C}$) occurs due to a strongly destabilizing (negative) enthalpic contribution at cold temperatures. Here, the opposing entropy contribution stabilizes the native fold ($-T\Delta S > 0$), but its stabilizing effect is weakened relative to the enthalpy term, as indicated by the shaded area in Fig. 1.1A. Since the protein entropy is always larger in the unfolded state compared to the native fold, the stabilizing but weakened effect of entropy at cold temperatures needs to be solvent-related. Understanding hydration thermodynamics is therefore essential to better understand the physics of protein folding and protein stability. Indeed, cold denaturation has been attributed to a temperature-induced weakening of the hydrophobic effect [20], which is known to be a major driving force in protein stability and folding [27–30] and is mainly driven by solvent entropy [29, 31–33].

The hydrophobic effect describes the propensity of non-polar solutes, such as some amino acids of a protein, to form aggregates that minimize the overall interface with the aqueous solution [28]. Its explanation on the molecular level has been the subject of long but controversial discussions [34], and many of its molecular aspects are still elusive [20, 35].

Frank and Evans [36] coined the well-known iceberg model, according to which hydrophobicity results from a cage of highly structured ice-like water that forms around a non-polar solute. The ice-like structure would have a low, unfavorable entropy, and thus the system would minimize its free energy by minimizing the solvated surface due to aggregation. Indeed, multiple studies found the hydration shell to be more ordered than bulk water [37, 38].

However, Yu and Karplus [39] and Ben-Naim [40] argued that the entropy loss from a more

ordered hydration shell, should it exist, is exactly compensated by the corresponding change of solvent-solvent interaction energies. They, therefore, concluded that solute-induced water re-orientations have a net-zero effect on the free energy and, therefore, cannot explain the hydrophobic effect. In particular, Yu and Karplus [39] showed analytically that the enthalpy and entropy contributions from ensemble averages over the solvent-solvent interaction energy $\langle U_{ss} \rangle_{\mathcal{H}}$ compensate each other exactly. Hence, solvent-solvent interactions contribute to the solvation free energy not directly but rather indirectly, as they are part of the Hamiltonian $\mathcal{H}$ that defines the thermodynamic ensemble itself. The proof by Yu and Karplus [39] alone is, therefore, not sufficient to disprove the iceberg hypothesis.

Ben-Naim [40] imagined the hydration as a two-step process. First, a water cavity is formed to make room for the solute, while a water order parameter is held fixed. Only then the order-constraint is released, and the water molecules are allowed to reorganize and adapt to the solute. Using thermodynamic and statistical mechanical arguments, he then showed that the solvation free energy is entirely determined by the cavity-creation sub-process, which is entropically unfavorable, whereas the solvent reorganization has a net-zero effect. He, therefore, inferred that the hydrophobic effect must be entirely explained by the cavity-creation entropy and that solvent reorganization processes, like the formation of "icebergs", are irrelevant for the solvation free energy.

This approach was criticized by Lee [41], who argued that a single arbitrary "order parameter" and even a finite set of such parameters are not sufficient to fully define the solvent organization and that the solvent generally also reorganizes during the first cavity-creation step. For instance, the probability density function of a solvent that contains a solute cannot be identical to that of a pure solvent system. Nevertheless, the order parameter can be forced to be identical in both cases. To rigorously identify solvation (sub-)processes with compensating free energy contributions, Lee [41] derived that exact enthalpy-entropy cancellation occurs if and only if the free energy change ΔG is identical to the average interaction energy between solute and solvent $\langle U_{ps} \rangle$. Since for cavity-creation, i.e., the solvation of a hard sphere, $\Delta G \neq 0$, but $\langle U_{ps} \rangle = 0$, it is a non-compensating process. For the second step, however, enthalpy and entropy contributions only approximately compensate each other for idealized solvents with only weakly attractive van-der-Waals-like solvent interactions.

For heterogeneous and typically charged proteins, however, such a dissection into physically meaningful compensating and non-compensating sub-processes is generally not possible. To better understand the solvent thermodynamics involved in protein-specific processes like protein folding, protein stability or ligand binding on the molecular level and the effects of each individual amino acid, a spatially resolved decomposition into physically meaningful enthalpy and entropy contributions is necessary.

2 Entropy Calculation

Although solvent free energies, enthalpies and entropies are fundamentally defined as non-local ensemble averages, spatially resolved maps can be obtained by exploiting the locality of the involved interaction types, Lennard-Jones interactions [42] and Coulomb interactions. The enthalpy can be straight-forwardly calculated on a three-dimensional grid as averages of the interaction energies, obtained directly from an MD force field [10, 11, 43]. The challenge, however, lies in the estimation of the entropy contribution, which is defined in terms of the solvent probability density function $\varrho(\{\boldsymbol{q}_i\}, \{\boldsymbol{p}_i\})$ as

$$S = -k_B \int \frac{d\boldsymbol{q}^N \boldsymbol{p}^N}{N! h^{6N}} \varrho(\{\boldsymbol{q}_i\}, \{\boldsymbol{p}_i\}) \log \varrho(\{\boldsymbol{q}_i\}, \{\boldsymbol{p}_i\}), \tag{2.1}$$

where the vectors $\{\boldsymbol{q}_i\}$ describe the configurational (3 translational, 3 orientational) coordinates of N water molecules, $\{\boldsymbol{p}_i\}$ are the respective (linear and angular) momenta, and h is the Planck constant. Whereas the kinetic contribution can be obtained analytically by carrying out the momentum-integral in eq. (2.1), the remaining configurational contribution

$$S_{\text{conf}} \propto -k_B \int \frac{d\boldsymbol{q}^N}{N!} \varrho(\{\boldsymbol{q}_i\}) \log \varrho(\{\boldsymbol{q}_i\}) \tag{2.2}$$

is notoriously difficult to sample. Here, the problem is twofold:

First, the configuration space for each, translational and rotational degrees of freedom is $3N$-dimensional, which implies that even for just 1000 water molecules, a 3000-dimensional space needs to be sampled to obtain translational or rotational entropies. Due to the "curse of dimensionality", which implies that for increasing dimensions an exponentially increasing number of data points (e.g., frames of an atomistic simulation) is needed to obtain a constant sampling density (same mean distance between data points), sufficient sampling poses a significant challenge. Albeit also large and high-dimensional, the configuration space of proteins is largely inaccessible due to stereochemical constraints. However, the configuration space for a diffusive solvent system is almost entirely accessible, resulting in even more severe sampling problems in comparison.

Secondly, volume of the configuration space scales with N^N, rather than with e^N, as it would be expected for an extensive thermodynamic quantity like entropy. The large volume renders an attempt to numerically sample the configuration space impossible, even if the system was small and the aforementioned curse of dimensionality would be overcome. The problem arises, because the concept of configuration space inherently treats particles as distinguishable, e.g., by assigning labels "1" to "N" to each of them. Due to the permutation symmetry, each microstate is counted redundantly $N!$ times, leading to a configuration space volume that is $N!$ times larger than physically necessary.

In an analytical treatment, the overcounting problem is, of course, solved by the Gibbs

factor $N!$ [44]. For numerical sampling, however, the Gibbs factor cannot be straightforwardly applied, as it requires sampling of the configuration space $\int d\boldsymbol{x}^N \propto N^N$ before the result is divided by $N!$.

In spite of these problems, several methods have been developed to calculate spatially resolved solvent entropies, of which the most established ones are the Grid Inhomogeneous Solvation Theory (GIST) [10, 43] (see section 12) and 3D 2-Phase Thermodynamics (2PT) [11] (see section 13).

GIST estimates the local entropy for each voxel of a three-dimensional grid using an expansion into a single-body entropy, which accounts for the mobility of the individual water molecules and multi-body water correlation terms. A particular advantage of this approach is, therefore, the fact that each term allows for an intuitive physical interpretation of the entropy change. As GIST is grid-based, it is agnostic towards the particle identity and, therefore, does not suffer from the overcounting problem, however, the "curse of dimensionality" results in severe sampling problems. Therefore, GIST is usually limited to the single-body term. Since water reorganization and, in particular, water-water correlations are expected to play a major role in hydrophobicity, this approach is not well-suited to study the solvation thermodynamics of protein folding and stability.

3D 2PT estimates the local entropy by treating the density of states of each water molecule as a superposition of a solid-like harmonic component and a hard-sphere gas, for each of which the entropy is known. The method was shown to yield accurate entropies for bulk water [45], but it is unclear how well surface water can be described as an interpolation between the two components.

To overcome these shortcomings, I will here develop and apply a nonparametric method ("Per|Mut") based on a **per**mutation reduction [46, 47], which solves the overcounting-problem and therefore increases the sampling density by the Gibbs factor $N!$, and a **mut**ual information entropy expansion [4, 48–50] (see section 16.2.1), which alleviates the "curse of dimensionality" and is akin to the expansion used in GIST but routinely captures up to three-body correlations.

In Per|Mut, the configurational entropy

$$S_{\text{conf}} = S_{\text{trans}} + S_{\text{rot}} - I_{\text{trans-rot}}$$

is split between the entropy contributions that arise from translational and rotational degrees of freedom and a mutual information (MI) term $I_{\text{trans-rot}}$ that accounts for the correlation between translational and rotational water motions. To calculate these entropy terms from an MD trajectory, first a permutation reduction is applied, which increases the sampling by the Gibbs factor $N!$ without changing the physics, as described in section 14. Here, the permutation symmetry of the identical water molecules is exploited by mapping the trajectory into a subregion of the configuration space, smaller by $N!$. Although the physics remains
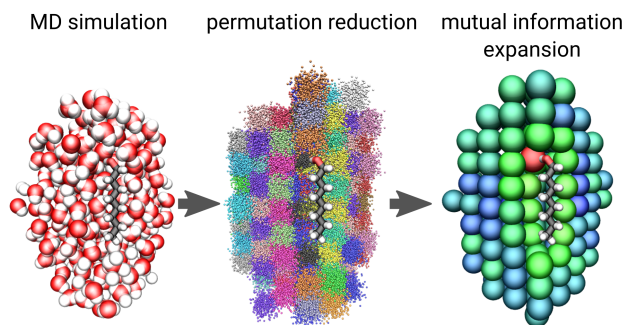
Figure 2.1: Flow chart of the method Per|Mut, visualized for the example of hydrated octanol. In the left image, water molecules are shown with red and white spheres. The center image shows the permutation-localized water molecules in different colors. The right image maps the per-molecule entropy on the water molecules. Note that the largest entropy loss occurs at the red water molecule, which forms a hydrogen bond to the octanol hydroxyl group.

unaltered, this results in a "localization" of the water molecules in real space, which then only fluctuate within a small region, as shown in Fig. 2.1, rather than diffusing through the simulation box.

In the next step, translational and rotational entropies are calculated from the permutationally reduced trajectory. Because the entropy integral in eq. (2.2) in $3N$ dimensions cannot be computed directly due to the "curse of dimensionality", a third-order mutual information expansion (MIE) [4, 48–50] (see section 16.2.1) is used to expand the full high-dimensional integral into multiple low-dimensional integrals over marginal distributions, which can be calculated numerically. The third-order expansion decomposes the entropy into the sum of single-molecule entropies (first order), which ignores all correlations and, therefore, overestimates the entropy, and second- and third-order terms, which account for the entropy loss due to correlations between molecule pairs and triples, respectively. To obtain these MI terms, the method uses a k-nearest-neighbor (kNN) estimator (see section 16.2.2), which estimates the density at each sample point by finding the k closest neighboring sample points and dividing by the volume of a ball that encloses the points.

Through the MI expansion, Per|Mut yields an entropy decomposition into contributions from single molecules, as well as from pair- and triple-correlations of permutationally localized water molecules. The method therefore allows for a spatial entropy representation, as shown in Fig. 2.1. Furthermore, it provides a physical interpretation of spatially resolved entropy changes on the molecular level by distinguishing between entropy losses due to reduced single-molecule mobility (which usually occurs close to protein charges) and due to increased water-water correlations (as is usually the case close to non-polar parts of a solute).

Per|Mut has been thoroughly tested on analytical test distributions and test systems, such as an argon gas and solvated octanol. When compared with reference entropies for solvated

alkanes, the results were on average accurate within $0.2\,\mathrm{J\cdot mol^{-1}\cdot K^{-1}}$ per water molecule for a solvation shell of 100 water molecules. The entropy error is equivalent to a free energy contribution of $0.06\,\mathrm{kJ\cdot mol^{-1}}$ or $0.024\,k_B T$ per water molecule at $300\,\mathrm{K}$. When applied to the hydration shells of the protein crambin, the solvent free energy effects of each amino acid as well as the effects of the protein shape on the solvent enthalpy and entropy were identified. Comparison of the solvent thermodynamics of the native crambin fold and that of an unfolded conformation revealed a strongly stabilizing effect of the solvent entropy, largely caused by induced water-correlations, which is consistent with the idea of solvent-driven protein stabilization discussed in the context of cold denaturation.

The presented method is, furthermore, relevant for the assessment of ligand binding affinities, as a ligand needs to displace surface water from a binding pocket in order to bind. In this context, Per|Mut might also prove useful in the realm of computational protein design.

My implementation of Per|Mut is available for download as a python package[2].

3 Membrane Thickening

The fusion of two lipid membranes is a universal process in living organisms and essential, e.g., for exocytosis during synaptic transmission, for the merger of sperm and oocyte, or for fusion of enveloped viruses into the host plasma membrane [51, 52]. The exact mechanism of membrane fusion, specifically its energetics, is matter of an ongoing debate, however, some intermediate states along the fusion pathway are universally agreed upon [51, 53–56].

First, both bilayers are brought into loose contact through the sequential assembly of SNARE proteins (soluble N-ethylmaleimide-sensitive factor activating protein receptor), which are anchored in the opposing leaflets and pull the bilayers together against repulsive electrostatic, hydration, and undulation forces. Next, a tight-docking state is reached, in which the proximal leaflets are potentially fully dehydrated, but still retain their integrity. In the actual fusion step, the two inner leaflets merge, forming a fusion stalk, which eventually extends into a merging of the full membranes.

As shown by Yavuz et al. [57], the tight-docking state is stable and SNARE-independent, which indicates that the repulsive forces have already been overcome or are compensated by an unknown attractive interaction. In a collaborative manuscript included below, Witkowska and Jahn have furthermore shown experimentally that the tight-docking state is characterized by an increased bilayer thickness. Although previous computational and experimental studies also reported a thickening of close or partially dehydrated lipid bilayers [58–63], no molecular-level explanation is available, to the best of my knowledge. In order to understand the energetics of the pre-fusion membrane states, it might thus be necessary to understand the associated bilayer thickening.

To reveal the underlying structural changes and molecular causes, I carried out all-atom

---

[2] https://gitlab.gwdg.de/lheinz/hydration_entropy

MD simulations of double-membrane systems at decreasing mutual distances. Indeed, the simulations reveal a thickening for distances below 1.5 nm. At close distances, dehydration leads to an electrostatically favorable tilt of the dipolar lipid head groups, which I identified as the main determinant for the observed thickening.

Structure of the Thesis

Chapter II summarizes fundamental concepts and methods that are relevant for the understanding of this thesis, but not explained in detail in a later chapter. To further provide an overview of the field of solvent entropy calculation and to characterize the requirements for a new approach, existing methods of entropy estimation are discussed in chapter III. The following chapters IV and V introduce and test Per|Mut for the calculation of rotational and translational entropies, respectively. To obtain spatially resolved solvent free energies, the calculation of enthalpy contributions is discussed in chapter VI, and the method is applied to the protein crambin. For the ease of the reader, some concepts and motivations are re-introduced at the beginning of chapters IV, V, and VI, which are written as stand-alone research articles that are either published or submitted for publication.

The dehydration-related thickening of close-contact pre-fusion lipid membranes is addressed in chapter VII.

II | Atomistic Simulations and Analysis

For a physically accurate description of the dynamics of individual atoms and their interactions through chemical bonds and non-bonded interactions, the time-dependent Schrödinger equation [64]

$$i\hbar\frac{d}{dt}|\psi\rangle = \mathcal{H}|\psi\rangle$$

needed to be solved, where $|\psi\rangle$ is the quantum mechanical state vector and $\mathcal{H}$ is the Hamiltonian of the system. A number of methods to approximate solutions of the Schrödinger equation for the electronic systems exist, of which the most important ones are the Hartree-Fock method [65–67] and density functional theory [68].

Despite decades of growing computer performance [69], it is still impossible to reach simulation timescales of nanoseconds to microseconds using quantum mechanical methods for even a small protein like crambin in solvent water (see chapter VI), which combined consists of ≈ 30000 atoms. Therefore, atomistic molecular dynamics (MD) simulations are carried out using several simplifications and approximations.

As shown by Born and Oppenheimer [70], the dynamics of nuclei and electrons can be approximately treated separately. Due to their much smaller mass, electrons move on faster timescales than nuclei, so that the electronic configuration of a molecule adapts within a typical simulation time step of 1 or 2 fs to a given set of nuclei positions. The electrons are therefore not explicitly simulated but instead give rise to an effective potential for the much slower dynamics of the atom nuclei. In MD simulations, this potential is approximated by a semi-empirical function, known as the force field (see section 4). Here, only the electronic ground state needs to be considered, since the thermal energy at $\approx 300\,\mathrm{K}$ is typically not sufficient to raise electrons into excited states.

Lastly, atomic nuclei are described by narrow wave packets relative to the length scale on which the effective interaction potential varies, because of their large masses and their short mean free paths. Therefore, Ehrenfest's theorem [71], which states that expectation values for sufficiently narrow wave-packets approximately follow classical equations of motion, can be used to switch from a quantum mechanical description to a purely classical picture.

Hence, the atoms are treated as point masses and Newton's equations of motion

$$\mathbf{M}\frac{\partial^2 \boldsymbol{r}}{\partial t^2} = F(\boldsymbol{r}) = -\nabla V(\boldsymbol{r}) \tag{3.1}$$

are solved in MD simulations. Here, $\boldsymbol{r} = (x_1, \ldots x_{3N})$ contains the positions of N atoms,

$V(\boldsymbol{r})$ is the effective potential energy, provided by a force field (see section 4), and $\mathbf{M}$ is a $3N \times 3N$ matrix containing the particle masses on its diagonal.

All atom MD simulations therefore require an initial structure $\boldsymbol{r}_0$, which for protein simulations is often obtained from the Protein Data Bank[3], and a force field that describes the interactions between the atoms. Equation (3.1) is then numerically integrated to obtain both positions and velocities in discrete time intervals using an integration algorithm, such as the leapfrog integrator described in section 5.

Under the assumption of ergodicity [72], thermodynamic quantities like entropies and free energies can be calculated from an MD simulation trajectory.

All MD simulations in this thesis were carried out using the software package GROMACS [73–77].

4 Force Fields

Force fields describe the forces $\boldsymbol{F}(\boldsymbol{r}) = -\nabla V(\boldsymbol{r})$ between the atoms of an MD system through an effective interaction potential $V(\boldsymbol{r})$ [78]. In the past, multiple force fields have been developed, of which to date, the most popular include the families of AMBER[4] [78, 79] and CHARMM[5] [80–84] force fields.

All simulations in this thesis were carried out using the CHARMM36m force field, which has the mathematical form

$$V(\boldsymbol{r}) = \sum_{\text{bonds}} \frac{1}{2}k_b(r - r_0)^2 + \sum_{\text{angles}} \frac{1}{2}k_a(\theta - \theta_0)^2 + \sum_{\text{torsions}} \frac{1}{2}V_n[1 + \cos(n\phi - \delta)]$$
$$+ \sum_{\text{improper}} \frac{1}{2}k_i(\omega - \omega_0)^2 + \sum_{\text{LJ}} 4\epsilon_{ij}\left[\left(\frac{\sigma_{ij}}{r_{ij}}\right)^{12} - \left(\frac{\sigma_{ij}}{r_{ij}}\right)^6\right] + \sum_{\text{Coulomb}} \frac{1}{4\pi\epsilon_0}\frac{q_i q_j}{r_{ij}}, \quad (4.1)$$

where the first four terms describe the interactions due to covalent bonds between the atoms, and last two terms are the unbonded Lennard-Jones [42] and Coulomb interactions, respectively. The individual terms are illustrated in table 4.1.
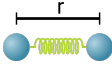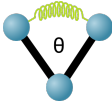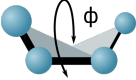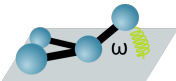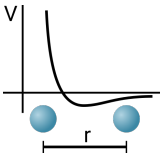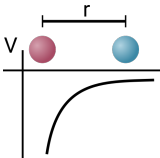
The bond stretching and angle bending, described by the first two terms, are typically very "stiff" degrees of freedom, such that the deviations from their equilibrium positions at $\approx 300\,\text{K}$ are small. Therefore, a second-order Taylor expansion is used to approximate the true potential around r_0 and θ_0, respectively. Because computers evaluate polynomials significantly faster than, e.g., the exponentials of the Morse potential [85], the approximation results in more efficient force and energy computations. Whereas the physically more accurate Morse potential is bound for large distances, the quadratic approximation grows to infinity

---

[3] https://www.rcsb.org
[4] https://ambermd.org/AmberModels.php
[5] https://www.charmm.org/charmm/resources/charmm-force-fields/

Table 4.1: Illustration of the force field terms, following [4].

Bonded potential between two atoms

$$V(r) = \frac{1}{2}k_b(r - r_0)^2$$

with a spring constant k_b and an equilibrium bond length r_0.

Angle potential between three bonded atoms

$$V(\theta) = \frac{1}{2}k_a(\theta - \theta_0)^2$$

with a spring constant k_a and an equilibrium angle θ_0.

Torsion potential across four bonded atoms

$$V(\phi) = \frac{1}{2}V_n[1 + \cos(n\phi - \delta)]$$

with strength V_n, multiplicity n, and offset δ.

Improper dihedral potential

$$V(\omega) = \frac{1}{2}k_i(\omega - \omega_0)^2$$

with a spring constant ω and equilibrium angle ω_0.

Lennard-Jones potential [42]

$$V(r) = 4\epsilon \left[\left(\frac{\sigma}{r}\right)^{12} - \left(\frac{\sigma}{r}\right)^6 \right]$$

between unbonded atoms, with strength ϵ and radius σ.

Coulomb interaction

$$V(r) = \frac{1}{4\pi\epsilon_0} \frac{q_1 q_2}{r}$$

between unbonded atoms with partial charges q_1 and q_2.

and therefore does not allow for bond ruptures or the formation of new bonds. This is a limitation, albeit not a fundamental one, that CHARMM shares with most other force fields.

The torsional dihedral angle describes the intertwist of four atoms bonded in a row, as sketched in table 4.1. Since, e.g., a methyl group typically rotates over the full range of $360°$ in less than a nanosecond, trigonometric functions rather than the quadratic approximation are used, as shown in equation (4.1).

The purely empirical improper dihedral potentials are used to maintain the chirality of molecules or the planarity of, e.g., aromatic rings [86]. As such, they merely need to restrain the dihedral angle around its target value ω_0 and hence a harmonic approximation can be used, as shown in equation (4.1).

The Lennard-Jones potential [42] in the next term describes van-der-Waals interactions and the Pauli repulsion [87] of atoms. The potential is therefore strongly repulsive for small distances and, due to cooperative polarization, slightly attractive for larger distances. ϵ defines the strength of the interaction, and σ describes the size of the atoms.

Whereas the attractive r^{-6}-part can be derived from first principles, the repulsive r^{-12}-part has no particular theoretical justification. However, the fact that r^{-12} can be calculated as the square of r^{-6}, which needs to be computed anyway, makes the Lennard-Jones potential computationally very efficient and therefore preferred over similar potentials [88].

The last term of equation (4.1) describes the electrostatic interactions between the atoms using Coulomb's law. Notably, all partial charges are fixed in classical force fields, and polarization effects are therefore neglected, which may create artifacts whenever strong electric fields are present [89, 90].

All force field parameters, consisting of force constants, bond lengths, bond angles, dihedral angles, Lennard-Jones parameters, partial charges and more can either be obtained from quantum mechanical calculations, deducted from experiments like spectroscopy, x-ray or electron diffraction, NMR or infrared analysis [91], or be indirectly inferred from other measurable quantities like diffusion constants or solvation free energies.

5 Integrator

The force field, described in section 4, yields the force $\boldsymbol{F}$ on any atom as a function of the atom position $\boldsymbol{x}$. To obtain a simulation trajectory through Newton's law

$$m\ddot{\boldsymbol{x}}(t) = m\boldsymbol{a}(t) = \boldsymbol{F}(\boldsymbol{x}(t)),$$

an integration algorithm is needed.

Out of the available algorithms, the leapfrog integrator [92, 93] is the most popular choice for MD simulations and was used in all simulations for this thesis.

The integrator reads

$$\boldsymbol{x}(t + \Delta t) = \boldsymbol{x}(t) + \boldsymbol{v}\left(t + \frac{\Delta t}{2}\right)\Delta t + \mathcal{O}(\Delta t^3) \tag{5.1}$$

$$\boldsymbol{v}\left(t + \frac{\Delta t}{2}\right) = \boldsymbol{v}\left(t - \frac{\Delta t}{2}\right) + \frac{\boldsymbol{F}(t)}{m}\Delta t + \mathcal{O}(\Delta t^3), \tag{5.2}$$

where the position $\boldsymbol{x}$ and velocity $\boldsymbol{v}$ are updated in an alternating manner every $\frac{\Delta t}{2}$ (which earned the algorithm its name). Although the leapfrog integrator is a second-order algorithm, it only requires a single force evaluation per time step, which makes it comparably efficient to, e.g., the Euler algorithm, but significantly more accurate.

6 Bond Constraints

To carry out the simulations as efficiently as possible, large time steps Δt are desirable. However, the leapfrog integrator is only numerically stable if the time step is significantly smaller than the period of the fastest degree of freedom [86]. Due to their small mass, this limit is given by the vibration frequency of covalent bonds to hydrogen atoms, which would require sub-femtosecond time steps.

To be able to use longer time steps, bonds to hydrogen atoms are constrained.

Due to the bond constraints, water molecules are treated as rigid bodies and therefore their bond-vibration entropy is neglected. The entropy of a quantum harmonic oscillator is

$$S = -k_B \log\left[1 - e^{-\beta\hbar\omega}\right] + \frac{\hbar\omega}{T}\frac{e^{-\beta\hbar\omega}}{1 - e^{-\beta\hbar\omega}}, \quad \omega = \sqrt{\frac{k}{\mu}},$$

which yields an estimation of the neglected vibrational entropy of $2.27 \cdot 10^{-5}$ J·mol^{-1}·K^{-1} per O$-$H bond, using $T = 300$ K, a reduced mass of $\mu \approx 1$ u, and a spring constant of $k = 376560$ kJ·mol^{-1}·nm^{-2} [84]. The result is much smaller than the entropy from the remaining 6 degrees of freedom (≈ 70 J·mol^{-1}·K^{-1}) and well within the margin of error of any hydration entropy estimator. Notably, a classical harmonic oscillator yields unphysical negative entropies for the same parameters. It has therefore been argued that constrained bonds are a better representation of their quantum mechanical nature than a classical harmonic potential [94].

In this thesis, 2 fs time steps were used. Bonds in water and bonds to hydrogen atoms in proteins were constrained using the SETTLE [95] and LINCS [96] algorithms, respectively.

7 Steepest Descent Minimization

Due to measurement errors in the PDB structure, imperfect force fields, homology modeling, or poorly placed water molecules, atomic clashes might be present in the starting structure of an MD simulation. In this case, an attempt to carry out the simulation would most likely

result in the system "blowing up", which could mean an immediate crash of the program or overly fast and unphysical movements of parts of the system. The effect is caused when atoms in the initial structure are too close, resulting in initial forces that are too large for the time step Δt. Such atoms are accelerated within the first time step according to equation (5.2) and "jump" too far, which can result in another clash that is possibly even more severe, triggering a chain reaction.

To avoid these problems, energy minimization needs to be carried out before the start of the actual MD simulation. Here, the potential energy $V(\boldsymbol{r})$ of the system is minimized until a local minimum, characterized by $\boldsymbol{F}(\boldsymbol{r}) = -\nabla V(\boldsymbol{r}) = 0$, is reached. In this thesis, the steepest descent (also known as gradient descent) implementation in GROMACS was used, which iteratively follows the direction of the steepest descent, i.e. the negative gradient $\boldsymbol{F}(\boldsymbol{r}) = -\nabla V(\boldsymbol{r})$ until the maximum scalar force on any atom $\max(\boldsymbol{F}_n)$ is below a threshold value [86].

The iterative algorithm is defined as

$$\boldsymbol{r}_{n+1} = \boldsymbol{r}_n + \frac{\boldsymbol{F}_n}{\max(\boldsymbol{F}_n)} h_n,$$

where $\boldsymbol{r}_0$ is the initial structure and h_n is a variable step size. In the GROMACS implementation [86], the new positions are accepted and the step size is updated to $h_{n+1} = 1.2 h_n$ if $V_{n+1} < V_n$. If $V_{n+1} > V_n$, the step size was too large. In this case, the new positions are rejected, and the step size is decreased to $h_{n+1} = 0.2 h_n$. The iterations stop, once $\max(\boldsymbol{F}_n) < \epsilon$.

A rough estimate of a safe threshold value ϵ can be obtained by considering the distance an accelerated atom with mass m travels during a time step according to equation (5.2) [97]

$$s = \frac{F}{m} \Delta t^2 \Rightarrow F = \frac{ms}{\Delta t^2}.$$

Using $s = 0.1\,\text{Å}$, which corresponds to the magnitude of inter-atom distances, $m = 1\,\text{u}$, the mass of a hydrogen atom, and the standard MD time step $\Delta t = 2\,\text{fs}$, a maximum force of $2.5 \cdot 10^3\,\text{kJ·mol}^{-1}\text{·nm}^{-1}$ is obtained.

8 Temperature and Pressure Coupling

If the leapfrog integrator (section 5) and the force field (section 4) were to be directly applied, a microcanonical NVE ensemble (constant number of particles, volume, and energy) would be simulated. In reality, however, most biophysically relevant processes and phenomena take place in the canonical NVT ensemble (constant number of particles, volume, and temperature) or in the isothermal-isobaric NPT ensemble, where, instead of volume, pressure is maintained. Hence, algorithms for temperature and pressure coupling are needed.

Temperature Coupling

At the beginning of an MD simulation, initial atom velocities are drawn from a Maxwell distribution

$$p(v_i) = \sqrt{\frac{m_i}{2\pi k_B T_0}} \exp\left(-\frac{m_i v_i^2}{2k_B T_0}\right),$$

according to the desired temperature T_0. However, in order to maintain this temperature during the simulation, the system needs to be coupled to a simulated heat bath.

For all simulations in this thesis, the V-rescale thermostat [86, 98] was used. In this algorithm, atom velocities are scaled at each simulation time step Δt by

$$\alpha = \sqrt{\frac{K_t}{K}}$$

to keep the desired temperature T_0 on average. Here, K is the total kinetic energy. To produce a canonic temperature distribution around T_0, the target kinetic energy K_t is redrawn from a stochastic process at every time step. Its distribution is

$$dK = \frac{K_0 - K}{\tau}dt + 2\sqrt{\frac{KK_0}{N_f}}\frac{dW}{\sqrt{\tau}},$$

where N_f is the number of degrees of freedom, $K_0 = \frac{N_f}{2}k_B T_0$, and W is a Wiener process. The first term is an exponential relaxation towards K_0, so that the temperature is always steered back to T_0. The second term is stochastic and produces the canonic temperature distribution. The time constant τ controls the coupling strength between a simulated system and a heat bath. Whereas small values yield a quickly thermalizing system, large values mean a slow temperature equilibration.

In this thesis, the standard value of $0.1\,\mathrm{ps}$ was used.

Pressure Coupling

In order to maintain a given pressure during an MD simulation, a pressure coupling algorithm, also known as a barostat, is needed. Two different algorithms, the Berendsen barostat [99] and the Parrinello-Rahman barostat [100, 101], were used in this thesis and are therefore going to be discussed here.

During an MD simulation, the pressure is calculated using the virial theorem [86, 97]. The pressure tensor reads

$$\boldsymbol{P} = \frac{2}{V}(\boldsymbol{E}_{\mathrm{kin}} - \boldsymbol{\Xi}),$$

where V is the current box volume,

$$\boldsymbol{E}_{\mathrm{kin}} = \frac{1}{2}\sum_{i=1}^{N} m_i \boldsymbol{v}_i \cdot \boldsymbol{v}_i^{\mathsf{T}}$$

is the kinetic energy tensor, and

$$\boldsymbol{\Xi} = -\frac{1}{2} \sum_{i=1}^{N} \boldsymbol{r}_i \cdot \boldsymbol{F}_i^{\mathsf{T}}$$

is the virial.

To maintain a target pressure of P_0 on average, the barostat iteratively rescales the simulation box dimensions, described by the box matrix $\boldsymbol{h} = (\boldsymbol{a}, \boldsymbol{b}, \boldsymbol{c})$ with the box vectors $\boldsymbol{a}$, $\boldsymbol{b}$, $\boldsymbol{c}$, which reads

$$\boldsymbol{h}' = \boldsymbol{h} \cdot \boldsymbol{\mu}.$$

Here, $\boldsymbol{\mu}$ is a diagonal 3×3 scaling matrix. In case of isotropic pressure coupling, which is normally used, the diagonal entries are identical. For membrane systems, however, it is useful to scale the membrane-orthogonal direction independently, in which case only two of the diagonal elements of $\boldsymbol{\mu}$ are identical.

Berendsen Barostat

In the Berendsen pressure coupling scheme, the pressure fluctuations are exponentially damped with a time constant τ. The scaling matrix update at every time step reads [86]

$$\mu_{ii} = 1 - \frac{1}{3} \frac{\Delta t}{\tau} \beta (P_{0ii} - \tilde{P}_{ii}),$$

where $\beta = 4.5 \cdot 10^{-10} \, \mathrm{bar}^{-1}$ is the isothermal compressibility of water. For isotropic scaling, the pressure tensor is substituted by its isotropic average $\tilde{P}_{ii} = \frac{1}{3} \mathrm{tr}(\boldsymbol{P})$. For semi-isotropic scaling $\tilde{\boldsymbol{P}}$ is defined as $\tilde{P}_{xx} = \tilde{P}_{yy} = \frac{1}{2}(P_{xx} + P_{yy})$, $\tilde{P}_{zz} = P_{zz}$.

The Berendsen barostat is very robust and stable. In particular, it is not prone to box oscillations, due to its exponential dampening. Since it strongly suppresses pressure fluctuations, it does not produce an exact thermodynamic NPT ensemble [86]. It has therefore only been used for equilibration simulations.

Parrinello-Rahman Barostat

An exact NPT ensemble can be produced by the Parrinello-Rahman barostat [102–104]. Here, the box size follows its own equation of motion [86]

$$\frac{d^2 \boldsymbol{h}^{\mathsf{T}}}{dt^2} = V \boldsymbol{W}^{-1} \boldsymbol{h}^{-1} (\boldsymbol{P} - \boldsymbol{P}_0),$$

where the mass-like matrix $\boldsymbol{W}$ is defined as

$$\boldsymbol{W}_{ij}^{-1} = \frac{4\pi^2 \beta}{3\tau^2 L} \delta_{ij}$$

with the largest box matrix element L. A large coupling timescale τ therefore results in a "heavy" box and a loose coupling.

To obtain the pressure fluctuations of a thermodynamic NPT ensemble, a friction term is added to Newton's equations of motion, which then reads

$$\frac{d^2\boldsymbol{r}_i}{dt^2} = \frac{\boldsymbol{F}_i}{m_i} - \boldsymbol{M}\frac{d\boldsymbol{x}_i}{dt}$$

$$\boldsymbol{M} = (\boldsymbol{h}^\intercal)^{-1}\left[\boldsymbol{h}^\intercal\frac{d\boldsymbol{h}}{dt}\frac{d\boldsymbol{h}^\intercal}{dt}\boldsymbol{h}\right]\boldsymbol{h}^{-1}.$$

$\boldsymbol{M}$ therefore plays the role of a friction coefficient.

Because the barostat is prone to oscillations, particularly when the initial pressure and the target pressure are very different, the algorithm is mainly suited for production runs after an initial equilibration phase with the Berendsen barostat.

9 Nearest Neighbor Search

In order to calculate rotational entropies (described in chapter IV), translational entropies as well as the translation-rotation correlation entropy (both described in chapter V), fast nearest neighbor search algorithms in Euclidean space (described by $\mathbb{R}^d$), the groups of rotations $SO(3)^d$, and the composite space $\mathbb{R}^3 \times SO(3)$ are required. Because entropy calculation via a third-order mutual information expansion (MIE, see section 16.2.1) requires k-nearest neighbor searches for each of the typically 10^5 data points for all molecules, all molecule pairs, and all molecule triples, a brute-force search approach is orders of magnitudes too slow, even when implemented on GPUs. Efficient nearest-neighbor search algorithms and their efficient implementation are therefore essential.

I have chosen two tree-based search algorithms, the k-d tree and the VP tree, which will be described in the following sections. Both algorithms aim to find the kth-nearest neighbor with respect to a metric $d_{\mathbb{S}}$ to each point from a set of points $\{\boldsymbol{x} \in \mathbb{S}\}$, where $\mathbb{S}$ is one of the spaces mentioned above.

k-d Tree

Before the k-dimensional tree [105] (k-d tree, k here not to be confused with the k-nearest neighbor) can be used to search for the nearest neighbor of a data point, the tree needs to be built. This process is referred to as indexing.

As shown in Fig. 9.1 for a two-dimensional example-data set, the k-d tree is built by recursively bisecting the data set through (hyper-)planes along the Cartesian axes.

First, the example data set is split along the x-axis. To this end, the center-most x-coordinate (median) is computed and used for bisection. In Fig. 9.1, the data set is therefore divided between points with $x < 5$, which form the left subtree, and points with $x \geq 5$, which

Figure 9.1: A two-dimensional k-d tree. The data set and the split-planes are shown on the left, the corresponding tree is shown on the right.

form the right subtree. Next, both subtrees are bisected along the y-axis using their median y-coordinates, which results in a split between $y < 7$ and $y \geq 7$ for the left subtree and a split between $y < 4$ and $y \geq 4$ for the right subtree. This schema of alternating splits along the axes is continued until each point in the data set is represented by a leaf node, as shown in Fig. 9.1 on the right.

The k-d tree can be directly generalized to data sets in $\mathbb{R}^d$, where instead of alternating splits along the x- and y-axes, the data set is bisected along the d Cartesian axes in a cyclic manner.

During nearest neighbor search, the tree is used to exclude entire subtrees from the search.

To find the nearest neighbor to (6|1) in the example from Fig. 9.2, the distances within the closest subtree are computed and (5|4) is identified as the "current best", as shown in Fig. 9.2A. Next, the search is continued one node-level higher. Because the sphere that includes the "current best" (shown in green) intersects the bisection plane at $x = 5$, it is possible that there is a closer point on the left of the plane. The next bisection-plane at $y = 7$, however, is further away than the "current best", hence the subtree $y \geq 7$ can be excluded, as shown in Fig. 9.2B. This leaves the point (1|1) as the only remaining candidate. Because it is further away from (6|1) than (5|4), the "current best" is not updated and (5|4) is correctly identified as the nearest neighbor to (6|1).

By repeating this process, the second-nearest, third-nearest, ..., k-nearest neighbor can be identified for all points of the data set.

If an $\mathcal{O}(n)$-algorithm is used to determine the (approximate) median for each of the $\mathcal{O}(\log n)$ bisection layers, indexing a k-d tree has a combined complexity of $\mathcal{O}(n \log n)$. In the best case, each layer is visited only once to find a nearest neighbor, thus finding a k-nearest neighbor is a $\mathcal{O}(k \log n)$-operation. Finding the k-nearest neighbor for each of n data points therefore

Figure 9.2: Nearest neighbor search schema using a k-d tree. The k-d tree is applied to find the nearest neighbor to $(6|1)$. In (A), the closest sub-tree to $(6|1)$ is searched and $(5|4)$ is stored as the "current best". In (B), the search is recursively continued to the higher nodes.
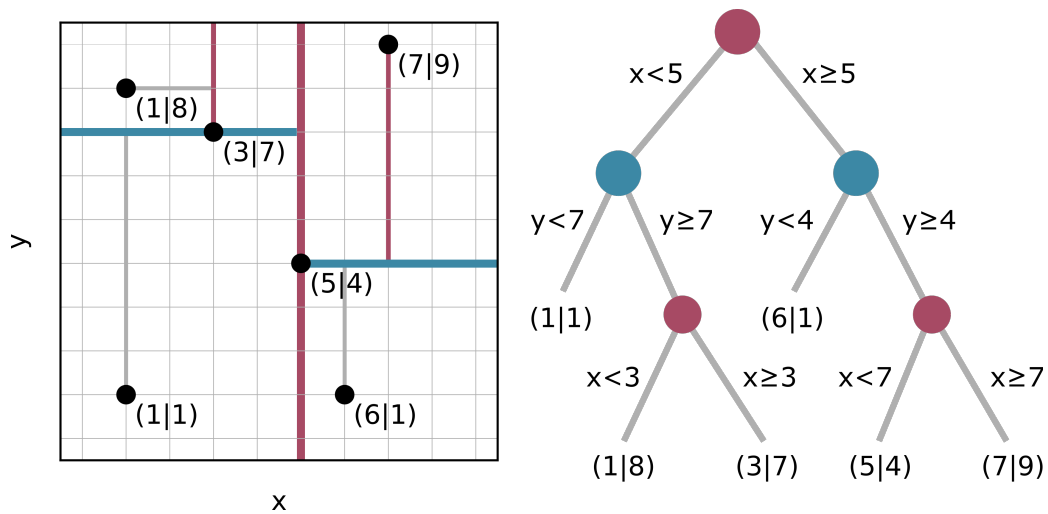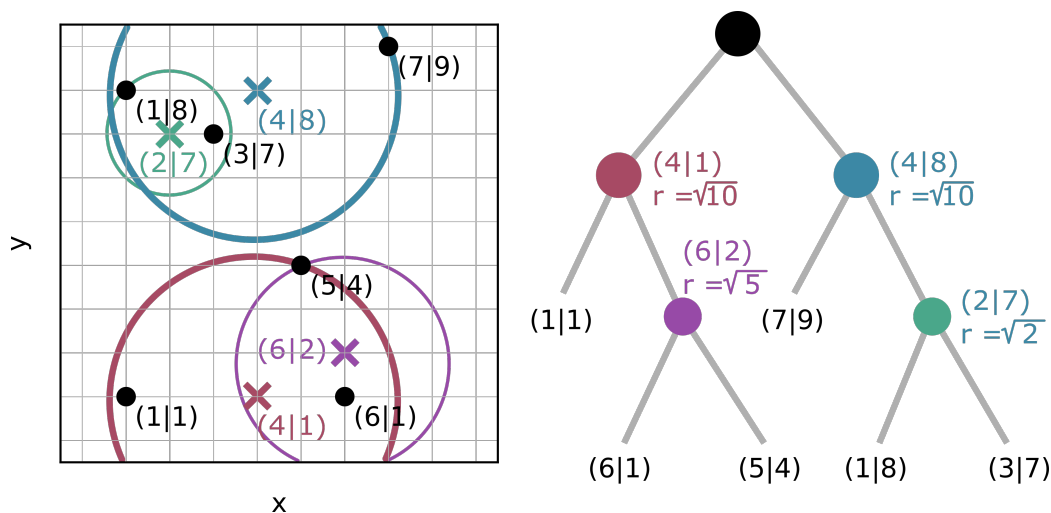
Figure 9.3: A two-dimensional VP tree. The data set and the balls are shown on the left, the corresponding tree is shown on the right.

has a complexity of $\mathcal{O}(kn\log n)$.

Because the algorithm uses axis-aligned splits, it can only be used with Cartesian coordinates in $\mathbb{R}^d$ with the Euclidean metric. Therefore, a different nearest-neighbor search algorithm needed to be used in $SO(3)^d$ and the composite space $\mathbb{R}^3 \times SO(3)$ for rotational entropies and translation-rotation correlation entropies, respectively.

VP Tree

Similar to the k-d tree, the vantage point tree [106–108] (VP tree, also known as ball tree) uses bisections to build a tree, which can then be used for nearest neighbor search. Most importantly, the VP tree only uses the triangle inequality and makes no further assumptions about the space or used coordinates. It can therefore be directly applied to $SO(3)^d$ or $\mathbb{R}^3 \times SO(3)$, as long as the used metrics fulfill the triangle inequality.

To build the tree, two pivot points ("vantage points") are randomly chosen[6], as shown in Fig. 9.3 in red and blue, respectively. The distances from the vantage points to all data points are calculated and the data points are assigned to their closest vantage point. For each vantage point, a radius, defined as the largest distance between the pivot point and its assigned data points, is stored. In the example from Fig. 9.3, both radii are $\sqrt{10}$ (in the Euclidean metric). This way, the data set is bisected into two "balls", which can, and generally will, overlap.

The bisection is recursively repeated until each data point can be represented as a leaf-node, as shown on the right of Fig. 9.3.

To find the nearest neighbor to the data point $\mathbf{Q} = (6|1)$, as exemplarily shown in Fig. 9.4A, the distances within the smallest enclosing ball (purple in Fig. 9.4) are computed and (5|4) is

---

[6]The implementation in the NMSLib [109, 110], which was used in this thesis, makes multiple attempts to choose the pivot points and accepts those with the largest variance of the data point distances.
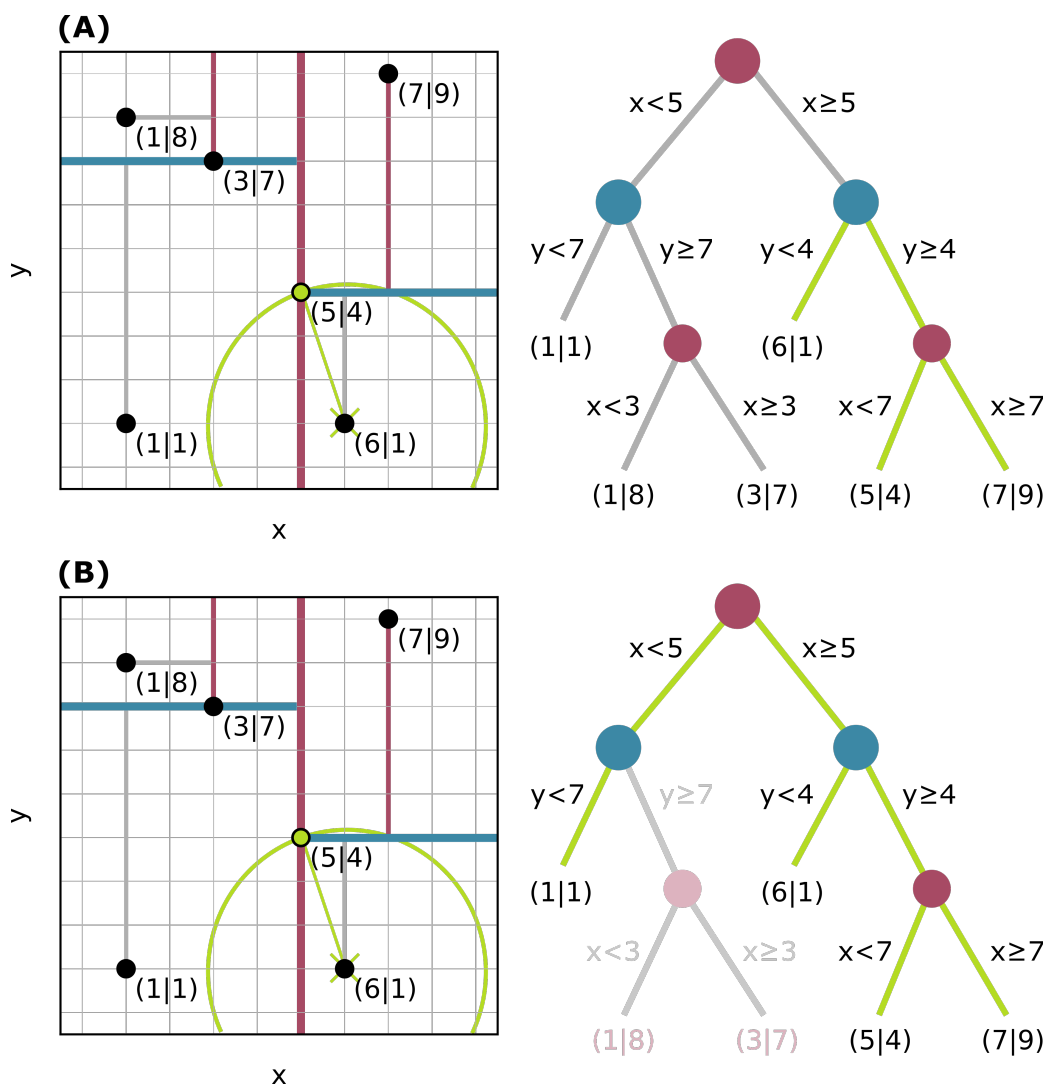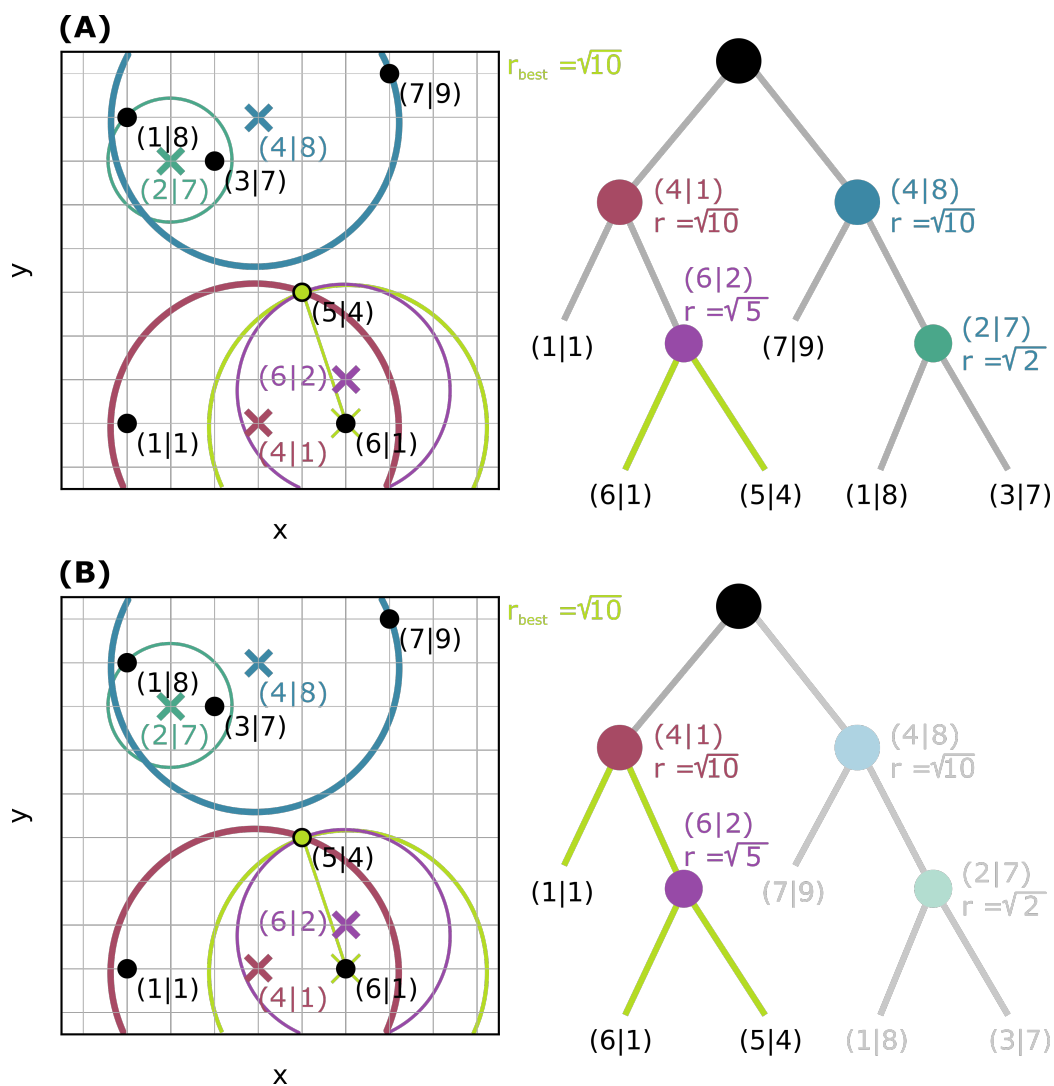
Figure 9.4: Nearest neighbor search schema using a VP tree. The VP tree is applied to find the nearest neighbor to (6|1). In (A), the closest sub-tree to (6|1) is searched and (5|4) is stored as the "current best". In (B), the search is recursively continued to the higher nodes, of which some can be excluded due to the triangle inequality.

stored as the "current best" with a distance of $r_{\text{best}} = \sqrt{10}$. Moving one tree-level upwards, the next ball (red) intersects with the "current best"-sphere (green). The node can therefore not be excluded, but since its leaf-node (1|1) is not closer than r_{best}, the "current best" is not updated. Moving further to the blue node, the triangle inequality yields

$$r_{\text{best}} < |d(\boldsymbol{Q}, \boldsymbol{P}_{\text{pivot}}) - r|,$$

with the pivot point $\boldsymbol{P}_{\text{pivot}} = (4|8)$ and the radius $r = \sqrt{10}$ of the VP-ball. The entire left subtree can therefore be excluded, as shown in Fig. 9.4B.

Following the same arguments as for the k-d tree, the complexity is $\mathcal{O}(n \log n)$ to build the tree and a k-nearest neighbor search for all n data points is a $\mathcal{O}(kn \log n)$-operation. The scaling behavior is therefore identical to the k-d tree, however, in practice the latter is usually faster. The k-d tree was therefore used to calculate translational entropies and the VP tree was used for the rotational entropies and the translation-rotation correlation terms, where the k-d tree cannot be applied.

For k-d trees, the implementation in the python module scikit-learn 0.20.3 [111] was used. For VP trees, the Non-Metric Space Library 1.7.3.6 [109, 112] (NMSLib) was used, for which I assessed and implemented metrics for the spaces of rotation $SO(3)^d$ (see chapter IV) and the composite space $\mathbb{R}^3 \times SO(3)$ (see chapter V).

III | Existing Methods for Entropy Estimation

In this chapter, an overview of existing methods for entropy calculation from MD simulations will be given. Methods that were used in this thesis, namely thermodynamic integration and permutation reduction, will be explained in more detail. Other methods will be explained briefly and their advantages and drawbacks, particularly with respect to calculating spatially resolved hydration shell entropies, will be discussed. Whenever possible within the scope of this work, I will sketch the basic idea and provide a short derivation.

10 Quasiharmonic Approximation

Introduced by Karplus and Kushick [113], the quasiharmonic (QH) method estimates entropies of macromolecules, such as proteins, by approximating the configuration space density by a multivariate Gaussian distribution, of which the entropy is known analytically [114].

The method is based on the observation that most folded globular proteins fluctuate around a single, well defined native fold configuration $\boldsymbol{r}_0$. The extremely complex and rugged potential energy landscape is then approximated by a quadratic function around $\boldsymbol{r}_0$, which reads

$$V(\boldsymbol{r}) \approx V(\boldsymbol{r}_0) + \frac{1}{2}(\Delta\boldsymbol{r})^{\mathsf{T}}(\beta\boldsymbol{C})^{-1}\Delta\boldsymbol{r},$$

where $\boldsymbol{r}$ is a $3N$-dimensional vector in configuration space, $\Delta\boldsymbol{r} = \boldsymbol{r} - \boldsymbol{r}_0$, and $(\beta\boldsymbol{C})^{-1}$ plays the role of a multivariate elastic force tensor. The system is hence approximated by a $3N$-dimensional harmonic oscillator. The approximation of the interaction energy implies a Gaussian configuration space density

$$\varrho(\boldsymbol{r}) \propto e^{-\beta V(\boldsymbol{r})} = e^{-\frac{1}{2}(\Delta\boldsymbol{r})^{\mathsf{T}}\boldsymbol{C}^{-1}\Delta\boldsymbol{r}}$$

with a covariance matrix $\boldsymbol{C}$, which is calculated from a long MD simulation trajectory. The analytic configurational entropy

$$S_{\mathrm{QH}} \approx \frac{k_B}{2}\left(3N + \log\left[(2\pi)^{3N}\det\boldsymbol{C}\right]\right)$$

of the Gaussian is then used as an entropy estimate.

As shown in figure 10.1, many eigenvalues $\{\sigma_i\}$ of the covariance matrix of a typical protein
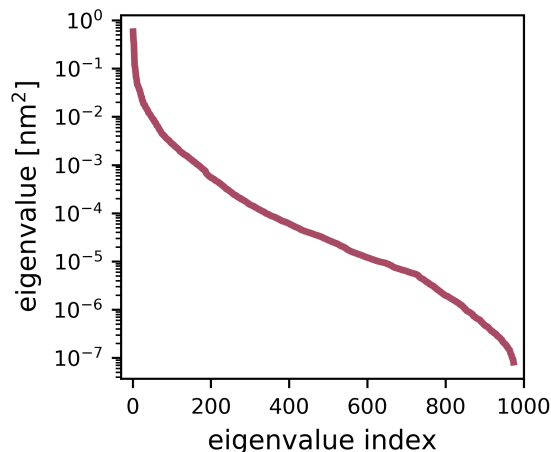
Figure 10.1: Eigenvalue spectrum of the covariance matrix of the crambin backbone. Obtained from $0.5\,\mu$s of unrestrained MD simulation of the protein in its native fold, as described in chapter VI.

are very small. Therefore, $\det \boldsymbol{C} = \prod_i \sigma_i$ takes on extremely small values, which may result in unphysical negative entropies. To solve this problem, a quantum mechanical version has been introduced by Schlitter [115], which uses the entropy of the quantum harmonic oscillator instead of the classical oscillator and reads

$$\tilde{S}_{\mathrm{QH}} \approx \frac{k_B}{2} \sum_{i=1}^{3N} \log\left[1 + \frac{k_B T e^2}{\hbar^2}\tilde{\sigma}_i^2\right].$$

Here, e is Euler's number, $\tilde{\sigma}_i^2$ are the eigenvalues of the mass-weighted coverance matrix $\tilde{\boldsymbol{C}} = \boldsymbol{M}^{\frac{1}{2}}\boldsymbol{C}\boldsymbol{M}^{\frac{1}{2}}$, and $\boldsymbol{M}$ has the atomic masses on its diagonal and is zero elsewhere.

Compared to other methods, the QH approximation is computationally efficient and yields good results for folded, globular proteins, which fluctuate around a single, native fold configuration in configuration space, so that the probability density can be well approximated by a Gaussian [114]. Since a Gaussian is the maximum-entropy distribution for a given variance, (see appendix section 36), the QH approximation results in a strict upper entropy limit if the configuration space is sufficiently sampled. The method therefore yields strongly overestimated entropies when the configuration space distribution is not Gaussian-like [115]. The QH approximation thus results in inaccurate entropy estimates for systems with multiple potential minima or diffusive systems like solvent water [46].

11 Thermodynamic Integration

Thermodynamic integration (TI) [116] is mostly used as a method to calculate free energy differences between an initial state, described by the Hamiltonian H_a, and a target state,

described by H_b [117–119]. A switching coordinate λ is used to interpolate between the two states, such that $H(\lambda = 0) = H_a$ and $H(\lambda = 1) = H_b$. The free energy difference ΔF is then calculated by integration of the absolute differential $\frac{dF}{d\lambda}$ along the interpolation pathway. Mathematically, this reads

$$
\begin{aligned}
\Delta F &= \int_{\lambda=0}^{\lambda=1} \frac{dF}{d\lambda} d\lambda \\
&= -k_B T \int_{\lambda=0}^{\lambda=1} \frac{\partial \log Z(\lambda)}{\partial \lambda} d\lambda \\
&= -k_B T \int_{\lambda=0}^{\lambda=1} \frac{1}{Z(\lambda)} \frac{\partial Z(\lambda)}{\partial \lambda} d\lambda \\
&= -k_B T \int_{\lambda=0}^{\lambda=1} \frac{1}{Z(\lambda)} \frac{\partial}{\partial \lambda} \left\{ \int e^{-\beta H(\lambda)} d\boldsymbol{x}^N d\boldsymbol{p}^N \right\} d\lambda \\
&= \int_{\lambda=0}^{\lambda=1} \frac{1}{Z(\lambda)} \left\{ \int \frac{\partial H(\lambda)}{\partial \lambda} e^{-\beta H(\lambda)} d\boldsymbol{x}^N d\boldsymbol{p}^N \right\} d\lambda \\
&= \int_{\lambda=0}^{\lambda=1} \left\langle \frac{\partial H(\lambda)}{\partial \lambda} \right\rangle_\lambda d\lambda,
\end{aligned}
$$

where Z is the partition function, and $\int d\boldsymbol{x}^N d\boldsymbol{p}^N$ denotes an integration over phase space.

In practice, the derivative of the Hamiltonian $H(\lambda)$ with respect to the switching coordinate λ is sampled along the interpolation path at multiple intermediate states $0 = \lambda_0 < \lambda_1 < \ldots < \lambda_n = 1$, e.g. by means of Monte Carlo or MD simulations. The integration is then carried out numerically, e.g., using the trapezoidal rule.

The same idea can also be applied to calculate entropy differences by integrating over the entropy differential $\frac{dS}{d\lambda}$. The entropy can be written as

$$
\begin{aligned}
S &= -\left(\frac{\partial F}{\partial T} \right)_{N,V} \\
&= k_B T \log Z(\lambda) + \frac{\int H(\lambda) e^{-\beta H(\lambda)} d\boldsymbol{x}^N d\boldsymbol{p}^N}{T \int e^{-\beta H(\lambda)} d\boldsymbol{x}^N d\boldsymbol{p}^N}.
\end{aligned}
$$

The differential with respect to λ reads [120]

$$
\left(\frac{dS}{d\lambda} \right)_{N,V,T} = \frac{1}{k_B T^2} \left\{ \left\langle \frac{\partial H(\lambda)}{\partial \lambda} \right\rangle_\lambda \langle H(\lambda) \rangle_\lambda - \left\langle \frac{\partial H(\lambda)}{\partial \lambda} H(\lambda) \right\rangle_\lambda \right\},
$$

and therefore the entropy difference ΔS is

$$
\Delta S = \frac{1}{k_B T^2} \int_{\lambda=0}^{\lambda=1} \left\{ \left\langle \frac{\partial H(\lambda)}{\partial \lambda} \right\rangle_\lambda \langle H(\lambda) \rangle_\lambda - \left\langle \frac{\partial H(\lambda)}{\partial \lambda} H(\lambda) \right\rangle_\lambda \right\} d\lambda.
$$

Since both free energy and entropy are state functions, the interpolation pathway is theoret-
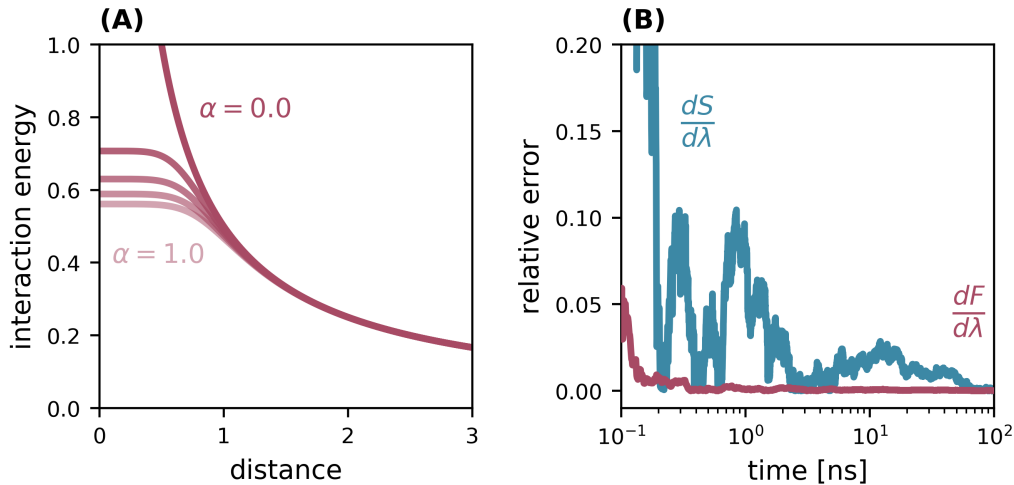
Figure 11.1: (A) shows an example of soft-core potentials with $\sigma = 1.0$ for α between 0.0 and 1.0. (B) shows the speed of convergence of $\frac{dF}{d\lambda}$ and $\frac{dS}{d\lambda}$ for 1728 argon atoms and $\lambda = 0.25$.

ically irrelevant. Therefore, a simple linear interpolation of the form $H(\lambda, r) = (\lambda - 1)H_a(r_a) + \lambda H_b(r_b)$ is commonly used. In practice, however, the singularities of the Lennard-Jones [42] and Coulomb interactions cause large fluctuations of $\frac{\partial H}{\partial \lambda}$ and H, slow convergence or even numerical problems, as particles in intermediate λ-states interact weakly enough to get extremely close [86]. To circumvent this problem, the simple linear interpolation is modified, and the singularities are removed using smoothed "soft-core" potentials. In the chosen Gromacs implementation, the soft-core potentials are shifted potentials such that their singularities at a distance of $r = 0$ are never reached [86]:

$$r_a = (\alpha \sigma_a^6 \lambda + r^6)^{\frac{1}{6}}$$
$$r_b = (\alpha \sigma_b^6 (1 - \lambda) + r^6)^{\frac{1}{6}},$$

where r is the actual distance between the atoms, σ_a and σ_b are the Lennard-Jones radii in states a and b, respectively, and α is a free parameter that controls the "smoothness" of the soft-core potential. Figure 11.1 (A) shows an example soft-core potential for α-values between 0 and 1 and $\lambda = 0.5$.

Contrary to other methods, TI does not rely on strong assumptions and therefore, given enough sampling, has no fundamental limit of accuracy other than those implied by MD simulations in general. However, as enthalpy and entropy fluctuations tend to compensate each other, TI usually converges significantly faster for free energy differences than when used for entropies. To demonstrate this, $\left\langle \frac{dF}{d\lambda} \right\rangle$ and $\left\langle \frac{dS}{d\lambda} \right\rangle$ were plotted in figure 11.1 (B) for the intermediate $\lambda = 0.25$ of 1728 ideal gas particles that were switched to argon atoms. Whereas the free energy derivative is converged within 1 ns, the entropy derivative takes 80 ns

to reach a relative error below 1%. Furthermore, because TI relies on integration over many λ-intermediates, errors accumulate for long integration paths. For both reasons, entropy TI is mainly suited for small systems and for small alterations to the system, i.e., when only few atoms need to be morphed or "grown". In addition, the spatial distribution of the entropy change cannot be calculated by TI, thus other methods need to be applied when a spatial resolution is necessary.

12 Inhomogeneous Solvation Theory

In the Inhomogeneous Solvation Theory (IST) [12, 13, 121, 122], the solvation entropy is approximated using a truncated series of multi-body correlation terms, akin to the mutual information expansion (MIE) used in this thesis.

Introduced by Lazaridis [12], the IST expansion of entropy reads

$$S = S^{(1)} + S^{(2)} + \dots.$$

The one- and two-body contributions

$$S^{(1)} = \frac{5}{2}Nk_B + Nk_B \log\left[\rho\left(\frac{2\pi m k_B T}{h^2}\right)^{\frac{3}{2}}\right] - k_B \int g^{(1)}(\boldsymbol{r})\log g^{(1)}(\boldsymbol{r})d\boldsymbol{r} \tag{12.1}$$

$$S^{(2)} = -\frac{1}{2}k_B\rho^2 \iint g^{(1)}(\boldsymbol{r})g^{(1)}(\boldsymbol{r'})[g^{(2)}(\boldsymbol{r},\boldsymbol{r'})\log g^{(2)}(\boldsymbol{r},\boldsymbol{r'}) - g^{(2)}(\boldsymbol{r},\boldsymbol{r'}) + 1]d\boldsymbol{r}d\boldsymbol{r'} \tag{12.2}$$

are expressed in terms of the correlation functions

$$g^{(1)}(\boldsymbol{r}) = \frac{\rho^{(1)}(\boldsymbol{r})}{\rho}$$

$$g^{(2)}(\boldsymbol{r},\boldsymbol{r'}) = \frac{\rho^{(2)}(\boldsymbol{r},\boldsymbol{r'})}{\rho^{(1)}(\boldsymbol{r})\rho^{(1)}(\boldsymbol{r'})}.$$

Here, ρ, $\rho^{(1)}(\boldsymbol{r})$, and $\rho^{(2)}(\boldsymbol{r},\boldsymbol{r'})$ are the number density, local number density, and local two-point number density, respectively. Higher-order terms $S^{(n)}$, $n > 2$ are usually neglected.

The IST can be generalized to also include rotational degrees of freedom [10, 43] but does not yield spatial resolution. Furthermore, since the correlation functions $g^{(1)}(\boldsymbol{r})$ and $g^{(2)}(\boldsymbol{r},\boldsymbol{r'})$ of the full system are dominated by the more abundant bulk water molecules, the effect of the crucial solvation shell is captured poorly.

To overcome these problems, Nguyen et al. [10, 43] developed the Grid Inhomogeneous Solvation Theory (GIST). Here, the entropy estimate is limited to the single-body term $S \approx S^{(1)}$, but the integral in eq. (12.1) is evaluated for every voxel of a three-dimensional grid. The spatial resolution of the approach captures the solvation shell, however, truncating the expansion after just the single-body term severely limits its accuracy.

Later, the two-body term has been added to GIST [123]. Here, the water-water correlation between each voxel and its neighboring voxels is evaluated as defined in eq. (12.2), which results in sampling problems of the high-dimensional space.

13 2-Phase Thermodynamics

The 2-Phase Thermodynamics (2PT) approach calculates solvation entropies by treading the density of states (DoS) $g(\omega)$ of the system as a superposition of solid-like and gas-like components [45, 124].

In the QH approximation, the partition function Z can be expressed as

$$\log Z = \int_0^\infty d\omega g_s(\omega) \log z_{\mathrm{HO}}(\omega),$$

where $z_{\mathrm{HO}}(\omega)$ is the partition function of a quantum harmonic oscillator with frequency ω and $g_s(\omega)$ is the DoS. As discussed above, the harmonic approximation describes diffusive systems poorly but is well-suited for crystalline and glassy solids.

The solid entropy reads

$$S_s = k_B \log Z + \beta^{-1} \left(\frac{\partial \log Z}{\partial T} \right)_{N,V} = k_B \int_0^\infty d\omega g_s(\omega) W_s(\omega),$$

with the weighting function

$$W_s(\omega) = \frac{\beta h \omega}{e^{\beta h \omega} - 1} - \log[1 - e^{\beta h \omega}].$$

To obtain better entropy estimates for liquids, 2PT introduces a second phase by approximating the total DoS as

$$g(\omega) = g_s(\omega) + g_g(\omega),$$

where $g_g(\omega)$ is the DoS of a gas-like component, accounting for $3fN$ degrees of freedom, and $g_s(\omega)$ is the remaining DoS, accounting for $3N(1 - f)$ degrees of freedom. Here, $f \in [0, 1]$ is a parameter that determines the proportion of solid-like and gas-like components.

The entropy reads

$$S = k_B \int_0^\infty d\omega g_s(\omega) W_s(\omega) + k_B \int_0^\infty d\omega g_g(\omega) W_g(\omega).$$

The gas-like component is modeled as a hard-sphere gas, for which the DoS and the weight-

ing function read

$$g_g(\omega) = \frac{g_0}{1 + \left[\frac{\pi g_0 \omega}{6fN}\right]^2}$$

$$W_g(\omega) = \frac{1}{3k_B} S_{\text{HS}},$$

respectively, where S_{HS} is the entropy of a hard-sphere gas at identical temperature and number density, and $g_0 = g_g(0) = g(0)$, because $g_s(0) = 0$ by definition.

In 2PT, the fit parameters g_0 and f are determined from the total velocity autocorrelation function

$$C(t) = \sum_{j=1}^{N} \sum_{k=1}^{3} m_j c_j^k(t),$$

with the atomic masses m_j and atomic velocity autocorrelation functions $c_j^k(t)$.

The DoS $g(\omega)$ can be obtained via the Wiener–Khintchine theorem [125, 126] as the Fourier transform

$$g(\omega) = \frac{2}{kT} \lim_{\tau \to \infty} \int_{-\tau}^{\tau} C(t) e^{-i2\pi\omega t} dt.$$

Therefore,

$$g_0 = g(0) = \frac{2}{kT} \int_{-\infty}^{\infty} C(t) dt.$$

The "fluidity" f is defined as

$$f = \frac{D(T, \rho)}{D_{\text{HS}}(T, \rho, \sigma_{\text{HS}})},$$

with the solvent diffusivity

$$D = \frac{1}{6mN} \int_{-\infty}^{\infty} C(t) dt$$

and the corresponding hard-sphere diffusivity

$$D_{\text{HS}}(T, \rho, \sigma_{\text{HS}}) = \frac{3}{8} \frac{1}{\rho \sigma_{\text{HS}}^2} \left(\frac{k_B T}{\pi m}\right)^{\frac{1}{2}}.$$

The hard-sphere radius σ_{HS} is determined from the hard-sphere packing density, as described by Lin et al. [124].

2PT can has been modified to also yield rotational entropies by substituting the hard-sphere model with a hindered rotor [45]. More recently, the 2PT-approach has been adapted to provide spatial resolution [11]. Here, the simulation box is divided into voxels, to which the method is applied individually. When a molecule is in a voxel at time t, it is followed until time $t + \tau_{\text{max}}$ (with, e.g., $\tau_{\text{max}} = 1.6\,\text{ps}$ [11]) to compute the autocorrelation function. A large value of τ_{max} is preferable to sample the full autocorrelation function and to obtain a good resolution in Fourier space, whereas small values yield a better spatial resolution, due

to the movement of the molecules during τ_{max}.

Although 2PT yields accurate absolute entropies for bulk water [45], it is unclear how well surface water features are captured by the method. In addition, the bulk water number density, rather than the local density, is used to determine the parameter f for three-dimensional voxels [11]. Overall, the authors report a solvation entropy contribution $-T\Delta S$ accuracy of $4.9\,\mathrm{kJ\cdot mol^{-1}}$ for monoatomic solutes or small molecules like benzene [11].

Fundamentally, 2PT is an interpolation between two idealized models — harmonic oscillators and a hard-sphere gas — which may not always allow for a good description of the local effects. A nonparametric approach is therefore preferred.

14 Permutation Reduction

As explained earlier, applying the QH approximation to a diffusive solvent system yields poor entropy estimates. The reason for this problem is twofold.

First, as discussed in the introduction, a huge amount of sampling is required to converge the entropy estimate, due to the very high dimensional configuration space.

Secondly, the configuration space distribution of a solvent system is highly non-Gaussian. As discussed in section 10, the QH approximation can thus not be safely applied. For an ideal gas in a cubic simulation box, the configuration space distribution is a uniform density in a $3N$-dimensional hypercube. For interacting particles, the distribution becomes even less Gaussian-like, because the excluded volumes of the particles cause "holes" in the configuration space density wherever two particles would intersect. Permutation reduction [46, 47] aims to solve both of these problems by mapping the full configuration space density into a subvolume, smaller by the Gibbs factor $N!$ without changing the physics. Permutation reduction thus improves sampling by $N!$ and makes the configuration space density more compact and therefore more Gaussian-like.

The method exploits the permutation symmetry of the identical solvent molecules. For each frame in a trajectory $\{\boldsymbol{x}_i(t)\}$, the particle permutation π is chosen so that minimizes the squared distance

$$\sum_{i=1}^{N} \|\boldsymbol{x}_{\pi(i)}(t) - \boldsymbol{r}_i\|^2 \tag{14.1}$$

in configuration space, where $\{\boldsymbol{r}_i\}$ is an arbitrary reference configuration. Because during the simulation, each particle carries a unique label (an index), the permutation of particles can also be interpreted as relabeling.

Figure 22.1A demonstrates the effect of relabeling on the two water molecules, where in the right panel the molecule labels are swapped since this minimizes the distance to the reference configuration. As shown in Fig. 22.1B for the simple example of two one-dimensional particles, the configuration space density is mapped into one of the two regimes ($x_1 < x_2$ or $x_1 > x_2$), depending on the chosen reference configuration. The volume to be sampled is
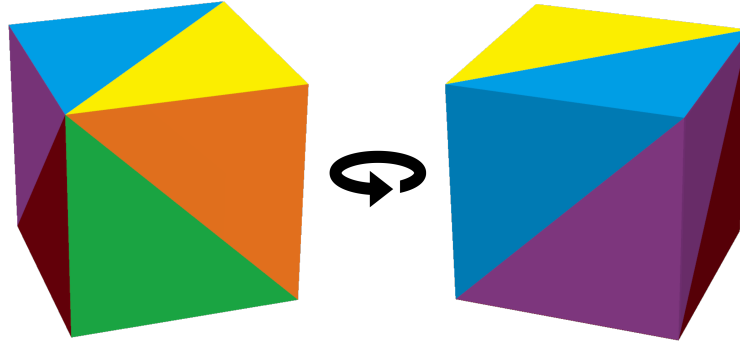
Figure 14.1: The configuration space of three one-dimensional particles split into $3! = 6$ "slices". For one-dimensional particles, the "slices" are simplices.

therefore reduced by $N!$ ($2! = 2$ in the given example), which increases the sampling density accordingly.

In 3D space, permutation reduction has the effect that each molecule is localized to a small region around its reference position, as shown in Fig. 22.1C.

Increased sampling

The definition of permutation reduction from eq. (14.1) can be reformulated as finding the permutation $\pi \in \mathcal{P}(N)$ out of all possible permutations of N molecules that fulfills

$$\sum_{i=1}^{N} (\boldsymbol{x}_{\pi(i)} - \boldsymbol{r}_i)^2 \leq \sum_{i=1}^{N} (\boldsymbol{x}_{\tau(i)} - \boldsymbol{r}_i)^2, \quad \forall \tau \in \mathcal{P}(N).$$

The permutation π is therefore accepted if the squared distance to the reference configuration $\{\boldsymbol{r}_i\}$ is smaller than that for any other permutation τ.

Writing translational configurations $\boldsymbol{X}_\pi = (\boldsymbol{x}_{\pi(1)}^\intercal, \ldots, \boldsymbol{x}_{\pi(N)}^\intercal)^\intercal \in \mathbb{R}^{3N}$ as single vectors, the definition can be further simplified to

$$\boldsymbol{X}_\pi \cdot \boldsymbol{R} \geq \boldsymbol{X}_\tau \cdot \boldsymbol{R}, \quad \forall \tau \in \mathcal{P}(N).$$

By setting $\pi = 1$ without loss of generality and by permuting the reference $\boldsymbol{R}$ rather than the configuration $\boldsymbol{X}$, the definition can be further simplified to

$$0 \leq (\boldsymbol{R} - \boldsymbol{R}_\tau) \cdot \boldsymbol{X}, \quad \forall \tau \in \mathcal{P}(N). \tag{14.2}$$

If the reference structure $\boldsymbol{R}$ is not degenerate, i.e., if $\boldsymbol{r}_i \neq \boldsymbol{r}_j, \quad \forall i \neq j$, this equation defines $N! - 1$ unique hyperplanes that divide the configuration space Γ into $N!$ "slices", as shown in Fig. 14.1.

Since all "slices" $\mathcal{S}(\boldsymbol{R})$ only differ by a permutation of the identical solvent particles, the

configuration space density

$$p(\boldsymbol{x}_1, \ldots, \boldsymbol{x}_N) = p(\boldsymbol{x}_{\pi(1)}, \ldots, \boldsymbol{x}_{\pi(N)}), \quad \forall \pi \in \mathcal{P}(N)$$

is invariant under particle permutation.

All configuration space integrals

$$\int_{\Gamma} f(p(\boldsymbol{X}))d\boldsymbol{X} = N! \cdot \int_{\mathcal{S}(\boldsymbol{R})} f(p(\boldsymbol{X}))d\boldsymbol{X},$$

can therefore be expressed as in terms of an integral over just one of the $N!$ slices $\mathcal{S}(\boldsymbol{R})$, independent of the reference $\boldsymbol{R}$ [127]. Here, $f(p(\boldsymbol{X}))$ is any arbitrary function of the configuration space density $p(\boldsymbol{X})$.

Furthermore, the sampled density reads

$$p_{\mathcal{S}(\boldsymbol{R})}(\boldsymbol{X}) = N! \cdot p(\boldsymbol{X}),$$

as the permutation reduction maps the entire trajectory into one of the $N!$ "slices". Sampling is therefore increased by $N!$.

Using the two identities above, the entropy reads [127]

$$\begin{aligned}
S_{\mathcal{S}(\boldsymbol{R})} &= -k_B \int_{\mathcal{S}(\boldsymbol{R})} d\boldsymbol{X} p_{\mathcal{S}(\boldsymbol{R})}(\boldsymbol{X}) \log p_{\mathcal{S}(\boldsymbol{R})}(\boldsymbol{X}) \\
&= -k_B \int_{\mathcal{S}(\boldsymbol{R})} d\boldsymbol{X} N! \cdot p(\boldsymbol{X}) \log[N! \cdot p(\boldsymbol{X})] \\
&= -k_B \int_{\Gamma} d\boldsymbol{X} p(\boldsymbol{X})[\log p(\boldsymbol{X}) + \log N!] \\
&= S - k_B \log N!,
\end{aligned}$$

which is therefore identical to the entropy of indistinguishable particles after the Gibbs factor has been applied.

Holes at the surface

The centers of the holes in the configuration space density are located where (at least) two particles i, j have identical positions, which can be written as

$$\boldsymbol{X}_{\text{hole}} = (\boldsymbol{x}_1^\intercal, \ldots, \underbrace{\boldsymbol{x}^\intercal}_{i}, \ldots, \underbrace{\boldsymbol{x}^\intercal}_{j}, \ldots, \boldsymbol{x}_N^\intercal)^\intercal.$$
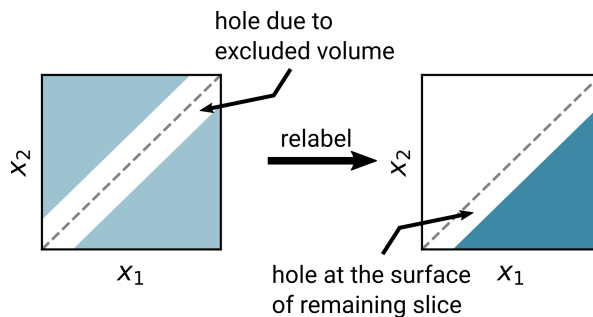
Figure 14.2: Configuration space density of two one-dimensional particles with excluded volume (left). After permutation reduction (right), the hole caused by the excluded volume is at the surface of the remaining density.

For an arbitrary reference $\boldsymbol{R}$ and the permutation $\tau = (i, j)$ that swaps particles i and j, the hyperplane normal vector in equation (14.2) reads

$$\boldsymbol{R} - \boldsymbol{R}_{(i,j)} = (0, \ldots, \underbrace{(\boldsymbol{r}_i - \boldsymbol{r}_j)^\mathsf{T}}_{i}, \ldots, \underbrace{(\boldsymbol{r}_j - \boldsymbol{r}_i)^\mathsf{T}}_{j}, \ldots, 0)^\mathsf{T}.$$

Since

$$(\boldsymbol{R} - \boldsymbol{R}_{(i,j)}) \cdot \boldsymbol{X}_{\text{hole}} = (\boldsymbol{r}_i - \boldsymbol{r}_j) \cdot \boldsymbol{x} + (\boldsymbol{r}_j - \boldsymbol{r}_i) \cdot \boldsymbol{x} = 0,$$

every hole center lies on one of the split-planes. Therefore, the holes in the configuration space density are located at the surface of the slice $\mathcal{S}(\boldsymbol{R})$, as sketched in Fig. 14.2.

Applying the quasiharmonic approximation

The subdivision of configuration space into $N!$ identical "slices" through permutation reduction alleviates the two main problems that arise when the QH approximation is applied to a solvent system. The sampling is increased by $N!$ and the configuration space density is mapped into a subspace that is more compact, without holes, and therefore presumably more Gaussian-like.

To test whether these nice properties of permutation reduction actually improve the entropy estimates, I applied the method to an argon test system, described in chapter V. Here, 1728 argon atoms were switched from an interaction-free ideal gas state to a fully interacting argon gas along the switching coordinate λ, and entropies were computed for intermediate λ values. Reference entropy values were obtained using TI, which is computationally significantly more expensive and can only be applied to small systems.

As shown in Fig. 14.3, permutation reduction (purple) indeed resulted in a significant improvement of the entropy estimates over the naive QH approximation (green), which fails to capture the effects of the particle-interactions entirely [46]. However, even with permutation reduction, the QH approximation still yields strongly overestimated entropies for interacting
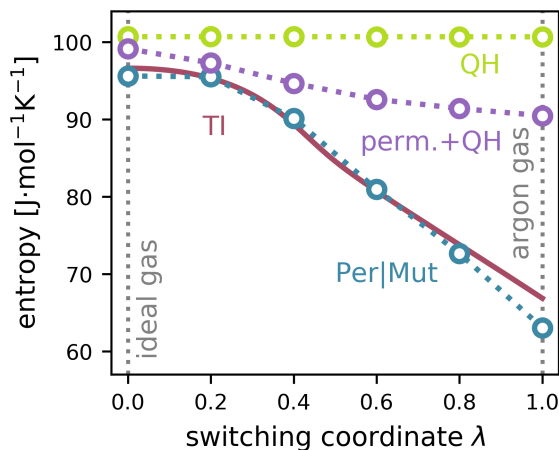
Figure 14.3: The QH approximation (green), permutation reduction + the QH approximation
(purple), and Per|Mut applied to an argon test system. TI (red) serves as an
entropy reference.

systems (larger λ).

This result demonstrates that even after permutation reduction, the configuration space
density is not sufficiently Gaussian-like to obtain accurate entropies from the QH approxima-
tion. Furthermore, the approach is limited to the translational entropies.

For these reasons, I developed the new method "Per|Mut" (blue in Fig. 14.3), which will
be presented in this thesis. The method still employs a permutation reduction to improve the
sampling but uses a mutual information expansion to calculate spatially resolved translational
and rotational entropies.

IV | Per|Mut: Rotational Entropy

The following text has been published[7] as

L.P. Heinz and H. Grubmüller, **"Computing spatially resolved rotational hydration entropies from atomistic simulations"**, *Journal of Chemical Theory and Computation*, vol. 16, no. 1, pp.108–118, 2019.

I carried out the research and wrote the manuscript. Helmut Grubmüller supervised the research and revised the manuscript.

Computing spatially resolved rotational hydration entropies from atomistic simulations

Leonard P. Heinz[1], Helmut Grubmüller[1]

[1] Department of Theoretical and Computational Biophysics, Max-Planck-Institute for Biophysical Chemistry, Göttingen, Germany
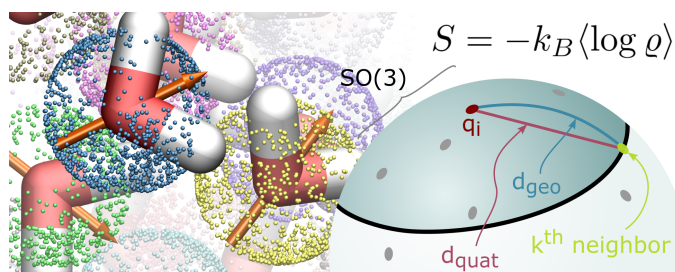
Abstract

For a first-principles understanding of macromolecular processes, a quantitative understanding of the underlying free energy landscape and in particular its entropy contribution is crucial. The stability of biomolecules, such as proteins, is governed by the hydrophobic effect, which arises from competing enthalpic and entropic contributions to the free energy of the solvent shell. While the statistical mechanics of liquids, as well as molecular dynamics simulations, have provided much insight, solvation-shell entropies remain notoriously difficult to calculate, especially when spatial resolution is required. Here, we present a method that allows for

---

the computation of spatially resolved rotational solvent entropies via a nonparametric k-nearest-neighbor density estimator. We validated our method using analytic test distributions and applied it to atomistic simulations of a water box. With an accuracy of better than $9.6\,\%$, the obtained spatial resolution should shed new light on the hydrophobic effect and the thermodynamics of solvation in general.

15 Introduction

Competing enthalpic and entropic contributions to the solvation free energies give rise to the hydrophobic effect [33], which is vital for protein function and folding [20, 28, 128]. Despite extensive theoretical work [33, 129], a quantitative understanding of the hydrophobic effect particularly at heterogeneous surfaces, such as of proteins and mixed bilayers, remains elusive.

Because surface water shows a significantly altered behavior compared to bulk [130, 131], it is essential for our understanding of the thermodynamics and energetics of protein solvation to better characterize, e.g., the relative contributions by different solvation shells or the effect of individual protein side chains on the solvent. Molecular dynamics (MD) simulations describe the hydrophobic effect at an atomic level [6, 132], but a deeper understanding of the molecular driving forces requires a quantitative and spatially resolved picture of solvation-shell thermodynamics, which poses considerable challenges.

Methods like thermodynamic integration (TI) [116, 120] allow for the calculation of solvation entropies based on MD simulations, but the lack of a spatial resolution precludes detailed analysis of how local features of the solvent-surface interface contribute and interact. Various order parameters [14–18] assess both the local translational and the local rotational order of water molecules but yield only a qualitative picture of the thermodynamic entropy.

Here, we limit our analysis to absolute rotational water entropies and present a method to reach a spatial resolution from atomistic simulations or Monte Carlo ensembles. Our method employs a mutual information expansion (MIE) to calculate the total entropy of N water molecules based on the contributions of each molecule individually and the entropy loss due to correlations between molecule pairs and triples. A similar approach was taken by, e.g., the grid inhomogeneous solvation theory (GIST) [10, 12, 13, 43, 121, 122, 133]. Rather than considering entropic contributions by correlations between individual molecules directly, GIST calculates discretized correlation integrals within voxels, which causes severe sampling problems for higher-order correlations. 3D-2-Phase-Thermodynamics (3D-2PT) [11, 134, 135] also uses voxels and approximates the system as a superposition of gas-like and solid-like components. Likewise, the Grid Cell Theory (GCT) [136] includes free energies and enthalpies, but it approximates rotational water correlation terms using a generalized Pauling's residual ice entropy model [137, 138]. Here, we address these correlations directly, convergence of which is challenging, as they require sampling and density estimates in high-dimensional configuration spaces.

In our approach, all MIE terms were calculated using a k-nearest-neighbor (kNN) density estimator, typically used in Euclidean spaces [139–141], which we modified and optimized for $SO(3)^n$, the Cartesian products of the group of rotations. We considered different metrics for the k-nearest neighbors in $SO(3)^n$, determined an optimal k-value, and provide a computationally efficient framework for rotational entropy calculation.

For easier notation, we will develop our method for water molecules, although it is general and applicable to any system with rotational degrees of freedom.

In the following sections, we will first provide the conceptual foundation and then describe our rotation entropy approach. Subsequently, we will apply it to analytical test distributions, as well as to MD water boxes.

16 Theory

16.1 Absolute entropy

Separating the entropy of water into rotational and translational contributions yields

$$S_{\text{total}} = S_{\text{rotation}} + S_{\text{translation}} - I_{\text{corr}},$$

where S_{rotation} is the entropy of the phase space distribution after projection onto the rotational degrees of freedom; $S_{\text{translation}}$, respectively, is the entropy arising from translational degrees of freedom; and the mutual information (MI) term I_{corr} quantifies the correlations between translation and rotation. In this paper, we focus on the rotational contribution S_{rotation}.

Note that some authors [10, 133, 142] define the rotational entropy as a conditional entropy, in which case it includes the MI term $-I_{\text{corr}}$.

Let the rotation of N water molecules of the simulation system be described by the Hamiltonian $\mathcal{H}(\{\boldsymbol{L}_i, \boldsymbol{\omega}_i\}) = \mathcal{T}(\{\boldsymbol{L}_i\}) + \mathcal{V}(\{\boldsymbol{\omega}_i\})$, with angular momenta $\boldsymbol{L}_i$, orientations $\boldsymbol{\omega}_i \in SO(3)$, the kinetic energy $\mathcal{T}$, and the potential energy $\mathcal{V}$, typically described by a molecular mechanics force field. The total entropy is

$$S_{\text{rotation}} = -k_B \int \frac{d\boldsymbol{L}^N d\boldsymbol{\omega}^N}{h^{3N}} \varrho \log \varrho,$$

with the Boltzmann constant k_B, Planck's constant h, and the normalized and dimensionless phase space density $\varrho = Z^{-1} \exp\left[-\frac{\mathcal{H}}{k_B T}\right] = \varrho_{\mathcal{T}} \varrho_{\mathcal{V}}$, with $\varrho_{\mathcal{T}} = Z_{\mathcal{T}}^{-1} \exp\left[-\frac{\mathcal{T}}{k_B T}\right]$, $\varrho_{\mathcal{V}} = Z_{\mathcal{V}}^{-1} \exp\left[-\frac{\mathcal{V}}{k_B T}\right]$, and the partiton function $Z = Z_{\mathcal{T}} Z_{\mathcal{V}}$. Because ϱ factorizes, the entropy can

be split into a kinetic and a configurational term

$$
\begin{aligned}
S_{\text{rotation}} = & -k_B \int \frac{d\boldsymbol{L}^N}{h_{\mathcal{T}}^{3N}} \varrho_{\mathcal{T}} \log \varrho_{\mathcal{T}} \\
& -k_B \int \frac{d\boldsymbol{\omega}^N}{h_{\mathcal{V}}^{3N}} \varrho_{\mathcal{V}} \log \varrho_{\mathcal{V}} \\
= & \underbrace{\frac{3Nk_B}{2} \log \left[\frac{2\pi e k_B T}{h_{\mathcal{T}}^2} \prod_{i=1}^{3} I_i^{\frac{1}{3}} \right]}_{S_{\text{kin}}} \\
& \underbrace{-k_B \int \frac{d\boldsymbol{\omega}^N}{h_{\mathcal{V}}^{3N}} \varrho_{\mathcal{V}} \log \varrho_{\mathcal{V}}}_{S_{\text{conf}}},
\end{aligned}
$$

where $h_{\mathcal{T}} > 0$ is arbitrary, $h_{\mathcal{V}} = h/h_{\mathcal{T}}$, and I_i are the eigenvalues of the moment-of-inertia tensor of a water molecule.

Because S_{kin} can be solved analytically, the challenge is to estimate S_{conf}.

16.2 Entropy estimation

Because the rotational entropy integral in $3N$ dimensions usually cannot be computed directly, we used a truncated mutual information expansion [4, 48–50] (see Section 16.2.1) to expand the full high-dimensional integral into multiple low-dimensional integrals over marginal distributions, which can be calculated numerically, similarly to the Inhomogeneous Solvation Theory (IST) [12, 13, 143], underlying GIST. To obtain these marginal entropies, a kNN estimator (see Section 16.2.2), which estimates the density at each sample point by finding the k closest neighboring sample points and dividing by the volume of a ball that encloses the points, was used. Here, the orientations of N water molecules in n_f different samples, e.g., frames of a computer simulation trajectory, were represented by a series of quaternions (see Section 16.2.3) $\{\boldsymbol{q}_{i,1}, \ldots, \boldsymbol{q}_{i,n_f}\}$ with $i = 1, \ldots, N$. We then defined suitable distance metrics, as required by the kNN algorithm, which are not trivial in curved spaces of rotations $SO(3)^n$ (see Section 16.2.4), and then calculated the volumes of balls, as induced by the metrics (see Section 16.2.5). We finally present a computationally efficient framework that allows finding k neighbors to each sample point (see Section 17.1).

16.2.1 Mutual information expansion

Figure 16.1A shows an example of an entropy expansion into mutual information (MI) terms of a system containing three subsystems, such as three water molecules, in a Venn diagram: The full entropy (S) is expanded into MI terms (I_m), of which the first term represents the entropies of each molecule individually and the further terms are correlation terms of second

(A)

(B)   **(C)**
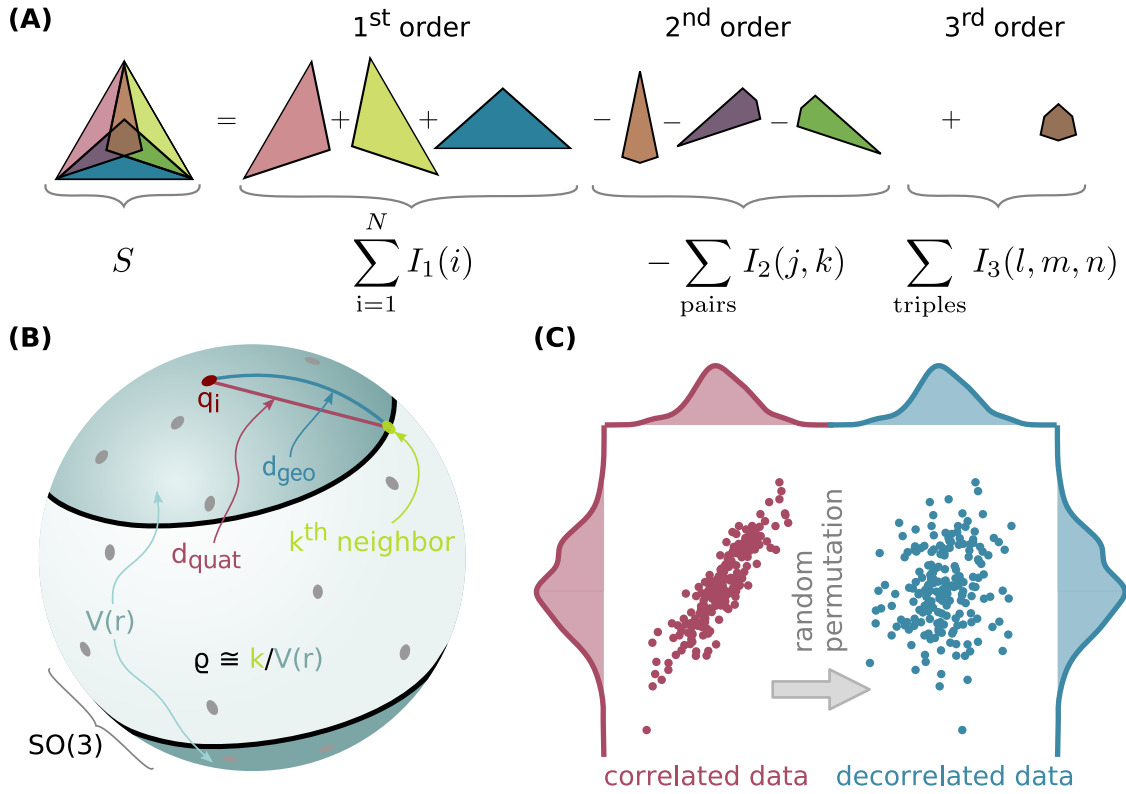
Figure 16.1: (A) Mutual information expansion illustrated for the entropy breakdown of three particles. (B) Sketch of density estimation on $SO(3)$ (here represented as a 2-sphere). Each dot on the sphere represents an orientation. For each point $\boldsymbol{x}_i$, the kth neighbor according to a distance metric (e.g., d_{quat} or d_{geo}) is found. The density is estimated via the volume $V(r)$ of a ball with radius $r = d(\cdot, \cdot)$. (C) Visualization of the fill mode approach: A correlated data set is shown on the left-hand side. The identical data is decorrelated by applying a random permutation along one axis, as shown on the right. The entropy of the decorrelated data is the sum of both "marginal entropies".

and third order, respectively,

$$I_1(i) = S(i) \tag{16.1a}$$

$$I_2(j,k) = S(j) + S(k) - S(j,k) \tag{16.1b}$$

$$I_3(l,m,n) = S(l) + S(m) + S(n) \tag{16.1c}$$
$$- S(l,m) - S(l,n) - S(m,n)$$
$$+ S(l,m,n).$$

In this notation, $S(\gamma_1, \ldots, \gamma_m)$ is the entropy of the marginal distribution with respect to molecules with indices $\gamma_1, \ldots, \gamma_m$.

For N water molecules, the expansion consists of N MI orders, of which the mth term in-

volves $(3m)$-dimensional integrals and takes all possible m-molecule correlations into account. Approximating the full entropy by a truncated expansion thus leads to lower dimensional integrals, which can be better sampled. Although there is no guarantee that truncated orders are small and can be neglected, it has been shown that a truncated expansion provides accurate entropy estimates if the correlations are short ranged [144], as for water in physiological conditions.

Here, we took up to 3-molecule correlations into account by truncating after the third order, hence,

$$
\begin{aligned}
S \approx & \sum_{i=1}^{N} I_1(i) \\
& - \sum_{(j,k)\in\mathrm{pairs}} I_2(j,k) \\
& + \sum_{(l,m,n)\in\mathrm{triples}} I_3(l,m,n),
\end{aligned}
\tag{16.2}
$$

where the first order includes the kinetic entropy contribution and a correction of $-N\log 2$, due to the 2-fold symmetry of the water molecule. The three terms are akin to the terms in IST [133]. In fact, closer analysis shows that in the thermodynamic limit, the second and third-order terms in the IST-expansion [12, 121, 122] of the molar entropy converge toward the respective terms in eq 16.2.

16.2.2 kNN entropy estimation

To evaluate eq 16.2 from a given sample of orientations $\{\boldsymbol{q}_1,\ldots,\boldsymbol{q}_{n_f}\}_i$ with $i = 1,\ldots,N$, the marginal entropies from eq 16.1 are calculated using a kNN entropy estimator [139–141, 145, 146]. For $SO(3)^1$, the kth nearest neighbor with respect to the sample point $\boldsymbol{q}_i$ is defined by a metric $d(\boldsymbol{q}_i, \boldsymbol{q}_j)$ (see Figure 16.1B), and $\varrho(\boldsymbol{q_i})$ is estimated as $(n_f-1)^{-1}k/V(r_{i,k})$, where k is a fixed integer, $V(r_{i,k})$ is the volume of a ball with radius $r_{i,k}$, the distance between $\boldsymbol{q}_i$ and its kth neighbor, and $(n_f - 1)^{-1}$ is a normalization constant. Results for $SO(3)^2$ and $SO(3)^3$ are obtained by generalizing the metric d and the volume $V(r_{i,k})$ to higher dimensions. The choice of metrices, on which the results may depend for finite sampling, and their corresponding volumes in $SO(3)^n$ will be discussed in section 16.2.4 and section 16.2.5. The entropy is

$$
\begin{aligned}
S &= -\langle \log \varrho \rangle \\
&\approx -\frac{1}{n_f} \sum_{i=1}^{n_f} \log\left(\frac{k}{(n_f-1)V(r_{i,k})}\right) - \gamma_k,
\end{aligned}
$$

where $\gamma_k = \psi(k) - \log k$ is a correction which accounts for the bias introduced by the kth neighbors being, by definition, on the edges of the balls [141]. ψ is the digamma function.

Because eqs 16.1b and 16.1c are sums and differences of integrals of different dimensional-

ities, biases are introduced: With increasing dimensionality and thus reduced sampling, the kNN estimator yields increasingly smoothed versions of the underlying true distributions. The estimator therefore overestimates entropies of distributions with higher-dimensional supports more than of those defined in lower-dimensional spaces, resulting in biases if entropies of different dimensionality are added or subtracted. To overcome this problem, the sampling space is expanded to equal dimensionality by using fill modes [4, 147]. I_2, defined in eq 16.1b as the sum of integrals in $SO(3)^1$ and $SO(3)^2$, can be rewritten as a sum of two $SO(3)^2$ integrals

$$I_2(j, k) = S(j, \hat{k}) - S(j, k)$$

if the corresponding joint distribution $\varrho(j, \hat{k})$ factorizes to $\varrho(j)\varrho(\hat{k}) = \varrho(j)\varrho(k)$. To achieve statistical independence, the sample points corresponding to index k were subjected to a random permutation $\{q_{\hat{k},1}, \ldots, q_{\hat{k},n_f}\} = \text{perm}\{q_{k,1}, \ldots, q_{k,n_f}\}$, which decorrelates $\{q_{j,1}, \ldots, q_{j,n_f}\}$ and $\{q_{k,1}, \ldots, q_{k,n_f}\}$, but leaves the marginal distributions unchanged, as sketched in Figure 16.1C. The joint entropy $S(j, \hat{k})$ is thus the sum of the initial marginal entropies $S(j) + S(k)$.

Similarly, the third-order MI term reads

$$\begin{aligned} I_3(l, m, n) = {} & 2S(\hat{l}, \hat{m}, \hat{n}) \\ & - S(l, m, \hat{n}) - S(l, \hat{m}, n) - S(\hat{l}, m, n) \\ & + S(l, m, n). \end{aligned}$$

16.2.3 Parametrization of orientations

From different parametrizations of orientations in 3D-space, such as Euler angles, Tait-Bryan angles, Hopf coordinates [148, 149], and spherical coordinates, we used quaternions [150], which, contrary to most other charts of $SO(3)$, do not suffer from Gimbal lock. They are defined as $q = (q_1, q_2, q_3, q_4)^\mathsf{T} = \pm(\cos\frac{\theta}{2}, u^\mathsf{T}\sin\frac{\theta}{2})^\mathsf{T}$, where u and θ are a normalized rotation axis and a rotation angle, respectively. q can thus be interpreted as an element of the 3-sphere, i.e., $\|q\|_2 = 1$. Because there is a one-to-one mapping of the 3-sphere to the Special Unitary group $SU(2)$, which in turn provides a double-covering of $SO(3)$, each orientation is described by two equivalent quaternions, which differ only by a sign [151].

16.2.4 Choice of metrics in $SO(3)^n$

We next considered the proper choice of metrics in $SO(3)^n$. At first sight, one might think that, of many possible metrics in $SO(3)$ [151, 152], only one, e.g., the geodesic metric $d_{\text{geo}}(q_1, q_2) = \arccos(|q_1 \cdot q_2|)$ shown in Figure 16.1B, yields the correct entropy. However, in the limit of infinite sampling, kNN entropy estimation with any metric is possible if used with its induced ball-volumes (see Section 16.2.5) [153]. Our choice was therefore guided by

the speed of convergence and computational efficiency.

We chose the quaternion metric [151, 154]

$$d_{\text{quat}}(\boldsymbol{q}_1, \boldsymbol{q}_2) = \min\{\|\boldsymbol{q}_1 - \boldsymbol{q}_2\|_2, \|\boldsymbol{q}_1 + \boldsymbol{q}_2\|_2\},$$

sketched in Figure 16.1B, which defines a metric between two rotations as the minimum Euclidean distance between unit quaternions, taking the sign ambiguity into account. In $SO(3)$, the quaternion metric and the more natural geodesic metric d_{quat} yield identical nearest neighbors. They are functionally equivalent because a positive continuous strictly increasing function h, such that $h \circ d_{\text{geo}} = d_{\text{quat}}$ (and vice versa), exists [151]. d_{quat} does not require evaluation of the inverse cosine function and thus is computationally more efficient; it was therefore preferred over d_{geo}.

Metrices in $SO(3)^2$ and $SO(3)^3$ were obtained by combining d_{quat} with the Euclidean norms in $\mathbb{R}^2$ and $\mathbb{R}^3$, respectively,

$$d_{\text{quat}^n}((\boldsymbol{q}_{1,1}, \ldots, \boldsymbol{q}_{1,n}), (\boldsymbol{q}_{2,1}, \ldots, \boldsymbol{q}_{2,n}))$$
$$= \sqrt{\sum_{i=1}^{n} d_{\text{quat}}(\boldsymbol{q}_{1,i}, \boldsymbol{q}_{2,i})^2},$$

with $(\boldsymbol{q}_{i,1}, \ldots, \boldsymbol{q}_{i,n}) \in SO(3)^n$. When combined with the Euclidean norms, the quaternion metric and the more natural geodesic metric are not functionally equivalent and hence yield, in general, different nearest neighbors. For small distances, i.e., for high sampling, the metrices are asymptotically identical.

To test whether this choice of metrics impacts the accuracy of the MI results, we compared our choice to the composite metric using d_{quat} and the maximum-norm in $\mathbb{R}^2$, which is functionally equivalent to the geodesic composite metric but slightly less efficient to evaluate than d_{quat^2}. For 10^5 frames, no significant difference between the MI values was seen.

16.2.5 Volumes of balls in $SO(3)^n$

The volumes $V(r) = \int_{d(\boldsymbol{q}_i, \boldsymbol{y}) < r} d\boldsymbol{y}$ (dark green in Figure 16.1B), enclosed by the kNN radius r, read

$$V_1(r) = 8\pi(r' - \sin r'),$$
$$r' = 2\arccos\left(1 - \frac{r^2}{2}\right),$$

for d_{quat} in $SO(3)$. The respective volumes for d_{quat^2} (in $SO(3)^2$) and d_{quat^3} (in $SO(3)^3$) reduce to

$$V_2(r) = 2^{10}\pi^2 \iint\limits_{\mathcal{V}_2} \sin^2\phi_A \sin^2\phi_B d\phi_A d\phi_B,$$

$$V_3(r) = 2^{15}\pi^3$$
$$\iiint\limits_{\mathcal{V}_3} \sin^2\phi_A \sin^2\phi_B \sin^2\phi_C d\phi_A d\phi_B d\phi_C,$$

respectively, with $\mathcal{V}_2 = \{2 - \cos\phi_A - \cos\phi_B \leq \frac{r^2}{2}\} \cap \{0 \leq \phi_A, \phi_B \leq \frac{\pi}{2}\}$ and $\mathcal{V}_3 = \{3 - \cos\phi_A - \cos\phi_B - \cos\phi_C \leq \frac{r^2}{2}\} \cap \{0 \leq \phi_A, \phi_B, \phi_C \leq \frac{\pi}{2}\}$. The integrals were solved numerically for 10^4 equally spaced values of r using the software Mathematica 10.0 [155] and the multidimensional rule; the results were stored in a lookup table. Cubic interpolation was used to obtain results from the stored values.

17   Methods

17.1   Nearest-neighbor search

Nearest-neighbor searches were performed using the Non-Metric Space Library 1.7.3.6 [109] (NMSLib)[8] and the above metrics. Each data set was indexed in a vantage-point tree [156, 157] (VP-tree) that rests on the triangle inequality. Our version of the NMSLib, modified to include the orientational metrices, is available online[9].

17.2   Accuracy assessment

17.2.1   Test distributions

To assess the accuracy of our method, we used analytical test-distributions $p^{(\mu)}$ in $SO(3)^1$, $SO(3)^2$, and $SO(3)^3$, derived from

$$p_1^{(\mu)}(\boldsymbol{q}) = \frac{1}{Z^{(\mu)}}\cos^\mu\phi_1 = \frac{1}{Z^{(\mu)}}q_1^\mu,$$

with a quaternion $\boldsymbol{q} \in SO(3)^1$, the first quaternion component q_1, the first azimuthal angle in spherical coordinates for the 3-sphere $\phi_1 \in [0, \pi/2)$, and the appropriate normalization constant $Z^{(\mu)}$ (Figure 18.1A). The analytical expression for the configurational entropy

---

[8]https://github.com/nmslib/nmslib
[9]https://gitlab.gwdg.de/lheinz/nmslib_quaternion

$\int d\boldsymbol{q} p_1^{(\mu)} \log p_1^{(\mu)}$ reads

$$S_1^{(\mu)} = \frac{1}{2} \left\{ \mu\psi\left(\frac{\mu+4}{2}\right) - \mu\psi\left(\frac{\mu+1}{2}\right) \right.$$
$$\left. +2\log\left(\frac{\Gamma(\frac{\mu+1}{2})}{\Gamma(\frac{\mu+4}{2})}\right) + \log(64\pi^3) \right\},$$

using the gamma function Γ and $k_B = h = 1$ for simpler notation. As illustrated in Figure 18.1A, the distribution depends on the localization parameter μ; a value of 0 yields a uniform distribution; larger μ values yield increasingly narrower distributions.

For $(\boldsymbol{q}_1, \boldsymbol{q}_2) \in SO(3)^2$ and $(\boldsymbol{q}_1, \boldsymbol{q}_2, \boldsymbol{q}_3) \in SO(3)^3$, probability distributions $p_2^{(\mu)}((\boldsymbol{q}_1, \boldsymbol{q}_2)) = p^{(\mu)}(\boldsymbol{q}_1)p^{(\mu)}(\boldsymbol{q}_2)$ and $p_3^{(\mu)}((\boldsymbol{q}_1, \boldsymbol{q}_2, \boldsymbol{q}_3)) = p^{(\mu)}(\boldsymbol{q}_1)p^{(\mu)}(\boldsymbol{q}_2)p^{(\mu)}(\boldsymbol{q}_3)$ were used to obtain uncorrelated distributions with entropies $S_2^{(\mu)} = 2S_1^{(\mu)}$, and $S_3^{(\mu)} = 3S_1^{(\mu)}$, respectively.

To also assess the accuracy for correlated distributions with $(\boldsymbol{q}_1, \boldsymbol{q}_2) \in SO(3)^2$, the test-distribution

$$p_{2,\text{corr}}^{(\mu)}((\boldsymbol{q}_1, \boldsymbol{q}_2)) = \frac{1}{8\pi^2 Z^{(\mu)}} \cos^{\mu}(d_{\text{geo}}(\boldsymbol{q}_1, \boldsymbol{q}_2))$$
$$= \frac{1}{8\pi^2 Z^{(\mu)}} |\boldsymbol{q}_1 \cdot \boldsymbol{q}_2|^{\mu}$$

was used, which was designed such that the marginals with respect to $\boldsymbol{q}_1$ and $\boldsymbol{q}_2$ are $p_1^{(\mu)}(\boldsymbol{q}_2)$ and $p_1^{(\mu)}(\boldsymbol{q}_1)$, respectively. The localization μ here controls the degree of correlation between $\boldsymbol{q}_1$ and $\boldsymbol{q}_2$, ranging from an uncorrelated uniform distribution ($\mu = 0$) on $SO(3)^2$ to strongly correlated distributions for larger values. The entropy of this distribution is

$$S_{2,\text{corr}}^{(\mu)} = S_1^{(\mu)} + \log(8\pi^2),$$

where $\log(8\pi^2)$ is the entropy of a free rotor.

Samples were obtained using a rejection method: First, a random point in $\boldsymbol{Q} = (\boldsymbol{q}_1, \ldots, \boldsymbol{q}_n) \in SO(3)^n$ was drawn from a uniform distribution by drawing n quaternions from the uniform distribution on the 3-sphere. Next, a random number a was drawn from a uniform distribution between 0 and $\max(p_n^{(\mu)})$. $\boldsymbol{Q}$ was accepted if $a < p_n^{(\mu)}(\boldsymbol{Q})$ and was rejected otherwise. This process was repeated until the desired number of samples was obtained.

The accuracy of our method was assessed for each test distribution for localization parameters between $\mu = 0$ and 50, nearest-neighbor k-values of $1, 5, 9$, and 13, and with 10^2 to 10^5 frames (n_f). The computed entropy and MI values were compared to the analytical results. To obtain statistical error estimates, the calculations for each parameter set was repeated 1000 times.

17.3 Molecular dynamics simulations

All MD simulations were carried out using a modified version[10] of the software package Gromacs 2018 [73–77] with an additional center of mass motion (COM) removal [158] method, used to individually constrain all oxygen atoms. We furthermore made small additional changes to apply COM removal to individual atoms and to overcome the limit of 254 COM removal groups[11]. The CHARMM36m force field [80–84] and the CHARMM-TIP3P water model [159] were used. All water molecules were subjected to SETTLE [95] constraints (i.e., rigid), and the leapfrog integrator with a time step of 2 fs was used. Electrostatic forces were calculated using the Particle-Mesh Ewald (PME) method [160] with a 1.2 nm real space cutoff; the same cutoff was used for Lennard-Jones potentials [42]. In all simulations, the V-rescale thermostat [98] with a time constant of 0.1 ps, and, if applicable, the Parrinello-Rahman barostat [100, 101] with a time constant of 1.0 ps and 1 bar pressure were used.

A total of 1728 water molecules were placed within a cubic simulation box, and the system was equilibrated for 1 ns at 300 K as a NPT ensemble. From the equilibrated system, resulting in a box size of approximately 3.7 nm, three 1 μs production runs were started, as shown in the first column of Figure 18.2. Run **m** ("mobile") was carried out as described above. To benchmark our method against the established method of thermodynamic integration (TI), a system with only rotational degrees of freedom was constructed. To this end, all oxygens were position-constrained using COM removal as shown in the first column of Figure 18.2 in run **p** ("pinned"), allowing only rotational movements around the oxygen atom under NVT conditions. The temperature was increased to 600 K, since the water molecules formed an almost rigid, ice-like hydrogen bond network at 300 K, showing only very little dynamics. Run **sp** ("sliced & pinned") was simulated like **p**, but all water molecules within a slice of 0.5 nm width were removed to create a water-vacuum interface.

17.3.1 Entropy calculation

For all three test systems, the entropy of rotation was calculated as described in section 16, each using a 1 μs trajectory with 10^5 frames. For the MI terms, a cutoff depending on the distance between average molecule positions was used. Whereas including the MIs of many molecule pairs by using a large cutoff distance gave rise to a more accurate MIE, it also introduced larger noise due to limited sampling. For pairwise MI terms, the cutoff was chosen as 1.0 nm, because for larger distances, the MI terms vanished within statistical errors (see Figure 18.2B, D). Similarly, triple MI contributions were cutoff at 0.45 nm.

Because the water molecules in system **m** were mobile, average positions across the obtained trajectory were unusable to define a cutoff. Therefore, the water molecules were relabeled in each frame, such that they remained as close as possible to a simple-cubic reference structure

---

[10]https://github.com/Tsjerk/gromacs/tree/rtc2018
[11]https://gitlab.gwdg.de/lheinz/gromacs-rtc2018_modif

using permutation reduction [46, 47], which left the physics of the system unchanged. In systems **p** and **sp**, the molecules were immobilized and the oxygen positions were used for applying the cutoff.

To quantify the precision of the method, the MD simulations and the subsequent entropy analyses were repeated in 10 independent calculations.

17.3.2 Thermodynamic integration reference

Reference entropy values for systems **p** and **sp** were obtained using thermodynamic integration [46, 116, 120] (TI). Interactions between water molecules were gradually switched off in a stepwise fashion to obtain the entropy difference between real water and noninteracting water. The absolute rotational entropy was obtained as the sum of the excess entropy, obtained via TI, and the ideal gas contribution,

$$\frac{3Nk_B}{2} \log \left[(4\pi^2)^{\frac{2}{3}} \cdot \frac{ek_BT}{2\pi\hbar^2} \prod_{i=1}^{3} I_i^{\frac{1}{3}} \right],$$

where I_i are the eigenvalues of the moment of inertia tensor of a water molecule.

Both TI calculations were performed using the soft-core [161] parameters $\alpha = 0.5$ and $\sigma = 0.3$. Coulomb interactions were linearly switched off in 80 windows of 20 ns each, and a further 10 windows were used to subsequently switch off the van-der-Waals interactions. The first nanosecond of each window was discarded.

18 Results and Discussion

18.1 Test distributions

We first assessed the accuracy of our method for three uncorrelated analytical test distributions (defined in $SO(3)^1$, $SO(3)^2$, and $SO(3)^3$) and one correlated analytical test distribution (defined in $SO(3)^2$), as described in Section 17.2.1. The distributions depend on the localization parameter μ, which, for the uncorrelated distributions $p_1^{(\mu)}, p_2^{(\mu)}$, and $p_3^{(\mu)}$, determines their width, as demonstrated in Figure 18.1A, and, for $p_{2,\text{corr}}^{(\mu)}$, controls the strength of the correlation.

As can be seen in Figure 18.1B, the kNN estimator largely agrees with the analytic results (dashed lines) for the uncorrelated distributions $p_1^{(\mu)}, p_2^{(\mu)}$, and $p_3^{(\mu)}$ for μ between 0 (uniform distribution) and 50 (strongly peaked) and the tested k-values between 1 and 13. The graphs for the three distributions are scaled and offset as indicated in the figure. We find that, for distributions $p_1^{(\mu)}$ and $p_2^{(\mu)}$, our method accurately reproduces the true entropy for all tested μ values within statistical errors, even for the small number of 100 frames. The statistical errors amount to 0.25 nats (natural units of information) for $p_2^{(\mu)}$, $k = 1$, $\mu = 50$ or less. Also for $p_3^{(\mu)}$ and $k = 1$, the analytical result is matched within statistical errors (0.28 nats
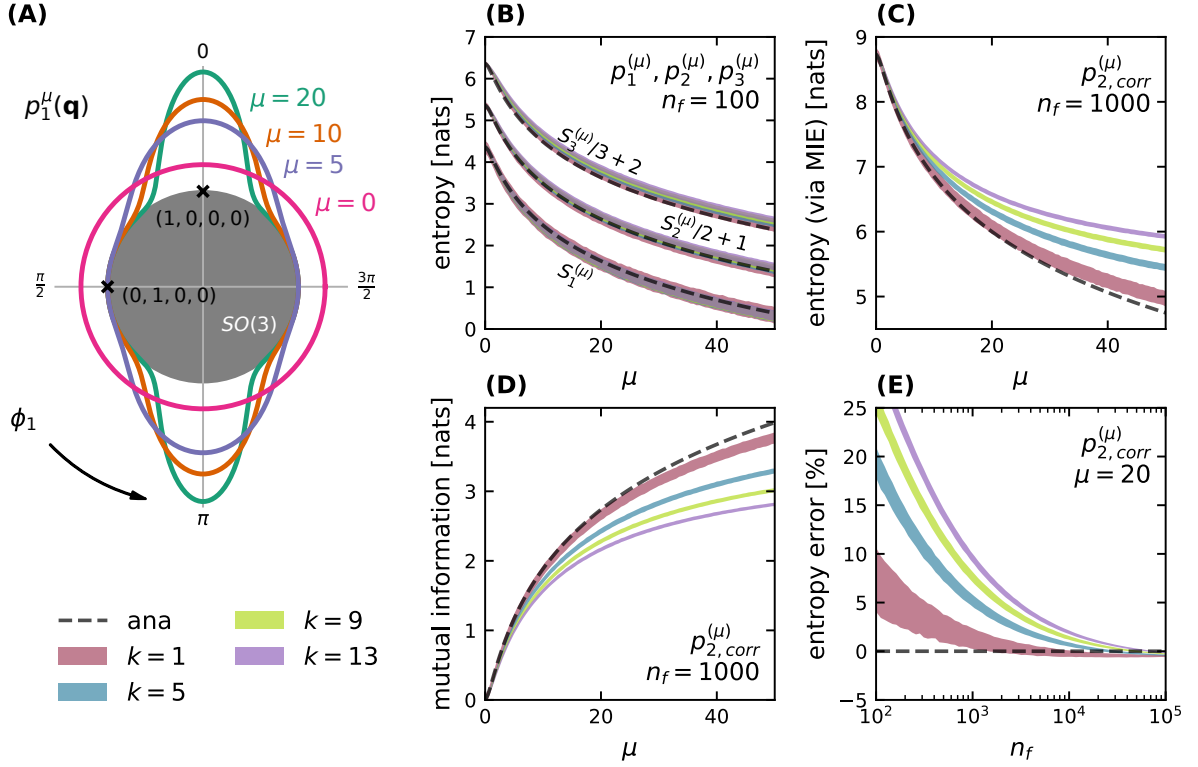
Figure 18.1:   Analytic test distribution compared to entropies and MI values obtained from density estimates. Panel (A) shows the distribution $p_1^{(\mu)}$ for increasing localizations μ, illustrated by different colors. Here, we represent $SO(3)$ as a 1-sphere and the distribution is renormalized accordingly for this 1d representation. The north pole $(1,0,0,0)$ and the quaternion $(0,1,0,0)$ are indicated with black crosses. Panels (B), (C), and (D) show entropies and MI values obtained using the test distributions $p_1^{(\mu)}, p_2^{(\mu)}, p_{2,\mathrm{corr}}^{(\mu)}$, and $p_3^{(\mu)}$ for varying coupling parameters μ. Panel (E) shows the convergence of the results for increasing sample sizes. In panels (B) to (E), the analytical result is shown by a dashed line; results for different k-values are colored according to the legend at the bottom left of the figure. Values that were fixed during the calculation, such as the choice of the test distribution, the number of frames n_f, or the coupling parameter μ, are stated in a corner of the respective panel. The shown errors denote 1σ regions.

at $\mu = 50$ at maximum), whereas larger values of k lead to overestimated entropies of up to 0.7 nats ($k = 13$, $\mu = 50$), caused by the limited sampling of just 100 frames and the increased dimensionality of $SO(3)^3$ compared to the other tested distributions.

Next, we assessed the accuracy for the correlated test distribution. The panels of Figure 18.1C, D show the entropy (calculated via the MIE as defined in eq 16.1b) and the MI of $p_{2,\text{corr}}^{(\mu)}$ for 1000 frames, respectively. For the uniform distribution ($\mu = 0$), the algorithm yields the analytic values of $2 \log(8\pi^2)$ and 0 for entropy and mutual information, respectively. With increasing correlation μ, the entropy is increasingly overestimated as MI is underestimated. Both effects are more pronounced for larger k-values: Whereas for $k = 1$, the algorithm yields accurate values within statistical errors up to a correlation of $\mu \approx 20$, the results deviate significantly for $k = 13$ even for very small μ-values. Overall, small k-values, such as $k = 1$, yield high accuracy but with reduced precision (i.e., larger statistical errors) compared to large k-values like 13, which gives rise to smaller statistical errors but reduced accuracy.

To further assess this trade-off and the convergence properties of our method, we calculated the relative entropy errors for $p_{2,\text{corr}}^{(20)}$ for sampling between 10^2 and 10^5 frames, shown in Figure 18.1E. For $k = 1$ and only 100 frames, the method overestimates the true entropy by 5 to 10%, which quickly drops to below 1% for more than $2 \cdot 10^3$ frames. For larger k-values, the entropy errors increase and the convergence becomes slower, e.g., $k = 13$ requires $2 \cdot 10^4$ frames to achieve an entropy error of less than 1%. The statistical errors at 10^5 frames are 0.11% and 0.05% for k-values of 1 and 13, respectively. Overall, $k = 1$ yields somewhat lower precision but significantly faster convergence compared to larger values, which becomes even more pronounced in higher dimensions. We therefore consider this value the optimal choice for the systems at hand and used it for all subsequent analyses.

The kNN entropy estimator rests on the assumption that the density is approximately constant and isotropic within each k-nearest-neighbor ball (see Figure 16.1B). This assumption implies that features of the true distribution that are smaller than the average distance between sample points are not resolved, which, in case of poor sampling, inevitably leads to an overestimated entropy, as seen for $p_3^{(\mu)}$ with large k or as shown in Figure 18.1E. The assumption of isotropy no longer holds for highly correlated data sets, such as $p_{2,\text{corr}}^{(\mu)}$ for large values of μ. In this case, also the k-nearest neighbors to each sample point are correlated and thus not isotropically distributed, which is not reflected by an isotropic kernel, i.e., a ball. For Euclidean spaces, this problem was addressed by using anisotropic kernels [162, 163]. Although this idea could also be applied in $SO(3)^n$, the correlation of water molecules at standard conditions is weak enough (Figure 18.2A) to allow for sufficiently accurate results under the isotropy assumption.

The trade-off between accuracy and precision with respect to the k-value is a general property of kNN entropy estimators, which has been characterized previously [163, 164], and is intuitively accessible: Whereas averaging over an increasing number of neighbors reduces statistical uncertainties and thus improves precision, the assumptions of approximately constant

isotropic densities are applied to increasingly larger balls, resulting in increasingly overestimated entropies for distributions with small scale features or strong correlations.

Overall, the kNN method with $k = 1$ yields most accurate results while being only slightly less precise than estimators with lager k. It retrieves the analytical entropies within statistical errors for the uncorrelated distributions, as well as for the correlated distribution with $\mu < 20$ using just 100 and 1000 frames, respectively.

18.2 Entropy calculated from MD simulations

Having assessed our rotational entropy method against analytic test distributions, we tested its accuracy for more realistic systems of up to 1728 interacting water molecules. To this end, we simulated three atomistic water MD systems (Figure 18.2, left column), as described in Section 17.3. For all systems, 10 independent MD simulations were performed, and for each system, entropies were calculated via a MIE as explained in Section 17.3.1.

System $\mathbf{m}$ ("mobile") comprises 1728 unconstrained water molecules. As shown in Figure 18.2A, an absolute rotational entropy of $(40.53 \pm 0.04)\,\mathrm{J \cdot mol^{-1} \cdot K^{-1}}$ per molecule is obtained, to which the first, second, and third MI orders contribute $(44.2349 \pm 0.0007)\,\mathrm{J \cdot mol^{-1} \cdot K^{-1}}$, $(-4.550 \pm 0.015)\,\mathrm{J \cdot mol^{-1} \cdot K^{-1}}$, and $(0.85 \pm 0.04)\,\mathrm{J \cdot mol^{-1} \cdot K^{-1}}$, respectively. Note that the provided values are averages and standard deviations of the 10 independent calculations and that the uncertainties are too small to be shown as error bars in Figure 18.2. The pair-mutual information terms I_2, shown in Figure 18.2B, reach a maximum of $0.8\,\mathrm{J \cdot mol^{-1} \cdot K^{-1}}$ for very close water molecules and vanish monotonically for molecules that are, after permutation reduction and on average, separated by more than $\approx 0.8\,\mathrm{nm}$. Note that the discrete nature of distances in Figure 18.2B is due to the choice of a simple cubic reference structure for permutation reduction.

To compare the obtained absolute entropies to TI [116, 120] (described in Section 17.3.2), the water movement was restricted to the rotational degrees of freedom in system $\mathbf{p}$ ("pinned") by pinning each molecule as described in Section 17.3. Here, the rotational entropy, shown in panel C, is reduced to $(29.53 \pm 0.03)\,\mathrm{J \cdot mol^{-1} \cdot K^{-1}}$. The second- and third-order mutual information terms contribute $(-18.47 \pm 0.01)\,\mathrm{J \cdot mol^{-1} \cdot K^{-1}}$ and $(4.21 \pm 0.02)\,\mathrm{J \cdot mol^{-1} \cdot K^{-1}}$, respectively. Compared to the results from TI shown in gray, the entropy is underestimated by $9.6\,\%$ due to the limited sampling of the strongly correlated system. Similar to what we observe for the analytical test case depicted in Figure 18.1D, the MI terms are underestimated for strong correlations, of which the third order is most severely affected due to the high dimensionality of the sampling space.

The I_2 terms, illustrated in Figure 18.2D, show a maximum of $7\,\mathrm{J \cdot mol^{-1} \cdot K^{-1}}$ and indicate that water molecules decorrelate beyond $\approx 0.4\,\mathrm{nm}$. The distribution shows secondary and tertiary peaks around 0.55 and 0.80 nm that arise from indirect coupling via one or two mediating water molecules, as indicated by the structures shown in Figure 18.2D. In this case, the correlations between the molecule pairs are not due to direct interactions; instead, mediating
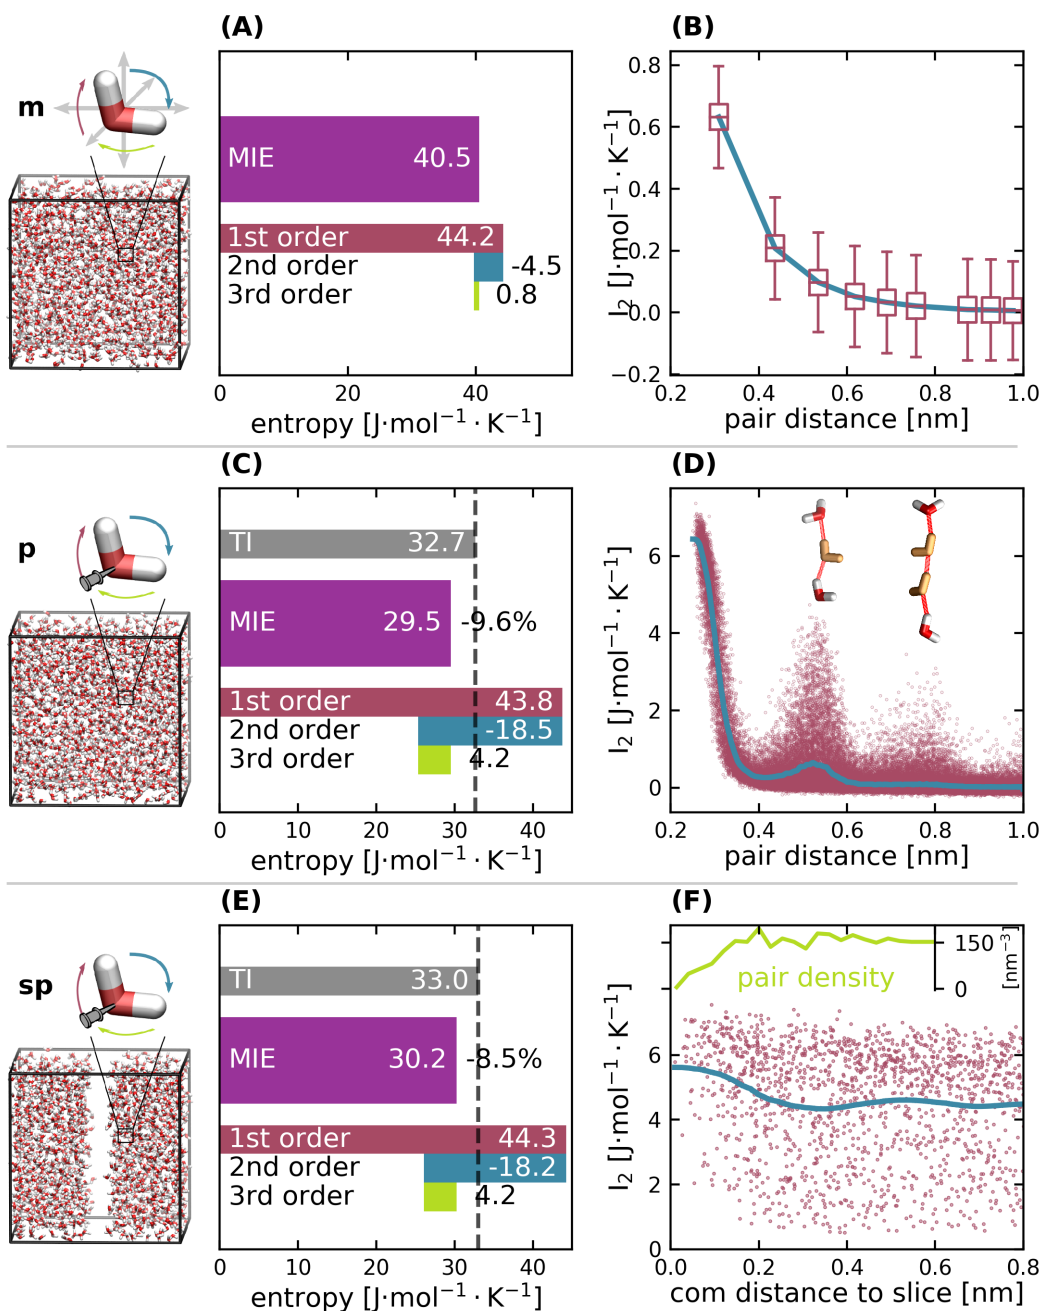
Figure 18.2: Entropies and MI contributions for the systems **m**, **p**, and **sp**. The first column
shows the three considered MD systems. Panels (A), (C), and (E) show the
rotational entropy computed using the MIE in purple; its breakdown into contri-
butions by first to third order is visualized underneath. For systems **p** and **sp**,
the result is drawn in comparison to the TI values in gray. Panels (B) and (D)
show the mutual information I_2 between all considered pairs of water molecules
depending on their distance. The blue lines correspond to running Gaussian
averages. Panel (F) displays I_2 between pairs of molecules that are closer than
0.33 nm in relation to the their distance to the vacuum slice in system **sp**. The
inset in green shows the molecule pair density with respect to the center of mass
distance to the slice.

water molecules (orange) enhance distant orientation correlations via short hydrogen-bonded chains (shown in red). This finding demonstrates that the method is able to identify regions of locally coupled water molecules and to quantify the resulting entropy losses, thus providing a spatially resolved picture of entropy changes.

To further assess and demonstrate the accuracy of the method for systems with spatial features, we included a $0.5\,\mathrm{nm}$ vacuum slice in system **sp** ("sliced & pinned", Figure 18.2), such that the dynamics of water molecules at the surface differs from those molecules in the bulk. For system **sp**, the accuracy of our entropy estimation relative to TI improves to $8.5\,\%$, whereas the contributions by higher MI orders remain almost identical (see Figure 18.2E). We assume that the improved accuracy is due to the smaller number of molecules (1728 vs 1493 with slice) and possibly because the vacuum slice limits the range of many-particle correlations that would not be captured by a third-order approximation.

Figure 18.2F shows the I_2 terms of molecule pairs that are closer than $0.33\,\mathrm{nm}$, i.e., those that are within their first hydration shells, relative to their distance to the slice. The correlations of pairs that are close to the vacuum interface are increased to $5.6\,\mathrm{J\cdot mol^{-1}\cdot K^{-1}}$ on average compared to $4.1\,\mathrm{J\cdot mol^{-1}\cdot K^{-1}}$ in bulk. Although the entropy per molecule increases compared to system **p**, mainly due to the dominating first-order term (see Figure 18.2C, E), the increased correlations at the surface and their associated entropy losses contribute to the thermodynamic unfavorability of water at a (hydrophobic) vacuum interface.

The MIE approaches the TI values for systems **p** and **sp** to $9.6\,\%$ and $8.5\,\%$, respectively, and additionally yields information about individual correlations and their associated entropy losses, thus providing spatial resolution. Remarkably, about 25-fold less computer time was required for the MIE compared to TI for the shown examples.

The large second and third-order contributions, illustrated in Figure 18.2C, E, show that both systems with pinned water exhibit strong correlations between water molecules. As for the test distributions illustrated in Figure 18.1, strong correlations result in systematically underestimated MI values. Due to their high dimensionality, and thus, low sampling density, we expect the third-order MIE contributions for systems **p** and **sp** to be mostly affected, contributing to their overall underestimated entropy. For the same reason, we expect entropies calculated from more loosely coupled mobile water to yield markedly more accurate results.

Although a direct comparison to TI is impossible for system **m**, we expect that the errors due to the truncation of higher-order MI terms, observed for the more tightly correlated systems **p** and **sp**, are larger than for unconstrained water. Therefore, the approximation of the truncated MIE yields more accurate results for realistic solute systems. These two effects combined, the performances obtained for the more correlated pinned water systems provide upper bounds for the expected errors.

19 Conclusion

We developed an estimator for spatially resolved rotational solvent entropies based on a truncated mutual information expansion and the k-nearest-neighbor algorithm on $SO(3)^n$. Accuracy and computational efficiency were assessed for both analytical test distributions and for systems of up to 1728 water molecules, described by atomistic MD simulations.

For the uncorrelated test distributions in $SO(3)^1$, $SO(3)^2$, and $SO(3)^3$, the estimator with $k = 1$ yields accurate entropies for as little as 100 sample points. For the correlated test distribution $p_2^{(\mu)}$, the entropies are overestimated for increasing coupling, caused by underestimating mutual information terms. The latter effect is especially pronounced for large k-values. Precision increased only marginally for larger k at the cost of decreased accuracy, which led us to conclude that $k = 1$ represents the best trade-off for the problem at hand. We furthermore demonstrated convergence within $2 \cdot 10^3$ frames for a correlated distribution ($\mu = 20$) and therefore expect our approach to accurately describe correlations of water molecules already in relatively short MD trajectories of $100\,\text{ns}$ to $1\,\mu\text{s}$.

For the considered MD systems, we find agreement within $9.6\,\%$ and $8.5\,\%$ with TI for pinned waters in systems **p** and **sp**, respectively, corresponding to energy deviations $(-T\Delta S)$ of $0.94\,\text{kJ·mol}^{-1}$ and $0.84\,\text{kJ·mol}^{-1}$ per water molecule at $300\,\text{K}$. The obtained rotational entropic contributions to the free energy are precise within $\pm 0.008\,\text{kJ·mol}^{-1}$ and $\pm 0.018\,\text{kJ·mol}^{-1}$, respectively. For the binding of a small ligand that displaces 10 water molecules at the binding pocket, we therefore expect to obtain absolute rotational entropy contributions corresponding to an accuracy of at least $10\,\text{kJ·mol}^{-1}$ and to resolve rotational entropy differences corresponding to at least $0.06\,\text{kJ·mol}^{-1}$. As seen in the second column of Figure 18.2, fully mobile water exhibits considerably smaller correlations than pinned water, rendering the tests using pinned water a tough benchmark compared to realistic solute systems. For a protein/water system, we would therefore expect markedly smaller error margins.

The algorithm provides spatial resolution by assessing the mutual information contributions on the level of individual molecules, distinguishing it from, e.g., GIST [10, 12, 13, 43, 121, 122]. For the hydrophobic vacuum interface, we calculated an entropy loss due to an increase in mutual information close to the surface. The ability to resolve the origin of entropy changes renders the method a promising tool to enhance our understanding of processes like the hydrophobic effect and the thermodynamics of solvated complex heterogeneous biomolecules in general.

Work on including the contributions by the translational entropy and the translation-rotation correlation to the overall entropy is in progress and will be published elsewhere. Also, our method can be extended to include intramolecular entropy contributions of flexible solvents, e.g., simulated water without SETTLE [95] constraints. In this case, additional correlation terms would arise from pairwise correlations between the internal degrees of freedom,

translational, and rotational, as well as the respective triple-correlation terms, which might be challenging to converge.

Although in this study we restricted the application and assessment of our approach to water, generalization to other solvents is straightforward. An implementation is available for download[12] as a python module [165, 166] with a C++ backend for fast neighbor search.

20 Supplement

20.1 Derivation of the test-distribution entropies

We describe elements of $SO(3)$ using quaternions that are, if treated as vectors, also elements of the 3-sphere by design. To carry out the integrations, it is therefore convenient to use the appropriate spherical coordinates:

$$q_1 = \cos\phi_1,$$
$$q_2 = \sin\phi_1\cos\phi_2,$$
$$q_3 = \sin\phi_1\sin\phi_2\cos\phi_3,$$
$$q_4 = \sin\phi_1\sin\phi_2\sin\phi_3,$$

with $\phi_1 \in [0, \pi/2)$, $\phi_2 \in [0, \pi)$, $\phi_2 \in [0, 2\pi)$. Because, for quaternions, $\boldsymbol{q}$ and $-\boldsymbol{q}$ denote the same rotation, integration over one hemisphere is sufficient. Therefore, ϕ_1 ranges from 0 to $\pi/2$ only. The surface element is given by $d^3S = 8\sin^2\phi_1\sin\phi_2 d\phi_1 d\phi_2 d\phi_3$, where the factor $8 = 2^3$ provides the necessary isotropic scaling by the factor of 2 per dimension [151]. As such, the total volume of $SO(3)$ reads

$$V = \int_{\phi_1=0}^{\pi/2}\int_{\phi_2=0}^{\pi}\int_{\phi_3=0}^{2\pi} 8\sin^2\phi_1\sin\phi_2 d\phi_1 d\phi_2 d\phi_3$$
$$= 32\pi\int_{\phi_1=0}^{\pi/2}\sin^2\phi_1 d\phi_1$$
$$= 8\pi^2,$$

which is also intuitively accessible, as the rotation of a principal axis provides 4π, i.e., the area of a 2-sphere, and a rotation around the principal axis contributes another factor of 2π. Using this result, the entropy of a free rotor in three dimensions is $\log(8\pi^2)$.

To obtain the entropy of $p_1^{(\mu)}(\boldsymbol{q}) = \frac{1}{Z^{(\mu)}}\cos^\mu\phi_1 = \frac{1}{Z^{(\mu)}}q_1^\mu$, it is first necessary to calculate

---

[12]https://gitlab.gwdg.de/lheinz/hydration_entropy

the normalization $Z^{(\mu)}$

$$Z^{(\mu)} = \int_{\phi_1=0}^{\pi/2} \int_{\phi_2=0}^{\pi} \int_{\phi_3=0}^{2\pi} 8 \cos^\mu \phi_1 \sin^2 \phi_1 \sin \phi_2 d\phi_1 d\phi_2 d\phi_3$$

$$= 32\pi \int_{\phi_1=0}^{\pi/2} \cos^\mu \phi_1 \sin^2 \phi_1 d\phi_1$$

$$= 32\pi \left[\int_{\phi_1=0}^{\pi/2} \cos^\mu d\phi_1 - \int_{\phi_1=0}^{\pi/2} \cos^{\mu+2} \phi_1 d\phi_1 \right]$$

$$= 32\pi \left[\frac{\sqrt{\pi}\Gamma\left(\frac{\mu+1}{2}\right)}{2\Gamma\left(\frac{\mu+2}{2}\right)} - \frac{\sqrt{\pi}\Gamma\left(\frac{\mu+3}{2}\right)}{2\Gamma\left(\frac{\mu+4}{2}\right)} \right]$$

$$= 8\pi^{\frac{3}{2}} \frac{\Gamma\left(\frac{\mu+1}{2}\right)}{\Gamma\left(\frac{\mu+4}{2}\right)},$$

where Γ is the gamma function. The entropy of $p_1^{(\mu)}(\boldsymbol{q})$ is

$$
\begin{aligned}
S_1^{(\mu)} &= -\langle \log p_1^{(\mu)}(\boldsymbol{q}) \rangle \\
&= \log Z^{(\mu)} - \frac{8}{Z^{(\mu)}} \int_{\phi_1=0}^{\pi/2} \int_{\phi_2=0}^{\pi} \int_{\phi_3=0}^{2\pi} \log(\cos^\mu \phi_1) \cos^\mu \phi_1 \sin^2 \phi_1 \sin \phi_2 d\phi_1 d\phi_2 d\phi_3 \\
&= \log Z^{(\mu)} - \frac{32\pi\mu}{Z^{(\mu)}} \int_{\phi_1=0}^{\pi/2} \log(\cos \phi_1) \cos^\mu \phi_1 \sin^2 \phi_1 d\phi_1 \\
&= \log Z^{(\mu)} - \frac{32\pi\mu}{Z^{(\mu)}} \int_{\phi_1=0}^{\pi/2} \frac{d}{d\mu} \cos^\mu \phi_1 \sin^2 \phi_1 d\phi_1 \\
&= \log Z^{(\mu)} - \frac{32\pi\mu}{Z^{(\mu)}} \frac{d}{d\mu} \int_{\phi_1=0}^{\pi/2} \cos^\mu \phi_1 \sin^2 \phi_1 d\phi_1 \\
&= \log Z^{(\mu)} - \frac{\mu}{Z^{(\mu)}} \frac{d}{d\mu} Z^{(\mu)} \\
&= \log Z^{(\mu)} - \mu \frac{\Gamma\left(\frac{\mu+4}{2}\right)}{\Gamma\left(\frac{\mu+1}{2}\right)} \frac{d}{d\mu} \frac{\Gamma\left(\frac{\mu+1}{2}\right)}{\Gamma\left(\frac{\mu+4}{2}\right)} \\
&= \log Z^{(\mu)} + \frac{\mu}{2} \left\{ \frac{\Gamma'\left(\frac{\mu+4}{2}\right)}{\Gamma\left(\frac{\mu+4}{2}\right)} - \frac{\Gamma'\left(\frac{\mu+1}{2}\right)}{\Gamma\left(\frac{\mu+1}{2}\right)} \right\} \\
&= \log Z^{(\mu)} + \frac{\mu}{2} \left\{ \psi\left(\frac{\mu+4}{2}\right) - \psi\left(\frac{\mu+1}{2}\right) \right\} \\
&= \log\left(8\pi^{\frac{3}{2}}\right) + \log\left(\frac{\Gamma\left(\frac{\mu+1}{2}\right)}{\Gamma\left(\frac{\mu+4}{2}\right)}\right) + \frac{\mu}{2} \left\{ \psi\left(\frac{\mu+4}{2}\right) - \psi\left(\frac{\mu+1}{2}\right) \right\} \\
&= \frac{1}{2} \left\{ \mu\psi\left(\frac{\mu+4}{2}\right) - \mu\psi\left(\frac{\mu+1}{2}\right) + 2\log\left(\frac{\Gamma\left(\frac{\mu+1}{2}\right)}{\Gamma\left(\frac{\mu+4}{2}\right)}\right) + \log\left(64\pi^3\right) \right\},
\end{aligned}
$$

with Γ as the gamma function and ψ as the digamma function.

The normalization of $p_{2,corr}^{(\mu)}((\boldsymbol{q}_1, \boldsymbol{q}_2)) \propto \cos^\mu(|\boldsymbol{q}_1 \cdot \boldsymbol{q}_2|)$ is obtained by using the translation-invariance of $\boldsymbol{q}_1 \cdot \boldsymbol{q}_2$ as $8\pi^2 Z^{(\mu)}$. In the same fashion, the entropy reads $S_{2,corr}^{(\mu)} = S_1^{(\mu)} + \log(8\pi^2)$.

V | Per|Mut: Translational Entropy and Correlation Term

The following text has been submitted as

L.P. Heinz and H. Grubmüller, **"Per|Mut: Spatially resolved hydration entropies from atomistic simulations"** to the *Journal of Chemical Theory and Computation* and is currently under review.

I carried out the research and wrote the manuscript. Helmut Grubmüller supervised the research and revised the manuscript.

Per|Mut: Spatially resolved hydration entropies from atomistic simulations

Leonard P. Heinz[1], Helmut Grubmüller[1]

[1] Department of Theoretical and Computational Biophysics, Max-Planck-Institute for Biophysical Chemistry, Göttingen, Germany

MD simulation   permutation reduction   mutual information expansion
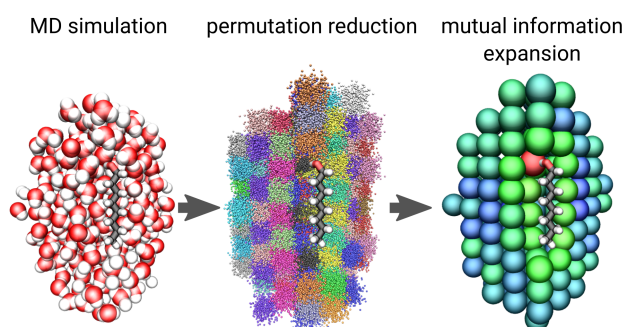
Abstract

The hydrophobic effect is essential for many biophysical phenomena and processes. It is governed by a fine-tuned balance between enthalpy and entropy contributions from the hydration shell. Whereas enthalpies can in principle be calculated from an atomistic simulation trajectory, calculating solvation entropies by sampling the extremely large configuration space is challenging and often impossible. Furthermore, to qualitatively understand how the balance

is affected by individual side chains, chemical groups, or the protein topology, a local description of the hydration entropy is required. In this study, we present and assess the new method "Per|Mut", which uses a permutation reduction to alleviate the sampling problem by a factor of $N!$ and employs a mutual information expansion to the third order to obtain spatially resolved hydration entropies. We tested the method on an argon system, a series of solvated n-alkanes, and solvated octanol.

21 Introduction

The thermodynamics of the hydration shell plays an important role in many biophyiscal processes, such as phase separation, membrane and micelle formation [6, 167, 168], or the function and folding of proteins [20, 28]. These processes are driven by the hydrophobic effect, [27, 33, 128, 169, 170] which results from a delicate balance between entropic and enthalpic contributions of the first few solvent layers but is quantitatively not yet fully understood [35]. A better understanding of the behavior of water molecules at heterogeneous surfaces is therefore necessary.

Atomistic molecular dynamics (MD) simulations have proven to reproduce the effects of hydrophobic interaction quantitatively [6, 132]. However, the lack of a straightforward way to quantify the hydration entropy contributions of specific side chains or functional groups of atoms precludes a deeper understanding of the molecular driving forces and the energetics of solvation. Further, the shallow energy landscape that governs the dynamics of the solvent molecules requires sampling of an extremely large configuration space and thus poses a severe challenge for entropy calculation.

The solvation entropy of a simple system can be calculated using thermodynamic integration (TI) [116, 120], which, however, lacks spatial resolution and a vast amount of sampling is needed for more complex systems.

3D-2-Phase-Thermodynamics (3D-2PT) [11, 134, 135] is a voxel-based approach and thus yields spatially resolved hydration entropies, but relies on the assumption that the system can be treated as a superposition of gas-like and solid-like components.

In the grid inhomogeneous solvation theory (GIST) [10, 12, 13, 43, 121, 122, 133], the entropy is approximated by a truncated expansion of single-body and multi-body correlation functions, which are calculated on a three-dimensional grid. Although the method also provides spatial resolution, the GIST expansion is usually truncated at the single-body, rarely at the two-body correlation term and therefore does not capture most multibody effects, which, as will be shown below, are important.

To address these issues, we developed an MD-based method to calculate the spatially resolved hydration entropy from atomistic simulations via permutation reduction and a mutual information expansion ("Per|Mut").

In a first step, sampling of the configuration space is enhanced by a factor of $N!$ by taking

advantage of the permutation symmetry of the identical solvent molecules. To calculate the entropy from the permutationally reduced trajectory, we use a mutual information expansion (MIE). The latter step is similar in spirit to GIST but here is based on particle positions rather than voxels, and all correlations up to three-body correlations are routinely included.

The method yields spatially resolved entropy contributions from translational and rotational degrees of freedom as well as from their correlation on a per-molecule level. The distinction between first-order entropy and contributions from many-body correlations furthermore provides an interpretation of the physical origins of solvent-driven free energy changes.

We have addressed the rotational part of solvent entropies in a previous publication [112], and will therefore focus the theory below on the translational part and the mutual information (MI) term that describes the correlation between translational and rotational degrees of freedom. Subsequently, we will assess the accuracy and convergence on 1728 argon atoms in section 24.1. In sections 24.2 and 24.3, we will apply Per|Mut to hydrated systems, namely to solvated alkanes and to octanol.

22 Theory

22.1 Absolute entropy

We separate the total entropy S_{tot} into the sum of contributions from the translational degrees of freedom S_{trans}, the rotational degrees of freedom S_{rot}, and a MI term $-I_{\text{trans-rot}}$, which quantifies the correlation between translational and rotational motions. Note that rotational entropy is also defined as a conditional entropy by some authors [10, 133, 142], in which case, it includes the MI term $-I_{\text{trans-rot}}$.

The translational entropy is

$$S_{\text{trans}} = -k_B \int \frac{d\boldsymbol{p}^N d\boldsymbol{x}^N}{h^{3N}} \varrho \log \varrho,$$

with the Boltzmann constant k_B, Planck's constant h, momenta $\{\boldsymbol{p}_i\}$, and positions $\{\boldsymbol{x}_i\}$ for N solvent molecules, the normalized and dimensionless phase space density $\varrho = Z^{-1} \exp\left[-\frac{\mathcal{H}}{k_B T}\right]$ with the Hamiltonian $\mathcal{H}$, and the partiton function Z. The translational entropy, in turn, separates into a kinetic part, which can be calculated analytically, and the configurational contribution is

$$S_{\text{trans}} = \underbrace{\frac{3Nk_B}{2} \log\left[\frac{2\pi e m k_B T}{h_{\mathcal{T}}^2}\right]}_{S_{\text{kin}}} \underbrace{-k_B \int \frac{d\boldsymbol{x}^N}{h_{\mathcal{V}}^{3N}} \varrho_{\mathcal{V}} \log \varrho_{\mathcal{V}}}_{S_{\text{conf}}},$$

where m is the mass of a water molecule, $h_{\mathcal{V}} = h/h_{\mathcal{T}}$ with arbitrary $h_{\mathcal{T}} > 0$, and $\varrho_{\mathcal{V}}$ is the probability density in $3N$-dimensional configuration space. We will therefore focus on S_{conf}.

We calculate S_{conf} by first carrying out an atomistic molecular dynamics (MD) simulation of solvent and solute. We then employ a permutation reduction of the identical solvent molecules as described in section 22.2 to increase sampling by a factor of $N!$. Finally, we calculate entropies from the permuted trajectory via a third-order MIE with a k-nearest neighbor density estimator, as described in section 22.3.

A MIE is also used in the same manner to estimate the correlation term $-I_{\text{trans-rot}}$.

22.2 Permutation reduction

The sampling of S_{conf}, and thus the calculation of solvent entropy from computer simulations, generally suffers from slow convergence. Contrary to the collective motion of macromolecules, which is usually highly coupled, the diffusive motion of solvent molecules leaves almost the entire configuration space accessible, thus rendering sufficient sampling practically impossible.

The volume of the full $3N$-dimensional translational configuration space of N water molecules with a constant density scales as N^N, which renders it impossible to numerically sample the configuration space even for a small system of ≈ 100 water molecules. This problem arises because the concept of phase space (or configuration space) inherently assigns unique labels to physically identical water molecules. Therefore, each microstate is counted redundantly $N!$ times, where the physically identical microstates only differ by a permutation of the indistinguishable water molecules, leading to a configuration space volume $N!$-times larger than physically necessary.

In an analytical treatment, this problem is, of course, solved by the Gibbs factor $N!$ [44]. For a numerical approach, however, the Gibbs factor cannot be straightforwardly applied. Permutation reduction [46, 47] solves this problem by relabeling (i.e., permuting) the solvent molecules in each frame of an atomistic trajectory such that the trajectory is mapped into a configurational subspace with a volume reduced by the Gibbs factor $N!$. Here, we summarize the aspects of permutation reduction that are relevant for the entropy estimate; an in-depth derivation is provided elsewhere [46, 47].

The alleviation of the sampling problem is achieved by choosing a permutation π for each frame of the trajectory $\{\boldsymbol{x}_i(t)\}$ that minimizes the distance

$$\sum_{i=1}^{N} \|\boldsymbol{x}_{\pi(i)}(t) - \boldsymbol{r}_i\|^2$$

with respect to an arbitrary but fixed set of reference positions $\{\boldsymbol{r}_i\}$.

Figure 22.1A demonstrates this approach for the simplest case of two water molecules. In the right panel, the water molecules have moved away from their reference positions (shown on the left), such that the deviation from the reference is minimized if the labels are swapped.

Demonstrating permutation reduction on ...

(A) ... two water molecules.

reference   for all frames

relabel

find permutation π
such that
$$\sum_{i=1}^{N} \|\vec{x}_{\pi(i)}(t) - \vec{r}_i\|^2$$
is minimized

(B) ... the configuration space of two 1D particles.

$x_1 < x_2$

$x_2 < x_1$

relabel

increase of
sampling by N!

$x_2 < x_1$

(C) ... a superimposed trajectory of 200 molecules.

original trajectory   relabeled trajectory
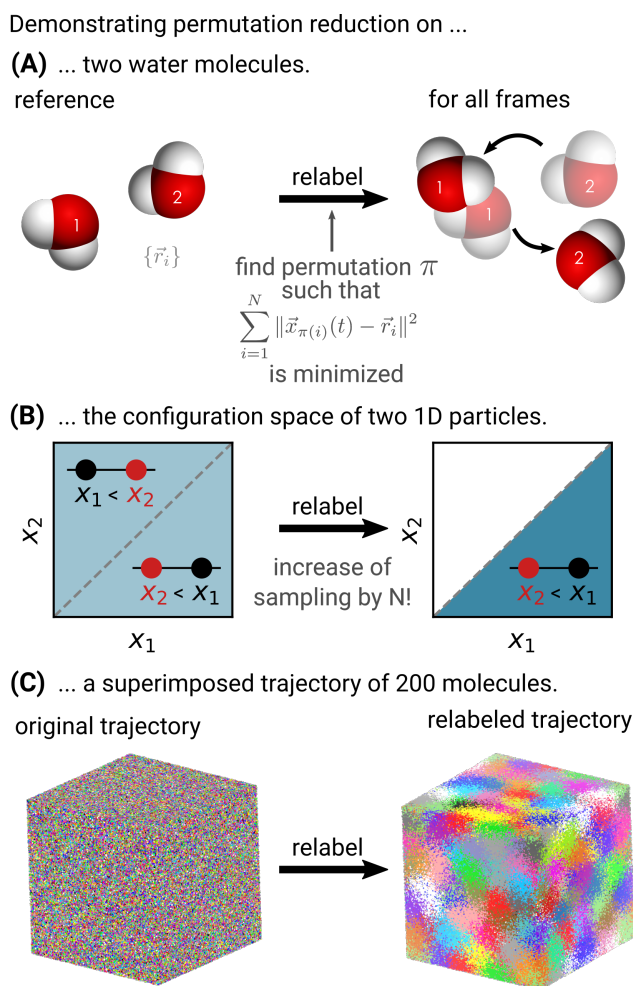

relabel

Figure 22.1: Concept of permutation reduction on (A) two water molecules, (B) two one-dimensional molecules, and (C) a superimposed trajectory of 200 molecules. In the last case, each color represents a single molecule.

Figure 22.1B visualizes the effect on the accessible configuration space for two one-dimensional molecules. Before permutation reduction, the system is either in a regime of $x_1 < x_2$ or $x_1 > x_2$, with the accessible configuration space marked in blue. After permutation reduction, the molecules are relabeled such that, depending on the reference configuration, the system remains in one of the two regimes. Thereby, the configuration space volume is reduced by a factor of 2!. Although this factor seems small, note that it is $N!$ for N molecules and hence represents an enormous alleviation of the sampling problem for larger N.

Figure 22.1C shows the effect of permutation reduction on 3D space for a box of 200 molecules. Now, each molecule is localized to a small region centered at the reference position. Importantly, the physics of the system is unchanged.

22.3 Entropy estimation

Even after permutation reduction, S_{conf} cannot be computed directly because it still requires sampling of a $3N$-dimensional space. Instead, we use a MIE [4, 48–50], and calculate its terms with a k-nearest neighbor estimator (kNN) [139–141, 145, 146], as previously described for rotational entropies [112].

The MIE

$$
\begin{aligned}
S \approx & \sum_{i=1}^{N} S_1(i) \\
& - \sum_{(j,k)\in\text{pairs}} I_2(j,k) \\
& + \sum_{(l,m,n)\in\text{triples}} I_3(l,m,n)
\end{aligned}
$$

is truncated after the 3-molecule correlation term, where the indices $1 \leq i,j,k,l,m,n \leq N$ label individual molecules. The first term is the sum of all single-molecule entropies, i.e., the entropies of the marginal distributions of individual molecules, and therefore does not take correlations between molecules into account. Pairwise and triplewise correlations are described by the second and third terms, respectively, for which the mutual information terms read

$$
I_2(j,k) = S(j) + S(k) - S(j,k) \tag{22.1a}
$$

$$
\begin{aligned}
I_3(l,m,n) = & \, S(l) + S(m) + S(n) \\
& - S(l,m) - S(l,n) - S(m,n) \\
& + S(l,m,n).
\end{aligned} \tag{22.1b}
$$

Here, the entropy terms again denote the entropies of the marginal distributions, where all degrees of freedom, except the ones of the molecules specified by the indices, are projected out. Because for a third-order expansion, a 9-dimensional space needs to be sampled, the expansion converges for $\approx 10^5$ samples. Higher-order terms are increasingly harder to sample, but typically have small contributions for short-ranged interactions and are therefore neglected [144].

The individual entropies of eq 22.1 are calculated using a kNN estimator [112]. The entropy of a given trajectory of n_f configurational coordinates $\{\boldsymbol{x}_1, \ldots, \boldsymbol{x}_{n_f}\}$ with $\boldsymbol{x}_i \in \mathbb{R}^d$, $d = 3$ (single-molecule term), $d = 6$ (pair term), or $d = 9$ (triple term) is given by

$$
\begin{aligned}
\frac{S}{k_B} &= -\langle \log \varrho \rangle \\
&\approx \frac{1}{n_f} \sum_{i=1}^{n_f} \log\left((n_f - 1) V^{(d)}(r_{i,k}) \right) - \psi(k),
\end{aligned}
$$

where k is a fixed positive integer, ψ is the digamma function, $r_{i,k}$ is the distance from $\boldsymbol{x}_i$ to

its kth neighbor using the Euclidean metric, and $V^{(d)}(r_{i,k}) = \pi^{\frac{d}{2}} r_{i,k}^d \Gamma(\frac{d}{2}+1)^{-1}$ is the volume of the $(d-1)$-sphere with radius $r_{i,k}$.

The correlation term $I_{\text{trans-rot}}$ is calculated as

$$I_{\text{trans-rot}} \approx \sum_{(j,\tilde{k}) \in \text{pairs}} I_2(j, \tilde{k}),$$

where the index j denotes the translational degrees of freedom of molecule j, and $\tilde{k}$ denotes the rotational degrees of freedom of molecule k.

To apply a kNN entropy estimator in the product space of $\mathbb{R}^3$ and the group of orientations $SO(3)$, we use the composite metric

$$d((\boldsymbol{x}_1, \boldsymbol{q}_1), (\boldsymbol{x}_2, \boldsymbol{q}_2))$$
$$= \sqrt{[\xi d_{\text{eucl}}(\boldsymbol{x}_1, \boldsymbol{x}_2)]^2 + d_{\text{quat}}(\boldsymbol{q}_1, \boldsymbol{q}_2)^2},$$

where the quaternions $\boldsymbol{q}_1$ and $\boldsymbol{q}_2$ describe molecular orientations, d_{eucl} is the Euclidean metric, and $d_{\text{quat}}(\boldsymbol{q}_1, \boldsymbol{q}_2) = \min\{\|\boldsymbol{q}_1 - \boldsymbol{q}_2\|_2, \|\boldsymbol{q}_1 + \boldsymbol{q}_2\|_2\}$ is the quaternion metric [112, 151]. The scaling factor ξ ensures equal units under the square root, where the distance in Euclidean space is measured in nanometers, and the distance in quaternion space is unitless. Its numerical value is chosen such that the typical distances in $\mathbb{R}^3$ and $SO(3)$ are of the same magnitude. For liquid water, a value of $\xi = 10 \, \text{nm}^{-1}$ is used, but tests with $20 \, \text{nm}^{-1}$ and $30 \, \text{nm}^{-1}$ did not yield significantly different results.

The volume of the ball, induced by the metric d in $\mathbb{R}^3 \times SO(3)$ reads

$$V^{(6)}(r) = 128 \xi^3 \pi^2 \int_{\mathcal{V}} R^2 \sin^2 \phi \, dR d\phi$$
$$\mathcal{V} = \{R^2 + 2(1 - \cos\phi) \leq r^2\}$$
$$\cap \{0 \leq \phi \leq \frac{\pi}{2}\} \cap \{0 \leq R \leq \infty\}$$

and was calculated numerically using the software Mathematica 10.0 [112, 155].

23 Methods

MD simulations were carried out using the software package Gromacs 2018 [73–77] with a leapfrog integrator (2 fs time step) and the CHARMM36m force field [81, 83, 84]. Bonds were constrained using SETTLE [95] (water molecules) and LINCS [96] (other bonds to hydrogen atoms). The V-rescale thermostat [98] with a time constant of 0.1 ps at 300 K was used in all simulations and NPT runs were performed using the Parrinello-Rahman barostat [100, 101] with a time constant of 1.0 ps and 1 bar pressure. Lennard-Jones potentials [42] were cutoff at 1.2 nm. The same value was used as the real-space cutoff of electrostatic interactions in

the Particle-Mesh Ewald (PME) method [160].

All production trajectories used for the entropy estimates were $1\,\mu s$ long; configurations were recorded every $10\,ps$.

23.1 Argon

To mimic the number density of liquid water, 1728 argon atoms were placed in a $(4\,nm)^3$ cubic box (equivalent to $\approx 10000\,bar$ pressure) and simulated under NVT conditions. Despite the large pressure, no crystallization occurred and the system remained diffusive. Translational entropies were calculated using permutation reduction and a mutual information expansion (Per|Mut), and its accuracy was assessed using reference entropies obtained via the more expensive thermodynamic integration (TI) [46, 112, 116, 120].

The TI was performed using 200 steps, during which the interactions were switched linearly from an ideal gas state ($\lambda = 0$), for which the entropy is known analytically, to full argon-argon interactions ($\lambda = 1$). The simulation runs for each step lasted $100\,ns$ and were carried out with soft-core [171] parameters $\alpha = 0.5$ and $p = 1$. Errors were estimated as the difference to a second TI with only $50\,ns$ per step.

Per|Mut entropies were calculated as described in section 22. Permutation reduction of the $1\,\mu s$ trajectory was carried out using a $12 \times 12 \times 12$ simple cubic reference configuration $\{\boldsymbol{r}_i\}$. Pairwise MI terms were calculated for atoms with an average distance of less than $1.0\,nm$ after permutation reduction. Similarly, a $0.45\,nm$ cutoff was used for third-order MI terms. A kNN value of $k = 1$ was used for all MI terms. Error bars were estimated from the standard deviation of the entropies of the individual atoms.

To compare the Per|Mut results for intermediate values of the TI switching-coordinate, Per|Mut was also applied to $1\,\mu s$ trajectories with λ-values of 0.8, 0.6, 0.4, 0.2, and 0.0 (ideal gas), as shown in Figure 24.1.

23.2 Alkanes and octanol

Per|Mut entropies were calculated for the n-alkanes between ethane and decane, as well as for octanol. Each solute molecule was solvated by 1728 water molecules in a cubic box and simulated under NPT conditions as described above. To prevent the solutes from deviating from their initial linear configurations, the positions of all atoms were fixed, such that only the water molecules remained mobile. Permutation reduction was carried out using a $12 \times 12 \times 12$ simple cubic reference configuration and the MI terms were calculated using $k = 1$, as well as $1.0\,nm$ and $0.33\,nm$ cutoffs for second-order and third-order terms, respectively. Errors were estimated from the standard deviation of entropies of bulk-phase water molecules, assuming that molecules close to the solute are subject to the same spread.

For alkanes, reference entropies were again obtained using TI for ethane, propane, pentane, octane, and decane, where the solutes were grown in a water box during 200 steps for the

Lennard-Jones interactions and an additional 50 steps for the Coulomb interactions. Each step lasted 50 ns and soft-core parameters identical to that of the argon TI were used.

For the spatially resolved entropy map of octanol (Figure 24.3), the entropy per permutation-localized water molecule (see Figure 22.1C) was calculated by splitting the contributions from pair (second MI order) and triple correlations (third MI order) equally between the involved molecules. The simulation box was divided into $128 \times 128 \times 128$ voxels, and for each voxel, the local entropy was given by the average of the entropies per water molecule, weighted according to the simulation trajectory.

23.3 Nearest-neighbor search

Nearest-neighbor searches for the kNN estimator in Euclidean space were performed using a KD-tree and the Python module scikit-learn 0.20.3 [111]. Nearest-neighbor searches in quaternion-space and the composite space were carried out using the Non-Metric Space Library 1.7.3.6 [109, 112].

24 Results and Discussion

24.1 Argon

To assess the accuracy of the translational entropies, we used a test system of 1728 argon atoms and compared entropies calculated with Per|Mut to entropies from thermodynamic integration (TI), as described in section 23.1. TI is computationally more expensive than Per|Mut, does not yield spatial resolution, and is unsuitable for more complex systems but can serve as a reference for the test system. During TI, we changed the interactions between the argon atoms from a noninteracting ideal gas state ($\lambda = 0$) to their normal interatom interactions at $\lambda = 1$ and calculated the entropy change along the switching coordinate λ. We subsequently used Per|Mut to calculate translational entropies along the switching coordinate for $\lambda = 0$ (ideal gas), 0.2, 0.4, 0.6, 0.8, and 1.0 (full argon gas) to compare with the respective reference values. As shown in Figure 24.1, Per|Mut closely follows the reference TI values for all λ-values with a maximum deviation of 5.7 %.

For the ideal gas state, the third-order Per|Mut expansion yields a quite accurate entropy of $95.6\,\mathrm{J\cdot mol^{-1}\cdot K^{-1}}$, and thus deviates by only 1 % from the reference value $96.6\,\mathrm{J\cdot mol^{-1}\cdot K^{-1}}$. In this state, the first order of Per|Mut contributes $100.4\,\mathrm{J\cdot mol^{-1}\cdot K^{-1}}$, the second-order term reduces the entropy by $5.4\,\mathrm{J\cdot mol^{-1}\cdot K^{-1}}$, and the third-order contributes a further $0.5\,\mathrm{J\cdot mol^{-1}\cdot K^{-1}}$.

As the atomic interactions are switched on, the reference TI entropy decreases by $29.8\,\mathrm{J\cdot mol^{-1}\cdot K^{-1}}$ to $66.7\,\mathrm{J\cdot mol^{-1}\cdot K^{-1}}$ for $\lambda = 1$. The Per|Mut entropies follow the same trend; for a full argon gas, the third-order Per|Mut yields an entropy of $63.0\,\mathrm{J\cdot mol^{-1}\cdot K^{-1}}$, which is within 5.7 % from the reference value. The reduction of the first-order term (light
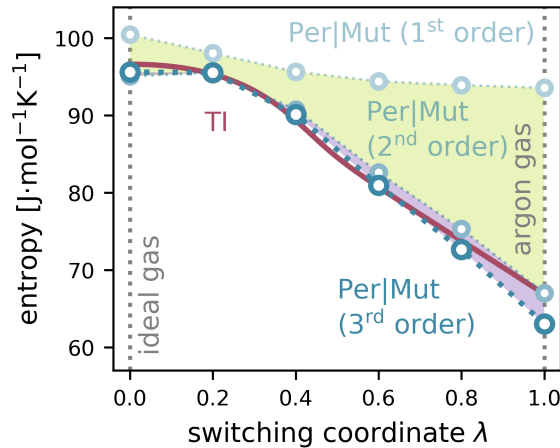
Figure 24.1: Entropy per particle of the argon system from TI (red) and Per|Mut (blue) along the switch from an ideal gas state (left) to a full argon gas (right). The MI expansions up to the first-, second-, and third order are shown with increasing opacity. The contribution by the second order, i.e., the difference between the first- and second-order expansion, is shaded in green. The contribution by the third order is shaded in magenta. Error bars are too small to be shown.

blue in Figure 24.1) reflects the effect of the decreased accessible volume due to the excluded volumes of the interacting particles and amounts to just 6.9 J·mol^{-1}·K^{-1}, which is 21.1 % of the total entropy loss. Significantly more entropy is lost due to correlated particle movement, as reflected by the second- and third-order contributions. Here, the pair correlations (second order) dominate by accounting for 21.2 J·mol^{-1}·K^{-1}, 65.0 % of the overall entropy loss, whereas the three-particle correlations (third order) contribute 4.5 J·mol^{-1}·K^{-1} (13.9 %).

Since the particles in an ideal gas are by definition uncorrelated, it might seem surprising that for $\lambda = 0$, the second- and third-order contributions of Per|Mut are nonzero, due to the permutation reduction. As shown in Figure 22.1B, even an uncorrelated input distribution (left) may become correlated once the trajectory is mapped into a permutation subspace (right). The argon test system demonstrates that an MI expansion to the third order is sufficient to compensate the effect within 1 % and that higher-order "pseudocorrelations" induced by permutation reduction are small. Furthermore, the second- and third-order "pseudocorrelations" for the ideal gas (5.4 J·mol^{-1}·K^{-1} and 0.5 J·mol^{-1}·K^{-1}, respectively) are small compared to the second- and third-order contributions for the full argon gas (21.2 J·mol^{-1}·K^{-1} and 4.5 J·mol^{-1}·K^{-1}, respectively). Interpreting the contributions by the second and third order as measures of physical two- and three-body correlations is therefore still warranted for sufficiently interacting systems.

The significant but small third-order contribution of 4.5 J·mol^{-1}·K^{-1} to the overall entropy loss shows that neglecting higher-order terms, which are expected to yield decreasing contributions, is justified.

As the interactions become stronger, the entropy is increasingly underestimated by up to 5.7 %, likely because higher-order expansion terms become more important at high pressures.

Argon at ≈ 10000 bar of pressure loses approximately 30 J·mol^{-1}·K^{-1} of entropy compared to its ideal gas state, which is significantly more than the ≈ 18 J·mol^{-1}·K^{-1} entropy loss of water compared to interactionsless water. Since kNN MI estimators are known to increasingly underestimate correlations the more correlated a system is [112, 163], argon at high pressure poses a harder benchmark than water at 1 bar. We therefore expect a smaller translational entropy error for water.

Overall, Per|Mut yields accurate solvation entropies for the argon test system. To test the accuracy and the ability to provide spatially resolved entropies, we applied Per|Mut to more complex systems, which will be discussed in sections 24.2 and 24.3.

24.2 Alkanes

Experimental and theoretical studies show that the solvation entropy of alkanes decreases approximately linearly with increasing alkane length [172–174]. To see if Per|Mut captures this linear relationship qualitatively and quantitatively, we calculated hydration entropies for the n-alkanes from ethane to decane. Here, we defined the hydration shell as the 100 closest water molecules to the solute after permutation reduction. The number was chosen such that even for the largest solute (decane), all water molecules that were affected by its presence were still included. Reference values were obtained by TI, as described in section 23.2.

As shown in Fig 24.2, Per|Mut indeed yields linear trends for the translational (blue symbols) and rotational (green symbols) entropies as well as for the translation-rotation correlation term $-\Delta I_{\text{trans-rot}}$ (purple symbols), which reduce the entropy by (5.6 ± 1.0) J·mol^{-1}·K^{-1}, (5.8 ± 0.9) J·mol^{-1}·K^{-1}, and (11.6 ± 0.6) J·mol^{-1}·K^{-1}, respectively, for each additional carbon atom. Combined, this results in a loss of (23.0 ± 1.1) J·mol^{-1}·K^{-1} per C atom (orange symbols). The result is in good agreement with the TI reference value of (21.0 ± 0.7) J·mol^{-1}·K^{-1} per C atom (red symbols) and more than the experimental [172] entropy loss of (13.3 ± 0.7) J·mol^{-1}·K^{-1}.

Since the difference between TI and experimental values is most likely due to force field errors, which equally affect Per|Mut, we consider TI to be the proper benchmark.

Notably, increased correlations between translational and rotational water motions for longer alkanes reduce the entropy by as much as the combination of translational and rotational modes. The correlation of translational and rotational modes of motion increases for molecules close to the solute. In this regime, the molecules likely experience an increased orientational bias by predominately forming hydrogen bonds facing away from the solute. By this reasoning, larger solutes result in a larger entropy loss from the correlation term.

Overall, Per|Mut accurately calculates the solvation entropy change for alkanes between ethane and decane. Furthermore, the method precisely captures the entropy change, induced by the addition of a chemical group as small as a methyl group to the solute for a hydration
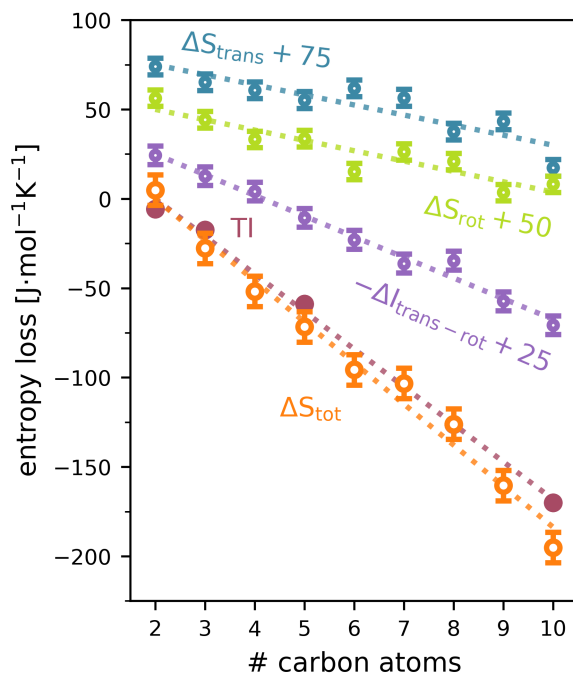
Figure 24.2: Hydration shell entropy loss relative to ethane with increasing alkane length. Translational and rotational entropies are illustrated with blue and green symbols, respectively. Entropy loss due to translation-rotation correlation is shown in purple. The total entropy change is shown in orange, and the TI reference entropies are shown in red. For easier visibility, the translational, rotational, and translation-rotation correlation data is offset by 75, 50, and 25 units, respectively. Dotted lines are from linear regression.

shell of 100 molecules.

24.3 Octanol

Our approach allows closer analysis of the molecular origin of entropy changes. To assess the spatial resolution yielded by Per|Mut, we calculated the hydration shell entropy of octanol. To this end, we simulated a fixed octanol molecule with 1728 water molecules in a similar manner as described in section 23, carried out the Per|Mut analysis, and calculated local entropies as described in section 23.2.

As shown in Figure 24.3, the spatial distribution of entropy differs significantly between the apolar tail and in the vicinity of the OH group of octanol.

The translational entropy (Figure 24.3A) is reduced by $15.2 \, \text{J·mol}^{-1}\text{·K}^{-1}$ per water molecule where the hydroxyl group acts as a proton donor (white arrow) and by $3.7 \, \text{J·mol}^{-1}\text{·K}^{-1}$ where it acts as a proton acceptor (cyan arrow). Close to the hydrophobic tail, the entropy reduction varies between 3.7 and $1.5 \, \text{J·mol}^{-1}\text{·K}^{-1}$.

The MI expansion allows for the entropy decomposition into contributions from the indi-
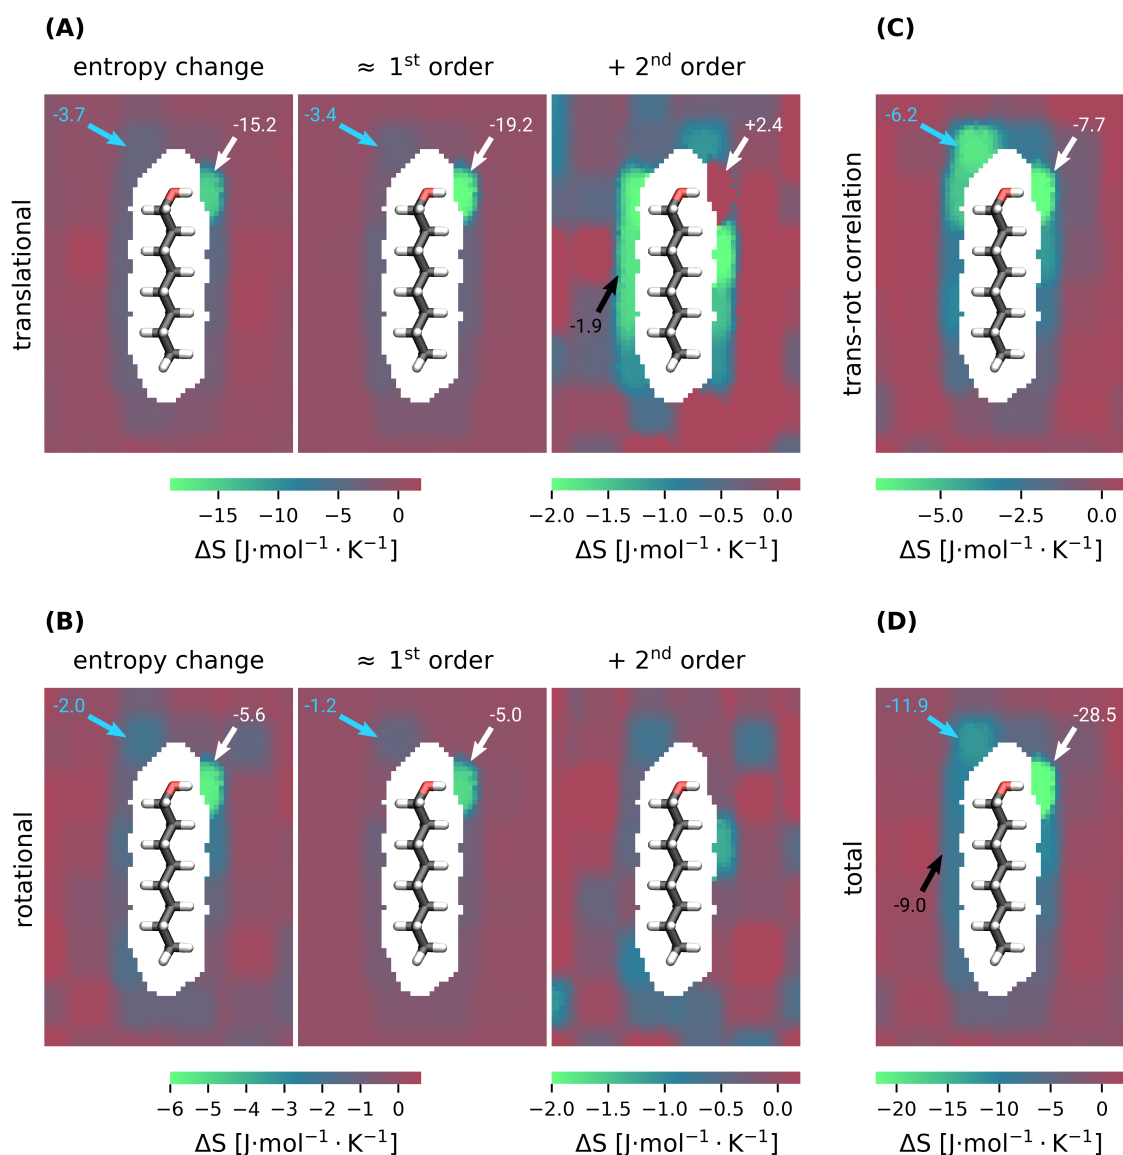
Figure 24.3: Spatially resolved octanol hydration shell entropy change per molecule relative to bulk quantities. (A) shows translational entropy values, and (B) shows rotational entropies. In both, (A) and (B), the three columns show total entropy change, first-order MI change, and second-order MI change, respectively. (C) shows the translation-rotation correlation, and (D) shows the total entropy change. Values at selected regions are highlighted with arrows.

vidual water molecules (first order) and their entropy loss due to pair correlations (second order). The first-order contribution (Figure 24.3A, center) shows that an entropy loss of 19.2 J·mol^{-1}·K^{-1} at the donor side of the hydroxyl group is due to the reduced mobility of a hydrogen-bonded water molecule. At the acceptor side, only 3.4 J·mol^{-1}·K^{-1} is lost. The entropy loss around the tail is mainly from the first-order contribution, where the presence

of the apolar hydrocarbon chain causes an entropy loss up to $3.5 \, \text{J·mol}^{-1}\text{·K}^{-1}$ due to the restricted mobility of individual water molecules.

As seen for the second-order contribution (Figure 24.3A, right), the entropy loss due to pair correlations of the water molecule that forms a hydrogen bond to the hydroxyl group is less than in bulk phase. The molecule therefore gains $2.4 \, \text{J·mol}^{-1}\text{·K}^{-1}$ relative to bulk (white arrow). A likely explanation is that the hydroxyl group replaces another water molecule as a hydrogen bond partner, which leaves the chemical environment almost unchanged but reduces the number of possible water-water correlation pairs in the vicinity. Furthermore, increased pairwise correlations decrease the entropy in a shell around the hydrophobic chain by additional $1.9 \, \text{J·mol}^{-1}\text{·K}^{-1}$ per molecule (black arrow). This result is unexpected, as water molecules close to the octanol molecule have fewer neighbors, thus showing that correlations with the remaining neighbors disproportionally increase at the hydrophobic tail.

The third-order correlation does not show significant spatial heterogeneity and, therefore, is not included in Figure 24.3.

The rotational entropy (Figure 24.3B) behaves similarly, albeit its contributions are smaller. $5.6 \, \text{J·mol}^{-1}\text{·K}^{-1}$ is lost by the water molecule for which the OH group of octanol acts as the proton donor (white arrow), to which the first order contributes $5.0 \, \text{J·mol}^{-1}\text{·K}^{-1}$. On the acceptor side, $2.0 \, \text{J·mol}^{-1}\text{·K}^{-1}$ is lost, to which the first order contributes $1.2 \, \text{J·mol}^{-1}\text{·K}^{-1}$ (cyan arrows). Near the hydrocarbon chain, the rotational entropy is reduced by $\approx 1.6 \, \text{J·mol}^{-1}\text{·K}^{-1}$ per molecule, of which $\approx 1.0 \, \text{J·mol}^{-1}\text{·K}^{-1}$ is due to a hindered rotational motion of the individual water molecules (first order). Again, there is a shell of more correlated water molecules around the hydrophobic part of octanol (second order), which is less pronounced than that for the translational entropy.

The mutual information of translational and rotational degrees of freedom (Figure 24.3C) shows strong correlations, equivalent to an entropy loss of $7.7 \, \text{J·mol}^{-1}\text{·K}^{-1}$ at the donor site (white arrow) and $6.2 \, \text{J·mol}^{-1}\text{·K}^{-1}$ at the acceptor site (cyan arrow). Furthermore, the translation-rotation correlations reduce the entropy of each molecule close to the apolar chain by 1.0 to $3.5 \, \text{J·mol}^{-1}\text{·K}^{-1}$.

As shown in Figure 24.3D, the total hydration entropy of octanol is mainly affected by the polar hydroxyl group, where entropy is reduced by $28.5 \, \text{J·mol}^{-1}\text{·K}^{-1}$ (white arrow) and $11.9 \, \text{J·mol}^{-1}\text{·K}^{-1}$ (cyan arrow) at the donor and acceptor sites, respectively.

The discrepancy between the two hydrogen binding sites is most likely caused by different bond strengths. Whereas hydrogen atoms of the hydroxyl group and of a water molecule carry almost identical partial charges, the oxygen atom of the hydroxyl group carries a partial charge of -0.65 elementary charges, significantly less than a water oxygen atom (-0.834 elementary charges).

In addition, the entropy of each water molecule close to the hydrocarbon chain is reduced by $9 \, \text{J·mol}^{-1}\text{·K}^{-1}$, which results from both a loss of mobility of individual molecules (first order) and increased correlations at the surface of the solute.

A quantitative comparison between the hydrophilic hydroxyl group and the apolar tail (or, equivalently, between octanol and octane) yields a solvent entropy difference of (36 ± 3) J·mol^{-1}·K^{-1}, which was determined using eight-molecule shells around the OH group and its hydrophobic counterpart, following the same rationale as for the alkanes. The result is significantly larger but comparable in magnitude to the TI reference estimate of (25.1 ± 0.1) J·mol^{-1}·K^{-1} (see Supporting Information). Aside from possible sampling issues, the deviation is likely the result of omitted higher-order correlations, which are affected differently by the polar hydroxyl group and the apolar chain.

25 Conclusions

We developed Per|Mut, a new method to calculate hydration entropies of water, and assessed its accuracy on argon, alkanes, and octanol test systems.

Our method rests on a permutation reduction [46, 47] (section 22.2), which alleviates the sampling problem by $N!$ and localizes the water molecules (Figure 22.1C), leaving the physics of the system unchanged. Due to the localization of the molecules, spatially resolved entropies can be calculated at the level of single water molecules. Further, a MIE is employed, which allows the absolute entropy to be decomposed into contributions from individual molecules, pair correlations, and triple correlations. The MIE reduces the dimensionality of the spaces that need to be sampled. By distinguishing between entropy contributions of individual molecules as well as pairwise and triple correlations, additional insight into the physical origin of entropy changes is provided.

We used the small argon test system to assess the accuracy of the translational entropy algorithm by comparing the obtained values with TI. Per|Mut yielded accurate entropy values for the full range of the switching coordinate within a maximum deviation of 5.7 % from the TI reference value.

To test the accuracy of Per|Mut as a whole, including the translation-rotation correlation term, we calculated the hydration entropies of n-alkanes from ethane to decane.

Indeed, we identified a linear entropy loss [174] of (23.0 ± 1.1) J·mol^{-1}·K^{-1} per additional carbon atom, as shown in Figure 24.2, which is in quantitative agreement with the reference entropy loss of (21.0 ± 0.7) J·mol^{-1}·K^{-1} per C atom, calculated by TI. Here, the increased correlation between the translational and rotational degrees of freedom for larger solutes was identified as the largest contribution to the entropy loss.

Because of its hydrophobic tail and its hydrophilic headgroup, we chose octanol as a test system to demonstrate how Per|Mut can characterize solvation entropies with a spatial resolution. Hydrogen bonding strongly reduces the local entropy by 11.9 J·mol^{-1}·K^{-1} and 28.5 J·mol^{-1}·K^{-1} for the water molecules to which the hydroxyl group of octanol acts as a proton acceptor or donor, respectively. The entropy loss at the donor site yields an entropic free energy contribution of 8.55 kJ·mol^{-1} at a temperature of 300 K, which is significantly less

than the $\approx 20\,\text{kJ·mol}^{-1}$ enthalpy loss due to the hydrogen bond [23]. The result shows that the solvation free energy difference of octanol and octane is enthalpy-driven.

Near the hydrophobic tail of octanol, the entropy is reduced for both the first- and second-order MI term. The losses of up to $1.9\,\text{J·mol}^{-1}\text{·K}^{-1}$ due to translational correlations, and up to $3.5\,\text{J·mol}^{-1}\text{·K}^{-1}$ due to correlations between translation and rotation show that the lack of strong interactions with the apolar octanol tail causes stronger interactions within the remaining water, an effect that is similarly discussed in previous publications [37, 38]. The finding does not necessarily imply an increased structural order in the hydration shell, as predicted by the controversial iceberg model [36], but identifies a reduced single-molecule mobility (first order) and increased water correlations (second order) as main causes for the entropy loss.

Our implementation of Per|Mut is available for download as a Python package[13].

26 Supplement

26.1 The mutual information expansion of hydration entropy

We describe the configuration of N water molecules using N translational degrees of freedom from $\mathbb{R}^3$, and N orientations from $SO(3)$. The set of translational degrees of freedom is $\mathcal{T} = \{1, \ldots, i, \ldots, N\}$, where the index i denotes the translational coordinates of molecule i. In the same fashion, the set of rotational degrees of freedom is $\mathcal{R} = \{1, \ldots, j, \ldots, N\}$.

The third-order mutual information expansion of the total system entropy is defined as [4, 48–50]

$$S_{\text{tot}} \approx \sum_{i \in (\mathcal{T} \cup \mathcal{R})} I_1(i) - \sum_{(j,k) \in \text{Pairs}(\mathcal{T} \cup \mathcal{R})} I_2(j,k) + \sum_{(l,m,n) \in \text{Triples}(\mathcal{T} \cup \mathcal{R})} I_3(l,m,n), \qquad (26.1)$$

where $(j,k) \in \text{Pairs}(\mathcal{T} \cup \mathcal{R})$ is a pair with $i \neq j$, and $(l,m,n) \in \text{Triples}(\mathcal{T} \cup \mathcal{R})$ is a triple of degrees of freedom with unique l, m, n.

The entropy from equation 26.1 is split into rotational, translational, and mixed parts:

$$\begin{aligned}
S_{\text{tot}} \approx & \sum_{i \in \mathcal{T}} I_1(i) + \sum_{\tilde{i} \in \mathcal{R}} I_1(\tilde{i}) \\
& - \sum_{(j,k) \in \text{Pairs}(\mathcal{T})} I_2(j,k) - \sum_{(\tilde{j},\tilde{k}) \in \text{Pairs}(\mathcal{R})} I_2(\tilde{j},\tilde{k}) - \sum_{j \in \mathcal{T},\, \tilde{k} \in \mathcal{R}} I_2(j,\tilde{k}) \\
& + \sum_{(l,m,n) \in \text{Triples}(\mathcal{T})} I_3(l,m,n) + \sum_{(\tilde{l},\tilde{m},\tilde{n}) \in \text{Triples}(\mathcal{R})} I_3(\tilde{l},\tilde{m},\tilde{n}) \\
& + \sum_{l \in \mathcal{T},\, (\tilde{m},\tilde{n}) \in \text{Pairs}(\mathcal{R})} I_3(l,\tilde{m},\tilde{n}) + \sum_{(l,m) \in \text{Pairs}(\mathcal{T}),\, \tilde{n} \in \mathcal{R}} I_3(l,m,\tilde{n}).
\end{aligned}$$

---

[13]https://gitlab.gwdg.de/lheinz/hydration_entropy

The translational entropy reads

$$S_{\text{trans}} \approx \sum_{i \in \mathcal{T}} I_1(i) - \sum_{(j,k) \in \text{Pairs}(\mathcal{T})} I_2(j,k) + \sum_{(l,m,n) \in \text{Triples}(\mathcal{T})} I_3(l,m,n),$$

and the rotational entropy reads

$$S_{\text{rot}} \approx \sum_{\tilde{i} \in \mathcal{R}} I_1(\tilde{i}) - \sum_{(\tilde{j},\tilde{k}) \in \text{Pairs}(\mathcal{R})} I_2(\tilde{j},\tilde{k}) + \sum_{(\tilde{l},\tilde{m},\tilde{n}) \in \text{Triples}(\mathcal{R})} I_3(\tilde{l},\tilde{m},\tilde{n}).$$

The translation-rotation correlation contribution is given by the remaining mixed terms and reads

$$I_{\text{trans-rot}} \approx \sum_{j \in \mathcal{T}, \ \tilde{k} \in \mathcal{R}} I_2(j,\tilde{k}) - \sum_{l \in \mathcal{T}, \ (\tilde{m},\tilde{n}) \in \text{Pairs}(\mathcal{R})} I_3(l,\tilde{m},\tilde{n}) - \sum_{(l,m) \in \text{Pairs}(\mathcal{T}), \ \tilde{n} \in \mathcal{R}} I_3(l,m,\tilde{n}).$$

In this Per|Mut implementation, we neglect the mixed three-body correlation terms due to their small contribution and slow convergence.

Using this separation, the full entropy is

$$S_{\text{tot}} = S_{\text{trans}} + S_{\text{rot}} - I_{\text{trans-rot}}.$$

VI | Application: Crambin

The following text has been submitted as

L.P. Heinz and H. Grubmüller, **"Spatially resolved free energy contributions of native fold and molten-globule-like crambin"** to the *Biophysical Journal* and is currently under review.

I carried out the research and wrote the manuscript. Helmut Grubmüller supervised the research and revised the manuscript.

Spatially resolved free energy contributions of native fold and molten-globule-like crambin

Leonard P. Heinz[1], Helmut Grubmüller[1]

[1] Department of Theoretical and Computational Biophysics, Max-Planck-Institute for Biophysical Chemistry, Göttingen, Germany

Abstract

The folding stability of a protein is governed by the free energy difference between its folded and unfolded states, which results from a delicate balance of much larger but almost compensating enthalpic and entropic contributions. The balance can therefore easily be shifted by an external disturbance, such as a mutation of a single amino acid or a change of temperature, in which case the protein unfolds. Effects like cold denaturation, in which a protein unfolds due to cooling, provide evidence that proteins are strongly stabilized by the solvent entropy contribution to the free energy balance. However, the molecular mechanisms behind this solvent-driven stability, their quantitative contribution in relation to other free energy contributions, and how the involved solvent thermodynamics is affected by individual amino acids, is largely unclear. Therefore, we addressed these questions using atomistic molecular dynamics simulations of the small protein crambin in its native fold and a molten-globule-like conformation, which here served as a model for the unfolded state. The free energy difference between these conformations was decomposed into enthalpic and entropic contributions from the protein and spatially resolved solvent contributions using the nonparametric method Per|Mut. From the spatial resolution, we quantified the local effects on the solvent free energy difference at each amino acid and identified dependencies of the local enthalpy and entropy

on the protein curvature. We identified a strong stabilization of the native fold by almost $500\,\text{kJ·mol}^{-1}$ due to the solvent entropy, revealing it as an essential contribution to the total free energy difference of $53\,\text{kJ·mol}^{-1}$. Remarkably, more than half of the solvent entropy contribution arose from induced water correlations.

Significance

The free energy difference between folded and unfolded states governs the thermodynamic stability of a protein. Effects like cold denaturation provide evidence that solvent-related contributions to the free energy difference strongly stabilize the native protein fold. Quantifying the solvent contribution and its dependency on the individual amino acids is therefore essential for a better understanding of the protein folding thermodynamics. The obtained spatial resolution of solvent free energy contributions might furthermore be relevant for protein design.

27 Introduction

Folding free energies of proteins at room temperature range typically in the order of a few $10\,\text{kJ·mol}^{-1}$ [19–22], which approximately corresponds to the interaction energy of just a few hydrogen bonds [23]. This small folding free energy results from a delicate balance between competing enthalpy and entropy contributions, each of which are much larger in magnitude, but compensate each other almost entirely [20]. Shifting the balance by external factors, e.g., due to a mutation of an amino acid or simply by changing the temperature, can lead to drastic and sometimes counter-intuitive consequences. For example, proteins can unfold at low temperatures [20, 24, 25], although the protein-internal interaction energies favor the folded state and the entropic protein free energy contribution $-T\Delta S$, which favors the unfolded state, should decrease in magnitude. This effect known as cold denaturation seems paradoxical only if the protein but not the solvent contributions are considered for the free energy difference estimation. Indeed, cold denaturation has been attributed to a temperature-induced weakening of the hydrophobic effect [20], which arises from the thermodynamics of the solvation shell. The hydrophobic effect is known to be a major driving force for protein folding and stability [20, 27–29].

Cold denaturation exemplarily illustrates the importance of solvation thermodynamics for protein stability. However, their quantitative role in relation to other free energy contributions, the molecular mechanisms of solvent-driven protein stability, and how the solvent thermodynamics is affected by individual amino acids are still largely unresolved. To study the effect of solvation on protein stability, a decomposition of the individual solvent- and protein-related free energy contributions of a protein native fold and the unfolded state is required. To furthermore characterize the effect of individual amino acids on the solvent

free energy contribution, a spatially resolved map of local solvent enthalpies and entropies is needed.

Here, we address these questions using all-atom molecular dynamics (MD) simulations of the small protein crambin. Because a representative ensemble of the extremely broad distribution of unfolded protein configurations is not sufficiently sampled within the timescales of MD simulations, we assessed the free energy difference between the native fold and a transiently stable molten-globule-like conformation.

We decomposed the free energy difference between the two conformations into enthalpy contributions from protein-protein, protein-solvent, and solvent-solvent interactions, and entropy contributions from both solvent and protein. To accurately compare the spatially resolved solvent entropy contributions to their respective enthalpic terms, we preferred a nonparametric method that also captures the entropy effects of multi-body correlations. We therefore selected Per|Mut [46, 47, 112, 175] to calculate the spatially resolved solvent entropy over alternative methods, such as the grid inhomogeneous solvation theory [10, 12, 13, 43, 121, 122, 133], which usually excludes multi-body correlations, the grid cell theory [136] or 3D two-phase thermodynamics [11, 134, 135], which both rely on strong model assumptions. Per|Mut provides a further solvent entropy decomposition into the entropy of the individual water molecules and the entropy changes arising from multi-body water-water correlations from both translational and rotational degrees of freedom. From the decomposition, we identified a substantial stabilizing effect of the solvent free energy contributions of almost $500\,\mathrm{kJ \cdot mol^{-1}}$. The further breakdown into different interaction energies and entropic water correlation terms allowed for an interpretation of the free energy changes on a molecular level. Here, a significant increase of water-water correlations around the molten-globule like conformation compared to the native fold was found, which corresponds to a stabilizing entropic free energy contribution. From the spatial resolution of both solvation enthalpy and entropy, we obtain the local free energy changes due to specific amino acid side chains and capture free energy effects of the protein shape.

28 Methods

28.1 Molecular dynamics simulations

All molecular dynamics (MD) simulations were carried out using the software package Gromacs 2018 [73–77], the CHARMM36m force field [81, 83, 84], and the leap frog integrator with a 2 fs time step. Unless stated otherwise, the temperature was kept at 300 K using the V-rescale thermostat [98] with a time constant of 0.1 ps. For equilibration simulations under NPT conditions, the Berendsen barostat [99] at 1 bar pressure with a time constant of 0.5 ps was used. All production runs were carried out using the Parrinello-Rahman barostat [100, 101] with a time constant of 1.0 ps and 1 bar pressure. During all simulations, Coulomb
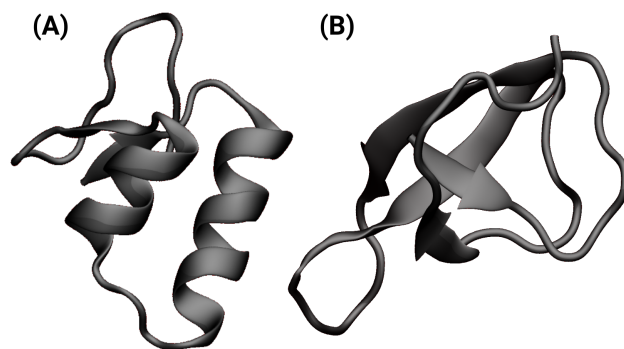
Figure 28.1: Ribbon-style visualization of crambin in its native fold (A) and in a molten-globule-like conformation (B). Images rendered using VMD [177].

and Lennard-Jones [42] cut-offs of $1.2\,\mathrm{nm}$ were used. Long-ranged electrostatic interactions were calculated using he Particle-Mesh Ewald (PME) method [160]. Bond-constraints were imposed on all water molecules using the SETTLE algorithm [95] and on all other bonds to hydrogen atoms using the LINCS algorithm [96]. During production runs, configurations were stored every $10\,\mathrm{ps}$ for subsequent analysis.

28.2 System setup

A high resolution structure of the protein crambin (46 amino acids) was retrieved from the Protein Data Bank (identifier 1CBN) [176]. The molecule was placed inside a cubic simulation box of size $(7\,\mathrm{nm})^3$ and solvated with 10950 water molecules. The system was then subjected to gradient descent energy minimization, which was terminated once the largest force on any atom decreased below $100\,\mathrm{kJ\cdot mol^{-1}nm^{-1}}$. To equilibrate the solvent water, a $1\,\mathrm{ns}$ MD simulation under NVT conditions was started, during which all heavy atoms of the protein were restrained using a force constant of $1000\,\mathrm{kJ\cdot mol^{-1}nm^{-2}}$. To relax the box size to $1\,\mathrm{bar}$, the system was subsequently simulated for $10\,\mathrm{ns}$ under NPT conditions, during which the restraints on the protein were maintained. Subsequently, the protein was equilibrated during a further $10\,\mathrm{ns}$ simulation under NVT conditions.

To obtain a molten-globule like confirmation of an unfolded state, the equilibrated protein was unfolded during $50\,\mathrm{ns}$ of heating and subsequent annealing. Here, the temperature was linearly increased to $600\,\mathrm{K}$ during the first $5\,\mathrm{ns}$, the temperature was then kept at $600\,\mathrm{K}$ for the following $40\,\mathrm{ns}$, and eventually decreased back to $300\,\mathrm{K}$ during the last $5\,\mathrm{ns}$. The unfolded structure was simulated for $2\,\mu\mathrm{s}$ under NPT conditions (Parrinello-Rahman barostat), during which a partial refolding occurred, resulting in the formation of two anti-parallel β-sheets, shown in Fig. 28.1B, whereas the rest of the molten-globule-like conformation remained unstructured. To ensure a proper equilibration of the folded system, it was also simulated for further $2\,\mu\mathrm{s}$.

Ensembles of four crambin structures were retrieved for both the native fold and the molten-

globule-like conformation by extracting equidistant frames from the last $3\,\mathrm{ns}$ of each of the equilibration runs. The ensembles therefore cover short-timescale processes like the sidechain reorientation, but do not capture larger configurational motions, particularly of the molten-globule-like conformation.

A production run of $1\,\mu\mathrm{s}$ under NPT conditions was carried out for each of the eight replicas. To restrict configurational changes of the protein, each atom was restrained using a force constant of $1000\,\mathrm{kJ\cdot mol^{-1}nm^{-2}}$.

28.3 Entropy calculation

Hydration entropies were calculated from the production runs of the eight replicas using the method Per|Mut [112, 175]. For each replica, a permutation reduction [46, 47] was carried out, which enhances sampling of the high-dimensional water configuration space by the Gibbs factor $N!$ and localizes the water molecules without changing their physical and statistical properties. The first trajectory frame was used as the reference configuration.

Spatially resolved entropies were calculated for the 4000 water molecules closest to the protein using a third-order mutual information (MI) expansion [4, 48–50]. Pairwise MI terms of the translational and rotational entropy, as well as for the translation-rotation correlation were calculated for water molecules with a maximum average distance of $1.0\,\mathrm{nm}$. For third-order terms, a $0.33\,\mathrm{nm}$ cut off was used. The individual entropy terms in the expansion were calculated using a k-nearest-neighbor algorithm with a value of $k = 1$ for all expansion orders. The entropy of the solvation shell consisting of the closest 1000 molecules was used for analysis. From those results, the entropy of the inner solvation shell, consisting of the closest 1000 molecules (equivalent to a thickness of $\approx 0.8\,\mathrm{nm}$), was calculated as $S_{\mathrm{internal}} - I$. Here, S_{internal} is the sum of all expansion terms concerning molecules from the inner solvation shell and I is the sum of all MI terms with a molecule inside and a molecule outside the inner solvation shell. Error bars were obtained as standard deviations of the four replicas per state.

Protein entropies were estimated for both states using Schlitter's method [115] on all non-hydrogen protein atoms from additional unrestrained MD simulations. Here, the molten-globule-like conformation partially re-folded after $1.5\,\mu\mathrm{s}$ (see Supplemental Information Fig. 31.1). To obtain values consistent with the restrained states, for which the remaining free energy terms were calculated, the analysis was limited to the first $1\,\mu\mathrm{s}$ of the trajectories. Errors were estimated from a second analysis using only $0.5\,\mu\mathrm{s}$ of the trajectories.

28.4 Enthalpy calculation

The solvation enthalpy was calculated as $U + pV$, where U is the sum of all interaction energies, averaged over the MD trajectory. The work term pV is the product of pressure ($1\,\mathrm{bar}$) and average box volume.

For each water molecule of the inner solvation shell, the interaction energies with the protein,

the outer solvation shell, and all other molecules of the inner shell were calculated. Lennard-Jones parameters and partial charges were taken from the CHARMM36m force field and a 2 nm cut off was used for Lennard-Jones interactions. The interaction energy of the protein was calculated directly with the software package Gromacs.

28.5 Visualization of entropy and enthalpy

The entropy and interaction energy contributions of each molecule were calculated by splitting the two-body MI terms, three-body MI terms, and the pairwise interaction energies equally between the respective molecules. The simulation box was then divided by a $129 \times 129 \times 129$ grid. The entropy and energy values of each voxel were calculated as an average, weighted by the time each localized water molecule spent in the voxel.

28.6 Error estimation

Errors of solvent entropy and solvent enthalpy contributions of each water molecule were estimated as standard deviation of bulk-phase water molecules, assuming that water molecules close to the solvent is subject to the same spread [175]. The estimated errors are listed in table 28.1.

Table 28.1: Estimated solvent enthalpy and solvent entropy errors per water molecule.

contribution	error [kJ·mol^{-1}]
U_{SS}	1.05
$-TS^{\mathrm{trans}}$	0.46
$-TS_1^{\mathrm{trans}}$	0.09
$-TS_{\geq 2}^{\mathrm{trans}}$	0.37
$-TS^{\mathrm{rot}}$	0.22
$-TS_1^{\mathrm{rot}}$	0.02
$-TS_{\geq 2}^{\mathrm{rot}}$	0.22
$-TS_{\geq 2}^{\mathrm{trans\text{-}rot}}$	0.18

Here, U_{SS} is the solvent-solvent interaction energy, $-TS^{\mathrm{trans}}$ and $-TS^{\mathrm{rot}}$ are the translational and rotational free energy contributions, respectively. Both are split into a single-body term $-TS_1$ and a multi-body term $-TS_{\geq 2}$. $-TS_{\geq 2}^{\mathrm{trans\text{-}rot}}$ corresponds to the entropy contribution from translation-rotation correlations.

28.7 Residue contributions and convexity

The local interaction energy and the local solvation entropy contribution around each amino acid was calculated as the average contribution of all water molecules within 0.4 nm of the amino acid. Here, the distance was measured as the minimum distance to any atom of the residue. The contribution of the pV term was small compared to the other contributions

(19.8 kJ·mol^{-1} for both the native fold and molten-globule-like conformations) and was therefore neglected. The local free energy was thus calculated as the sum of the interaction energy and the entropy contribution.

As a measure of the local topology, the convexity of the protein-solvent surface was calculated for each residue as

$$c = \frac{N_d}{\left(\frac{4}{3}\pi d^3 - V_{\text{res}}\right) n_W},$$

where N_d is the number of water molecules within a radius of $d = 1\,\text{nm}$ around the amino acid, V_{res} is the volume of the residue itself, and n_W is the number density of water.

Accordingly, a convexity of $c = 0$ indicates that the entire volume within the radius d around the residue is occupied by other protein atoms, i.e., that the residue is buried. Conversely, a value of $c = 1$ shows that the residue is fully exposed to the solvent.

To quantify the hydrophobicity of a residue and for comparison with our calculations, the empirical hydrophobicity scale by Zhao and London [178] was used: A small value indicates a hydrophilic amino acid, whereas a large value denotes hydrophobicity.

29 Results and Discussion

29.1 Spatially resolved solvation free energy contributions

To quantify and characterize the solvent contribution to the stability of the protein crambin, we carried out and compared MD simulations of the native fold and a molten-globule-like protein conformation, as described in the Methods section. We will first address the contribution of individual amino acids and the protein shape to the solvent free energy contributions for a number of exemplary residues and subsequently provide a more systematic assessment. To this end, we calculated spatially resolved solvent interaction energies and corresponding entropies for both the native fold and the molten-globule-like conformation, respectively.

Because any spatially resolved picture of the solvation-shell thermodynamics depends on and changes with the particular conformation of the protein, we chose a representative conformation for each, the folded and the molten-globule-like ensemble. To obtain better statistics, we carried out simulations of four similar replicas for each conformation. As the molten-globule-like conformation is structurally unstable, we restrained the structures as described in the Methods section.

For each of the 2×4 replicas, spatially resolved solvent interaction energies ΔU were calculated directly from the interaction energies between protein and solvent (ΔU_{PS}) and between the solvent molecules (ΔU_{SS}). Solvent entropy contributions were calculated using the method Per|Mut [112, 175], which yields a spatially resolved entropy decomposition into terms S_1^{trans}, S_1^{rot}, $S_{\geq 2}^{\text{trans}}$, $S_{\geq 2}^{\text{rot}}$, and $S_{\geq 2}^{\text{trans-rot}}$. Here, S_1^{trans} and S_1^{rot} are the translational and rotational entropies of the individual water molecules, ignoring any correlations. $S_{\geq 2}^{\text{trans}}$, $S_{\geq 2}^{\text{rot}}$ capture the entropy loss due to two- and three-body correlations between water molecules for

translational and rotational degrees of freedom, respectively. A positive value of $-T\Delta S_{\geq 2}$ denotes an unfavorable free energy increase relative to bulk water due to increased water-water correlations. Conversely, a negative value of $-T\Delta S_{\geq 2}$ indicates a favorable free energy contribution due to weakened correlations relative to bulk water. Similarly, $\Delta S_{\geq 2}^{\text{trans-rot}}$ is the entropy change due to correlations between translational and rotational degrees of freedom of water molecules.

Figure 29.1 compares these spatially resolved free energy contributions for the native fold (A) and the molten-globule-type conformation (B), respectively. As shown in the upper left of Fig. 29.1A, the interaction energy ΔU of water molecules near the protein is generally more negative (stronger) than in bulk water, particularly close to charged residues such as R17 and D43, with interaction energy differences of $-48\,\text{kJ·mol}^{-1}$ and $-27\,\text{kJ·mol}^{-1}$, respectively. Similarly, in the vicinity of the polar hydroxyl groups of the tyrosine residues Y29 and Y44, the energy is lower than in bulk water ($\Delta U = -22\,\text{kJ·mol}^{-1}$), and close to residue C40, the interaction energy difference is $-27\,\text{kJ·mol}^{-1}$. Residues T2, F13, and R10 form a "groove" in the native-fold structure, with strong interaction energies below $-30\,\text{kJ·mol}^{-1}$ relative to bulk water.

Figure 29.1A (top row, center and right) reveals the individual interaction energy contributions from the interactions between protein and solvent ΔU_{PS}, as well as from interactions between solvent molecules ΔU_{SS}. The interaction energy of hydration is dominated by the protein-solvent interactions, where charged and polar residues show particularly favorable interactions. However, the solvent-solvent interactions (top row, right) partially counteract this effect. Here, water molecules that bind to charged residues perturb the hydrogen bond network of the surrounding water, thereby weakening the solvent-solvent interactions. Accordingly, around the charged residues R10, R17, and D43, the solvent-solvent interaction energy becomes strongly unfavorable with values between $+17$ and $+30\,\text{kJ·mol}^{-1}$ relative to bulk water. Similarly, but to a lesser extent, solvent-solvent interactions are weakened for water molecules around the hydroxyl groups of residues Y29 and Y44, resulting in unfavorable energy changes ranging between $+10$ and $+15\,\text{kJ·mol}^{-1}$.

Next, we quantified the impact of these strong interactions on the local entropy. For the charged residues R10, R17, and D43, as well as for the polar residues Y29 and Y44, both the translational and rotational solvent entropies are reduced (Fig. 29.1A, center rows). At these sites, the free energy contributions $-T\Delta S$ vary from $+5.5$ to $+12\,\text{kJ·mol}^{-1}$ (translational) and from $+3.5$ to $+8.2\,\text{kJ·mol}^{-1}$ (rotational) and therefore weaken the otherwise strong solvation. The entropy loss is dominated by the first-order terms $\Delta S_1^{\text{trans}}$ and ΔS_1^{rot}, revealing that the reduced mobilities of the individual water molecules due to the protein (partial) charges, and not water-water correlations, are the main cause of the entropy loss.

Translational entropy free energy contributions from multi-body correlations $-T\Delta S_{\geq 2}$ very close to the charged and polar protein parts are indeed relatively weak, but not negligibly so, with values between -2.0 and $-1.5\,\text{kJ·mol}^{-1}$ relative to bulk water. Remarkably, for second-
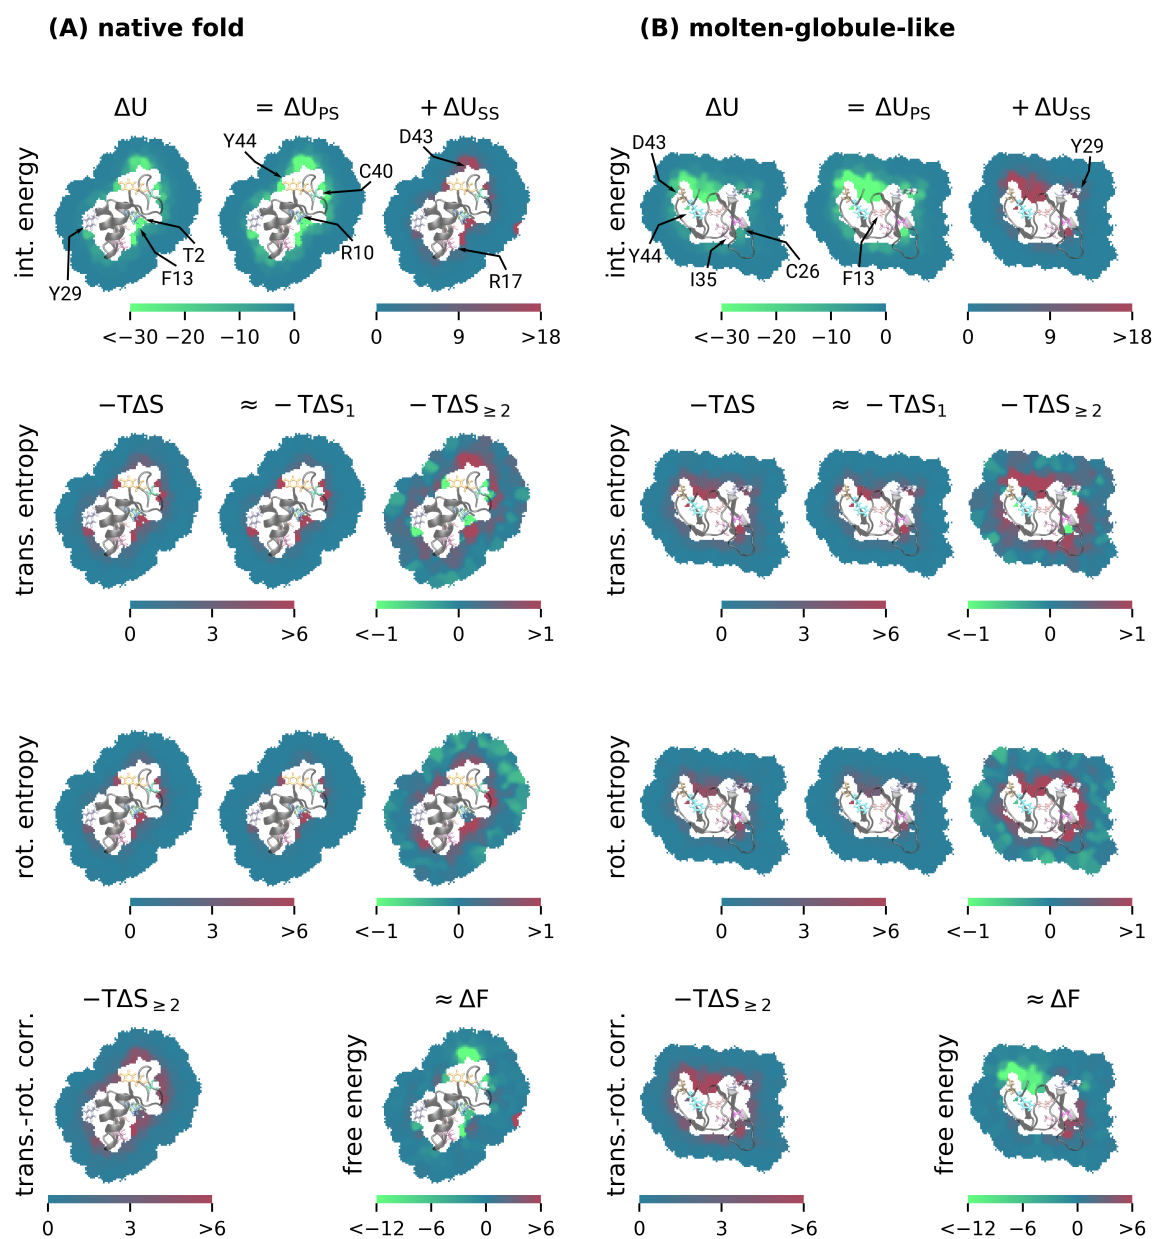
Figure 29.1: Spatially resolved hydration free energy decomposition of crambin in its native fold (A) and the molten-globule-like conformation (B) relative to bulk water, visualized on a representative two-dimensional slice through the center of the molecules. The total interaction energy ΔU (top row, left) is split into protein-solvent interactions ΔU_{PS} (top row, center), and solvent-solvent interactions ΔU_{SS} (top row, right). In the same manner, the translational and rotational entropy contributions $-T\Delta S$ (center rows) are split into the single-molecule entropy $-T\Delta S_1$, and multi-body correlations $-T\Delta S_{\geq 2}$. The entropy contribution from the translation-rotation correlation is shown in the bottom left. The spatially resolved free energy change (sum of the first column) is shown in the bottom right. All numerical values are given in kJ·mol^{-1} and important residues are highlighted by arrows.

shell water molecules around charged residues, these correlations again change sign, e.g., for residues R10 ($+1.0$ kJ·mol^{-1}), D43 ($+1.2$ kJ·mol^{-1}), and Y44 ($+0.6$ kJ·mol^{-1}), thus being less favorable compared to bulk water. A similar effect is seen for the rotational multi-body correlations where the free energy contribution almost vanishes at the first hydration shell (e.g., $+0.2$ kJ·mol^{-1} at R10, $+0.4$ kJ·mol^{-1} at Y29 and Y44, each with an estimated error of 0.22 kJ·mol^{-1}), but increases to values from $+0.7$ kJ·mol^{-1} to $+1.0$ kJ·mol^{-1} within the second hydration shell.

A possible explanation for this effect is that the strong binding of the first-shell water molecules (for rotational degrees of freedom the strong directionality due to the water dipole moment [179]) results in reduced fluctuations, which limits the correlations with the remaining water molecules (for a more detailed analysis, see *Supplemental Information*). As a result, the multi-body correlation entropy contribution is more favorable for the bound molecules, compared to bulk water. In contrast, the fluctuations of the second-shell water molecules are larger due to the reduced electrostatic interactions with the charged and polar residues, as shown by bulk-like single-body entropy terms. Therefore, the correlations with the adjacent water molecules reduce the entropy by a larger amount.

As the last remaining entropy term, we also quantified the entropy reduction due to translation-rotation correlations (Fig. 29.1A, lower left), which reveals increased correlations close to the protein surface, particularly in the vicinity of charged residues and polar chemical groups. Here, the respective solvent free energy contribution is increased by $+3.2$ kJ·mol^{-1} (R10), $+5.9$ kJ·mol^{-1} (R17), $+2.7$ kJ·mol^{-1} (Y29), $+5.0$ kJ·mol^{-1} (D43), and $+4.0$ kJ·mol^{-1} (Y44) compared to bulk water. The correlations and the corresponding unfavorable free energy changes arise because the strong directionality of the water molecules close to a protein (partial) charge quickly decreases with distance. A water molecule close to a charged amino acid is therefore on average more strongly oriented than a water molecule at a larger distance.

To characterize the local solvent free energy change relative to bulk water, we considered the sum of all free energy contributions (Fig. 29.1A, bottom right). Close to hydrophilic residues, enthalpic contributions dominate, such that the total free energy change is favorable, as expected. The water molecules that form hydrogen bonds to the hydroxyl groups of Y29 and Y44 contribute -5.3 and -6.2 kJ·mol^{-1} to the hydration free energy difference of the protein, respectively. Around the charged residues R17 and D43, the free energy per water molecule is reduced by -17.5 and -32 kJ·mol^{-1}, respectively. Close to the residue C40, the free energy is reduced by -5.4 kJ·mol^{-1} and in the "groove", formed by residues T2, R10, F13, the free energy contribution per water molecule is -6.6 kJ·mol^{-1} compared to bulk water. Here, the average interaction energy of the closest water molecule to the three amino acids is $+4$, -57, and 0 kJ·mol^{-1}, respectively. Although the entropic contributions cannot be assigned to specific residues, their total contribution to the free energy (≈ 30 kJ·mol^{-1}) is smaller than the enthalpic contribution of R10. In this sense, the favorable free energy is enthalpically dominated due to the charged residue R10. Due to the water-water correlation

effects discussed above, second-shell water molecules around protein charges contribute unfavorably to the hydration free energies, with values ranging from $+1.8$ to $3.2\,\mathrm{kJ\cdot mol^{-1}}$ per molecule.

Next, we carried out the same analysis for a molten-globule-like conformation to identify possible differences in the solvation thermodynamics between the two states. As shown in Fig. 29.1B, the enthalpic and entropic changes for water molecules close to charged and polar residues, e.g., for Y29, D43, and Y44, are similar to the native fold. However the spatial distribution of the individual free energy contributions is different. Whereas in the folded state, prominent differences of the local free energy contributions are isolated and can be well attributed to individual (charged or polar) residues, a particularly large region of favorable interaction energies and unfavorable entropies is seen around residues D43, Y44, and F13 of the molten-globule-like fold. Here we considered two possibilities: First, the (coincidental) co-localization of specific amino acids in the molten-globule-like state could give rise to the spread-out distribution. Alternatively, the locally concave shape of the molten-globule-like conformation could provide a possible explanation of the large volume of the affected region. In the latter case, the semi-cavity exposes a markedly larger surface area to the solvent, which strengthens the solvent-protein interactions ($-50\,\mathrm{kJ\cdot mol^{-1}}$ per molecule), but weakens the solvent-solvent interactions ($+20\,\mathrm{kJ\cdot mol^{-1}}$) as the number of neighboring water molecules is reduced due to the protein geometry. Correspondingly, the concave shape restricts the mobility of the water molecules, resulting in an unfavorable entropy change compared to bulk water ($+5.8\,\mathrm{kJ\cdot mol^{-1}}$ translational, $+5.0\,\mathrm{kJ\cdot mol^{-1}}$ rotational). Indeed, similar entropy and enthalpy effects were also seen to a smaller extent at the concave "groove", formed by residues T2, R10, and F13 in the native fold. To probe whether this observation is anecdotal, we systematically compared the local curvature to the local free energy contributions of all amino acids in the next subsection.

In addition to the single-body entropy reduction in the concavity, we also observed unfavorable entropy changes due to the multi-body correlation terms, which contribute $+2.2\,\mathrm{kJ\cdot mol^{-1}}$ (translational), $+2.1\,\mathrm{kJ\cdot mol^{-1}}$ (rotational), and $+9.5\,\mathrm{kJ\cdot mol^{-1}}$ (translation-rotation correlation). In light of the geometry-induced weaker solvent-solvent interactions, such high correlations are unexpected. To check whether this observation is anecdotal or, alternatively, more general, we systematically compared the multi-body entropy terms to the local convexity for each amino acid of the 2×8 replicas. As shown in *Supplemental Information* Fig. 31.2, there are indeed stronger correlations (i.e., unfavorable free energy contributions) between rotational degrees of freedom ($-T\Delta S_{\geq 2}^{\mathrm{rot}}$) and between translational and rotational degrees of freedom ($-T\Delta S_{\geq 2}^{\mathrm{trans\text{-}rot}}$). Translational correlations ($-T\Delta S_{\geq 2}^{\mathrm{trans}}$) show an inverted dependence on the local convexity. So far, we are unable to provide an explanation for these effects.

For the native fold, as well as for the molten-globule-like conformation, the well known tug of war between enthalpic and entropic free energy contributions results in a partial compensation of the two contributions [41]. For the native fold, this compensation also applies to the

multi-body correlation entropies and the solvent-solvent interaction energies. However, this compensation of solvent-solvent terms is not seen for water molecules in the concavity of the molten-globule-like conformation, which might be an effect of the concave shape of the protein. Despite the partial enthalpy-entropy compensation, the solvent free energy is mainly affected by the protein-solvent interaction energies.

The entropy contributions in the native-fold state mainly originate from water molecules that are bound to protein charges and mostly affect the first hydration shell. To the contrary, the entropy contributions in the molten-globule-like state are more spread out and also the second hydration shell is significantly affected. This observation might be a direct consequence of the hydrophobic driving forces of protein folding [28], where hydrophobic residues, which result in increased multi-body correlation entropies [37, 38, 175], are predominately packed into the protein interior of the native fold.

We furthermore note that the water entropy seems to become more unfavorable in concave parts of the molecules, for example as seen in the "groove" at residues T2, R10, and F13 for the native fold and around residues R13, D43, and Y44 for the molten-globule-like conformation. The water interaction energy seems to show a compensating effect.

29.2 Residue contributions

So far, we focused on a few illustrative example residues. For a more comprehensive and systematic assessment of the protein-shape effect on each of the local free energy contributions, we compared the free energies close to all residues to the local convexity, calculated as described in methods section *Residue contributions and convexity*. To distinguish protein topology-effects from those of the solvation properties of the individual amino acids, we furthermore compared the local free energy changes around each residue to its hydrophobicity index [178].

Figure 29.2, reveals a clear correlation between the local interaction energy differences with respect to bulk water and the convexity for both the native fold (Fig. 29.2A) and the molten-globule-like states (Fig. 29.2B), reflected in Pearson correlation coefficients of 0.50 and 0.42, respectively. Here, the interaction energy of the water molecules at a concave surface (convexity ≈ 0.3) is reduced by up to $40\,\text{kJ·mol}^{-1}$ for the native fold and by up to $33\,\text{kJ·mol}^{-1}$ for the molten-globule-like fold, whereas the interaction energy of water molecules at convex surfaces of the protein (convexity > 0.8) only differs by a small amount ($\Delta U > -5\,\text{kJ·mol}^{-1}$) relative to bulk water. This effect is to be expected, as more solvent-exposed residues interact with a larger number of water molecules.

An even stronger but inverse correlation (correlation coefficients of -0.76 and -0.71 for the native-fold and the molten-globule-like states, respectively) is seen in Fig. 29.2 for the solvation entropy contribution $-T\Delta S$: Here, the water entropy becomes strongly unfavorable at concave surfaces with contributions of up to $28\,\text{kJ·mol}^{-1}$ for the native-fold and of $24\,\text{kJ·mol}^{-1}$ for the molten-globule-like configurations. Correspondingly, the entropy contributions at convex
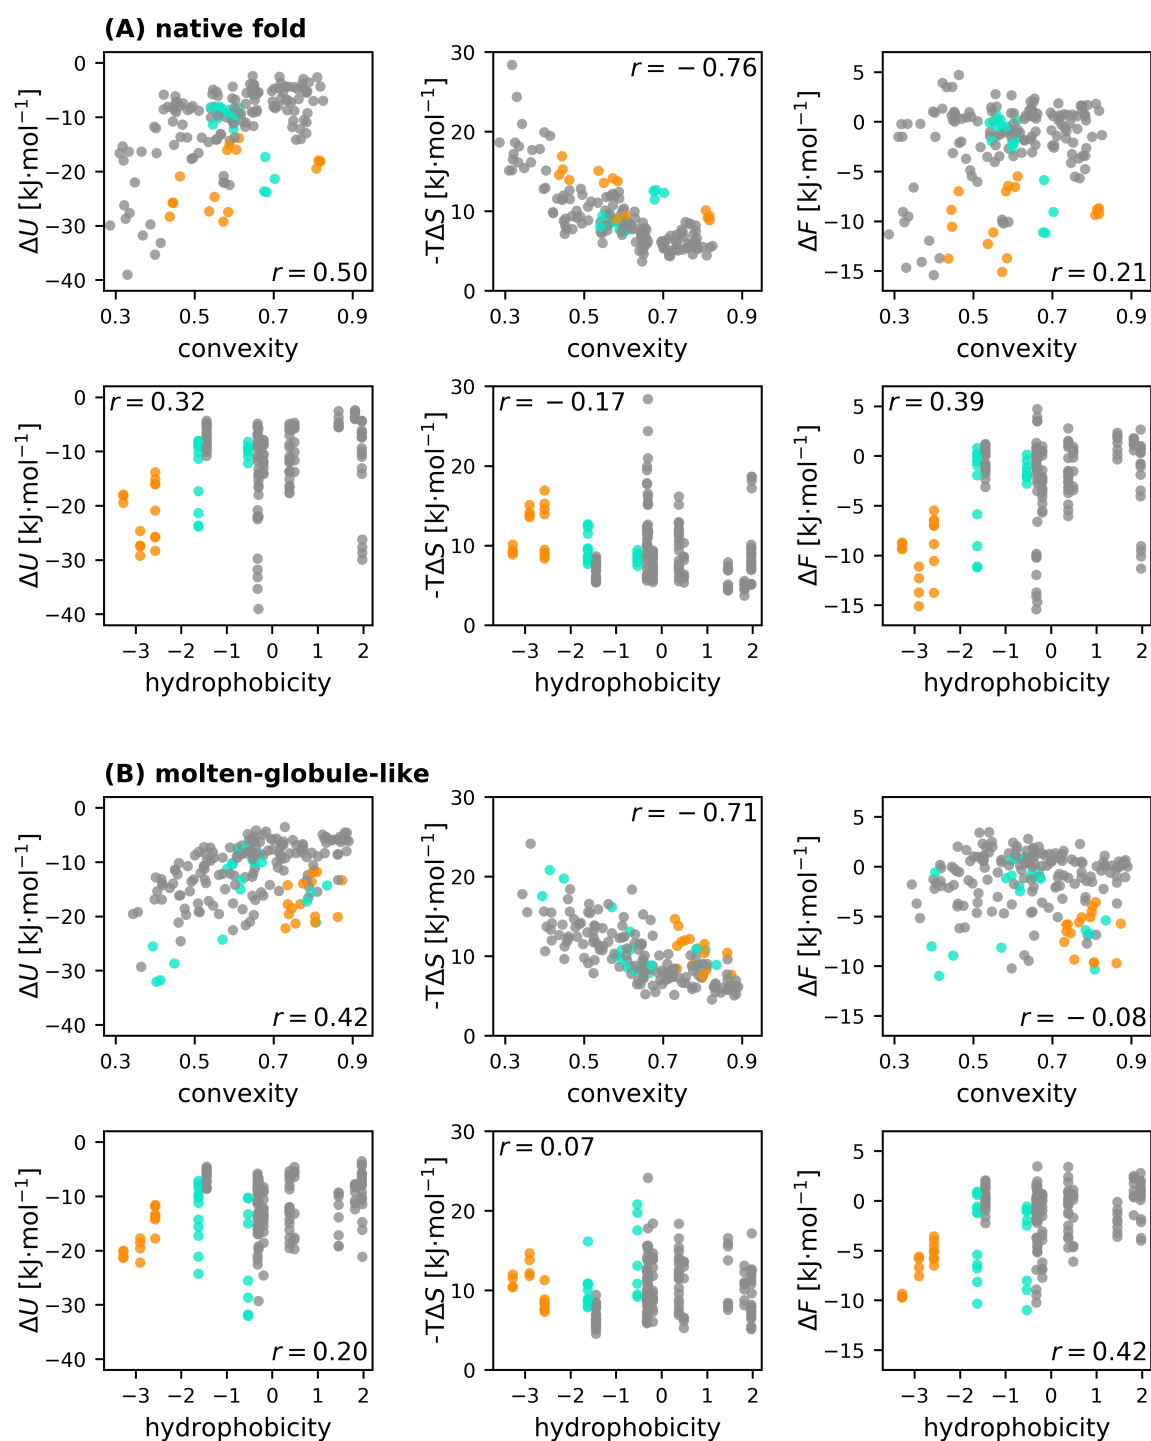
Figure 29.2: Changes in interaction energy (ΔU), entropy ($-T\Delta S$), and free energy (ΔF) of the water molecules around each residue of the eight replicas, relative to bulk water. The free energy change, and its energetic and entropic contributions, are shown relative to the local convexity and relative to the amino acid hydrophobicity. Results for the native fold are shown in (A), results for the molten-globule-like conformation are shown in (B). Charged amino acids are depicted as orange circles, polar amino acids are shown in cyan. Apolar amino acids are colored grey. Pearson correlation coefficients are stated in the corners of each plot.

surfaces is small ($-T\Delta S < 5\,$kJ·mol^{-1}). This observation is in line with previous reports about unfavorable water entropies and favorable water enthalpies in cavities [180, 181] and altered water behavior at concave surfaces [182].

Interestingly, the convexity dependencies of enthalpy and entropy almost compensate each other, such that the local free energy, shown on the top right of Fig. 29.2A and B, shows no significant dependency on the surface convexity ($r = 0.21$ and -0.08 for the native fold and the molten-globule-like state, respectively). Furthermore, the side chain polarity, as indicated by the color code in Fig. 29.2, has no significant impact on the convexity effects of interaction energy and entropy.

As shown in the bottom rows of Fig. 29.2A and B, there is only a modest correlation of the local interaction energy and entropy contributions with the hydrophobicity index of each amino acid ($r = 0.32$ and 0.20 for native fold and molten-globule-like state, respectively). However, as expected, the local free energy correlates more strongly with the hydrophobicity index ($r = 0.39$ and 0.42 for native fold and molten-globule-like state, respectively), where the most hydrophilic residues show favorable local solvation free energies of -16 to $-5\,$kJ·mol^{-1}, whereas local free energy changes of -11 to $+3\,$kJ·mol^{-1} are attributed to hydrophobic residues. Unexpectedly, for the non-polar residues, no correlation between their solvation free energies in the protein context and their hydrophobicity index is seen. As shown in the bottom right of Fig. 29.2A and B, the most favorable free energy is observed around charged residues (colored orange), with a free energy change of $(-8.1 \pm 2.3)\,$kJ·mol^{-1}. Around polar residues (colored cyan), the average free energy change is $(-3.14 \pm 3.72)\,$kJ·mol^{-1}; around apolar residues (colored grey) an average contribution of $(-1.53 \pm 4.45)\,$kJ·mol^{-1} is seen.

Comparing the native fold with the molten-globule-like conformation, the most striking difference is seen for the charged residues, for which the local free energy contribution is more favorable for the native fold ($(-9.64 \pm 2.58)\,$kJ·mol^{-1}) compared to the molten-globule-like conformation ($(-6.53 \pm 1.97)\,$kJ·mol^{-1}). This finding suggests that an important mechanism to stabilize a folded state is an optimized geometry, which maximally exposes hydrophilic charged residues.

29.3 Free energy decomposition

To obtain a quantitative understanding of the overall hydration contribution to protein stability, we decomposed the total free energy differences $\Delta G = G_{\mathrm{m}} - G_{\mathrm{f}}$ of the molten-globule-like conformation and the native fold into individual free energy contributions. In addition to solvent contributions, we also calculated protein interaction energies U_{PP} directly from the ensemble of the eight restrained replicas, as well as protein entropies S_{P} from unrestrained MD simulations of the native fold and the molten-globule-like conformation, respectively. To ensure that the unrestrained simulation of the molten-globule-type only system samples the conformation of the four similar molten-globule-like ensemble members used for the calcula-
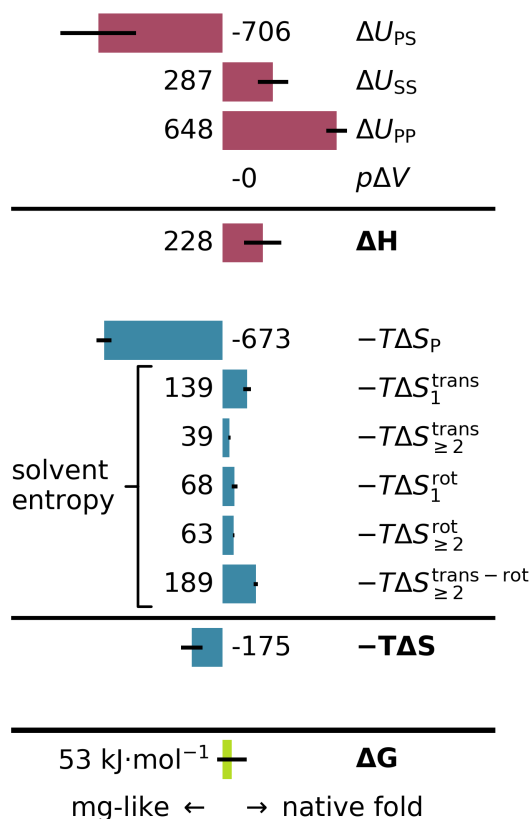
Figure 29.3: Decomposition of the folding free energy (green) into enthalpy (red) and entropy terms (blue). Positive values favor the native fold, negative values favor the molten-globule-like conformation. Values in kJ·mol^{-1}.

tion of the other free energy contributions, the simulation time was restricted to $1\,\mu$s. During this time, the conformation did not change markedly, with an RMSD below $0.7\,$nm. Accordingly, our free energy budget includes only part of the conformational entropy of the whole molten-globule-type ensemble.

Figure 29.3 summarizes the calculated differences between the native fold and the molten-globule-like conformation for the various free energy contributions. Positive values indicate a stabilization of the native fold, negative values favor the molten-globule-like conformation. As expected, we obtained a total free energy difference that favors the folded state ((53 ± 84) kJ·mol^{-1}). Although, to the best of our knowledge, a measured value is not available, this result agrees with the expected range of a few $10\,$kJ·mol^{-1} [20, 21].

Likely because the molten-globule-like conformation has a larger solvent accessible surface area ((36.9 ± 0.6) nm^2 vs. (31.3 ± 0.3) nm^2 for the native fold), the total interaction energy between the protein and inner solvent shell U_{PS} destabilizes the native fold by (-706 ± 216) kJ·mol^{-1} with respect to the molten-globule-like conformation. In contrast, the

interaction energies within the inner solvation shell (closet 1000 molecules) favor the native fold ($\Delta U_{\mathrm{SS}} = 287 \pm 86$) kJ·mol^{-1}). The work term $p\Delta V = (-0.002 \pm 0.002)$ kJ·mol^{-1} shows no significant difference between the two conformations. As expected, the internal interaction energies of the protein strongly stabilize the folded state ($\Delta U_{\mathrm{PP}} = (648 \pm 59)$ kJ·mol^{-1}). Overall, the enthalpy contributions strongly favor the folded state ($\Delta H = (228 \pm 106)$ kJ·mol^{-1}), in line with textbook thermodynamics of protein folding.

Next we quantified all entropic contributions to the free energy difference, which are expected to largely compensate the strongly stabilizing enthalpy difference. Indeed, the protein entropy strongly favors the molten-globule-like state even without parts of the configurational entropy ($-T\Delta S_{\mathrm{P}} = (-673 \pm 42)$ kJ·mol^{-1}). Crucially, however, all solvation entropy terms stabilize the native fold, adding up to 498 kJ·mol^{-1}. One of the largest contributions is the singe-particle translational entropy $S_1^{\mathrm{trans}} = (139 \pm 22)$ kJ·mol^{-1}. In contrast, the respective rotational entropy contribution, where many studies focus on, yields only a difference of (68 ± 17) kJ·mol^{-1}. Also intriguingly, the second and third order MI terms $S_{\geq 2}^{\mathrm{trans}}$ and $S_{\geq 2}^{\mathrm{rot}}$, which denote the two- and three-body correlations, contribute large differences of (39 ± 6) kJ·mol^{-1} and (63 ± 4) kJ·mol^{-1}, respectively. Strikingly, the largest solvent entropy contribution stems from the translation-rotation correlation entropy, which favors the native fold by (189 ± 12) kJ·mol^{-1}. Obviously, these entropies, which arise from the correlated motion of water molecules, contribute markedly to the folding thermodynamics of crambin.

In summary, the largest contribution to solvent-induced entropic stabilization (498 kJ·mol^{-1}) of the native fold is not due to the mobility of single water molecules but due to their correlated motion (291 kJ·mol^{-1}). Conversely, the protein-solvent interaction energies ΔU_{PS} and the protein entropy ΔS_{P} are the only destabilizing terms.

30 Conclusion

To quantitatively assess the role of the protein structure and individual side chains in the solvation thermodynamics and stability of the native fold, we compared spatially resolved enthalpic and entropic solvation free energy contributions of the native fold and the molten-globule-like state of the small globular protein crambin. A systematic analysis of the solvation thermodynamics of all residues revealed that their local enthalpies and entropies are mainly affected by the protein charges and the local curvature of the protein surface. Close to charged or polar residues, a more favorable interaction energy and a less favorable entropy contribution were observed. Interaction energies are here dominated by protein-solvent terms, whereas solvent-solvent interactions are weakened and therefore become less favorable around protein charges.

Single body entropy terms are strongly reduced close to protein charges, as the water molecules bind to the protein and therefore lose translational and orientational mobility. For solvent multi-body correlation contributions, a more complex picture emerged. Whereas cor-

relations with the remaining water molecules are weakened for the bound water molecules, the second hydration shell water molecules show an increased correlation and therefore an unfavorable multi-body entropy contribution. The translation-rotation correlation entropy becomes unfavorable close to the protein surface, particularly close to protein charges, indicating that the correlation between translational and rotational degrees of freedom increases close to the protein surface.

The local contributions from enthalpy and entropy largely compensate each other, as expected, so that the local free energy values are smaller in magnitude. The spatial distribution of the local free energy changes is mainly dominated by favorable contributions close to polar or charged residues. This result suggests that, at least for crambin, the solvation of hydrophilic residues contributes to a larger extent to the stability of the folded protein than burying hydrophobic residues into the protein interior. We furthermore observed that concave protein surfaces seem to cause more favorable hydration enthalpies and less favorable entropies.

To distinguish between a possible curvature effect and the influence of individual amino acids, we compared the local convexity of the protein surface and the hydrophobicity index of each protein residue to the local enthalpies, entropies, and free energies. Indeed, strong correlations between surface convexity and solvation interaction energies and entropies are seen. A more concave shape of the protein surface generally correlates with more favorable interaction energies and less favorable solvation entropy contributions. Due to the opposing effects on the two quantities, the local free energy of solvation remains almost unaffected by the protein curvature. As expected, we found more favorable local solvent free energies close to hydrophilic residues (i.e., those with a small hydrophobicity index) than close to hydrophobic residues. The finding indicates that the physico-chemical properties of the respective amino acids have a larger effect on the local free energy than the protein curvature, although curvature strongly affects both entropy and enthalpy.

To see how much the solvation contributions weigh in the total free energy balance, we decomposed the total folding free energy difference between the native fold and molten-globule-like conformation of crambin into individual enthalpy and entropy contributions from both protein and solvent. Although other computational methods, such as Replica Exchange Molecular Dynamics (REMD) [132, 183–185], allow to calculate ΔG with better precision and higher accuracy, our decomposition into individual energy and entropy terms provides additional insights into the thermodynamics of protein stability, even on globular protein energy budget level. Specifically, the protein-solvent interaction energies and the protein entropy turn out to be the only destabilizing contributions to the free energy, whereas the solvent-related entropy contributions favor the folded state of crambin by almost $500\,\mathrm{kJ\cdot mol^{-1}}$. Furthermore, the solvent entropy difference is not dominated by water mobilities, but by induced multi-body water correlations.

As the full ensemble of all potential molten-globule-like conformations is not sampled by our MD simulations, all differences of the free energy contributions were calculated with respect

to four ensemble members of the same molten-globule-like conformation. Particularly the calculated protein entropy is therefore underestimated. For a rough estimate of the entropy error, we assume an effective number of 100 to 1000 metastable conformations [186, 187] with similar conformational entropies. In this case, a configurational entropy contribution between $k_B T \log 100 = 11\,\text{kJ·mol}^{-1}$ and $k_B T \log 1000 = 17\,\text{kJ·mol}^{-1}$ would be missed, thereby not changing the budget qualitatively. As the first transition occurred after $1.5\,\mu s$ (see *Supplemental Information* Fig. 31.1), 1000 metastable states would correspond to an approximate but reasonable equilibration time of 1.5 ms.

With this caveat in mind, our calculations are in line with the known delicate balance between enthalpy and entropy, resulting a free energy difference of $(53 \pm 84)\,\text{kJ·mol}^{-1}$, which corresponds to the enthalpy change due to just a few hydrogen bonds. The solvent entropy (free energy contribution of $\approx 500\,\text{kJ·mol}^{-1}$) was found to be one of the most important stabilizing contributions to the free energy difference. Our study therefore also provides quantitative mechanistic insights into cold denaturation, driven by temperature-induced weakening of hydrophobic interactions [20]. Indeed, the obtained mean entropy and enthalpy contributions would yield cold and heat denaturation temperatures of crambin within the range of $-42\,°\text{C}$ to $-22\,°\text{C}$, and $60\,°\text{C}$ to $68\,°\text{C}$, respectively, assuming protein-typical heat capacity differences ΔC_P between 5 and $15\,\text{kJ·mol}^{-1}\text{·K}^{-1}$ [20, 26] (see *Supplemental Information*). These temperature ranges agree with typical unfolding temperatures observed for globular proteins [20, 24, 25]. The spatial resolution of the solvent free energy terms identified favorable free energy contributions due to solvated hydrophilic residues (charged or polar) as significantly larger than unfavorable contributions of solvated hydrophobic amino acids. The methodology might be relevant for computational in the context of computational protein design.

As the local free energy contributions around each residue also depend on the stereochemical environment formed by adjacent residues, the results presented here are specific to crambin in the assessed native fold and the molten-globule-like configuration. Nevertheless, because crambin exhibits a number of stereochemical motifs also present in most other globular proteins, it seems quite likely that the observed effects are - at least qualitatively - features of most globular proteins in general.

31 Supplemental Information

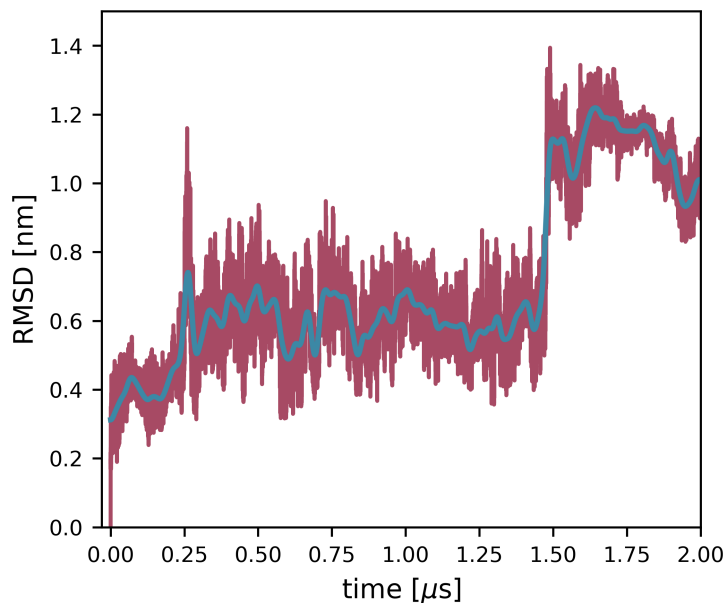## 31.1 Unrestrained MD simulation for the protein entropy calculation

Figure 31.1: RMSD of the molten-globule-like conformation during an unrestrained MD simulation for the protein entropy calculation or crambin. The $10\,\mathrm{ns}$ running average is shown in blue.

31.2 Decrease of water-water correlations upon binding to a protein charge

To show that the mutual information between a bound water molecule and a second-shell water molecule is smaller compared to bulk water, we considered two molecules with positions x and y.

In bulk, the molecules have single-body entropies $S(X)$, $S(Y)$ and the joint entropy $S(X,Y)$. Their mutual information (correlation) is $I(X,Y) = S(X) + S(Y) - S(X,Y)$.

Upon binding of the first molecule (X) to a protein charge, its entropy $\tilde{S}(X) < S(X)$ is reduced compared to bulk water (see main text Fig. 29.1A). As shown by the distribution of $-T\Delta S_1$ in the main text Fig. 29.1A, the second-shell molecule is almost unaffected by the binding of X, i.e., its single-body entropy $\tilde{S}(Y) = S(Y)$ remains unchanged. We furthermore assume that the conditional entropy $\tilde{S}(X|Y) \leq S(X|Y)$ is also reduced compared to bulk water, as the mobility of X is hindered upon binding.

The mutual information

$$\begin{aligned}
\tilde{I}(X,Y) &= \tilde{S}(X) + \tilde{S}(Y) - \tilde{S}(X,Y) \\
&= \tilde{S}(X) - \tilde{S}(X|Y) \\
&< S(X) - S(X|Y) = I(X,Y)
\end{aligned}$$

is therefore smaller compared to bulk.

31.3 Comparison of the local convexity with the entropic multi-body contributions
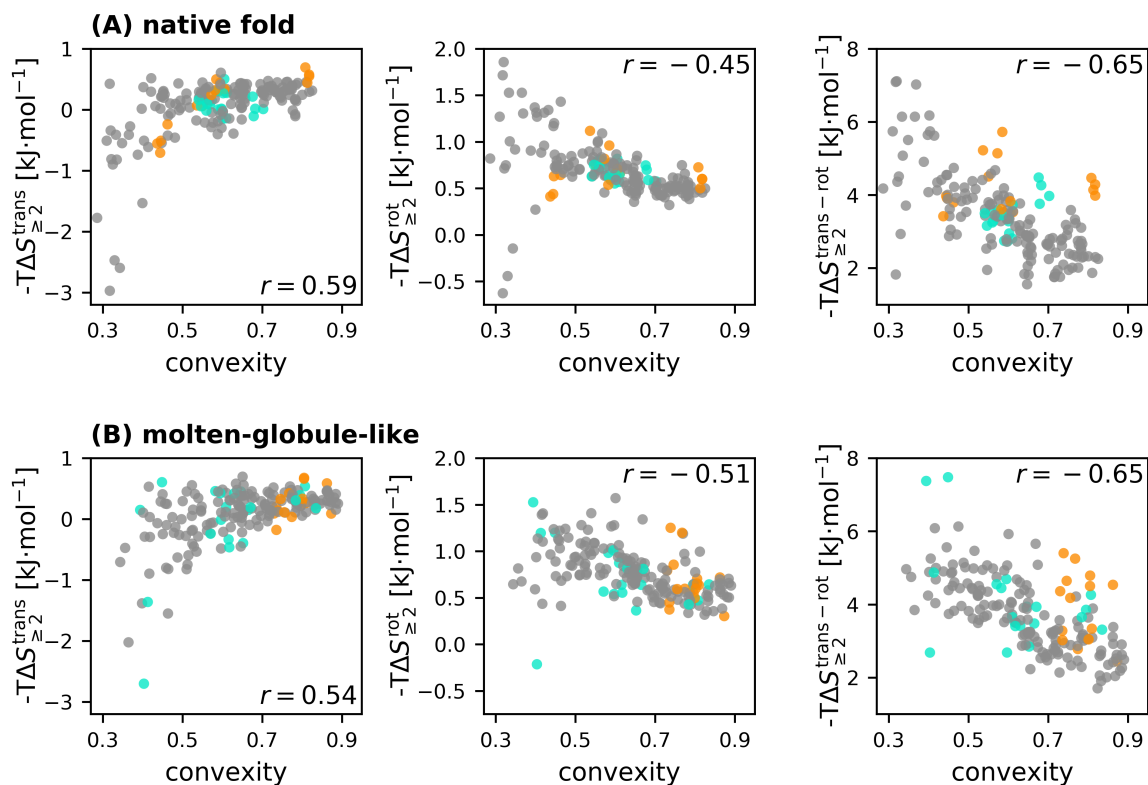

Figure 31.2: Dependence of the multi-body correlation contributions $-T\Delta S_{\geq 2}$ (relative to bulk water) on the local convexity for the native fold (A) and the molten-globule-like conformations (B). Charged amino acids are colored orange, polar amino acids are shown in cyan. Apolar amino acids are colored grey. Pearson correlation coefficients are stated in the corners of each plot.

31.4 Estimation of the denaturation temperatures from the simulation results

The temperature dependence of the enthalpy difference between folded and unfolded states of a protein can be approximated as Dias et al. [20]

$$\Delta H(T) = \Delta H(T_c) + (T - T_c)\Delta C_P,$$

where T_c is the heat denaturation temperature and ΔC_P is the heat capacity difference between folded and unfolded states, which here is assumed to be constant. In the same manner, the entropy difference reads Dias et al. [20]

$$\Delta S(T) = \frac{\Delta H(T_c)}{T_c} + \Delta C_P \log\left(\frac{T}{T_c}\right).$$

Using the simulation results $\Delta H_{300} = \Delta H(300\,\text{K}) \approx 228\,\text{kJ·mol}^{-1}$ and $-300\,\text{K} \cdot \Delta S_{300} = -300\,\text{K} \cdot \Delta S(300\,\text{K}) \approx -175\,\text{kJ·mol}^{-1}$, the above expressions can be rewritten as

$$\Delta H(T) = \Delta H_{300} + (T - 300\,\text{K})\Delta C_P$$
$$\Delta S(T) = \Delta S_{300} + \Delta C_P \log\left(\frac{T}{300\,\text{K}}\right).$$

Assuming typical values of ΔC_P between 5 and $15\,\text{kJ·mol}^{-1}\text{·K}^{-1}$ Dias et al. [20], Robertson and Murphy [26], the free energy difference $\Delta G(T) = \Delta H(T) - T\Delta S(T)$ has a concave shape, as illustrated in Fig. 31.3.
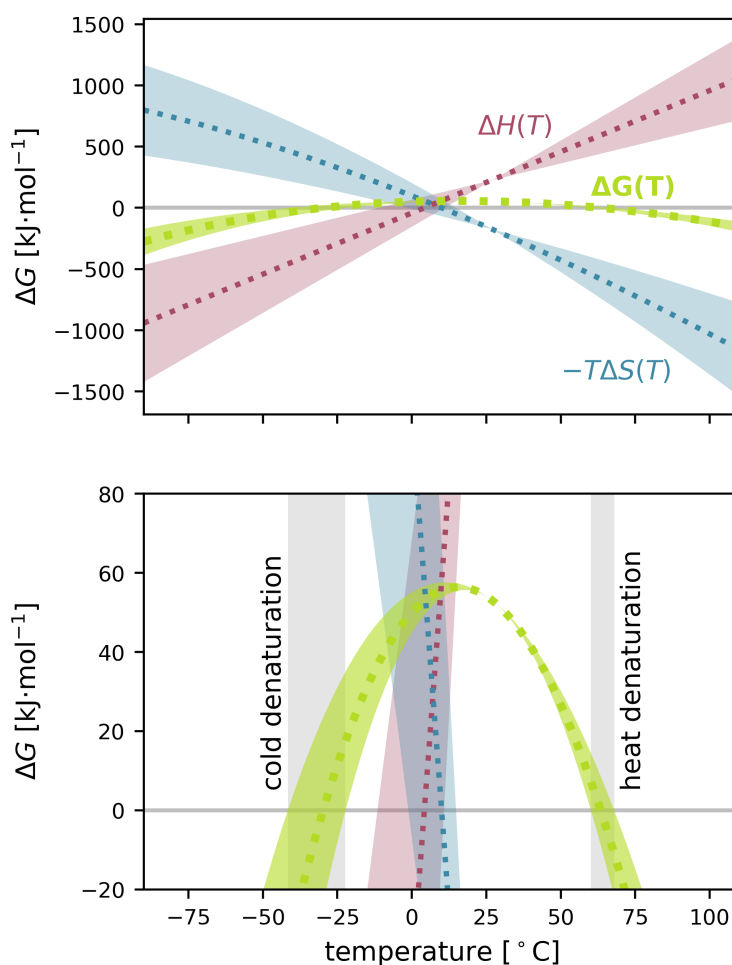

Figure 31.3: Temperature dependence of $\Delta H(T)$, $\Delta S(T)$, and $\Delta G(T)$ for ΔC_P between 5 and $15\,\text{kJ·mol}^{-1}\text{·K}^{-1}$. The bottom panel is a zoomed-in version. Cold and heat denaturation temperature ranges are marked in grey.

The roots of $\Delta G(T)$, correspond to the cold (between $-42\,°\text{C}$ and $-22\,°\text{C}$) and heat denaturation (between $60\,°\text{C}$ and $68\,°\text{C}$) temperatures, respectively.

31.5 Spatially resolved thermodynamic quantities on off-center slices

The following figures show the spatially resolved solvent enthalpies, entropies, and free energies on slices across the three dimensional space, as visualized in Fig. 29.1 of the main text, but offset by 1.0, 0.5, −0.5, and −1.0 nm in depth relative to the protein center.

Figure 31.4: Spatially resolved solvent enthalpies, entropies, and free energies of crambin in the native fold (A) and a molten-globule-like conformation (B), as shown in Fig. 29.1 of the main text. Visualization-slice across the three-dimensional space offset by $\approx 1.0\,\mathrm{nm}$ relative to the center of the protein.

Figure 31.5: Spatially resolved solvent enthalpies, entropies, and free energies of crambin in the native fold (A) and a molten-globule-like conformation (B), as shown in Fig. 29.1 of the main text. Visualization-slice across the three-dimensional space offset by $\approx 0.5\,\mathrm{nm}$ relative to the center of the protein.
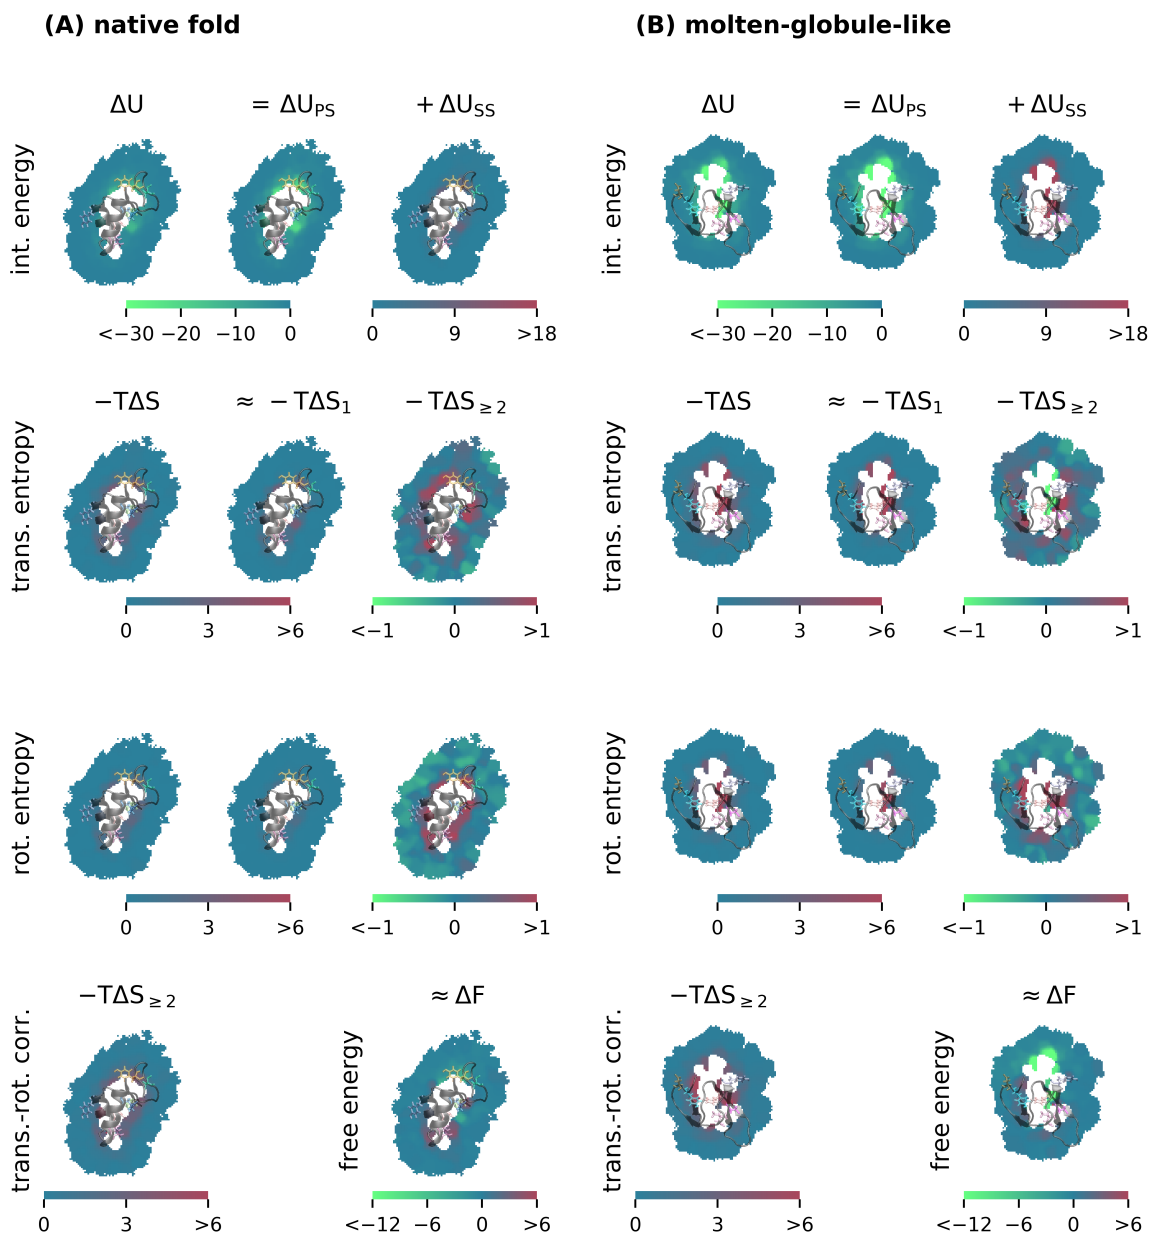
Figure 31.6: Spatially resolved solvent enthalpies, entropies, and free energies of crambin in the native fold (A) and a molten-globule-like conformation (B), as shown in Fig. 29.1 of the main text. Visualization-slice across the three-dimensional space offset by $\approx -0.5\,\text{nm}$ relative to the center of the protein.
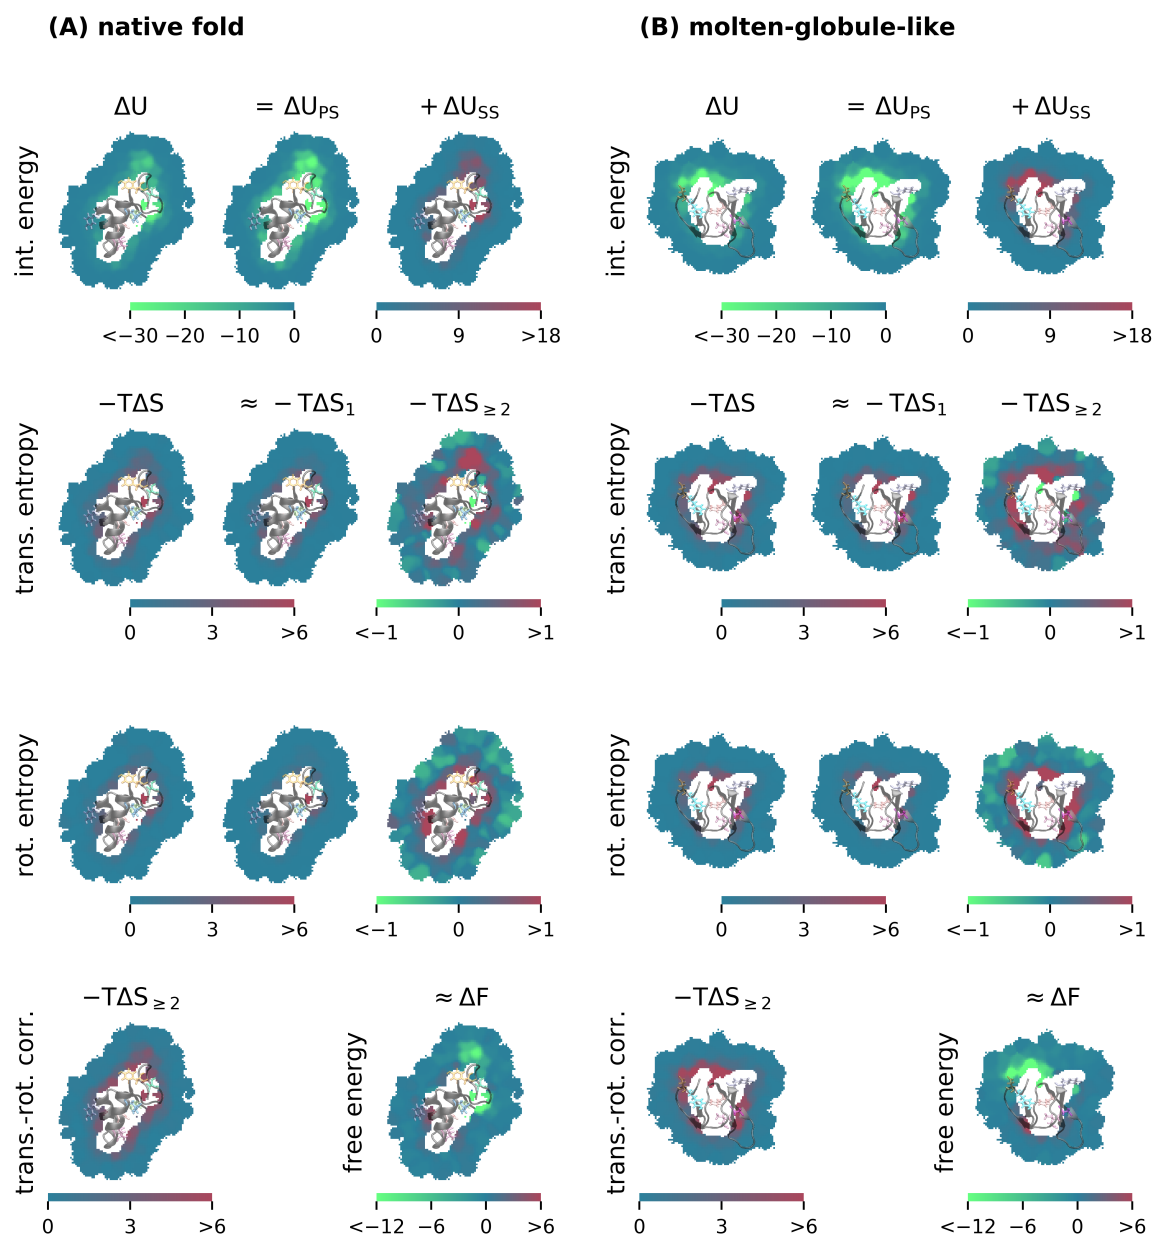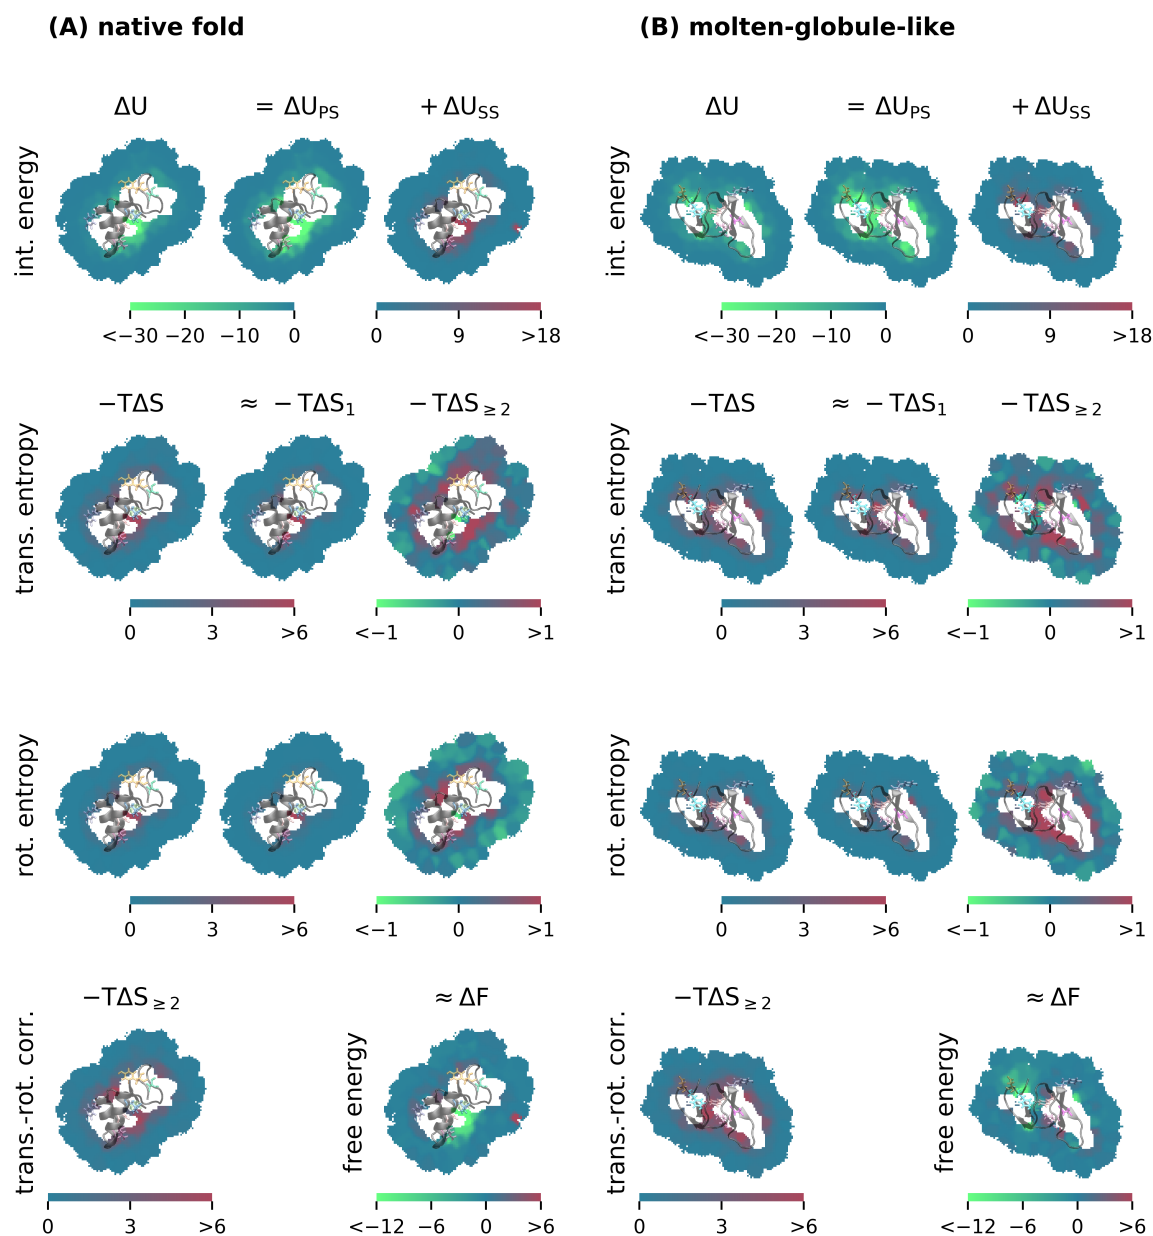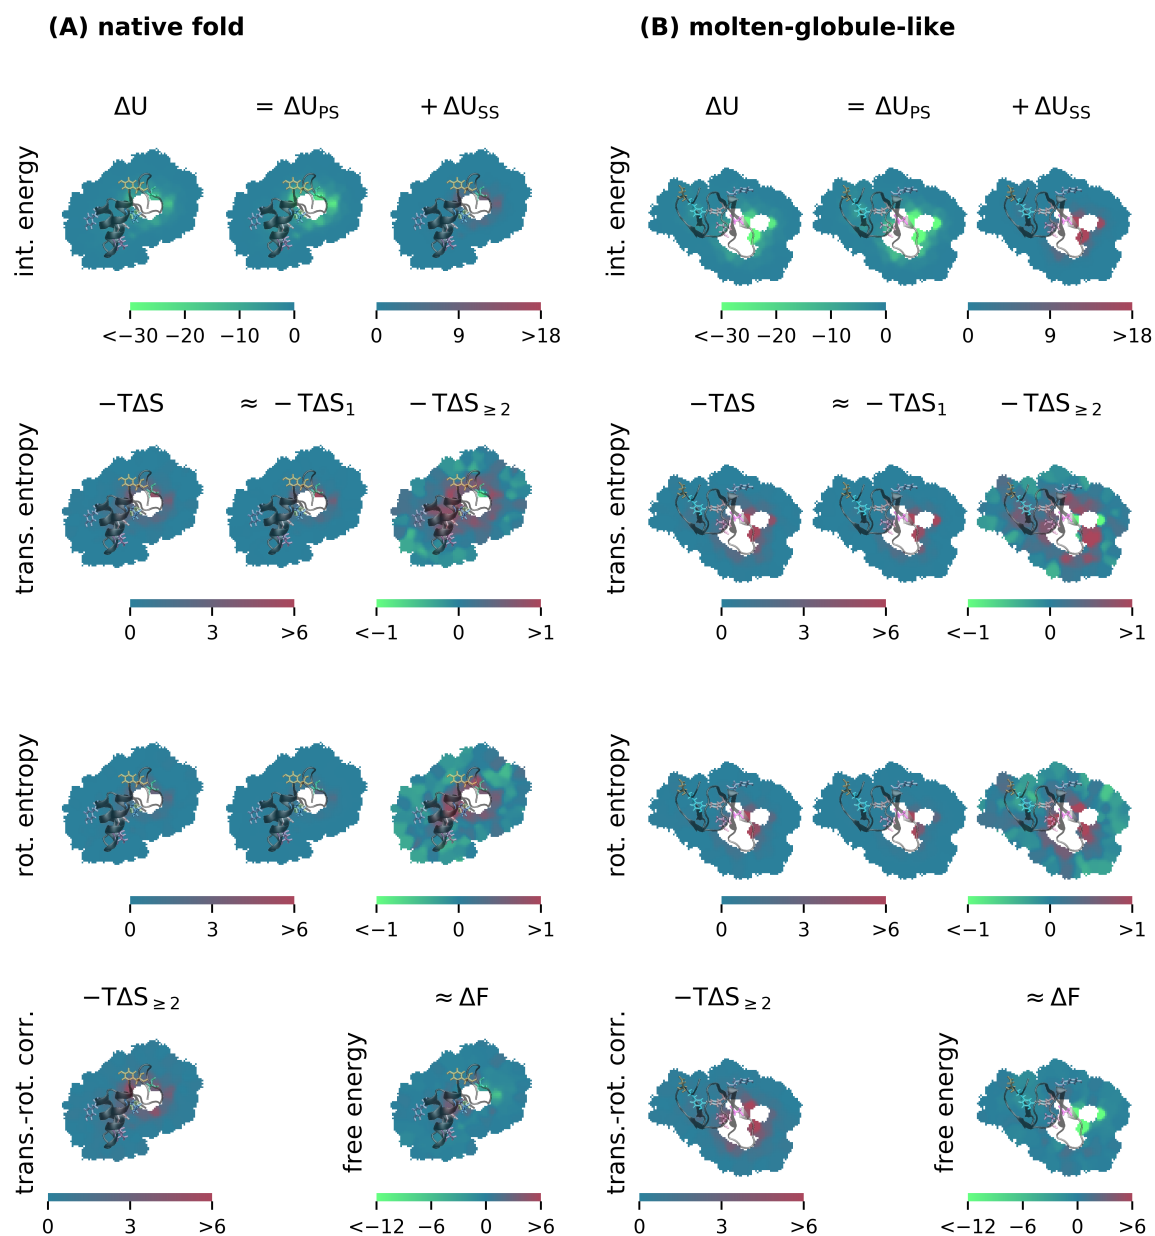
Figure 31.7: Spatially resolved solvent enthalpies, entropies, and free energies of crambin in the native fold (A) and a molten-globule-like conformation (B), as shown in Fig. 29.1 of the main text. Visualization-slice across the three-dimensional space offset by ≈ -1.0 nm relative to the center of the protein.

VII | Membrane Thickening

The following text resulted from a collaboration with the "Laboratory of Neurobiology", lead by Prof. Dr. Reinhard Jahn, and is currently in preparation for submission as
A. Witkowska, L.P. Heinz, H. Grubmüller and R. Jahn **"Tight docking of membranes before fusion represents a novel, metastable state with unique properties"**.

The project was initiated and lead be Reinhard Jahn and Agata Witkowska, who also carried out all experiments presented below. All MD-related research was supervised by Helmut Grubmüller and the MD simulations were carried out by myself. The manuscript was written by Agata Witkoswka, except for the MD-related sections, which were written by myself and revised by Helmut Grubmüller. The entire manuscript was revised by Reinhard Jahn.

Tight docking of membranes before fusion represents a novel, metastable state with unique properties

Agata Witkowska[1,†], Leonard P. Heinz[2], Helmut Grubmüller[2], Reinhard Jahn[1,3]

[1] Laboratory of Neurobiology, Max-Planck-Institute for Biophysical Chemistry, Göttingen, Germany

[2] Department of Theoretical and Computational Biophysics, Max-Planck-Institute for Biophysical Chemistry, Göttingen, Germany

[3] University of Göttingen, Göttingen, Germany

[†] Current address: Department of Molecular Pharmacology & Cell Biology, Leibniz-Forschungsinstitut für Molekulare Pharmakologie (FMP), Berlin, Germany

Abstract

Membrane fusion is a multi-step process leading to merger of two membranes that is essential in various biological processes such as membrane trafficking or viral infection. Proteins that catalyze membrane fusion need to overcome energy barriers to induce temporal loss of integrity of single bilayers. Here, we investigate structural changes in membranes arrested in tight docking intermediates that precede hemifusion. We use lipid vesicles docked with SNARE proteins as well as atomistic molecular dynamics simulations to probe properties of such docking interfaces. We demonstrate divalent cation-independent formation of a metastable, protein-free intermediate characterized by local membrane thickening due to profound lipid rearrangements induced by membrane surface dehydration.

One Sentence Summary

Before fusion, membranes form a tight and protein-free connection and become thicker due to dehydration and the resulting changed electrostatic interactions.

32 Main Text

Fusion of biological membranes is fundamental for the functioning of all living organisms ranging from the cell entry of enveloped viruses to the exocytotic release of neurotransmitters. While the proteins mediating fusion are evolutionarily unrelated and structurally diverse, the merger of two bilayers appears to follow a common pathway involving a sequence of structurally distinct intermediates. These begin by loose protein-mediated membrane contact and is followed by tight apposition of the membranes while still maintaining bilayer integrity. Then, membrane structure is disrupted by the merger of the proximal leaflets, resulting in a fusion stalk or a hemifusion diaphragm. This is followed by the rupture of the diaphragm (fusion pore) that then expands, thus completing membrane merger [55, 56].

While the pathway outlined above is supported by an increasing body of experimental evidence and theoretical modeling [54–56, 188, 189], there are still crucial gaps in knowledge, in particular with respect to the steps immediately before the first bilayer disruption. Importantly, these are also the steps on which proteins regulating final progression to fusion operate. Thus, we need to clarify the structural and energetic details of these intermediates to precisely understand the molecular mechanisms of such regulatory proteins (e.g. synaptotagmins) that are controversially discussed since more than 20 years [189–191].

Here we focus on the tightly docked intermediate in which membranes are apposed to each other with a distance of $< 1\,\mathrm{nm}$. While this state is well known from cryo-electron microscopy (cryoEM) studies as an intermediate in events as diverse as fusion of trafficking organelles [57, 192, 193], fusion of mitochondria [194], or cell entry of influenza virus [195], its biophysical features have remained largely enigmatic. To reach such state, repulsive forces between negatively charged lipid headgroups and the energy barrier involved in dehydration need to be overcome. It is unclear, which forces are counteracting (headgroup chelation by divalent cations, van der Waals and hydrophobic forces, protein "clamps"), if and how proteins are cleared from these contact sites, and how the subsequent transition to stalk formation is facilitated [55, 56, 196].

Here we have used an in vitro system in which fusion of artificial membranes is mediated by SNARE (soluble N-ethylmaleimide-sensitive-factor attachment receptor) proteins involved in neuronal exocytosis. These proteins are known to fuse membranes by a consecutive assembly of complementary SNAREs between the membrane that is initiated at the N-terminal ends and progresses towards the C-terminal membrane anchors, referred to as SNARE zippering [189, 190]. The in vitro system was optimized for studying fusion intermediates by omitting
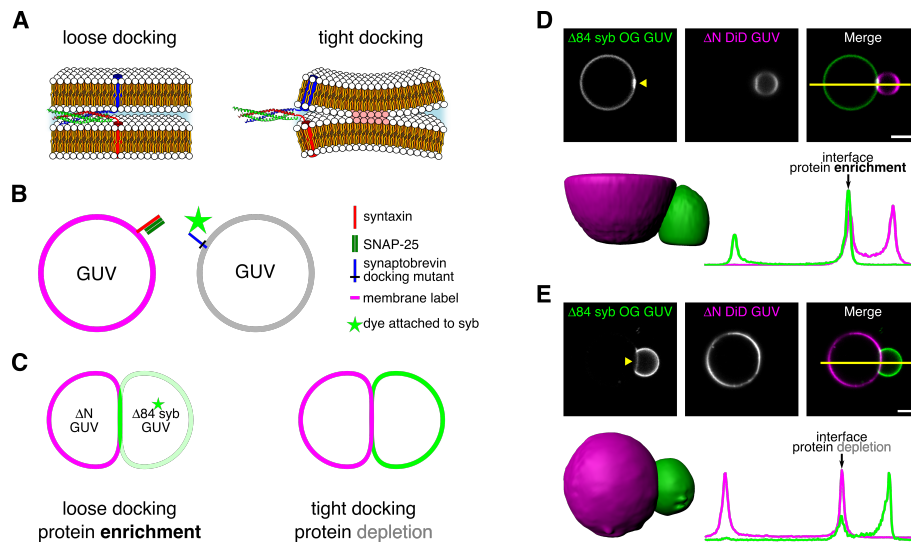
Figure 32.1: Protein inclusion or depletion at the interface differentiates loose and tight dock-
ing. (A) Schematic illustration of loose and tight docking. (B) Schematic illus-
tration of a GUV-GUV docking as an experimental system to probe membrane
and protein interactions at loose and tight docking interfaces with one GUV set
containing membrane label (magenta) and Q-SNARE acceptor complex, and a
second set with labelled syb docking mutant, such as $\Delta 84$ syb (green). (C) In
experimental system presented in panel B, at loose docking interfaces trapping
and enrichment of assembled SNARE complexes occurs over time (accumulation
of green label at the interface). At tight docking interfaces due to interaction of
lipids from opposing membrane bilayers, exclusion of proteins occurs resulting in
depletion of protein (green label) signal from the docking interface. (D and E)
Example of a loosely (D) and tightly (E) docked GUV pair with syb signal ($\Delta 84$
syb labeled with Oregon Green 488 (OG488)) accumulation (D) or depletion (E)
at the docking interface. Position of a line profile (graph below the images) is
marked with yellow line. Q-SNARE GUVs are labeled with a membrane dye
DiD. Scale bars $5\,\mu$m.

all upstream steps regulating initial assembly [197] and consists of an activated Q-SNARE
complex in the acceptor membrane and the R-SNARE in the donor membrane. We have
shown previously that in this system the tightly docked state can be stably reproduced as a
fusion intermediate and furthermore, that it can be arrested in this state by a point mutation
in the R-SNARE synaptobrevin [57, 198]. As shown in Fig. 32.1A, membrane contact results
in N-terminal assembly of the SNARE proteins. In this state (referred to as loosely docked
state), membranes are still separated by a hydrated gap, and disruption of SNARE assembly
by the disassembly ATPase NSF or by competition with soluble synaptobrevin fragments
leads to vesicle dissociation [57]. The system then progresses to the tightly docked state that
cannot be reverted anymore by SNARE disassembly suggesting the involvement of attractive
forces of unknown nature that only operate at subnanometer distances.

To allow for using light microscopic techniques for the characterization of the intermediates,

we adapted this system by using GUVs (giant unilamellar vesicles) instead of large unilamellar vesicles. Two sets of GUVs were reconstituted with complementary sets of SNARE proteins — synaptobrevin 2 (syb) that in neurons resides on synaptic vesicles and a stabilized complex containing syntaxin 1A and SNAP25a [197] that are naturally present on the plasma membrane. In this system, syb contains a single residue deletion (Δ84) that is known to trap fusing membranes at the tightly docked state [57, 192, 198]. In one vesicle set, membrane of a GUV was labelled allowing to control localization of lipids as well as making sure that no hemifusion has occurred. In the other set of vesicles, syb was labelled to allow for monitoring of protein behavior at the interface (Fig. 32.1B). When mixed together, these vesicles would interact with each other, fuse, or get stalled at one of the fusion intermediates such as docking or hemifusion. The low curvature of the GUV membrane leads to a higher energy barrier for membrane merger [192], thereby increasing the probability of obtaining arrested docked states.

If our previous assumptions that SNAREs are excluded from tight docking interfaces are correct, we should be able to observe in such experimental setup decreased protein intensity at docking interfaces (Fig. 32.1C left). Conversely, at arrested loose docking interfaces, fusion complexes in trans will get trapped and thus accumulated over time, visible as increased protein signal at the interface (Fig. 32.1C right). Indeed, with our experimental system we were able to demonstrate that there are two distinct docking states different from hemifusion, that are characterized by either protein accumulation at the docking interface (loose docking, Fig. 32.1D) or protein depletion from the membrane-membrane contact site (tight docking, Fig. 32.1E). Moreover, in some rare cases, we were also able to observe transition of these trapped intermediates to full fusion as shown in Fig. 32.2A and B providing evidence that tight docking can be productive fusion intermediate.

To better understand the forces that initiate and stabilize the tightly docked state we probed the occurrence probability of loose and tight docking interfaces in various electrostatic conditions by varying concentrations of divalent cations — Ca^{2+} or Mg^{2+} — and negatively charged lipids PS and PI(4,5)P2 (Fig. 32.2C). We find that increasing concentration of divalent cations (both Ca^{2+} and Mg^{2+}) increases the probability of tight docking interface formation, also within the physiologically relevant range of $\approx 50\,\mu$M. Moreover, addition of negatively charged lipid (PI(4,5)P2) on one of the docking membranes (as present on the plasma membrane) reduces the number of tight docking interfaces indicating a higher energy barrier due to electrostatic repulsion (Fig. 32.2C right). Following this result, we also probed localization of negatively charged lipids in docked vesicles with fluorescently labelled lipids (PS-TMR or PI(4,5)P2-TMR, Fig. 32.2D and E). We did not observe depletion of these lipid species from docking interfaces, however in case of loose docking, PI(4,5)P2 seemed to be enriched (as visible on loose docking example in Fig. 32.2E), probably due to its interaction with SNARE proteins (specifically syntaxin 1) that accumulate at the interface.

Previously [57] different membrane-membrane distances were observed by cryoEM. Specifi-
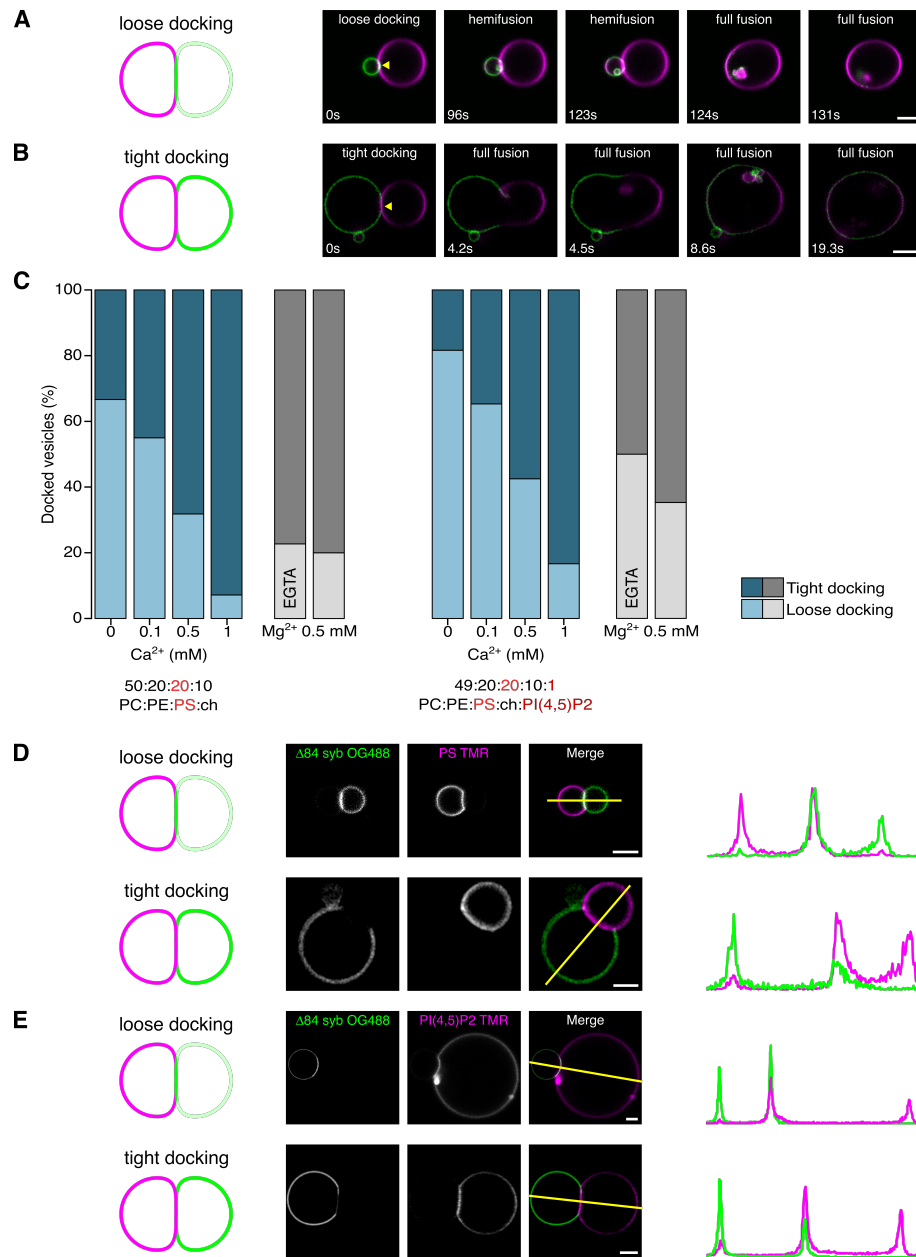
Figure 32.2: Loose and tight docking are productive membrane fusion intermediates promoted but not exclusively caused by lipid charge compensation from divalent cations included in the surrounding buffer. (A and B) Time series of GUVs docked in loose (A) or tight (B) way transiting to the fully fused state. Membrane and protein labeling as in Fig. 32.1. (C) Occurrence frequency of loose and tight docking in various concentrations of either Ca^{2+} or Mg^{2+} and in the presence of membranes that contain various amounts of negatively charged lipids (PS and PI(4,5)P2) (D and E) Example of a loosely and tightly docked GUV pairs with $\Delta 84$ syb labeled with OG488 and membranes containing TMR labelled PS (D) or PI(4,5)P2 (E). PS distribution remains unaffected in both docking types while there is an enrichment of PI(4,5)P2 at loose docking interfaces. Scale bars $5\,\mu m$.

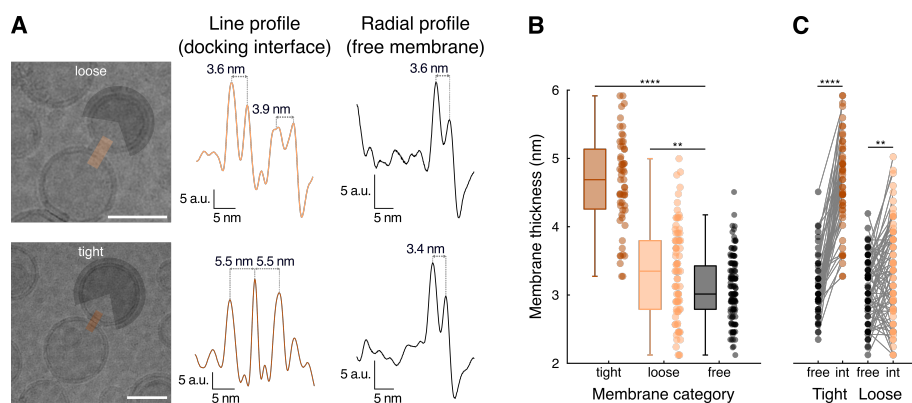Figure 32.3: Membrane thickening at tight docking interfaces. (A) Example liposome membrane radial profiles (membrane outside docking interfaces) and line scans (membranes at the interface) used for comparing membrane thickness. Scale bars 50 nm. (B) Comparison of membrane thicknesses at tight, loose, and docking-free interfaces obtained by analyzing cryoEM images as shown in A. Boxes represent interquartile range, and whiskers below and above indicate full data range. Line in a box represents median and square point represents the mean. (C) Membrane thickness variations within single membranes being involved in docking interfaces (int) or free.

cally, in tightly docked membranes with extended, flat docking interfaces (up to 100 nm-long for LUVs [57]) signals coming from lipid headgroups of proximal leaflets would blend together into one, thin line (see also Fig. 32.3A). We wanted to further characterize structural rearrangements of the membranes trapped in the tightly docked state and analyzed membrane thicknesses of at docking interfaces as well as membranes not involved interactions (Fig. 32.3A–C). Strikingly, we observed profound membrane thickening localized directly at the tight docking interfaces (Fig. 32.3). These structural changes were absent in membranes not involved in the docking process and were not observed in vesicle pairs classified as loosely docked.

To confirm that membrane thickening at close distances is due to membrane intrinsic properties, and to reveal the underlying molecular mechanism, we performed atomistic simulations of two opposing membranes at varying distances (Fig. 33.1A), mimicking experimental conditions but without proteins. Indeed, for distances below 1.5 nm, significant thickening of the membranes was observed also in the simulations, independent of lipid composition (Fig. 32.4A). We also observed shrinkage of the membrane area (Fig. 32.4B), which raises the question which of the two changes is the primary cause. To answer this question, Fig. 32.4C shows that the volume of the membrane decreases with the membrane distance. Because the membrane behaves as a nearly incompressible fluid, this finding suggests that area shrinkage drives membrane thickening, as otherwise the volume should increase. Indeed, inducing a similar area shrinkage by increasing the lateral membrane pressure in additional simulations (Fig. 32.4D, purple) quantitatively reproduces the thickening (green). Our simulations also
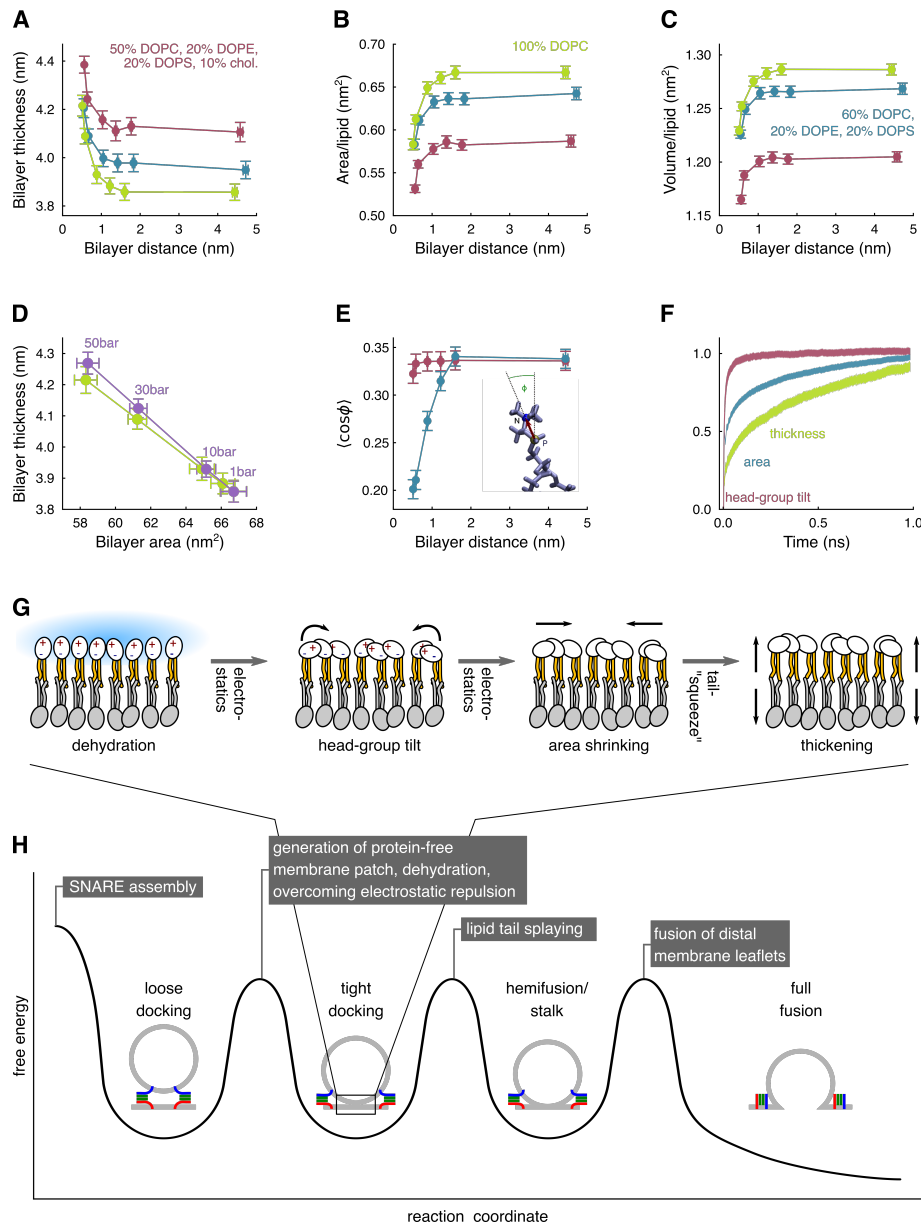
Figure 32.4: Membrane thickening at close distances is caused by dehydration-driven electrostatically-induced headgroup tilt and area compression. Atomistic molecular dynamics simulations show changes in membrane thicknesses (A), area (B), and volume (C) when two membranes with three different lipid compositions (see plots) are apposed to each other at various distances. (D) Enforced membrane area compression reproduces membrane thickening at close distances. (E) Lipid headgroups within the inner membrane leaflet at the docking interface exhibit higher tilt ($\cos\phi$) upon close approach. (F) Change of head-group tilt, area, and thickness over time (average over all 500 simulation runs; for easier comparison, all values were normalized to a range between 0 and 1). (G) Scheme proposing order of events resulting in membrane thickening. (H) Redefined energy landscape for membrane fusion.

showed that with decreasing distance, lipid chain order (Fig. 33.1B) as well as headgroup tilt of the opposing membrane leaflets increase (Fig. 32.4E, blue), whereas the outwards facing leaflets (Fig. 32.4E, red) are nearly unaffected. The strong electrostatic dipole moment of the lipid head groups suggests [58] electrostatic interactions as the primary cause. In particular, the larger tilt of the lipid head groups should allow for an electrostatically more favorable antiparallel arrangement. Indeed, the electrostatic interaction energy between the inner leaflets decrease upon approach, whereas the outer leaflets show only little change (Fig. 33.1C).

To confirm that lipid headgroup tilting is the actual cause of area shrinkage, in further simulations we enforced them to tilt, which indeed reduced the membrane area (Fig. 33.1D). We attribute these electrostatic effects to reduced shielding caused by dehydration, which was also seen in our simulations (Fig. 33.1C). Independent support comes form the timing of events, which follows the proposed causal chain (Fig. 32.4F and Fig. 33.1E-F). Indeed, lipid headgroup tilting precedes area shrinkage, which precedes membrane thickening. Figure 32.4G summarizes the causal chain of events revealed by our atomistic simulations: Due to reduced electrostatic shielding, dehydration of the inner leaflet surfaces tilts the lipid headgroups, which laterally contracts the membrane. This contraction causes increases lipid ordering and, ultimately, drives thickening of the membrane.

Taken together, we now provide a coherent description of a metastable docking state that constitutes a novel intermediate in the pathway leading to membrane fusion with unique properties. Despite growing experimental evidence it has so far been difficult to reconcile such a tightly docked intermediate with standard concepts in membrane biophysics. Current fusion models agree on the notion that due to electrostatic repulsion membrane contact destabilizes the membranes before the formation of a fusion stalk [56]. To minimize the energetic penalty of close membrane contact, many models imply that the contact area is limited by the formation of point-like membrane protrusions with fusion being facilitated by lipid packing defects at the apex of a highly deformed membrane [188]. Accordingly, tightly docked membrane-membrane interfaces reported in cryoEM studies [57, 192, 193] were considered to be stalled, off-pathway states.

Our results significantly contribute to our understanding of the energy landscape governing the early steps of membrane fusion (Fig. 32.4H). After exergonic assembly of SNARE proteins [199] a metastable state is reached where the membranes are still several nm apart from each other. This is followed by an energy barrier: Water and ions must be removed from the membrane contact zone, and the electrostatic repulsion between opposite lipid headgroups needs to be overcome [55, 56]. We now demonstrate that a metastable state follows that is characterized by a partially dehydrated and tight adhesion between membranes. The contact zone is free of proteins, does not require, albeit being stabilized by, divalent cations, and is associated with dehydration and a change in lipid organization, resulting in membrane thickening. Progression towards stalk formation probably requires tail splaying of membrane lipids [54] which may be facilitated by the increased hydrophobicity and the changed lipid

head geometry at the contact site. Probably, stalks are initiated at the edge of the membrane contact zone where increased curvature may cause lipid packing defects facilitating such transitions [192]. Furthermore, we propose that in fast secretory cells such as neurons fusion may be arrested at this state. It is conceivable that the energy barrier separating this state from stalk formation and subsequent fusion is moderate and can easily be overcome by accessory proteins perturbing lipid packing at the membrane surface, such as synaptotagmins.

33 Supplementary Materials

33.1 Materials and methods

Materials

Lipids: DOPC (1,2-dioleoyl-sn-glycero-3-phosphocholine), DOPE (1,2-dioleoyl-sn-glycero-3-phosphoethanolamine), DOPS (1,2-dioleoyl-sn-glycero-3-phospho-L-serine), 18:1 biotinyl cap PE (1,2-dioleoyl-sn-glycero-3-phosphoethanolamine-N-(cap biotinyl)), DOPS (1,2-dioleoyl-sn-glycero-3-phospho-L-serine), cholesterol (ovine wool), PI(4,5)P2 L-α-phosphatidylinositol-4,5-bisphosphate (Brain, Porcine), TopFluor TMR (1-oleoyl-2-(6-((4,4-difluoro-1,3-dimethyl-5-(4-methoxyphenyl)-4-bora-3a,4a-diaza-s-indacene-2-propionyl)amino)hexanoyl)-sn-glycero-3-) PI(4,5)P2 and PS were purchased from Avanti Polar Lipids. Lipophilic tracer DiD, NeutrAvidin, biotinylated bovine serum albumin, were from Thermo Fisher Scientific.

Protein expression, purification, and labelling

SNARE proteins (syntaxin-1A (183–288) [200], SNAP-25 (cysteine free) [201], synaptobrevin-2 Δ84 C28 [202], and synaptobrevin-2 fragment (49—96) [197] were derived from Rattus norvegicus. Proteins were expressed in Escherichia coli strain BL21 (DE3) and purified via nickel-nitrilotriacetic acid affinity chromatography (Qiagen) followed by ion exchange chromatography on an Äkta system (GE Healthcare) as described before [57, 203]. The activated Q-SNARE complex [203] consisting of syntaxin, SNAP-25, and synaptobrevin fragment 49-–96 was obtained by overnight mixing at 4 °C in a buffer containing CHAPS, followed by ion exchange chromatography (MonoQ) as described before [197, 203]. Fluorescence labeling of synaptobrevin was carried out according to the manufacturer's instructions using Oregon Green 488 (OG488) iodoacetamide (Molecular Probes) [57].

Liposome and GUV preparation

For basic liposome mixtures lipids PC, PE, PS, and cholesterol were mixed in a ratio of 5:2:2:1, respectively. In GUVs 1 mol% of DOPE was replaced with biotinyl-cap-PE for immobilization on neutravidin-functionalized coverslips [203], and for fluorescent labelling 1 mol% of DOPC

was replaced with a lipophilic tracer DiD. For localizing PS and PI(4,5)P2 1 mol% was replaced with TMR-labelled versions of this lipids. Giant unilamellar vesicles were prepared by electroformation from small proteoliposomes (diameter ≈ 40 nm) containing either activated Q-SNARE complex (protein:lipid molar ratio 1:1000) or synaptobrevin (protein:lipid molar 1:500), as described before [203].

Microscopy imaging and analysis

Imaging was performed on scanning confocal microscopes Zeiss LSM 780 and 880. Classification of docked vesicles was performed on docked GUV pairs (recognized based on labels in each of the GUV type). Pair was classified as loosely docked if there was physical contact between vesicles, and fluorescent signal from syb labelled with OG488 was detectable and not reduced in comparison to protein signal on the free membrane of syb vesicle at the interface. Pair was classified as tightly docked if syb signal was decreased or absent at the GUV-GUV docking interface. Enrichment or decrease of syb signal at the interface was assessed by line scans perpendicular to docking interface. Line scans and images for some visualizations were rotated using bicubic extrapolation rotation from Fiji.

Cryo-EM image analysis

Docked liposome pairs identified in cryo-EM images with nominal magnification of 20,000–40,000x were analyzed with a custom written script executed in Fiji [204] and GNU Octave [205]. Briefly, membrane thickness was extracted from line (docking interfaces) or radial (free membrane) profiles as distance between peaks denoting headgroups of both membrane leaflets. Radial profiles were obtained with the "Radial Profile Extended" plugin for ImageJ.

Moleculad dynamics simulations

All molecular dynamics simulations were carried out using the software package Gromacs 2018 [73–77] with the CHARMM36m force field [81, 83, 84]. SETTLE [95] and LINCS [96] constraints were applied to all water molecules and bonds to hydrogen atoms respectively, and the system was time-propagated using the leap frog integrator with a time step of 2 fs. Electrostatic forces were calculated using the Particle-Mesh Ewald (PME) method [160] with a 1.2 nm real space cut-off; the same cut-off was used for Lennard-Jones potentials [42]. The V-rescale thermostat [98] with a time constant of 0.1 ps was used. Unless stated otherwise, all simulations were performed at 1 bar pressure, independently applied along the membrane and the lateral directions (semi-isotropic pressure coupling). Equilibration simulations were carried out using the Berendsen barostat with a time constant of 5.0 ps and production runs were performed using the Parrinello-Rahman barostat [100, 101] with a time constant of 1.0 ps.
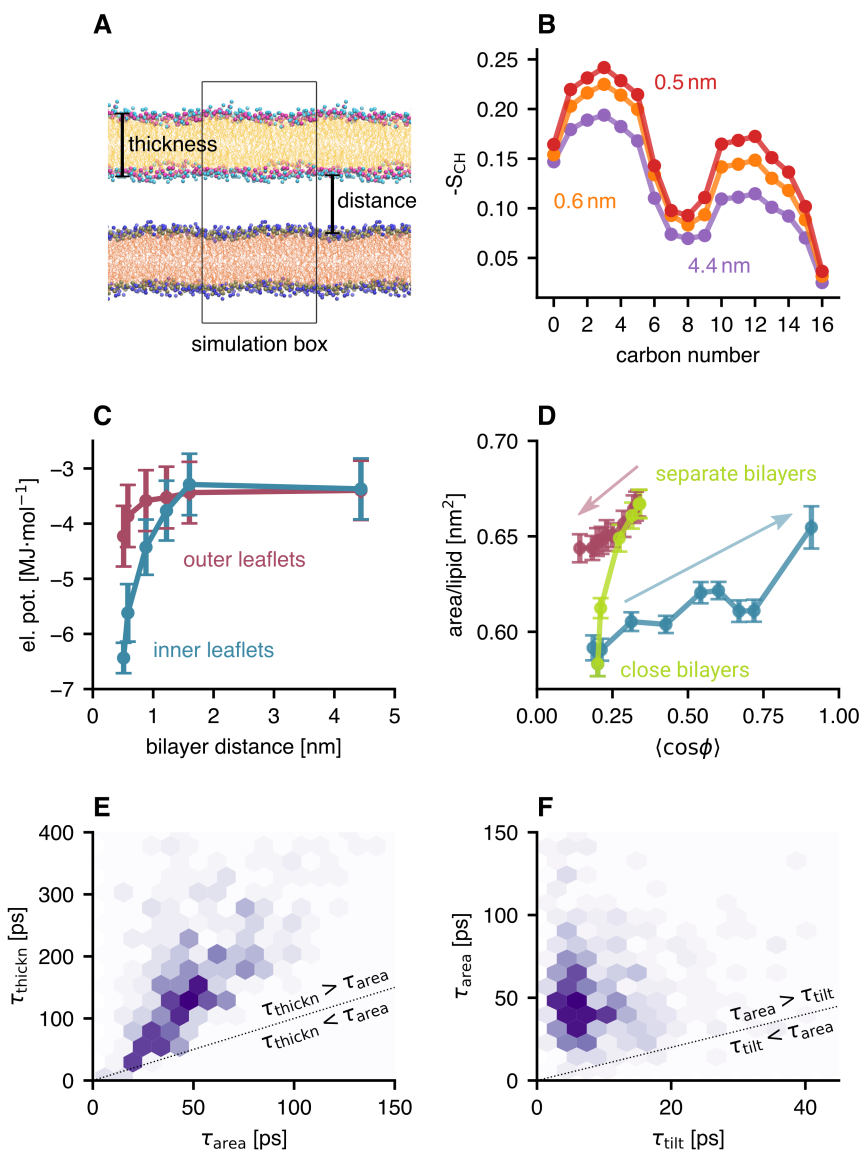
Figure 33.1: (A) Double-membrane simulation setup. The two bilayers are shown in yellow and orange; P-atoms are golden and pink, N-atoms are shown in blue and cyan color. (B) Deuterium order parameters for membrane consisting of $100\,\%$ DOPC at selected distances. (C) Average leaflet-internal electrostatic potential for inner (blue) and outer (red) leaflets depending on the bilayer distance. (D) Relationship between headgroup tilt $(\cos\phi)$ and area per lipid. The data of membrane consisting of $100\,\%$ DOPC at various distances is shown in green. Results from in-plane biased simulations of separated bilayers and out-of-plane biased simulations of bilayers at close distance are shown in red and blue, respectively. (E, F) Histograms of the response times of area shrinkage, bilayer thickening, and head group tilt, showing that the tilt precedes the area shrinkage, which precedes the thickening.

Simulation setup

For each of the tested lipid compositions, a symmetric bilayer of 100 lipids per leaflet was created using the MemBuilder II [206] webserver[14]. Subsequently the systems were solvated using the Gromacs tool gmx solvate with approximately 18000 water molecules; if necessary, charge-neutrality was achieved by adding K-ions, and further ions equivalent to $75\,\mathrm{mmol \cdot L^{-1}}$ KCl were added, mirroring the experimental conditions. Gradient descent energy minimization was performed until a maximum force of $100\,\mathrm{kJ \cdot mol^{-1} \cdot nm^{-1}}$ was reached, and the systems were equilibrated for $50.5\,\mathrm{ns}$.

The equilibrated bilayers were doubled to obtain double-membrane systems with inter-membrane distances varying between 0.4 and $4.7\,\mathrm{nm}$, as shown in Fig. 33.1A. Again, the systems were solvated and KCl ions were added to each compartment separately, as described above. All systems were again energy-minimized and equilibrated for a further $1\,\mathrm{ns}$, before data-acquisition runs, each lasting $200\,\mathrm{ns}$, were started.

Mean distances and mean thicknesses were calculated from the last $190\,\mathrm{ns}$ of each simulation run based on the distances between the average P-atom layers.

Response-time analysis

In order to determine the time sequence of headgroup tilt, area shrinkage and bilayer thickening, we carried out 500 independent non-equilibrium simulations, of the transition of two equilibrated single membranes in close contact to an equilibrated, i.e., thickened, double-membrane.

500 start structures were obtained by taking snapshots every $1\,\mathrm{ns}$ from a seeding trajectory of a single bilayer. The bilayers were then doubled as described above to obtain non-equilibrated double-membranes at a distance of $0.5\,\mathrm{nm}$, and each replica was simulated for $1\,\mathrm{ns}$. To allow analysis of the equilibration process, trajectory coordinates were stored every $0.1\,\mathrm{ps}$.

To determine the timescales of headgroup tilt, area, and thickness changes, exponential relaxations of the form $f(t) = x_\infty + (x_0 - x_\infty)e^{-t/\tau}$ were fitted to the data of each replica. Here, x_0 denotes the initial value of the observable (e.g. bilayer thickness at the beginning of the simulation), x_∞ is the mean in the double-membrane equilibrium, and τ is the relaxation time. The values obtained for τ_{tilt}, τ_{area}, and τ_{thickn} yield the time sequence of headgroup tilt, area shrinkage, and bilayer thickening, as shown in Fig. 33.1E and F.

Enforced headgroup tilt

Headgroup tilts, measured using the angle ϕ between N and P atoms of each lipid and the membrane normal (z-axis), were enforced by adding the angle-restraint term $k(1-\cos(\phi-\phi_0))$ to the force field. To bias separated bilayers towards more tilted headgroups, the minimum

---

[14]http://bioinf.modares.ac.ir/software/mb2/builder.php

of the potential ϕ_0 was placed at 90°. Similarly, the potential minimum was placed at 0° to bias double bilayers at close distance (0.5 nm) towards less tilted head groups. In both cases, the force constant was varied between 0.5 and 30 kJ·mol^{-1}, depending on the desired biasing strength, and each simulation lasted 200 ns.

VIII | Conclusion

34 Entropy Calculation

Obtaining sufficient sampling poses a significant challenge to any method for solvent entropy calculation, as the $6N$-dimensional configuration space of N water molecules requires an exponentially growing number of sample points for the same sampling quality due to the "curse of dimensionality". Furthermore, attempts to naively compute the configurational solvent entropy as

$$S_{\text{conf}} \propto -k_B \int \frac{d\boldsymbol{q}^N}{N!} \varrho(\{\boldsymbol{q}_i\}) \log \varrho(\{\boldsymbol{q}_i\})$$

suffer from an additional — and substantial — decrease in sampling quality, as the Gibbs factor $N!$ cannot be straightforwardly applied in numerical sampling schemes. In an analytical treatment, the Gibbs factor [44] fixes the overcounting problem, which arises, because each (physically equivalent) permutation of the identical solvent molecules in a microstate is counted separately, leading to an overcounting of microstates by $N!$. In a numerical approach, however, the full configuration space $\int d\boldsymbol{q}^N$ needs to be sampled first, including all $N!$ physically identical molecule permutations, and only then, the result is multiplied by $1/N!$.

Although the overcounting problem can be solved by switching from a molecule-based to a grid-based entropy calculation scheme that is agnostic about the particle identity, the "curse of dimensionality" has historically limited the accuracy of hydration shell entropy estimates. For instance, the accuracy of the method GIST [10, 43] is limited, as it usually approximates the full water entropy as the sum of single-body entropy terms, which are defined as integrals in just three dimensions and are thus easy to sample. This approximation is particularly critical when addressing the thermodynamics of the hydrophobic effect, where, as also this work has shown, multi-body correlation effects are crucial. Conversely, the method 3D 2-Phase Thermodynamics [11], in principle, captures the full high-dimensional configuration-entropy integral but does so at the cost of depending on strong model assumptions, which also results in limited accuracy.

To obtain solvation-shell entropies at a higher accuracy, I developed the new method Per|Mut, which employs a permutation reduction to address the overcounting problem and uses a third-order mutual information expansion to alleviate the "curse of dimensionality".

The permutation reduction maps all sample points into a smaller configuration space subvolume ("slice"), and thereby increases the sampling density by the Gibbs factor $N!$ without changing the physics. In (three-dimensional) real space, permutation reduction localizes the

water molecules. Instead of undergoing normal diffusive motion, each water molecule only samples a small region around a reference position.

The mutual information expansion decomposes the high-dimensional entropy integral into multiple low-dimensional integrals that can be sampled sufficiently. The MI terms denote entropies due to the fluctuations of individual water molecules (first order) and the entropy losses due to two- and three-body correlations (second and third order, respectively). The decomposition into first-order entropies and higher-order correlations between localized water molecules not only provides a local description of entropy and thereby a spatial resolution but, crucially, also provides a physical interpretation of entropy changes.

All MI terms were numerically calculated using a k-nearest neighbor (kNN) density estimator, which calculates the local probability density at every sample point by finding its kth neighbor and dividing by the volume of the ball that encloses the k neighbors. For Euclidean space and translational entropies, nearest neighbor search algorithms are implemented in popular machine learning libraries, such as scikit-learn [111] and the non-metric space library (NMSLib) [109, 110]. The algorithm has also been used to calculate rotational entropies [152, 207], however, these attempts used the volume of an Euclidean n-sphere, rather than the ball volume induced by the respective orientational metric, to calculate probability densities, which results in an entropy bias. I, therefore, generalized the kNN algorithm for the group of orientations $SO(3)^n$ by defining appropriate distance metrics for molecular orientations based on quaternions and calculating their respective induced ball volumes. For fast orientational nearest neighbor searches, I implemented the quaternion distances in the NMSLib. The rotational entropy method was subsequently tested on analytical test distribution and "pinned" water molecules, for which only rotational degrees of freedom remained. Here, I identified a kNN value of $k = 1$ as the best trade-off between accuracy (expected to be better for smaller k) and precision (expected to be better for larger k) within the realm of water entropy calculation. For the "pinned" water test systems, I furthermore obtained values within $9.6\,\%$ from reference entropies. As "pinned" water exhibits significantly stronger correlations than real (mobile) water, I expect the entropy contributions from neglected higher-order correlations to be larger, and therefore the test system to be a harder benchmark than unperturbed water at ambient conditions.

To obtain a full picture of the solvent entropy, I subsequently included translational entropies and the entropy loss due to correlations of translational and rotational degrees of freedom, calculated by the same means, into Per|Mut. To assess the accuracy of the translational entropy algorithm, I used a small argon test system. Here, reference entropies were obtained by performing a thermodynamic integration (TI) from an ideal gas, for which the entropy is known analytically, to an interacting argon gas. Per|Mut yielded accurate entropy values for all tested TI switching-coordinate intermediates within a maximum deviation of $5.7\,\%$ from the reference values.

To furthermore test the accuracy of Per|Mut as a whole, including the translation-rotation

correlation term, I selected hydrated n-alkanes as test systems. From both theory and experiments, it is known that the hydration entropy of alkanes decreases approximately linearly with the alkane length [172–174]. Indeed, Per|Mut identified a linear entropy loss of (23.0 ± 1.1) J·mol^{-1}·K^{-1} per additional carbon atom for alkanes from ethane to decane, which is in quantitative agreement with the reference entropy loss of (21.0 ± 0.7) J·mol^{-1}·K^{-1} per C-atom from TI. Remarkably, the increased correlation between the translational and rotational degrees of freedom was identified as the largest contribution to the entropy loss.

To demonstrate how Per|Mut can characterize solvation entropies with spatial resolution, hydrated octanol was chosen as a further test system. As expected, Per|Mut identified an entropy loss for the water molecules that form hydrogen bonds to the hydroxyl group of octanol, equivalent to an unfavorable free energy difference of 8.55 kJ·mol^{-1} at 300 K relative to bulk water. Since the enthalpy change due to a hydrogen bond is ≈ 20 kJ·mol^{-1} [23], the result indicates that the solvation free energy difference of octanol and octane is enthalpy-dominated.

Close to the hydrophobic tail of octanol, the entropy losses originate from both first- and second-order MI terms. Whereas the single-body entropy loss simply originates from the reduced water mobility close to the solute, the entropy loss due to increased translational and rotational correlations shows that the lack of strong interactions with the apolar octanol tail causes stronger interactions within the remaining water. This effect has been similarly discussed in previous publications [37, 38] and adds a new flavor to the controversial iceberg hypothesis of hydrophobicity [36], as it attributes parts of the entropy loss not to the increased order of "ice-like cathedrals" of solvation shell water but to increased water correlations.

To assess the role of the solvent on the stability of proteins as well as to study how the solvent thermodynamics is affected by individual amino acids, I calculated spatially resolved solvent enthalpies and entropies for the globular protein crambin in its native fold and in a molten-globule-like conformation, which served as a model for the unfolded state.

For both conformations, the spatial distribution of the free energy contributions was mainly affected by charged and polar residues. For these hydrophilic residues, the method identified favorable interaction energies, which dominate the favorable free energy, and a less favorable entropy that partially compensates the enthalpic contribution. Interestingly, the solvation of charged amino acids was associated with favorable local solvent free energy differences of up to 15 kJ·mol^{-1} relative to bulk water, whereas the solvation of the most hydrophobic residues resulted in unfavorable free energy changes by maximally +5 kJ·mol^{-1}. The result provides evidence that the hydrophilic effect is stronger and arguably more impactful to protein stability than the hydrophobic effect, which has also been argued previously [208–210]. However, both effects work in concert to achieve protein stability, which typically results from folding free energy differences of only a few 10 kJ·mol^{-1} [19–22]. As such, both effects are vital for protein stability, and weakening of either of them would likely result in protein unfolding.

To quantify the role of the solvent in this tug of war, I calculated the remaining contributions to the free energy difference between the native fold and the molten-globule-like conformation. The free energy balance revealed an overall difference of $(53 \pm 84)\,\mathrm{kJ\cdot mol^{-1}}$, which falls within the expected range of folding free energy differences [19–22] and corresponds to the enthalpy change due to just a few hydrogen bonds [23]. Protein-solvent interaction energies $((-706 \pm 216)\,\mathrm{kJ\cdot mol^{-1}})$ and the protein entropy $((-673 \pm 42)\,\mathrm{kJ\cdot mol^{-1}})$ were identified as the only destabilizing contributions, whereas the protein-internal interaction energies $((648 \pm 59)\,\mathrm{kJ\cdot mol^{-1}})$ showed the largest stabilizing effect. In comparison, the solvent entropy contributions were revealed to be strongly stabilizing and almost equally as large in magnitude $(\approx 500\,\mathrm{kJ\cdot mol^{-1}})$, which is consistent with the picture of solvent-driven stability, discussed in the context of cold denaturation.

Indeed, the simulation results are consistent with typical cold-denaturation temperatures around $-20\,°\mathrm{C}$ or lower [20, 24, 25]. Remarkably, more than half of the solvent entropy contribution originated from strong induced water correlations for the molten-globule-like conformation.

Taken together, Per|Mut reveals and quantifies a mechanistic picture of solvent entropy contributions to protein stability. Although the discussed results are, of course, specific to the protein crambin and the selected molten-globule-like conformation, I expect similar mechanisms to also hold qualitatively for other proteins, as they are comprised of the same building blocks, identical structural elements, and are subject to similar evolutionary pressures.

Whereas some protein free energy contributions, such as the protein-internal interaction energies, can be straightforwardly calculated from atomistic simulations, the solvent contributions, specifically the solvent entropy, has often remained elusive [20, 34, 35]. To this end, the new method Per|Mut provides spatial resolution of the solvent entropy, as well as a physically interpretable decomposition into single-body entropies and contributions from multi-body correlations. To my knowledge, no other method currently allows for an entirely nonparametric, spatially resolved entropy decomposition while capturing up to three-body correlations. As, e.g., demonstrated by solvent-driven protein stability from induced water correlations, the method provides new insight into the mechanics of solvent thermodynamics.

A further potentially important application of Per|Mut, which so far has not been explored, is ligand binding. Since ligands, e.g., drugs or neurotransmitters, need to displace water molecules from the binding site before binding to a receptor, the water thermodynamics is believed to profoundly affect the binding affinity [211–214]. Commercial programs, such as WaterMap[15], an implementation of the Inhomogeneous Solvation Theory [12, 13] that also underlines the method GIST, are currently used to predict and optimize binding affinities of ligands. However, higher accuracies in computational protein and drug design could potentially be reached with Per|Mut.

A problem that, so far, remains unsolved is the optimal choice of the permutation reduc-

---

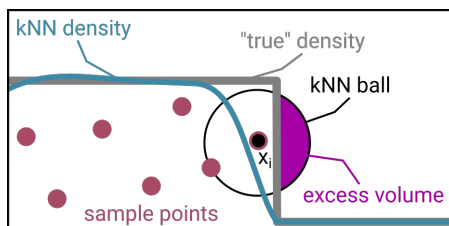[15] https://www.schrodinger.com/watermap

Figure 34.1: Sketch of the configuration space probability density at the boundary of a permutation "slice". The "true" underlying density is shown in grey, the density perceived by the kNN density estimator is shown in blue. Sample points are drawn in red. The kNN ball of sample point $\boldsymbol{x}_i$ is drawn for $k = 1$; its excess volume, i.e., the volume that extends across the permutation "slice" boundary, is shown in purple.

tion reference configuration. The reference configuration determines the shape of the permutational configuration space subregion ("slice") but has no effect on entropy in the limit of infinite sampling. At finite sampling, however, the Per|Mut entropies depend on the chosen reference due to smoothing of the probability density at the boundaries of the permutation "slice". Here, the "true" underlying configuration space probability density vanishes abruptly, as sketched in Fig. 34.1. Close to this "hard edge", the kNN method yields a smoothed density, as the kNN balls extend across the boundary of the permutation "slice" (Fig. 34.1, purple), which results in entropy errors. Attempts to calculate the excess volumes of each kNN ball, either using a Monte Carlo integration scheme or a second part of the MD trajectory, have proven to solve the problem. However, these solutions are computationally inefficient and, therefore, cannot be practically applied.

For small solutes like alkanes or octanol, the entropy error is alleviated by choosing a simple cubic reference conformation, which minimizes the surface area of the permutation "slice" — and therefore also the entropy error. For larger and irregularly shaped solutes like crambin, which themselves exclude parts of the configuration space, a scheme to find the optimal permutation reference configuration is still needed.

A further limitation of Per|Mut is that it neglects correlations between four or more molecules from the entropy estimate. Although a third-order expansion is an advancement compared to existing methods, higher-order correlations could potentially be significant under certain conditions. For instance, "pinned" water molecules exhibit strong orientational correlations (see chapter IV). Although normal water molecules are never strictly immobilized, it seems plausible that higher-order correlations significantly contribute to the entropy of bound first-hydration-shell water molecules. Indeed, data obtained after the defense of this thesis indicates a significant effect of higher-order correlations, particularly close to solvated charges.

In the future, Per|Mut could hence be improved further by including, or at least approximating, contributions from four-body water correlations. In addition, further optimized nearest neighbor search algorithms could significantly speed up MI calculation, which would allow for

increased sampling to be processed within the same amount of time.

In spite of the mentioned limitations, Per|Mut is a significant step forward compared to existing methods and will shed light onto the solvent thermodynamics of, e.g., hydrophobic interaction and ligand binding. My current implementation of Per|Mut is available for download as a python package with a C++ backend for fast nearest neighbor search[16].

35 Membrane Thickening

The solvent contributions to the energetics of membrane fusion constitute a further potentially relevant application of Per|Mut. Here, two membranes are first brought into loose contact through the assembly of SNARE proteins. Subsequently, dehydration of the proximal leaflets occurs as the membranes transition into a tight-docking state, in which the bilayers are closely aligned but retain their integrity.

It is thus expected that the displacement of hydration-shell water molecules upon dehydration contributes to the largely debated free energy landscape of membrane fusion [51, 53–56] in a similar manner as the displacement of water molecules from binding pockets affects ligand binding affinities. A systematic analysis of the solvent's role in membrane fusion is still ongoing work, however, another hydration-related question arose in this context.

As shown by Witkowska and Jahn in chapter VII, the protein-free and calcium-independent tight-docking regime of pre-fusion membranes in close contact represents a productive metastable intermediate in a membrane fusion pathway. Tight docking thus represents a new local minimum of the membrane fusion free energy landscape, a better understanding of which is vital to unravel membrane fusion as a whole. Using cryo-electron microscopy, they furthermore showed that the tight-docking state is characterized by an increased membrane thickness.

To assess the mechanics of the observed membrane thickening, I carried out MD simulations of double-membrane setups at decreasing mutual distances, defined by a given number of water molecules between the bilayers. The simulations reproduced the experimentally observed thickening of lipid bilayers at close distance for three different lipid compositions. In addition, the area per lipid and the membrane volume decreased as the two bilayers were moved closer together. To identify the cause of the thickening, I demonstrated that an enforced headgroup tilt leads to an area-shrinkage, and an enforced area-shrinkage, in turn, leads to thickening. However, this finding alone is not sufficient to conclude that the three events — headgroup tilt, area shrinkage, and thickening — also naturally occur in this sequence. To determine the correct sequence of events, I carried out non-equilibrium simulations. Here, two equilibrated single membranes were instantaneously brought into close contact and simulated until the new double-membrane equilibrium was reached. Following the rationale that the cause should precede the effect, a response-time analysis resulted in the mechanism shown in Fig. 32.4G.

---

[16]https://gitlab.gwdg.de/lheinz/hydration_entropy

The lipid headgroups tilt in-plane upon dehydration, which is likely caused by their electrostatic interactions, as the level of shielding by water molecules decreases (first arrow). A more antiparallel arrangement of the headgroups is electrostatically favorable, which creates a contraction force and therefore leads to a smaller area per lipid (second arrow). As a result, the lipid tails are "squeezed" into a more upright position: the bilayer thickness and the lipid tail order increase (third arrow).

It should, however, be noted that the thickening-effects at close distances strongly depend on the way the mutual membrane distance is controlled. In fact, simulations in mimicked grand canonical conditions (by adding a pore) with harmonic restraints between the two inner layers of headgroups showed no thickness-effect. Harmonic restraints applied between the two bilayer centers of mass even revealed a reversed distance-thickness relationship, as also observed previously by Smirnova et al. [62]. It is therefore intriguing to speculate that the isothermal-isobaric (NPT) ensemble better describes the experimental data than a grand canonical ensemble, because the exchange of water molecules between the inter-membrane water layer and bulk is (partially) blocked.

In summary, the collaboration has characterized a new metastable tight-docking state that is associated with a change in lipid organization, resulting in membrane thickening. Remarkably, neither divalent cations nor proteins are required. Although the role of lipid bilayer thickening in membrane fusion is not yet clear, the heavily tilted headgroups could potentially facilitate hemifusion stalk-formation, which is triggered by a lipid tail "flip" into the adjacent proximal leaflet [54]. Understanding the molecular mechanism that underlie the thickening provides an additional piece in the puzzle of close-contact, pre-fusion membrane physics.

IX | Appendix

36 The Maximum Entropy Distribution for a Fixed Variance

Out of all distributions with given mean μ, variance σ^2 on the support $(-\infty, \infty)$, the Gaussian distribution has the highest entropy. This statement can be easily proven using Lagrange multipliers.

The objective function reads

$$\mathcal{F}[p] = -\int_{-\infty}^{\infty} p(x)\log p(x)dx + \lambda_0 \left(\int_{-\infty}^{\infty} p(x)dx - 1 \right)$$
$$+ \lambda_1 \left(\int_{-\infty}^{\infty} xp(x)dx - \mu \right) + \lambda_2 \left(\int_{-\infty}^{\infty} x^2 p(x)dx - \sigma^2 \right),$$

where the first term is the entropy of the distribution $p(x)$ and the remaining terms constrain the unit-normalization, the mean μ, and the variance σ^2, respectively, using Langrange multipliers λ_0, λ_1, and λ_2.

The derivative with respect to the distribution p reads

$$\frac{\partial \mathcal{F}}{\partial p dx} = -\log p(x) - 1 + \lambda_0 + \lambda_1 x + \lambda_2 x^2.$$

Setting the derivative to zero yields

$$p(x) = e^{\lambda_0 - 1 + \lambda_1 x + \lambda_2 x^2}.$$

From re-substituting the constraints follows the Gaussian distribution

$$p(x) = \frac{1}{\sqrt{2\pi}\sigma} e^{-\frac{x^2}{2\sigma^2}}.$$

The second derivative

$$\frac{\partial^2 \mathcal{F}}{\partial p^2 dx} = -\frac{1}{p(x)} < 0$$

confirms that the Gaussian distribution $p(x)$ is indeed an entropy maximum.

X | References

[1] L. J. Rothschild and R. L. Mancinelli, "Life in extreme environments," *Nature*, vol. 409, no. 6823, pp. 1092–1101, 2001.

[2] M. Chaplin, "Do we underestimate the importance of water in cell biology?" *Nature Reviews Molecular Cell Biology*, vol. 7, no. 11, p. 861, 2006.

[3] D. Laage, T. Elsaesser, and J. T. Hynes, "Water dynamics in the hydration shells of biomolecules," *Chemical Reviews*, vol. 117, no. 16, pp. 10 694–10 725, 2017.

[4] M. Fengler, "Estimating orientational water entropy at protein interfaces," Ph.D. dissertation, Georg-August-Universität Göttingen, 2011.

[5] D. R. Canchi, D. Paschek, and A. E. Garcia, "Equilibrium study of protein denaturation by urea," *Journal of the American Chemical Society*, vol. 132, no. 7, pp. 2338–2344, 2010.

[6] A. H. De Vries, A. E. Mark, and S. J. Marrink, "Molecular dynamics simulation of the spontaneous formation of a small dppc vesicle in water in atomistic detail," *Journal of the American Chemical Society*, vol. 126, no. 14, pp. 4488–4489, 2004.

[7] J. Kleinjung and F. Fraternali, "Design and application of implicit solvent models in biomolecular simulations," *Current opinion in structural biology*, vol. 25, pp. 126–134, 2014.

[8] F. Fraternali and W. Van Gunsteren, "An efficient mean solvation force model for use in molecular dynamics simulations of proteins in aqueous solution," *Journal of molecular biology*, vol. 256, no. 5, pp. 939–948, 1996.

[9] S. Rauscher, V. Gapsys, M. J. Gajda, M. Zweckstetter, B. L. de Groot, and H. Grubmüller, "Structural ensembles of intrinsically disordered proteins depend strongly on force field: a comparison to experiment," *Journal of chemical theory and computation*, vol. 11, no. 11, pp. 5513–5524, 2015.

[10] C. N. Nguyen, T. Kurtzman Young, and M. K. Gilson, "Grid inhomogeneous solvation theory: hydration structure and thermodynamics of the miniature receptor cucurbit [7] uril," *The Journal of chemical physics*, vol. 137, no. 4, p. 044101, 2012.

[11] R. A. Persson, V. Pattni, A. Singh, S. M. Kast, and M. Heyden, "Signatures of solvation thermodynamics in spectra of intermolecular vibrations," *Journal of chemical theory and computation*, vol. 13, no. 9, pp. 4467–4481, 2017.

[12] T. Lazaridis, "Inhomogeneous fluid approach to solvation thermodynamics. 1. theory," *The Journal of Physical Chemistry B*, vol. 102, no. 18, pp. 3531–3541, 1998.

[13] ——, "Inhomogeneous fluid approach to solvation thermodynamics. 2. applications to simple fluids," *The Journal of Physical Chemistry B*, vol. 102, no. 18, pp. 3542–3550, 1998.

[14] J. R. Errington and P. G. Debenedetti, "Relationship between structural order and the anomalies of liquid water," *Nature*, vol. 409, no. 6818, p. 318, 2001.

[15] J. R. Errington, P. G. Debenedetti, and S. Torquato, "Cooperative origin of low-density domains in liquid water," *Physical review letters*, vol. 89, no. 21, p. 215503, 2002.

[16] Z. Yan, S. V. Buldyrev, P. Kumar, N. Giovambattista, P. G. Debenedetti, and H. E. Stanley, "Structure of the first-and second-neighbor shells of simulated water: Quantitative relation to translational and orientational order," *Physical Review E*, vol. 76, no. 5, p. 051201, 2007.

[17] A. Godec, J. C. Smith, and F. Merzel, "Increase of both order and disorder in the first hydration shell with increasing solute polarity," *Physical review letters*, vol. 107, no. 26, p. 267801, 2011.

[18] M. Heyden, "Heterogeneity of water structure and dynamics at the protein-water interface," *The Journal of chemical physics*, vol. 150, no. 9, p. 094701, 2019.

[19] P. Privalov and N. Khechinashvili, "A thermodynamic approach to the problem of stabilization of globular protein structure: a calorimetric study," *Journal of molecular biology*, vol. 86, no. 3, pp. 665–684, 1974.

[20] C. L. Dias, T. Ala-Nissila, J. Wong-ekkabut, I. Vattulainen, M. Grant, and M. Karttunen, "The hydrophobic effect and its role in cold denaturation," *Cryobiology*, vol. 60, no. 1, pp. 91–99, 2010.

[21] E. E. Lattman and G. D. Rose, "Protein folding–what's the question?" *Proceedings of the National Academy of Sciences*, vol. 90, no. 2, pp. 439–441, 1993.

[22] V. R. Agashe and J. B. Udgaonkar, "Thermodynamics of denaturation of barstar: evidence for cold denaturation and evaluation of the interaction with guanidine hydrochloride," *Biochemistry*, vol. 34, no. 10, pp. 3286–3299, 1995.

[23] K. Wendler, J. Thar, S. Zahn, and B. Kirchner, "Estimating the hydrogen bond energy," *The Journal of Physical Chemistry A*, vol. 114, no. 35, pp. 9529–9536, 2010.

[24] D. Sanfelice and P. A. Temussi, "Cold denaturation as a tool to measure protein stability," *Biophysical chemistry*, vol. 208, pp. 4–8, 2016.

[25] S. B. Kim, J. C. Palmer, and P. G. Debenedetti, "Computational investigation of cold denaturation in the trp-cage miniprotein," *Proceedings of the National Academy of Sciences*, vol. 113, no. 32, pp. 8991–8996, 2016.

[26] A. D. Robertson and K. P. Murphy, "Protein structure and the energetics of protein stability," *Chemical reviews*, vol. 97, no. 5, pp. 1251–1268, 1997.

[27] W. Kauzmann, "Some factors in the interpretation of protein denaturation," in *Advances in protein chemistry*. Elsevier, 1959, vol. 14, pp. 1–63.

[28] D. Chandler, "Interfaces and the driving force of hydrophobic assembly," *Nature*, vol. 437, no. 7059, p. 640, 2005.

[29] K. A. Dill, T. M. Truskett, V. Vlachy, and B. Hribar-Lee, "Modeling water, the hydrophobic effect, and ion solvation," *Annu. Rev. Biophys. Biomol. Struct.*, vol. 34, pp. 173–199, 2005.

[30] C. N. Pace, "Contribution of the hydrophobic effect to globular protein stability," *Journal of molecular biology*, vol. 226, no. 1, pp. 29–35, 1992.

[31] J. H. Hildebrand, "Is there a "hydrophobic effect"?" *Proceedings of the National Academy of Sciences of the United States of America*, vol. 76, no. 1, p. 194, 1979.

[32] N. T. Southall, K. A. Dill, and A. Haymet, "A view of the hydrophobic effect," 2002.

[33] A. Ben-Naim, "Hydrophobic interaction and structural changes in the solvent," *Biopolymers: Original Research on Biomolecules*, vol. 14, no. 7, pp. 1337–1355, 1975.

[34] A. Ben-Naim, "Myths and verities in protein folding theories: From frank and evans iceberg-conjecture to explanation of the hydrophobic effect," *The Journal of Chemical Physics*, vol. 139, no. 16, p. 10B626_1, 2013.

[35] G. Hummer, S. Garde, A. Garcıa, and L. Pratt, "New perspectives on hydrophobic effects," *Chemical Physics*, vol. 258, no. 2-3, pp. 349–370, 2000.

[36] H. S. Frank and M. W. Evans, "Free volume and entropy in condensed systems iii. entropy in binary liquid mixtures; partial molal entropy in dilute solutions; structure and thermodynamics in aqueous electrolytes," *The Journal of Chemical Physics*, vol. 13, no. 11, pp. 507–532, 1945.

[37] N. Galamba, "Water's structure around hydrophobic solutes and the iceberg model," *The Journal of Physical Chemistry B*, vol. 117, no. 7, pp. 2153–2159, 2013.

[38] S. Y. Noskov, G. Lamoureux, and B. Roux, "Molecular dynamics study of hydration in ethanol- water mixtures using a polarizable force field," *The Journal of Physical Chemistry B*, vol. 109, no. 14, pp. 6705–6713, 2005.

[39] H.-A. Yu and M. Karplus, "A thermodynamic analysis of solvation," *The Journal of chemical physics*, vol. 89, no. 4, pp. 2366–2379, 1988.

[40] A. Ben-Naim, "Hydrophobic interaction and structural changes in the solvent," *Biopolymers: Original Research on Biomolecules*, vol. 14, no. 7, pp. 1337–1355, 1975.

[41] B. Lee, "Enthalpy-entropy compensation in the thermodynamics of hydrophobicity," *Biophysical chemistry*, vol. 51, no. 2-3, pp. 271–278, 1994.

[42] J. E. Jones, "On the determination of molecular fields.—ii. from the equation of state of a gas," *Proc. R. Soc. Lond. A*, vol. 106, no. 738, pp. 463–477, 1924.

[43] C. N. Nguyen, T. K. Young, and M. K. Gilson, "Erratum: "grid inhomogeneous solvation theory: Hydration structure and thermodynamics of the miniature receptor cucurbit [7] uril" [j. chem. phys. 137, 044101 (2012)]," *The Journal of chemical physics*, vol. 137, no. 14, p. 149901, 2012.

[44] J. W. Gibbs, "On the equilibrium of heterogeneous substances," 1879.

[45] S.-T. Lin, P. K. Maiti, and W. A. Goddard III, "Two-phase thermodynamic model for efficient and accurate absolute entropy of water from molecular dynamics simulations," *The Journal of Physical Chemistry B*, vol. 114, no. 24, pp. 8191–8198, 2010.

[46] F. Reinhard and H. Grubmüller, "Estimation of absolute solvent and solvation shell entropies via permutation reduction," *The Journal of chemical physics*, vol. 126, no. 1, p. 014102, 2007.

[47] F. Reinhard, O. F. Lange, J. S. Hub, J. Haas, and H. Grubmüller, "g_permute: Permutation-reduced phase space density compaction," *Computer Physics Communications*, vol. 180, no. 3, pp. 455–458, 2009.

[48] H. Matsuda, "Physical nature of higher-order mutual information: Intrinsic correlations and frustration," *Physical Review E*, vol. 62, no. 3, p. 3096, 2000.

[49] V. Hnizdo, E. Darian, A. Fedorowicz, E. Demchuk, S. Li, and H. Singh, "Nearest-neighbor nonparametric method for estimating the configurational entropy of complex molecules," *Journal of computational chemistry*, vol. 28, no. 3, pp. 655–668, 2007.

[50] V. Hnizdo, J. Tan, B. J. Killian, and M. K. Gilson, "Efficient calculation of configurational entropy from molecular simulations by combining the mutual-information expansion and nearest-neighbor methods," *Journal of computational chemistry*, vol. 29, no. 10, pp. 1605–1614, 2008.

[51] R. Jahn, T. Lang, and T. C. Südhof, "Membrane fusion," *Cell*, vol. 112, no. 4, pp. 519–533, 2003.

[52] S. Martens and H. T. McMahon, "Mechanisms of membrane fusion: disparate players and common principles," *Nature reviews Molecular cell biology*, vol. 9, no. 7, pp. 543–556, 2008.

[53] R. Jahn and R. H. Scheller, "Snares—engines for membrane fusion," *Nature reviews Molecular cell biology*, vol. 7, no. 9, pp. 631–643, 2006.

[54] H. J. Risselada, C. Kutzner, and H. Grubmüller, "Caught in the act: visualization of snare-mediated fusion events in molecular detail," *ChemBioChem*, vol. 12, no. 7, pp. 1049–1055, 2011.

[55] J. M. Hernández and B. Podbilewicz, "The hallmarks of cell-cell fusion," *Development*, vol. 144, no. 24, pp. 4481–4495, 2017.

[56] S. L. Leikin, M. M. Kozlov, L. V. Chernomordik, V. S. Markin, and Y. A. Chizmadzhev, "Membrane fusion: overcoming of the hydration barrier and local restructuring," *Journal of theoretical biology*, vol. 129, no. 4, pp. 411–425, 1987.

[57] H. Yavuz, I. Kattan, J. M. Hernandez, O. Hofnagel, A. Witkowska, S. Raunser, P. J. Walla, and R. Jahn, "Arrest of trans-snare zippering uncovers loosely and tightly docked intermediates in membrane fusion," *Journal of Biological Chemistry*, vol. 293, no. 22, pp. 8645–8655, 2018.

[58] M. Kanduc, E. Schneck, and R. R. Netz, "Hydration interaction between phospholipid membranes: insight into different measurement ensembles from atomistic molecular dynamics simulations," *Langmuir*, vol. 29, no. 29, pp. 9126–9137, 2013.

[59] O. G. Mouritsen, *Life-as a matter of fat.* Springer, 2005.

[60] F. Chen and W. Hung, "Structural changes of lipid membrane induced by dehydration," *Chinese Journal of Physics*, vol. 34, no. 6, pp. 1363–1372, 1996.

[61] S. Aeffner, T. Reusch, B. Weinhausen, and T. Salditt, "Energetics of stalk intermediates in membrane fusion are controlled by lipid composition," *Proceedings of the National Academy of Sciences*, vol. 109, no. 25, pp. E1609–E1618, 2012.

[62] Y. G. Smirnova, S. Aeffner, H. J. Risselada, T. Salditt, S.-J. Marrink, M. Müller, and V. Knecht, "Interbilayer repulsion forces between tension-free lipid bilayers from simulation," *Soft Matter*, vol. 9, no. 45, pp. 10 705–10 718, 2013.

[63] V. A. Parsegian, N. Fuller, and R. P. Rand, "Measured work of deformation and repulsion of lecithin bilayers," *Proceedings of the National Academy of Sciences*, vol. 76, no. 6, pp. 2750–2754, 1979.

[64] E. Schrödinger, "An undulatory theory of the mechanics of atoms and molecules," *Physical review*, vol. 28, no. 6, p. 1049, 1926.

[65] D. Hartree, "The wave mechanics of an atom with a non-coulomb central field. part i. theory and methods," in *Mathematical Proceedings of the Cambridge Philosophical Society*, vol. 24, no. 1. Cambridge University Press, 1928, pp. 89–110.

[66] J. C. Slater, "The self consistent field and the structure of atoms," *Physical Review*, vol. 32, no. 3, p. 339, 1928.

[67] V. Fock, "Näherungsmethode zur lösung des quantenmechanischen mehrkörperproblems," *Zeitschrift für Physik*, vol. 61, no. 1-2, pp. 126–148, 1930.

[68] R. G. Parr, "Density functional theory of atoms and molecules," in *Horizons of quantum chemistry*. Springer, 1980, pp. 5–15.

[69] G. E. Moore *et al.*, "Cramming more components onto integrated circuits," 1965.

[70] M. Born and R. Oppenheimer, "Zur quantentheorie der molekeln," *Annalen der physik*, vol. 389, no. 20, pp. 457–484, 1927.

[71] P. Ehrenfest, "Bemerkung über die angenäherte gültigkeit der klassischen mechanik innerhalb der quantenmechanik," *Zeitschrift für Physik*, vol. 45, no. 7-8, pp. 455–457, 1927.

[72] R. D. Skeel, "What makes molecular dynamics work?" *SIAM Journal on Scientific Computing*, vol. 31, no. 2, pp. 1363–1378, 2009.

[73] H. J. Berendsen, D. van der Spoel, and R. van Drunen, "Gromacs: a message-passing parallel molecular dynamics implementation," *Computer Physics Communications*, vol. 91, no. 1-3, pp. 43–56, 1995.

[74] D. Van Der Spoel, E. Lindahl, B. Hess, G. Groenhof, A. E. Mark, and H. J. Berendsen, "Gromacs: fast, flexible, and free," *Journal of computational chemistry*, vol. 26, no. 16, pp. 1701–1718, 2005.

[75] B. Hess, C. Kutzner, D. Van Der Spoel, and E. Lindahl, "Gromacs 4: algorithms for highly efficient, load-balanced, and scalable molecular simulation," *Journal of chemical theory and computation*, vol. 4, no. 3, pp. 435–447, 2008.

[76] S. Pronk, S. Páll, R. Schulz, P. Larsson, P. Bjelkmar, R. Apostolov, M. R. Shirts, J. C. Smith, P. M. Kasson, D. van der Spoel, B. Hess, and E. Lindahl, "Gromacs 4.5: a high-throughput and highly parallel open source molecular simulation toolkit," *Bioinformatics*, pp. 845–854, 2013.

[77] S. Pall, M. J. Abraham, C. Kutzner, B. Hess, and E. Lindahl, "Tackling exascale software challenges in molecular dynamics simulations with gromacs," in *International Conference on Exascale Applications and Software*. Springer, 2014, pp. 3–27.

[78] J. W. Ponder and D. A. Case, "Force fields for protein simulations," in *Advances in protein chemistry*. Elsevier, 2003, vol. 66, pp. 27–85.

[79] J. A. Maier, C. Martinez, K. Kasavajhala, L. Wickstrom, K. E. Hauser, and C. Simmerling, "ff14sb: improving the accuracy of protein side chain and backbone parameters from ff99sb," *Journal of chemical theory and computation*, vol. 11, no. 8, pp. 3696–3713, 2015.

[80] B. R. Brooks, R. E. Bruccoleri, B. D. Olafson, D. J. States, S. Swaminathan, and M. Karplus, "Charmm: a program for macromolecular energy, minimization, and dynamics calculations," *Journal of computational chemistry*, vol. 4, no. 2, pp. 187–217, 1983.

[81] A. D. MacKerell Jr, D. Bashford, M. Bellott, R. L. Dunbrack Jr, J. D. Evanseck, M. J. Field, S. Fischer, J. Gao, H. Guo, S. Ha, D. Joseph-McCarthy, L. Kuchnir, K. Kuczera, F. T. K. Lau, C. Mattos, S. Michnick, T. Ngo, D. T. Nguyen, B. Prodhom, W. E. Reiher III, B. Roux, M. Schlenkrich, J. C. Smith, R. Stote, J. Straub, M. Watanabe, J. Wiórkiewicz-Kuczera, D. Yin, and M. Karplus, "All-atom empirical potential for molecular modeling and dynamics studies of proteins," *The journal of physical chemistry B*, vol. 102, no. 18, pp. 3586–3616, 1998.

[82] B. R. Brooks, C. L. Brooks, A. D. MacKerell, L. Nilsson, R. J. Petrella, B. Roux, Y. Won, G. Archontis, C. Bartels, S. Boresch, A. Caflisch, L. S. D. Caves, Q. Cui, A. R. Dinner, M. Feig, S. Fischer, J. Gao, M. Hodoš ček, W. Im, K. Kuczera, T. Lazaridis, M. Jun, V. Ovchinnikov, E. Paci, R. W. Pastor, C. B. Post, J. Z. Pu, M. Schaefer, B. Tidor, R. M. Venable, H. L. Woodcock, X. Wu, W. Yang, D. M. York, and M. Karplus, "Charmm: the biomolecular simulation program," *Journal of computational chemistry*, vol. 30, no. 10, pp. 1545–1614, 2009.

[83] J. Huang and A. D. MacKerell, "Charmm36 all-atom additive protein force field: Validation based on comparison to nmr data," *Journal of computational chemistry*, vol. 34, no. 25, pp. 2135–2145, 2013.

[84] J. Huang, S. Rauscher, G. Nawrocki, T. Ran, M. Feig, B. L. de Groot, H. Grubmüller, and A. D. MacKerell Jr, "Charmm36m: an improved force field for folded and intrinsically disordered proteins," *nature methods*, vol. 14, no. 1, p. 71, 2017.

[85] P. M. Morse, "Diatomic molecules according to the wave mechanics. ii. vibrational levels," *Physical review*, vol. 34, no. 1, p. 57, 1929.

[86] GROMACS development team, "Gromacs documentation, release 2020.3," 2020. [Online]. Available: http://manual.gromacs.org/current/manual-2020.3.pdf

[87] W. Pauli, "Über den zusammenhang des abschlusses der elektronengruppen im atom mit der komplexstruktur der spektren," *Zeitschrift für Physik A Hadrons and Nuclei*, vol. 31, no. 1, pp. 765–783, 1925.

[88] R. A. Buckingham, "The classical equation of state of gaseous helium, neon and argon," in *Proceedings of the Royal Society of London A: Mathematical, Physical and Engineering Sciences*, vol. 168, no. 933. The Royal Society, 1938, pp. 264–283.

[89] G. Lamoureux and B. Roux, "Modeling induced polarization with classical drude oscillators: Theory and molecular dynamics simulation algorithm," *The Journal of chemical physics*, vol. 119, no. 6, pp. 3025–3039, 2003.

[90] G. Lamoureux, A. D. MacKerell Jr, and B. Roux, "A simple polarizable model of water based on classical drude oscillators," *The Journal of chemical physics*, vol. 119, no. 10, pp. 5185–5197, 2003.

[91] D. C. Rapaport, *The art of molecular dynamics simulation.* Cambridge university press, 2004.

[92] R. W. Hockney, S. Goel, and J. Eastwood, "Quiet high-resolution computer models of a plasma," *Journal of Computational Physics*, vol. 14, no. 2, pp. 148–158, 1974.

[93] H. J. Berendsen and W. F. Van Gunsteren, "Practical algorithms for dynamic simulations," *Molecular-dynamics simulation of statistical-mechanical systems*, pp. 43–65, 1986.

[94] M. Christen, C. D. Christ, and W. F. van Gunsteren, "Free energy calculations using flexible-constrained, hard-constrained and non-constrained molecular dynamics simulations," *ChemPhysChem*, vol. 8, no. 10, pp. 1557–1564, 2007.

References

[95] S. Miyamoto and P. A. Kollman, "Settle: An analytical version of the shake and rattle algorithm for rigid water models," *Journal of computational chemistry*, vol. 13, no. 8, pp. 952–962, 1992.

[96] B. Hess, H. Bekker, H. J. Berendsen, and J. G. Fraaije, "Lincs: a linear constraint solver for molecular simulations," *Journal of computational chemistry*, vol. 18, no. 12, pp. 1463–1472, 1997.

[97] L. P. Heinz, "Molecular dynamics simulation of the ryanodine receptor 1," 2017.

[98] G. Bussi, D. Donadio, and M. Parrinello, "Canonical sampling through velocity rescaling," *The Journal of chemical physics*, vol. 126, no. 1, p. 014101, 2007.

[99] H. J. Berendsen, J. v. Postma, W. F. van Gunsteren, A. DiNola, and J. R. Haak, "Molecular dynamics with coupling to an external bath," *The Journal of chemical physics*, vol. 81, no. 8, pp. 3684–3690, 1984.

[100] H. C. Andersen, "Molecular dynamics simulations at constant pressure and/or temperature," *The Journal of chemical physics*, vol. 72, no. 4, pp. 2384–2393, 1980.

[101] M. Parrinello and A. Rahman, "Polymorphic transitions in single crystals: A new molecular dynamics method," *Journal of Applied physics*, vol. 52, no. 12, pp. 7182–7190, 1981.

[102] ——, "Polymorphic transitions in single crystals: A new molecular dynamics method," *Journal of Applied physics*, vol. 52, no. 12, pp. 7182–7190, 1981.

[103] S. Nosé and M. Klein, "Constant pressure molecular dynamics for molecular systems," *Molecular Physics*, vol. 50, no. 5, pp. 1055–1076, 1983.

[104] R. Martoňák, A. Laio, and M. Parrinello, "Predicting crystal structures: The parrinello-rahman method revisited," *Physical review letters*, vol. 90, no. 7, p. 075503, 2003.

[105] J. L. Bentley, "Multidimensional binary search trees used for associative searching," *Communications of the ACM*, vol. 18, no. 9, pp. 509–517, 1975.

[106] P. N. Yianilos, "Data structures and algorithms for nearest neighbor search in general metric spaces," in *Soda*, vol. 93, no. 194, 1993, pp. 311–21.

[107] T. Bozkaya and M. Ozsoyoglu, "Indexing large metric spaces for similarity search queries," *ACM Transactions on Database Systems (TODS)*, vol. 24, no. 3, pp. 361–404, 1999.

[108] S. Brin, "Near neighbor search in large metric spaces," 1995.

References

[109] L. Boytsov and B. Naidan, "Engineering efficient and effective non-metric space library," in *International Conference on Similarity Search and Applications*. Springer, 2013, pp. 280–293.

[110] B. Naidan, L. Boytsov, Y. Malkov, and D. Novak, "Non-metric space library (nmslib) manual," 2019. [Online]. Available: https://raw.githubusercontent.com/nmslib/nmslib/master/manual/latex/manual.pdf

[111] F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, J. Vanderplas, A. Passos, D. Cournapeau, M. Brucher, M. Perrot, and Édouard Duchesnay, "Scikit-learn: Machine learning in python," *Journal of Machine Learning Research*, vol. 12, no. 85, pp. 2825–2830, 2011. [Online]. Available: http://jmlr.org/papers/v12/pedregosa11a.html

[112] L. Heinz and H. Grubmüller, "Computing spatially resolved rotational hydration entropies from atomistic simulations," *Journal of Chemical Theory and Computation*, vol. 16, no. 1, pp. 108–118, 2019.

[113] M. Karplus and J. N. Kushick, "Method for estimating the configurational entropy of macromolecules," *Macromolecules*, vol. 14, no. 2, pp. 325–332, 1981.

[114] C.-E. Chang, W. Chen, and M. K. Gilson, "Evaluating the accuracy of the quasiharmonic approximation," *Journal of Chemical Theory and Computation*, vol. 1, no. 5, pp. 1017–1028, 2005.

[115] J. Schlitter, "Estimation of absolute and relative entropies of macromolecules using the covariance matrix," *Chemical physics letters*, vol. 215, no. 6, pp. 617–621, 1993.

[116] J. G. Kirkwood, "Statistical mechanics of fluid mixtures," *The Journal of Chemical Physics*, vol. 3, no. 5, pp. 300–313, 1935.

[117] S. Spicher and S. Grimme, "Efficient computation of free energy contributions for association reactions of large molecules," *The Journal of Physical Chemistry Letters*, 2020.

[118] N. T. Mai, N. T. Lan, T. Y. Vu, P. T. M. Duong, N. T. Tung, and H. T. T. Phung, "Estimation of the ligand-binding free energy of checkpoint kinase 1 via non-equilibrium md simulations," *Journal of Molecular Graphics and Modelling*, vol. 100, p. 107648, 2020.

[119] M. Reinhardt and H. Grubmüller, "Determining free energy differences through variationally-derived intermediates," *Journal of Chemical Theory and Computation*, 2020.

[120] C. Peter, C. Oostenbrink, A. van Dorp, and W. F. van Gunsteren, "Estimating entropies from molecular dynamics simulations," *The Journal of chemical physics*, vol. 120, no. 6, pp. 2652–2661, 2004.

[121] D. C. Wallace, "On the role of density fluctuations in the entropy of a fluid," *The Journal of chemical physics*, vol. 87, no. 4, pp. 2282–2284, 1987.

[122] A. Baranyai and D. J. Evans, "Direct entropy calculation from computer simulation of liquids," *Physical Review A*, vol. 40, no. 7, p. 3817, 1989.

[123] C. N. Nguyen, T. Kurtzman, and M. K. Gilson, "Spatial decomposition of translational water–water correlation entropy in binding pockets," *Journal of chemical theory and computation*, vol. 12, no. 1, pp. 414–429, 2016.

[124] S.-T. Lin, M. Blanco, and W. A. Goddard III, "The two-phase model for calculating thermodynamic properties of liquids from molecular dynamics: Validation for the phase diagram of lennard-jones fluids," *The Journal of chemical physics*, vol. 119, no. 22, pp. 11 792–11 805, 2003.

[125] N. Wiener *et al.*, "Generalized harmonic analysis," *Acta mathematica*, vol. 55, pp. 117–258, 1930.

[126] A. Khintchine, "Korrelationstheorie der stationären stochastischen prozesse," *Mathematische Annalen*, vol. 109, no. 1, pp. 604–615, 1934.

[127] F. Reinhard, "Computation of solvent entropies from molecular synamics simulations," Diploma thesis, Georg-August-Universität Göttingen, 2005.

[128] B. J. Berne, J. D. Weeks, and R. Zhou, "Dewetting and hydrophobic interaction in physical and biological systems," *Annual review of physical chemistry*, vol. 60, 2009.

[129] J. Grdadolnik, F. Merzel, and F. Avbelj, "Origin of hydrophobicity and enhanced water hydrogen bond strength near purely hydrophobic solutes," *Proceedings of the National Academy of Sciences*, vol. 114, no. 2, pp. 322–327, 2017.

[130] Y.-K. Cheng and P. J. Rossky, "Surface topography dependence of biomolecular hydrophobic hydration," *Nature*, vol. 392, no. 6677, p. 696, 1998.

[131] M. Tarek and D. J. Tobias, "The dynamics of protein hydration water: a quantitative comparison of molecular dynamics simulations and neutron-scattering experiments," *Biophysical journal*, vol. 79, no. 6, pp. 3244–3257, 2000.

[132] Y. Sugita and Y. Okamoto, "Replica-exchange molecular dynamics method for protein folding," *Chemical physics letters*, vol. 314, no. 1-2, pp. 141–151, 1999.

[133] C. N. Nguyen, T. Kurtzman, and M. K. Gilson, "Spatial decomposition of translational water–water correlation entropy in binding pockets," *Journal of chemical theory and computation*, vol. 12, no. 1, pp. 414–429, 2015.

[134] S.-T. Lin, M. Blanco, and W. A. Goddard III, "The two-phase model for calculating thermodynamic properties of liquids from molecular dynamics: Validation for the phase diagram of lennard-jones fluids," *The Journal of chemical physics*, vol. 119, no. 22, pp. 11 792–11 805, 2003.

[135] S.-T. Lin, P. K. Maiti, and W. A. Goddard III, "Two-phase thermodynamic model for efficient and accurate absolute entropy of water from molecular dynamics simulations," *The Journal of Physical Chemistry B*, vol. 114, no. 24, pp. 8191–8198, 2010.

[136] G. Gerogiokas, G. Calabro, R. H. Henchman, M. W. Southey, R. J. Law, and J. Michel, "Prediction of small molecule hydration thermodynamics with grid cell theory," *Journal of chemical theory and computation*, vol. 10, no. 1, pp. 35–48, 2013.

[137] L. Pauling, "The structure and entropy of ice and of other crystals with some randomness of atomic arrangement," *Journal of the American Chemical Society*, vol. 57, no. 12, pp. 2680–2684, 1935.

[138] R. H. Henchman and S. J. Irudayam, "Topological hydrogen-bond definition to characterize the structure and dynamics of liquid water," *The Journal of Physical Chemistry B*, vol. 114, no. 50, pp. 16 792–16 810, 2010.

[139] L. Kozachenko and N. N. Leonenko, "Sample estimate of the entropy of a random vector," *Problemy Peredachi Informatsii*, vol. 23, no. 2, pp. 9–16, 1987.

[140] H. Singh, N. Misra, V. Hnizdo, A. Fedorowicz, and E. Demchuk, "Nearest neighbor estimates of entropy," *American journal of mathematical and management sciences*, vol. 23, no. 3-4, pp. 301–321, 2003.

[141] A. Kraskov, H. Stögbauer, and P. Grassberger, "Estimating mutual information," *Physical review E*, vol. 69, no. 6, p. 066138, 2004.

[142] T. Lazaridis and M. Karplus, "Orientational correlations and entropy in liquid water," *The Journal of chemical physics*, vol. 105, no. 10, pp. 4294–4316, 1996.

[143] B. J. Killian, J. Yundenfreund Kravitz, and M. K. Gilson, "Extraction of configurational entropy from molecular simulations via an expansion approximation," *The Journal of chemical physics*, vol. 127, no. 2, p. 024107, 2007.

[144] M. Goethe, I. Fita, and J. M. Rubi, "Testing the mutual information expansion of entropy with multivariate gaussian distributions," *The Journal of chemical physics*, vol. 147, no. 22, p. 224102, 2017.

References

[145] A. B. Tsybakov and E. Van der Meulen, "Root-n consistent estimators of entropy for densities with unbounded support," *Scandinavian Journal of Statistics*, pp. 75–83, 1996.

[146] D. Evans, "A computationally efficient estimator for mutual information," in *Proceedings of the Royal Society of London A: Mathematical, Physical and Engineering Sciences*, vol. 464, no. 2093.   The Royal Society, 2008, pp. 1203–1215.

[147] U. Hensen, O. F. Lange, and H. Grubmüller, "Estimating absolute configurational entropies of macromolecules: the minimally coupled subspace approach," *PloS one*, vol. 5, no. 2, p. e9179, 2010.

[148] H. Hopf, "Über die abbildungen der dreidimensionalen sphäre auf die kugelfläche," in *Selecta Heinz Hopf.*   Springer, 1964, pp. 38–63.

[149] A. Yershova, S. Jain, S. M. Lavalle, and J. C. Mitchell, "Generating uniform incremental grids on so (3) using the hopf fibration," *The International journal of robotics research*, vol. 29, no. 7, pp. 801–812, 2010.

[150] C. F. Karney, "Quaternions in molecular modeling," *Journal of Molecular Graphics and Modelling*, vol. 25, no. 5, pp. 595–604, 2007.

[151] D. Q. Huynh, "Metrics for 3d rotations: Comparison and analysis," *Journal of Mathematical Imaging and Vision*, vol. 35, no. 2, pp. 155–164, 2009.

[152] D. J. Huggins, "Comparing distance metrics for rotation using the k-nearest neighbors algorithm for entropy estimation," *Journal of computational chemistry*, vol. 35, no. 5, pp. 377–385, 2014.

[153] S. Singh and B. Póczos, "Analysis of k-nearest neighbor distances with application to entropy estimation," *arXiv preprint arXiv:1603.08578*, 2016.

[154] B. Ravani and B. Roth, "Motion synthesis using kinematic mappings," *Journal of mechanisms, Transmissions, and Automation in Design*, vol. 105, no. 3, pp. 460–467, 1983.

[155] W. R. Inc., "Mathematica, Version 10.0," Champaign, IL, 2014.

[156] J. K. Uhlmann, "Satisfying general proximity/similarity queries with metric trees," *Information processing letters*, vol. 40, no. 4, pp. 175–179, 1991.

[157] P. N. Yianilos, "Data structures and algorithms for nearest neighbor search in general metric spaces," in *SODA*, vol. 93, no. 194, 1993, pp. 311–321.

[158] A. Amadei, G. Chillemi, M. Ceruso, A. Grottesi, and A. Di Nola, "Molecular dynamics simulations with constrained roto-translational motions: theoretical basis and statistical mechanical consistency," *The Journal of Chemical Physics*, vol. 112, no. 1, pp. 9–23, 2000.

References

[159] W. L. Jorgensen, J. Chandrasekhar, J. D. Madura, R. W. Impey, and M. L. Klein, "Comparison of simple potential functions for simulating liquid water," *The Journal of chemical physics*, vol. 79, no. 2, pp. 926–935, 1983.

[160] T. Darden, D. York, and L. Pedersen, "Particle mesh ewald: An n log (n) method for ewald sums in large systems," *The Journal of chemical physics*, vol. 98, no. 12, pp. 10 089–10 092, 1993.

[161] T. C. Beutler, A. E. Mark, R. C. van Schaik, P. R. Gerber, and W. F. Van Gunsteren, "Avoiding singularities and numerical instabilities in free energy calculations based on molecular simulations," *Chemical physics letters*, vol. 222, no. 6, pp. 529–539, 1994.

[162] U. Hensen, H. Grubmüller, and O. F. Lange, "Adaptive anisotropic kernels for nonparametric estimation of absolute configurational entropies in high-dimensional configuration spaces," *Physical Review E*, vol. 80, no. 1, p. 011913, 2009.

[163] S. Gao, G. Ver Steeg, and A. Galstyan, "Efficient estimation of mutual information for strongly dependent variables," in *Artificial Intelligence and Statistics*, 2015, pp. 277–286.

[164] S. Khan, S. Bandyopadhyay, A. R. Ganguly, S. Saigal, D. J. Erickson III, V. Protopopescu, and G. Ostrouchov, "Relative performance of mutual information estimation methods for quantifying the dependence among short and noisy data," *Physical Review E*, vol. 76, no. 2, p. 026209, 2007.

[165] R. J. Gowers, M. Linke, J. Barnoud, T. J. E. Reddy, M. N. Melo, S. L. Seyler, J. Domanski, D. L. Dotson, S. Buchoux, I. M. Kenney *et al.*, "Mdanalysis: a python package for the rapid analysis of molecular dynamics simulations," Los Alamos National Lab.(LANL), Los Alamos, NM (United States), Tech. Rep., 2019.

[166] N. Michaud-Agrawal, E. J. Denning, T. B. Woolf, and O. Beckstein, "Mdanalysis: a toolkit for the analysis of molecular dynamics simulations," *Journal of computational chemistry*, vol. 32, no. 10, pp. 2319–2327, 2011.

[167] J. N. Israelachvili, D. J. Mitchell, and B. W. Ninham, "Theory of self-assembly of hydrocarbon amphiphiles into micelles and bilayers," *Journal of the Chemical Society, Faraday Transactions 2: Molecular and Chemical Physics*, vol. 72, pp. 1525–1568, 1976.

[168] L. Maibaum, A. R. Dinner, and D. Chandler, "Micelle formation and the hydrophobic effect," *The Journal of Physical Chemistry B*, vol. 108, no. 21, pp. 6778–6781, 2004.

[169] C. Tanford, "The hydrophobic effect and the organization of living matter," *Science*, vol. 200, no. 4345, pp. 1012–1018, 1978.

[170] T. Lazaridis, "Hydrophobic effect," *eLS*, 2013.

[171] T. T. Pham and M. R. Shirts, "Identifying low variance pathways for free energy calculations of molecular transformations in solution phase," *The Journal of chemical physics*, vol. 135, no. 3, p. 034114, 2011.

[172] S. Dec and S. Gill, "Heats of solution of gaseous hydrocarbons in water at 25°C," *Journal of solution chemistry*, vol. 13, no. 1, pp. 27–41, 1984.

[173] W. L. Jorgensen, J. Gao, and C. Ravimohan, "Monte carlo simulations of alkanes in water: Hydration numbers and the hydrophobic effect," *The Journal of Physical Chemistry*, vol. 89, no. 16, pp. 3470–3473, 1985.

[174] E. Gallicchio, M. Kubo, and R. M. Levy, "Enthalpy- entropy and cavity decomposition of alkane hydration free energies: numerical results and implications for theories of hydrophobic solvation," *The Journal of Physical Chemistry B*, vol. 104, no. 26, pp. 6271–6285, 2000.

[175] L. Heinz and H. Grubmüller, "Per|mut: Spatially resolved hydration entropies from atomistic simulations," submitted to the *Journal of Chemical Theory and Computation*, 2020.

[176] M. M. Teeter, S. M. Roe, and N. H. Heo, "Atomic resolution (0· 83 å) crystal structure of the hydrophobic protein crambin at 130 k," *Journal of molecular biology*, vol. 230, no. 1, pp. 292–311, 1993.

[177] W. Humphrey, A. Dalke, and K. Schulten, "VMD: visual molecular dynamics," *Journal of molecular graphics*, vol. 14, no. 1, pp. 33–38, 1996.

[178] G. Zhao and E. London, "An amino acid "transmembrane tendency" scale that approaches the theoretical limit to accuracy for prediction of transmembrane helices: relationship to biological hydrophobicity," *Protein science*, vol. 15, no. 8, pp. 1987–2001, 2006.

[179] K. Tielrooij, S. Van Der Post, J. Hunger, M. Bonn, and H. Bakker, "Anisotropic water reorientation around ions," *The Journal of Physical Chemistry B*, vol. 115, no. 43, pp. 12 638–12 647, 2011.

[180] V. P. Denisov, K. Venu, J. Peters, H. D. Hörlein, and B. Halle, "Orientational disorder and entropy of water in protein cavities," *The Journal of Physical Chemistry B*, vol. 101, no. 45, pp. 9380–9389, 1997.

[181] D. J. Huggins, "Quantifying the entropy of binding for water molecules in protein cavities by computing correlations," *Biophysical journal*, vol. 108, no. 4, pp. 928–936, 2015.

[182] M. von Domaros, D. Bratko, B. Kirchner, G. Hummer, and A. Luzar, "Multifaceted water dynamics in spherical nanocages," *The Journal of Physical Chemistry C*, vol. 123, no. 10, pp. 5989–5998, 2019.

[183] Y. M. Rhee and V. S. Pande, "Multiplexed-replica exchange molecular dynamics method for protein folding simulation," *Biophysical journal*, vol. 84, no. 2, pp. 775–786, 2003.

[184] A. E. García and K. Y. Sanbonmatsu, "Exploring the energy landscape of a β hairpin in explicit solvent," *Proteins: Structure, Function, and Bioinformatics*, vol. 42, no. 3, pp. 345–354, 2001.

[185] K. Y. Sanbonmatsu and A. E. García, "Structure of met-enkephalin in explicit aqueous solution using replica exchange molecular dynamics," *Proteins: Structure, Function, and Bioinformatics*, vol. 46, no. 2, pp. 225–234, 2002.

[186] F. Noé, I. Horenko, C. Schütte, and J. C. Smith, "Hierarchical analysis of conformational dynamics in biomolecules: Transition networks of metastable states," *The Journal of chemical physics*, vol. 126, no. 15, p. 04B617, 2007.

[187] F. Noé, C. Schütte, E. Vanden-Eijnden, L. Reich, and T. R. Weikl, "Constructing the equilibrium ensemble of folding pathways from short off-equilibrium simulations," *Proceedings of the National Academy of Sciences*, vol. 106, no. 45, pp. 19 011–19 016, 2009.

[188] M. M. Kozlov and L. V. Chernomordik, "Membrane tension and membrane fusion," *Current opinion in structural biology*, vol. 33, pp. 61–67, 2015.

[189] R. Jahn and D. Fasshauer, "Molecular machines governing exocytosis of synaptic vesicles," *Nature*, vol. 490, no. 7419, pp. 201–207, 2012.

[190] J. Rizo and J. Xu, "The synaptic vesicle release machinery," *Annual review of biophysics*, vol. 44, pp. 339–367, 2015.

[191] S. Chang, T. Trimbuch, and C. Rosenmund, "Synaptotagmin-1 drives synchronous ca 2+-triggered fusion by c 2 b-domain-mediated synaptic-vesicle-membrane attachment," *Nature neuroscience*, vol. 21, no. 1, pp. 33–40, 2018.

[192] J. M. Hernandez, A. Stein, E. Behrmann, D. Riedel, A. Cypionka, Z. Farsi, P. J. Walla, S. Raunser, and R. Jahn, "Membrane fusion intermediates via directional and full assembly of the snare complex," *Science*, vol. 336, no. 6088, pp. 1581–1584, 2012.

[193] C. Imig, S.-W. Min, S. Krinner, M. Arancillo, C. Rosenmund, T. C. Südhof, J. Rhee, N. Brose, and B. H. Cooper, "The morphological and molecular nature of synaptic vesicle priming at presynaptic active zones," *Neuron*, vol. 84, no. 2, pp. 416–431, 2014.

[194] T. Brandt, L. Cavellini, W. Kühlbrandt, and M. M. Cohen, "A mitofusin-dependent docking ring complex triggers mitochondrial fusion in vitro," *Elife*, vol. 5, p. e14618, 2016.

[195] L. Gui, J. L. Ebner, A. Mileant, J. A. Williams, and K. K. Lee, "Visualization and sequencing of membrane remodeling leading to influenza virus fusion," *Journal of virology*, vol. 90, no. 15, pp. 6948–6962, 2016.

[196] B. R. Shrestha and X. Banquy, "Hydration forces at solid and fluid biointerfaces," *Biointerphases*, vol. 11, no. 1, p. 018907, 2016.

[197] A. V. Pobbati, A. Stein, and D. Fasshauer, "N-to c-terminal snare complex assembly promotes rapid membrane fusion," *Science*, vol. 313, no. 5787, pp. 673–676, 2006.

[198] A. Witkowska, S. Spindler, R. G. Mahmoodabadi, V. Sandoghdar, and R. Jahn, "Differential diffusional properties in loose and tight docking prior to membrane fusion," *bioRxiv*, 2020.

[199] K. Wiederhold and D. Fasshauer, "Is assembly of the snare complex enough to fuel membrane fusion?" *Journal of Biological Chemistry*, vol. 284, no. 19, pp. 13 143–13 152, 2009.

[200] C. G. Schuette, K. Hatsuzawa, M. Margittai, A. Stein, D. Riedel, P. Küster, M. König, C. Seidel, and R. Jahn, "Determinants of liposome fusion mediated by synaptic snare proteins," *Proceedings of the National Academy of Sciences*, vol. 101, no. 9, pp. 2858–2863, 2004.

[201] D. Fasshauer, W. Antonin, M. Margittai, S. Pabst, and R. Jahn, "Mixed and noncognate snare complexes characterization of assembly and biophysical properties," *Journal of Biological Chemistry*, vol. 274, no. 22, pp. 15 440–15 446, 1999.

[202] T. J. Siddiqui, O. Vites, A. Stein, R. Heintzmann, R. Jahn, and D. Fasshauer, "Determinants of synaptobrevin regulation in membranes," *Molecular biology of the cell*, vol. 18, no. 6, pp. 2037–2046, 2007.

[203] A. Witkowska, L. Jablonski, and R. Jahn, "A convenient protocol for generating giant unilamellar vesicles containing snare proteins using electroformation," *Scientific reports*, vol. 8, no. 1, pp. 1–8, 2018.

[204] J. Schindelin, I. Arganda-Carreras, E. Frise, V. Kaynig, M. Longair, T. Pietzsch, S. Preibisch, C. Rueden, S. Saalfeld, B. Schmid *et al.*, "Fiji: an open-source platform for biological-image analysis," *Nature methods*, vol. 9, no. 7, pp. 676–682, 2012.

[205] J. W. Eaton, D. Bateman, and S. Hauberg, *GNU Octave version 3.0. 1 manual: a high-level interactive language for numerical computations.* SoHo Books, 2007.

[206] M. M. Ghahremanpour, S. S. Arab, S. B. Aghazadeh, J. Zhang, and D. van der Spoel, "Membuilder: a web-based graphical interface to build heterogeneously mixed membrane bilayers for the gromacs biomolecular simulation program," *Bioinformatics*, vol. 30, no. 3, pp. 439–441, 2014.

[207] E. P. Raman and A. D. MacKerell Jr, "Rapid estimation of hydration thermodynamics of macromolecular regions," *The Journal of chemical physics*, vol. 139, no. 5, p. 08B609_1, 2013.

[208] A. Ben-Naim, *Molecular Theory of Water and Aqueous Solutions: The role of water in protein folding, self-assembly and molecular recognition.* World Scientific, 2009, vol. 2.

[209] A. Ben-Naim *et al.*, "The rise and fall of the hydrophobic effect in protein folding and protein-protein association, and molecular recognition," *Open Journal of Biophysics*, vol. 1, no. 01, p. 1, 2011.

[210] A. Y. Ben-Naim, *Hydrophobic interactions.* Springer Science & Business Media, 2012.

[211] T. Beuming, Y. Che, R. Abel, B. Kim, V. Shanmugasundaram, and W. Sherman, "Thermodynamic analysis of water molecules at the surface of proteins and applications to binding site prediction and characterization," *Proteins: Structure, Function, and Bioinformatics*, vol. 80, no. 3, pp. 871–883, 2012.

[212] J. E. Ladbury, "Just add water! the effect of water on the specificity of protein-ligand binding sites and its potential application to drug design," *Chemistry & biology*, vol. 3, no. 12, pp. 973–980, 1996.

[213] T. Young, R. Abel, B. Kim, B. J. Berne, and R. A. Friesner, "Motifs for molecular recognition exploiting hydrophobic enclosure in protein–ligand binding," *Proceedings of the National Academy of Sciences*, vol. 104, no. 3, pp. 808–813, 2007.

[214] R. A. Friesner, R. B. Murphy, M. P. Repasky, L. L. Frye, J. R. Greenwood, T. A. Halgren, P. C. Sanschagrin, and D. T. Mainz, "Extra precision glide: Docking and scoring incorporating a model of hydrophobic enclosure for protein- ligand complexes," *Journal of medicinal chemistry*, vol. 49, no. 21, pp. 6177–6196, 2006.

XI | Acknowledgments