

The relative entropy is fundamental to multiscale and inverse thermodynamic problems

M. Scott Shell^{a)}

Department of Chemical Engineering, University of California, Santa Barbara, California 93106, USA

(Received 7 May 2008; accepted 10 September 2008; published online 10 October 2008)

We show that the relative entropy, $S_{\text{rel}} \equiv \sum p_T \ln(p_T/p_M)$, provides a fundamental and unifying framework for multiscale analysis and for inverse molecular-thermodynamic problems involving optimization of a model system (M) to reproduce the properties of a target one (T). We demonstrate that the relative entropy serves as a generating function for principles in variational mean-field theory and uniqueness and gives intuitive results for simple case scenarios in model development. Moreover, we suggest that the relative entropy provides a rigorous framework for multiscale simulations and offers new numerical techniques for linking models at different scales. Finally, we show that S_{rel} carries physical significance by using it to quantify the deviations of a three-site model of water from simple liquids, finding that the relative entropy, a thermodynamic concept, even predicts water's kinetic anomalies. © 2008 American Institute of Physics.
[DOI: 10.1063/1.2992060]

I. INTRODUCTION

Many physical problems of interest entail interactions at multiple length and time scales that are challenging to model in a bottom-up manner using atomic physical principles. Such problems abound in protein folding and assembly, for example, where small-scale single-residue modifications can have profound effects on large-scale macromolecular structure and association. Detailed all-atom models, while rooted in molecular physics, are not generally able to reproduce the large length and time scale properties of such systems due to immense computational demands in simulating them. On the other hand, simplified coarse-grained (CG) models have been widely used to study such properties, but it has remained challenging to derive these exclusively from atomic physics, with a more common practice being hand-tuning models to reproduce experimentally observed bulk behavior. What is missing here is a basic multiscale theory for linking models at different resolutions in a seamless manner.

Closely related to multiscale analyses are so-called inverse simulation methods. While simulations are routinely used to study the properties of systems with specific interaction potentials and at prescribed thermodynamic conditions, inverse methods aim to do exactly the opposite of this procedure: given an observed set of properties, what are the molecules, energy functions, and/or thermodynamic conditions that give rise to them? In multiscale studies, inverse methods provide a way to develop coarser models from the results of simulations of finer ones. The aim of these methods is accurate CG models that faithfully reproduce the properties of interest of the atomistic ones, while permitting studies of much larger systems and at longer time scales.

An early inverse algorithm was the reverse Monte Carlo method.^{1,2} This approach determines pair potentials that give

rise to measured pair distribution functions (PDFs) by way of iterations between Monte Carlo simulations with test potentials and updates based on deviations of the computed PDFs. On the other hand, Izvekov and Voth,³ building on the earlier work of Ercolessi and Adams,⁴ refined a “force-matching” method to fit potentials by minimizing the differences in predicted atomic forces with respect to a reference system. Both methods are now routinely used in CG model development. However, it has been challenging to identify general theoretical approaches for arbitrary systems, energy functions, and ensembles.

In this work we introduce a new unifying framework for linking systems at different scales that provides a basic theory for multiscale analysis and the development of inverse algorithms. We argue that a quantity called the relative entropy is of fundamental physical significance to multiscale problems. We consider the basic task at hand to be the identification of one thermodynamic ensemble (i.e., a particular system, interaction potential, and set of state conditions) that best reproduces the features of an existing ensemble. We term the former system the *model* and the latter the *target*. We show that the model most optimally represents the target when the relative entropy is minimized, given by

$$S_{\text{rel}} = \sum_i p_T(i) \ln \frac{p_T(i)}{p_M(i)}, \quad (1)$$

where the summation is over all configurations, $p(i)$ is the probability of configuration i in an ensemble, and T and M denote target and model, respectively. This expression bears some resemblance to the usual statistical-thermodynamic entropy, except that it entails *two* sets of ensemble probabilities. The relative entropy stems from information theory, but recently Wu and Kofke,⁵ in their work on bias in free energy measurements, made the key discovery that S_{rel} is also useful for quantifying the overlap between two molecular ensembles. Earlier work by Hummer *et al.*⁶ also used the rela-

^{a)}Electronic mail: shell@engineering.ucsb.edu.

tive entropy to develop a model of hydrophobic effects.

Here we provide a basic derivation of the relative entropy in the context of molecular ensembles, showing that it stems from the log likelihood that a model ensemble mimics a target. Subsequently, we show that its minimization unifies some established results in statistical mechanics and ensemble optimization, and on that basis, we argue that it is a fundamentally important quantity in multiscale problems. We also show that the relative entropy suggests several strategies for the optimization of CG models using numerical methods. Finally, we perform a case study in which we compute the relative entropy between liquid water and the Lennard-Jones (LJ) liquid, demonstrating that S_{rel} captures the degree of anomalous thermodynamic and kinetic behavior in water relative to that of simple liquids.

II. A LIKELIHOOD-BASED DERIVATION OF S_{rel}

Here we provide a rough derivation of the relative entropy based on computing the likelihood that a sampling of configurations from one system reproduces the expected distribution in the other. Consider a model and target system, each with known ensemble probabilities p_M and p_T . For the sake of simplicity, we will first take the number of degrees of freedom to be the same such that a one-to-one mapping can be performed between model and target configurations. We then perform a simple test operation of drawing a large number n of random configurations from the model according to p_M . We want to calculate the likelihood that these configurations appear with the *expected* frequencies we would have obtained if instead we had drawn them from the target ensemble, the latter given by

$$\bar{n}(i) = n \times p_T(i). \quad (2)$$

In other words, we want to calculate the probability that this test of the model ensemble, consisting of n random configurations, could be used to reconstruct accurately the target probability distribution p_T . The likelihood is given by a simple multinomial expression:

$$\begin{aligned} L(T|M) &= n! \prod_i \frac{1}{\bar{n}(i)!} p_M(i)^{\bar{n}(i)} \\ &= n! \prod_i \frac{1}{[n p_T(i)]!} p_M(i)^{n p_T(i)}. \end{aligned} \quad (3)$$

Of interest is the limiting behavior of the likelihood in an exhaustive test where $n \rightarrow \infty$. Taking the logarithm of this expression and applying Stirling's approximation,

$$\ln L(T|M) = -n \sum_i p_T(i) \ln \frac{p_T(i)}{p_M(i)}. \quad (4)$$

Here, the log likelihood scales linearly with n . The constant of proportionality is the negative relative entropy given by Eq. (1). Thus a physical interpretation of S_{rel} might be that it conveys the log probability that one test configuration of the model ensemble is representative of the target.

Though this treatment considered only systems with the same number of degrees of freedom, the probabilistic approach enables a straightforward application to the case in

which the model system has less degrees of freedom than that target, i.e., in the case that it is coarse-grained. Required is a way to map any configuration of the target system to the model. This may be expressed using a mapping function M that converts a set of coordinates in the more detailed target system, \mathbf{r}_T , to a set of coordinates in the model system, \mathbf{r}_M :

$$\mathbf{r}_M = M(\mathbf{r}_T). \quad (5)$$

In shorthand notation, we will write $j=M(i)$, where i represents a target configuration \mathbf{r}_T and j a model configuration \mathbf{r}_M . The mapping function follows directly from the choice of the form of the CG model, which depends on outside considerations as to the degree of coarse-graining desired.

Note that a single model configuration can be generated by multiple target ones. The number of target configurations mapping onto a model configuration k is given by the degeneracy

$$\Omega_{\text{map}}(k) = \sum_{\text{all } i} \delta_{k,M(i)}, \quad (6)$$

where the summation is performed over all target configurations and the delta function filters for those for which $k=M(i)$. Given a single model conformation j , with no more information, it is impossible to assign it exclusively to any one of the $\Omega_{\text{map}}(j)$ target conformations. Returning to the test scenario above, therefore, the probability that a specific target conformation i is produced by the model ensemble is given by the probability of the corresponding model configuration in the model ensemble decreased by this degeneracy:

$$p'_M(i) = \frac{p_M(M(i))}{\Omega_{\text{map}}(M(i))}, \quad (7)$$

where $p_M(j)$ gives the normalized probability of a configuration j in the model ensemble. In other words, the sum $p_M(j)$ over all model configurations j is unity, while the sum $p'_M(i)$ over all target conformations i is 1. With these considerations in mind, and considering the form of the likelihood in Eq. (3), the relative entropy can be expressed as

$$\begin{aligned} S_{\text{rel}} &= \sum_i p_T(i) \ln \frac{p_T(i)}{p'_M(i)} \\ &= \sum_i p_T(i) \ln \frac{p_T(i)}{p_M(M(i))} + \sum_i p_T(i) \ln \Omega_{\text{map}}(M(i)) \\ &= \sum_i p_T(i) \ln \frac{p_T(i)}{p_M(M(i))} + \langle S_{\text{map}} \rangle_T, \end{aligned} \quad (8)$$

where the sums are performed for all possible configurations (i.e., over all degrees of freedom) in the target ensemble and $\langle S_{\text{map}} \rangle_T$ is the average entropy that results from degeneracies in the target-model mapping. It is particularly important to note that this "mapping entropy" does not depend on properties of the model ensemble beyond the specification of the mapping function M . In other words, $\langle S_{\text{map}} \rangle_T$ is independent of the model Hamiltonian.

This line of reasoning, based on a log likelihood, shows that S_{rel} measures the extent to which a model overlaps with a target ensemble. As a result, it has several basic properties.

First, it is always positive since L is a probability. Second, at model optimality when the likelihood is a maximum, S_{rel} is at a minimum. It is zero in the case that the model perfectly reproduces the target. Finally, the relative entropy is directional, that is, $S_{\text{rel}}(T|M) \neq S_{\text{rel}}(M|T)$. This asymmetry stems from the need for a *basis* for weighing differences between ensemble probabilities at different configurations; that basis is given by the target ensemble.

III. S_{rel} IN THE CANONICAL ENSEMBLE

Though the relative entropy is generic to any ensemble, we now specialize to the canonical one. Substituting the canonical configurational probabilities into Eq. (1) gives

$$S_{\text{rel}} = \beta \langle U_M - U_T \rangle_T - \beta (A_M - A_T) + \langle S_{\text{map}} \rangle_T, \quad (9)$$

where $\beta = 1/k_B T$, U is the potential energy, A is the configurational part of the Helmholtz free energy, and the average over the potential energies is performed in the target ensemble:

$$\langle U_M - U_T \rangle_T = \sum_i p_T [U_M(M(i)) - U_T(i)], \quad (10)$$

where the mapping function M converts a target configuration i to a model one, as before.

In light of this result, it is first useful to consider the instructive case scenario in which the target ensemble consists of a single configuration. Such situations routinely occur in efforts to design protein sequences or energy functions that stabilize a unique given structure. In this case, the target ensemble is characterized by a single nonzero $p_T(i^*)$ for some structure i^* . This simplifies the expression for the relative entropy in Eq. (9) to $S_{\text{rel}} = U_M(i^*) - A_M$, neglecting any mapping degeneracies (which do not depend on the model energy function). Thus, minimization here leads directly to maximization of the model probability for the target configuration, $p_M(i^*) = \exp[\beta A_M - \beta U_M(i^*)]$, an intuitive result for this simple one-structure case.

Based on the strict positivity of the relative entropy, $S_{\text{rel}} \geq 0$, Eq. (9) immediately shows that

$$A_M \leq A_T + \langle U_M - U_T \rangle_T. \quad (11)$$

This expression recapitulates the venerable Gibbs–Bogoulibov–Feynmann (GBS) inequality, central to variational mean-field theory.⁷ The relative entropy generates this result intuitively on the basis of its connection to a log likelihood. In this sense, variational mean-field theory can be seen as a maximum likelihood approach to model development. A strength of the relative entropy approach is that it serves as a generating function for similar inequalities in arbitrary ensembles. Interestingly, in the context of variational theory using the GBS inequality, the mean-field system corresponds to the target ensemble. That means the variational principle optimizes the target rather than the model. This seeming inversion of roles occurs for an important reason: the target provides the basis for weighing ensemble differences, which is only calculable in the simpler mean-field system. A better mean-field solution would require full solution of the original system as the target.

An example of this inversion of roles is the following: in the one dimensional Ising model, the usual variational solution gives the self-consistent mean field as $h_{\text{MF}} = 2J \langle s \rangle_{\text{MF}}$, where J is the coupling constant and $\langle s \rangle_{\text{MF}}$ is the average spin. Alternatively, treatment of the mean-field system as the model gives a value of h_{MF} that is the solution to the implicit equation $\langle s \rangle_{\text{MF}} = \langle s \rangle_{\text{orig}}$ and that is therefore dependent on temperature. This latter case, which returns a more optimal h_{MF} , ultimately requires complete evaluation of the original system to determine its average spin as a function of temperature.

In the canonical ensemble, a simple relationship describes the conditions at which a model system is optimal to a target. Consider a model potential energy function U_M that depends on a collection of adjustable parameters $\{\lambda_1, \lambda_2, \dots\}$. Finding the minimum of Eq. (9) with respect to one of these yields

$$\left\langle \frac{\partial U_M}{\partial \lambda_i} \right\rangle_M = \left\langle \frac{\partial U_M}{\partial \lambda_i} \right\rangle_T. \quad (12)$$

Thus at model optimality, the average derivatives of the trial energy function, with respect to the parameters, are equal in the model and target ensembles. This expression is reminiscent of force-matching techniques for coarse-graining, although here the energy derivatives are not matched between two energy functions but between two *averages* of one energy function.

Consider the simple case in which U_M consists of pairwise interactions characterized by single adjustable length and energy scales, $U_M = \sum_{i < j} \epsilon f(r_{ij}/\sigma)$, where f is a dimensionless function. This kind of functionality characterizes a broad range of typical force field components. One finds, after substitution into Eq. (12), that the optimal values of ϵ and σ , which reproduce a target are given by the conditions

$$\langle U_M \rangle_M = \langle U_M \rangle_T, \quad (13a)$$

$$\langle W_M \rangle_M = \langle W_M \rangle_T, \quad (13b)$$

where W_M is the virial of the model, $W = \sum \mathbf{F}_i \cdot \mathbf{r}_i$. In physical terms, the values of ϵ and σ most appropriate to the target are those that achieve the same average model potential energies and configurational pressures. Note that the optimal ϵ and σ depend on the state conditions of the target ensemble (e.g., T and ρ).

For continuously deformable Hamiltonians, it is possible to introduce a variational version of Eq. (12) for the minimum relative entropy condition

$$\left\langle \frac{\delta U_M[\mathbf{r}'_M]}{\delta U_M[\mathbf{r}_M]} \right\rangle_M = \left\langle \frac{\delta U_M[\mathbf{r}'_M]}{\delta U_M[\mathbf{r}_M]} \right\rangle_T, \quad (14)$$

which simplifies to

$$\langle \delta[\mathbf{r}_M - \mathbf{r}'_M] \rangle_M = \langle \delta[\mathbf{r}_M - \mathbf{r}'_M] \rangle_T, \quad (15)$$

where δ is a multidimensional delta function selecting for those conformations for which $\mathbf{r}_M = \mathbf{r}'_M$, and the averages involve permutations of the coordinates \mathbf{r}'_M for fixed values \mathbf{r}_M . In the target ensemble, the average involves a sum over target degrees of freedom \mathbf{r}'_T , with $\mathbf{r}'_M = M(\mathbf{r}'_T)$. Recognizing that

the averages over these delta functions yield probability distributions, we finally arrive at

$$\varphi_M(\mathbf{r}_M) = \varphi_T(\mathbf{r}_M), \quad (16)$$

where $\varphi(\mathbf{r}_M)$ gives the probability of conformation \mathbf{r}_M in the given ensemble. The condition described by Eq. (16) is that the model energy function reproduces the true free energy surface of the target system when projected through the mapping function onto the model degrees of freedom. That is, we have the following relations:

$$\varphi_M(\mathbf{r}_M) \propto \exp[-\beta U_M(\mathbf{r}_M)], \quad (17a)$$

$$\begin{aligned} \varphi_T(\mathbf{r}_M) &\propto \sum_i \exp[-\beta U_T(\mathbf{r}_T^i)] \delta[\mathbf{r}_M - M(\mathbf{r}_T^i)] \\ &\propto \exp[-\beta F_T(\mathbf{r}_M)], \end{aligned} \quad (17b)$$

where the summation i proceeds over all target ensemble conformations and $F_T(\mathbf{r}_M)$ is a generalized free energy (potential of mean force) in the target ensemble along the reduced degrees of freedom \mathbf{r}_M . Thus, unconstrained relative entropy minimization of a model with CG degrees of freedom rigorously returns the correct free energy surface in the corresponding target system within an arbitrary additive constant.

This variational approach can also be used to show that a measured PDF is produced by a unique pair potential, the so-called uniqueness theorem developed by Henderson.⁸ Let $v_M(r)$ and $v_T(r)$ be the pair interaction potentials of a model and target system, respectively, and $g_M(r)$ and $g_T(r)$ their corresponding PDFs. We want to show that if $g_M = g_T$, then the pair potentials must be the same. We proceed by finding the optimal v_M . In analogy with Eq. (14),

$$\left\langle \frac{\delta U_M[v_M(r)]}{\delta v_M(r)} \right\rangle_M = \left\langle \frac{\delta U_M[v_M(r)]}{\delta v_M(r)} \right\rangle_T, \quad (18)$$

or, after simplification of the functional derivatives,

$$\left\langle \sum_{i < j} \delta[r - r_{ij}] \right\rangle_M = \left\langle \sum_{i < j} \delta[r - r_{ij}] \right\rangle_T. \quad (19)$$

By multiplying each side by $V/4\pi r^2 N^2$, this expression can be converted to the equality

$$g_M(r) = g_T(r). \quad (20)$$

That is, the optimal model system, when subject to variational modification of the pair potential, yields the same PDF as the target. If both systems' energy functions consist entirely of pair potentials, then the specific v_M of the target is fully contained within the solution space of the model, and hence the bound $S_{\text{rel}}=0$ can be achieved. Since there is no more optimal model than this case, we must conclude that Eq. (20) corresponds to $S_{\text{rel}}=0$. Finally, it follows that v_M at most can differ from v_T by a constant value since any other differences would alter the configurational probabilities p_M and p_T , and hence increase the value of S_{rel} . Thus, the condition of equality of PDFs implies that two systems are optimal to one another, which in turn demands equality of pair potentials. It is important to note that if three-body or other terms are present in the energy function, a similar derivation

shows that higher-order correlations are necessary to constrain uniqueness of the potential.

The relative entropy also suggests a general framework for numerical methods for model optimization. Commonly, for example, it is desired to parametrize a CG model based on simulation results of an atomistic one. Here, numerical minimization of S_{rel} with respect to the energy function parameters of the CG model enables its optimization. Two approaches can be used to compute the relative entropy for this purpose.

- (i) Minima of S_{rel} can be located directly by its calculation over a broad range of model parameter space. For example, Eq. (9) requires only the computation of the quantities $\langle U_M \rangle_T$ and A_M for all values of the parameters to be optimized, as the remaining terms in Eq. (9) are constant (provided also the mapping function is fixed). Any one of many current free energy methods can be used to compute changes in A_M with parameters here, including thermodynamic integration, multiple histogram reweighting, and flat histogram techniques.⁹ When many parameters are to be simultaneously optimized, it would be efficient to perform "line searches" in parameter space, whereby a gradient-based minimization procedure is coupled with on-the-fly computation of free energy changes along parameter vectors. Such an approach is increasingly feasible with modern flat histogram methods.
- (ii) For continuous parameters, one can find local minima of S_{rel} using standard root-finding methods applied to its derivative, readily implemented as an iterative simulation protocol. Here we derive a simple update rule based on a Newton–Raphson iteration for single parameter λ in the canonical ensemble:

$$\begin{aligned} \lambda_{i+1} &= \lambda_i - (\partial S_{\text{rel}}/\partial \lambda)/(\partial^2 S_{\text{rel}}/\partial \lambda^2) \\ &= \lambda_i - \left[\left\langle \frac{\partial U}{\partial \lambda} \right\rangle_T - \left\langle \frac{\partial U}{\partial \lambda} \right\rangle_M \right] \left[\left\langle \frac{\partial^2 U}{\partial \lambda^2} \right\rangle_T \right. \\ &\quad \left. - \left\langle \frac{\partial^2 U}{\partial \lambda^2} \right\rangle_M + \beta \left\langle \left(\frac{\partial U}{\partial \lambda} \right)^2 \right\rangle_M - \beta \left\langle \frac{\partial U}{\partial \lambda} \right\rangle_M^2 \right]^{-1}. \end{aligned} \quad (21)$$

Here, U is the model potential function U_M , i is an iteration index for the model version, the model averages are computed from a test model simulation using the current parameter λ_i , and the target averages stem from a single prior target simulation. Alternatively, if the second term in brackets is negative, the current value of λ_i does not sit near a local minimum of S_{rel} , and the Newton–Raphson update will fail. In this case, this term can be replaced by a step factor to perform a steepest descent procedure.

Equation (21) provides a simple way to refine models: for each candidate parameter set, the target averages are evaluated using measured distributions from the target ensemble (e.g., PDFs), and canonical simulations are performed to evaluate averages in the model ensemble. Then, the values of the model parameters are updated using Eq.

(21), new model simulations are performed, and the process repeats until the parameters converge. For the multiscale task of parametrizing a CG model, this provides an efficient approach to model refinement using a single atomistic simulation coupled with iterative CG ones.

In the special case that the model energy function is linear in λ , $U = \lambda f(\mathbf{r}^N) + \dots$, Eq. (21) simplifies to

$$\lambda_{i+1} = \lambda_i - k_B T \frac{\langle f \rangle_T - \langle f \rangle_M}{\langle f^2 \rangle_M - \langle f \rangle_M^2}. \quad (22)$$

Note here that the second derivative of the relative entropy with respect to λ is formally positive and relates to the variance of f in the model ensemble. Thus, for linear terms such as these, the relative entropy has a single minimum, and the Newton–Raphson procedure converges to a minimum for any starting value of λ . Many common potential energy functional forms entail linear prefactors, such as atomic charges and LJ coefficients; moreover, the parameters of fitted cubic spline potentials are also linear. These features bode well for the numerical stability of S_{rel} minimization in model development.

IV. A STUDY OF S_{rel} IN LIQUID WATER

To demonstrate the physical significance of the relative entropy, we use it to quantify the anomalous behavior of liquid water, using the SPC/E model.¹⁰ In recent years, there has been considerable interest in characterizing, from a molecular point of view, many of the unusual properties that water exhibits when compared to “simple” liquids, such as expansion upon cooling and compression induced increases in diffusivity.¹¹ Simple molecular-geometric order parameters have emerged as insightful predictors of bulk anomalous behavior. A study by Errington and Debenedetti¹² examined the SPC/E model and found that order metrics based on translational and tetrahedral orientational correlations predicted anomalous behavior in both water’s thermodynamic and kinetic properties. That work interpreted bulk anomalies in terms of preferential microscopic interactions that favor open tetrahedral geometries via hydrogen bonding.

Here we take a different approach that *directly* compares water to a simple fluid. Namely, we ask the following: how much overlap does water have with a reference simple fluid, across different state points, as measured by the relative entropy? For the reference fluid, we take a pure LJ system and find those values of effective LJ parameters ϵ_{eff} and σ_{eff} that best reproduce the properties of water at a given state point, using relative entropy minimization.

We measure the relative entropy between SPC/E water (the target) and a pure LJ system with the same number of molecules (the model) to quantify water’s differences with simple liquids. This procedure consists of two steps: (1) first identifying the LJ parameters ϵ_{eff} and σ_{eff} that minimize the relative entropy with respect to water and (2) evaluating the value of S_{rel} for the optimum effective parameters. In order to compute the relative entropy, we use the expression in Eq. (9). For this equation, we compute $A_{\text{SPCE}}(\rho, T)$ and $A_{\text{LJ}}(\epsilon, \sigma)$ to absolute values using a flat histogram Monte Carlo technique,¹³ integrating to an ideal gas state in two parts:




FIG. 1. (Color online) S_{rel} between a pure LJ system (model) and SPC/E water (target) for a range of effective LJ parameters ϵ and σ . The relative entropy is computed on a grid of 50 values for ϵ and 30 for σ .

first, through an isochore to high T and, second, along an isotherm to low density. We then use molecular dynamics simulations to compute the average SPC/E potential energies $\langle U_{\text{SPCE}} \rangle_{\text{SPCE}}$ and the SPC/E center-of-mass PDFs. We then integrate the SPC/E PDFs, using the pairwise LJ potential, to find $\langle U_{\text{LJ}} \rangle_{\text{SPCE}}$. These two free energies and two potential energy averages enable the determination of S_{rel} via Eq. (9).

Simulation details are as follows. All simulations are performed with 216 molecules. For SPC/E, electrostatics are treated with the Ewald summation, with $\alpha L = 5.6$ and k -vectors with magnitude $\leq 5 \times 2\pi/L$ included in the reciprocal space summation. Non-Coulombic interactions are cut and shifted at 2.5σ . Molecular dynamics simulations are performed using the velocity Verlet integrator with a Nosé–Hoover thermostat. A time step of 1 fs is used, with total run lengths varying between 0.1 ($T=350$ K) and 20 ns (240 K). Diffusion constants are computed using the Einstein relation with 20 time origins spaced 1 ps apart. Bonds are constrained using the RATTLE protocol. For pure LJ, the potential is cut and shifted at 3.0σ .

Figure 1 shows the value of the relative entropy at the single water state point $\rho=950$ kg/m³ and $T=300$ K as a function of the effective LJ parameters ϵ_{eff} and σ_{eff} . These results reveal a minimum at the values $\epsilon_{\text{eff}}=0.277$ kJ/mol and $\sigma_{\text{eff}}=2.88$ Å, indicating that these parameters best reproduce the thermodynamic ensemble of SPC/E water at that particular state point. Interestingly, the minimum is much broader in the energy than in the distance parameter. This is most likely due to the fact that changes to σ_{eff} affect sensitively the effective packing fraction of the liquid through modifications to the pairwise potential’s repulsive core.

We have repeated the calculation in Fig. 1 for a range of state points, and the effective LJ parameters corresponding to the relative entropy minimum for each are shown in Fig. 2.




FIG. 2. (Color online) Effective LJ parameters as a function of state conditions, stemming from relative entropy minimization with SPC/E water as target.

Here, we find a slight dependence of these parameters on the temperature and density. These deviations, which would be absent if water were perfectly modeled by the LJ potential, demonstrate important departures from simple liquid behavior. Both ϵ_{eff} and σ_{eff} appear to increase slightly with temperature, which could be an indication that water maintains a more structured molecular environment as it is heated (e.g., retains more pronounced correlation functions) than would be expected for a simple liquid under similar conditions. Furthermore, while the optimal energy parameter seems fairly insensitive to density, the same is not true for σ_{eff} , which grows slightly with density. Thus, the variation in water's effective packing fraction ($\rho\sigma^3$) is more pronounced than that of an equivalent simple fluid.

Overall, however, the average variation in ϵ_{eff} and σ_{eff} across the state conditions is relatively small at 15% and 2%, respectively. Therefore, for the remainder of our calculations, we hold these two parameters constant at their values for the state point $\rho=950 \text{ kg/m}^3$ and $T=300 \text{ K}$, as shown in Fig. 1. This enables us to determine the relative entropy at higher water densities than those shown in Fig. 2, where otherwise the S_{rel} minimization procedure tends to locate effective parameters that correspond to a crystallized LJ system (results not shown).

Figure 3 shows the values of S_{rel} for the LJ system targeted to SPC/E water over a range of ρ and T . Here, S_{rel} is clearly effective as a metric for water's deviations from the LJ system. Its higher values at lower temperatures and densities readily demonstrate water's anomalous behavior at those conditions, consistent with the previous results of Errington and Debenedetti¹² and the idea that hydrogen bonding promotes transient, open, and tetrahedrally coordinated local geometries at these state points. It is important to note, however, that the relative entropy generates these results without directly considering microscopic structure, and it is com-




FIG. 3. (Color online) Top: S_{rel} between a pure LJ system (model) and SPC/E water (target). The effective parameters of the pure LJ system are found by minimization of S_{rel} at one state point, as described in the text. Bottom: The log ratio of the self-diffusivity of SPC/E water with that of a pure LJ system. The inset shows the nonmonotonic behavior of water's diffusivity with compression at lower temperatures. All lines are polynomial fits to guide the eyes.

pletely general so far as a reference simple fluid can be chosen. Moreover, the nonzero absolute value of S_{rel} , even at the highest temperature studied, is consistent with the fact that the LJ liquid lacks any orientational degrees of freedom, which results in a residual mapping entropy of $\langle S_{\text{map}} \rangle_T = \ln(8\pi^2) \approx 4.4$ per molecule.

Remarkably, the relative entropy also tracks differences in the kinetic behavior of the two systems, despite its purely thermodynamic construction. In the bottom of Fig. 3, the logarithm of the ratio of the self-diffusion constants is shown to closely follow the behavior of the relative entropy: at points where water's self-diffusion is substantially slower than would be predicted for a LJ fluid at similar conditions (i.e., with water-optimized parameters), the relative entropy is high, indicating a departure of water from simple behavior.




FIG. 4. (Color online) Comparison of the relative entropy and the diffusion constants of a simple LJ reference fluid and SPC/E water.

Importantly, the relative entropy's ability to track kinetic differences persists even within the region where water's self-diffusivity shows anomalous nonmonotonic behavior upon compression, as shown in the inset of Fig. 3. Figure 4 shows more closely the connection between the relative entropy and the diffusion constants of these two systems. These results reflect ongoing ideas about the role of excess entropies in condensed-phase dynamics,^{14,15} although it is important to note that the relative entropy cannot be strictly separated into purely excess entropy quantities.

V. CONCLUSIONS

In summary, we have shown that the relative entropy provides a useful and fundamental framework for multiscale analysis by way of measuring the overlap between two molecular-thermodynamic ensembles. Of much interest may be its future use in new coarse-graining simulation techniques. In particular, the close connection found here between S_{rel} and diffusion in water is encouraging that the relative entropy approach can simultaneously optimize *both* the thermodynamic and kinetic properties of CG models.

ACKNOWLEDGMENTS

The author gratefully acknowledges T. M. Truskett for helpful comments, and a Dreyfus Foundation New Faculty Award and Santa Barbara Cottage Hospital Special Research Award for support.

- ¹R. L. McGreevy and L. Pusztai, *Mol. Simul.* **1**, 359 (1988).
- ²A. P. Lyubartsev and A. Laaksonen, *Phys. Rev. E* **52**, 3730 (1995).
- ³S. Izvekov and G. A. Voth, *J. Phys. Chem. B* **109**, 2469 (2005); *J. Chem. Phys.* **123**, 134105 (2005).
- ⁴F. Ercolessi and J. B. Adams, *Europhys. Lett.* **26**, 583 (1994).
- ⁵D. Wu and D. A. Kofke, *J. Chem. Phys.* **123**, 054103 (2005); **123**, 084109 (2005).
- ⁶G. Hummer, S. Garde, A. E. Garcia, A. Pohorille, and L. R. Pratt, *Proc. Natl. Acad. Sci. U.S.A.* **93**, 8951 (1996).
- ⁷H. Callen, *Thermodynamics and Thermostatistics* (Wiley, New York, 1985).
- ⁸R. L. Henderson, *Phys. Lett.* **49A**, 197 (1974).
- ⁹*Free Energy Calculations: Theory and Applications in Chemistry and Biology*, edited by C. Chipot and A. Pohorille (Springer, Berlin, 2007).
- ¹⁰H. J. C. Berendsen, J. R. Grigera, and T. P. Straatsma, *J. Phys. Chem.* **91**, 6269 (1987).
- ¹¹C. A. Angell, R. D. Bressel, M. Hemmati, E. J. Sare, and J. C. Tucker, *Phys. Chem. Chem. Phys.* **2**, 1559 (2000).
- ¹²J. R. Errington and P. G. Debenedetti, *Nature (London)* **409**, 318 (2001).
- ¹³M. S. Shell, P. G. Debenedetti, and A. Z. Panagiotopoulos, *J. Chem. Phys.* **119**, 9406 (2003).
- ¹⁴Y. Rosenfeld, *Phys. Rev. A* **15**, 2545 (1977).
- ¹⁵J. Mittal, J. R. Errington, and T. M. Truskett, *J. Chem. Phys.* **125**, 076102 (2006).