

*Identifying Nonadditive Contributions to the Hydrophobicity of Chemically Heterogeneous Surfaces via Dual-Loop Active Learning*

Atharva S. Kelkar, Bradley C. Dallin, and Reid C. Van Lehn\*

Department of Chemical and Biological Engineering, University of Wisconsin – Madison, 1415 Engineering Drive, Madison, WI, 53706 USA

\*send correspondence to: [vanlehn@wisc.edu](mailto:vanlehn@wisc.edu)

## **Abstract**

Hydrophobic interactions drive numerous phenomena involving surfaces that are chemically heterogeneous at the nanoscale. Nonadditive contributions to the hydrophobicity of such surfaces depend on the chemical identities and spatial patterns of polar and nonpolar groups in ways that remain poorly understood. Here, we develop an active learning framework that utilizes molecular dynamics (MD) simulations, enhanced sampling, and a convolutional neural network to predict the hydration free energy (a thermodynamic descriptor of hydrophobicity) for nearly 200,000 chemically heterogeneous self-assembled monolayers (SAMs). Analysis of this data set reveals that SAMs with distinct polar groups exhibit substantial variations in hydrophobicity as a function of their composition and patterning, but the clustering of nonpolar groups is a common signature of highly hydrophobic patterns. Further MD analysis relates such clustering to the perturbation of interfacial water structure. These results provide new insight into the influence of chemical heterogeneity on hydrophobicity via quantitative analysis of a large set of surfaces, enabled by the active learning approach.

## Introduction

The hydrophobicity of an interface dictates the magnitude of attractive, water-mediated hydrophobic interactions that drive diverse processes, including molecular recognition,<sup>1-3</sup> protein folding,<sup>2, 4, 5</sup> colloidal aggregation,<sup>3, 6</sup> and assembly at the nano-bio interface.<sup>7-9</sup> While the hydrophobicity of spherical, uniformly nonpolar solutes can be related to the molecular rearrangement of water,<sup>10</sup> the hydrophobicity of solutes with chemically heterogeneous surfaces—characterized by polar and nonpolar groups in close ( $\sim$ nm) proximity—is less well-understood.<sup>11, 12</sup> Approaches to predict quantitative metrics of hydrophobicity for these surfaces often assume that contributions from polar and nonpolar groups are additive; example approaches include predicting water contact angles based on area-weighted sums of the contact angles of polar and nonpolar surface regions (*i.e.*, the Cassie equation),<sup>13, 14</sup> predicting transfer free energies by multiplying the nonpolar solvent-accessible surface area of a solute by a constant fitting parameter,<sup>15-19</sup> or predicting partitioning into nonpolar media by summing hydrophobicity scale values<sup>20</sup> or octanol-water partition coefficients.<sup>21-24</sup> However, experiments have demonstrated that the hydrophobicity of chemically heterogeneous surfaces can deviate from these additive approximations depending on the spatial arrangement of polar and nonpolar groups at the nanoscale (*i.e.*, surface patterning) and the chemical identity of polar groups. For example, Woodward *et al.* attributed deviations from the Cassie equation in binary self-assembled monolayers (SAMs) to a “boundary region” around hydrophobic patches.<sup>25</sup> Wang *et al.* and Ma *et al.* found that the hydrophobic force between a nonpolar SAM and a  $\beta$ -peptide containing polar and nonpolar amino-acid side chains was only measurable if the polar side chains were adjacent to a well-defined nonpolar patch.<sup>11, 26</sup> These same studies revealed the impact of the adjacent polar groups: lysine side chains (with an amine group) had no measurable effect on the hydrophobic force, whereas glutamine side chains (with an amide group) eliminated the hydrophobic force completely.<sup>11, 26</sup> This interplay between surface patterning and polar group chemistry on the hydrophobicity of chemically heterogeneous surfaces remains poorly understood.

To complement experimental studies, atomistic molecular dynamics (MD) simulations have been employed to analyze interfacial water structure at chemically heterogeneous interfaces and examine the effect of surface patterning on hydrophobicity.<sup>27-34</sup> For example, Xi *et al.* studied the impact of surface patterning by systematically varying the distance between hydroxyl groups in a binary SAM and found a nonmonotonic relationship between this distance and interfacial

water density fluctuations.<sup>12</sup> Monroe and Shell leveraged a genetic algorithm to show that hydroxylated surfaces with highly clustered polar domains had the highest surface water diffusivity as a measurement of hydrophobicity.<sup>35</sup> This work seeks to build upon these and related studies<sup>36, 37</sup> by exhaustively mapping surface patterning to hydrophobicity, thereby enabling the quantitative analysis of pattern features that impact hydrophobicity. A key challenge is the combinatorial complexity associated with surface patterning; for example, a binary SAM with polar and nonpolar ligands arranged on an  $n \times n$  lattice has  $\sim 2^{n^2}$  arrangements of polar and nonpolar groups. Enhanced sampling methods that calculate the hydration free energy (HFE), a thermodynamic descriptor of hydrophobicity,<sup>27, 28</sup> are too computationally expensive to tractably explore all such patterns. We recently developed a data-driven method for predicting HFEs utilizing a convolutional neural network (CNN) that requires minimal simulation time.<sup>32</sup> However, the CNN's reduced accuracy combined with the combinatorial complexity of patterning necessitates more efficient methods to study the effect of patterning on hydrophobicity, particularly when considering multiple types of polar groups.


In this work, we develop a new computational approach to efficiently map surface patterning to hydrophobicity for chemically heterogeneous surfaces. Our approach uses a Gaussian Process Regression (GPR) model to predict the HFE of a patterned SAM surface using a binary matrix representation of the pattern as input. To efficiently train the GPR model, we develop a dual-loop active learning scheme to label a small subset of patterned SAMs with HFEs. Active learning accelerates model training by choosing training data using an information gain criterion<sup>38</sup> and has been recently employed for materials discovery.<sup>39-41</sup> Our scheme combines a fast, reduced-accuracy method (a pre-trained CNN) with a slow, higher-accuracy method (Indirect Umbrella Sampling) to label selected training patterns with HFEs. Once trained, the GPR model can rapidly predict HFEs for tens of thousands of patterns without additional simulations. We leverage this capability to identify patterns that display large deviations from additive behavior (*i.e.*, the HFE cannot be approximated by additive contributions from polar and nonpolar groups) for nearly 200,000 binary SAMs containing either amine, hydroxyl, or amide polar end groups. Analysis of this large dataset reveals that SAMs with different polar groups exhibit substantial variations in hydrophobicity as a function of surface composition and patterning, but highly hydrophobic patterns are always distinguished by the clustering of nonpolar groups which we attribute to the disruption of interfacial water molecules. These results thus provide new physical insight into the

role of chemical heterogeneity on hydrophobicity via quantitative analysis of a substantially larger set of surfaces than would be accessible by simulations alone.

## Results

### Set of patterned surfaces and descriptor for hydrophobicity

Identifying patterned surfaces with interfacial hydrophobicities that deviate from additive behavior requires defining (i) the set of patterned surfaces, (ii) a descriptor of interfacial hydrophobicity, and (iii) additive behavior. We chose a set of patterned binary SAMs consisting of alkanethiol ligands terminated with either methyl groups (referred to as nonpolar ligands) or one of three polar end groups – amine, hydroxyl, or amide (referred to as polar ligands). Ligand and end group chemical structures are shown in Figure 1A. SAMs were constructed from a hexagonal lattice of 64 ligands with a 4×4 patch of polar and nonpolar ligands embedded within a background of 48 nonpolar ligands. The number and spatial positioning of polar and nonpolar ligands within the patch were permuted to generate different patterns as shown in Figure 1B. In total, this set included



**Figure 1: Space of self-assembled monolayers and labels.** (A) Ligands used to create patterned self-assembled monolayers (SAMs) consist of an alkane backbone with either a nonpolar (methyl) or polar (amine, hydroxyl, or amide) end group. (B) Representative pattern obtained from varying the number and spatial positioning of nonpolar and polar ligands within a central 4×4 ligand patch. Nonpolar ligands within the pattern are blue, polar ligands within the pattern are red, and nonpolar ligands surrounding the pattern are gray. Beneath each pattern is a corresponding Voronoi diagram to illustrate the partitioning of the 2×2 nm<sup>2</sup> area covered by the cavity (described in C) into polar and nonpolar cells following the same color scheme. Each diagram is labeled by the corresponding polar area fraction. (C) Simulation snapshot of a 2×2×0.3 nm<sup>3</sup> cavity, indicated by  $v$ , placed at the SAM-water interface. The snapshot illustrates the cavity when emptied of water during an indirect umbrella sampling simulation.

$3 \times 2^{16} = 196,608$  SAMs, permitting broad exploration of the relationship between spatial patterning and hydrophobicity for different polar end groups.

We quantified SAM interfacial hydrophobicity via the hydration free energy (HFE), which measures the magnitude of water density fluctuations.<sup>27,28</sup> Equation 1 defines the HFE in terms of the probability,  $P_\nu(0)$ , that a cavity (indicated by subscript  $\nu$ ) near the SAM is occupied by zero water molecules:<sup>27</sup>

$$\text{HFE} = -k_B T \ln(P_\nu(0)) \quad (1)$$

Water density fluctuations are enhanced near hydrophobic surfaces, increasing  $P_\nu(0)$  and decreasing the HFE. The HFE is an appropriate descriptor of hydrophobicity for chemically heterogeneous surfaces because it has been shown to vary for surfaces with hand-selected patterns.<sup>12,42</sup> Previous studies have also shown HFEs to correlate with experimental measurements of surface contact angles<sup>43,44</sup> and hydrophobic forces.<sup>33,42</sup> Because the HFE is sensitive to the placement and size of the cavity, all HFE calculations used a  $2.0 \times 2.0 \times 0.3$  nm<sup>3</sup> cavity (Figure 1C) positioned with its base on a constant water (number) density isosurface.<sup>45</sup> SI Section S1.1 further discusses the cavity placement. HFEs were obtained either by Indirect Umbrella Sampling (INDUS),<sup>27,28</sup> an enhanced sampling technique that biases the removal of water molecules from the cavity during independent MD simulations, or by using a previously developed 3D convolutional neural network (CNN),<sup>32</sup> a machine learning model that maps interfacial water positions sampled from a short MD simulation to an HFE. INDUS is a high accuracy approach but requires substantial MD simulation time (~65 ns) whereas the CNN is less accurate but requires minimal simulation time (~2 ns). Details on the usage of these two methods are provide below.

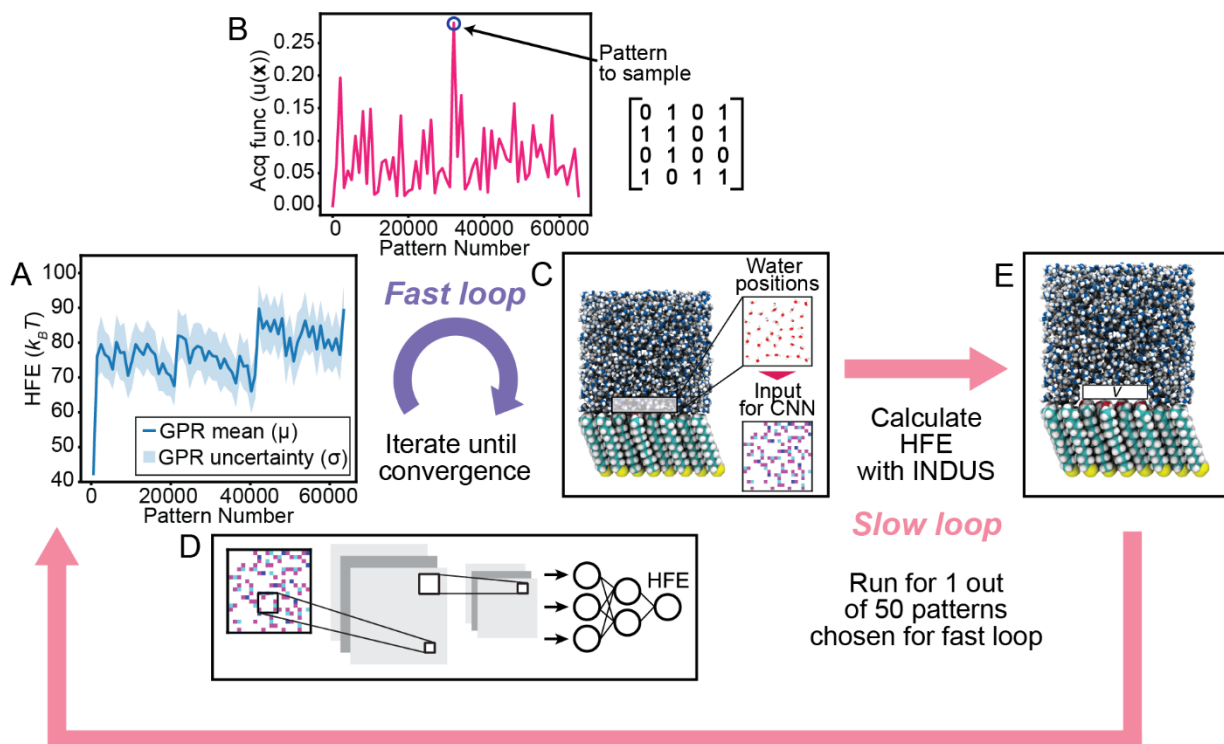
We expect the HFE of a surface that exhibits additive behavior to depend on the relative amounts of polar and nonpolar area covered by the cavity.<sup>10,15-19</sup> Accordingly, we computed Voronoi diagrams to partition the total area covered by the cavity into polar and nonpolar cells determined by the closest ligand end group at equilibrium (Figure 1B). Cells were generated from a subset of MD simulations for each polar end group and were found to be nearly equal hexagons in all simulations. We thus define the polar area fraction as the fraction of covered area that is occupied by polar cells, assuming that cell areas are independent of the pattern and polar end group, and define a pattern as exhibiting additive behavior if its HFE is a linear combination of the HFEs

of SAMs containing only polar and only nonpolar ligands weighted by its polar area fraction. Because the cavity has a square 2D footprint while the SAM lattice is hexagonal, the polar area fraction can vary by up to  $\sim 0.05$  for patterns containing the same number of polar groups and reaches a maximum of  $\sim 0.82$  for a pattern with only polar groups due to the boundary of nonpolar groups (Figure 1B). Details on the calculation of the polar area fraction are provided in the Supplementary Information.

### **Dual-loop active learning scheme to efficiently map patterns to hydrophobicity**

The large number of possible patterned SAMs inhibits the direct calculation of HFEs for all possible SAMs. Instead, we trained a Gaussian Process Regression (GPR) model to map patterns (encoded as  $4 \times 4$  binary matrices) to HFE labels; the model is trained on HFEs calculated from simulation data, but after training the GPR model can predict HFEs for all patterns without additional simulations. To minimize the amount of training data while maximizing GPR model accuracy, we designed a dual-loop active learning scheme that exploits the advantages of both INDUS (high accuracy, more computationally expensive) and the CNN (lower accuracy, less computationally expensive) during GPR model training. This scheme uses accurate INDUS calculations to label selected patterns with large deviations from additive behavior and faster, but less accurate, CNN calculations to explore large portions of pattern space.

Figure 2 illustrates the active learning scheme; individual elements are detailed in the Methods. The GPR model is initially trained using a “seed” set of 52 SAMs for each polar end group with INDUS-calculated HFE labels. Two loops are then performed in parallel to iteratively label additional SAMs with HFEs: a “fast loop” utilizing the CNN and a “slow loop” utilizing INDUS. 50 iterations of the fast loop require approximately the same computational time as one iteration of the slow loop. A pattern is chosen for each iteration of the fast loop by maximizing an acquisition function that optimizes the tradeoff between choosing patterns with high deviations from additive behavior and sampling patterns in uncertain regions of pattern space. The chosen pattern is then labeled with the CNN. The pattern that was labeled with the largest deviation from expected additive behavior during the 50 iterations of the fast loop is then chosen for the slow loop; the fast loop continues for an additional 50 iterations while this pattern is labeled with INDUS. The parallel execution of the fast and slow loops is repeated until convergence of GPR predictions,<sup>46</sup> which required only 30 iterations of the slow loop and 1448 iterations of the fast



**Figure 2: Illustration of the dual-loop active learning scheme.** (A) A Gaussian Process Regression (GPR) model is used to map SAM patterns ( $\mathbf{x} \in \mathbb{R}^d$ ) to HFE labels ( $y \in \mathbb{R}$ ). The GPR model predicts a mean HFE value ( $\mu$ ) and uncertainty in the predicted mean ( $\sigma$ ) for all possible patterns given training data (patterns labeled with HFEs). (B) The pattern which maximizes the acquisition function  $u(\mathbf{x})$ , is chosen as the next pattern to sample during the fast loop. (C) An unbiased MD simulation of a SAM with the chosen pattern is performed. Interfacial water molecules are converted to an input representation for the pre-trained CNN following previous work.<sup>32</sup> (D) Schematic of the CNN, which receives the interfacial water molecule representation as input and outputs an HFE label. (E) The pattern with the highest deviation from expected additive behavior identified during the previous 50 iterations of the fast loop is chosen for the slow loop. This pattern is labeled using accurate but computationally expensive Indirect Umbrella Sampling (INDUS) simulations. The HFE label calculated from the INDUS simulation is then added to the GPR training data.

loop (detailed in the Supplementary Information). Table 1 summarizes the scheme, which we repeated once for each polar end group.

### **Predictions from the Gaussian Process Regression model before and after active learning**

We first trained the GPR model using the same set of randomly selected seed patterns for all three polar end groups, then initiated the dual-loop active learning scheme. Figure 3A depicts envelopes of GPR-predicted HFEs after being trained only on seed patterns. Each envelope illustrates the range of GPR-predicted HFEs for different patterns as a function of polar area

**Table 1: Summary of the dual-loop active learning algorithm.**

<b>Dual-Loop Active Learning Algorithm</b>	
<b>1</b>	Generate a set of 52 seed patterns and label with INDUS simulations
<b>2</b>	Use GPR model to map patterns to HFE labels
<b>3</b>	<i>while</i> (convergence criterion is not met):
<b>4</b>	Execute 50 iterations of the <i>fast loop</i> and 1 iteration of the <i>slow loop</i> in parallel
<b>5</b>	<i>end</i>

---

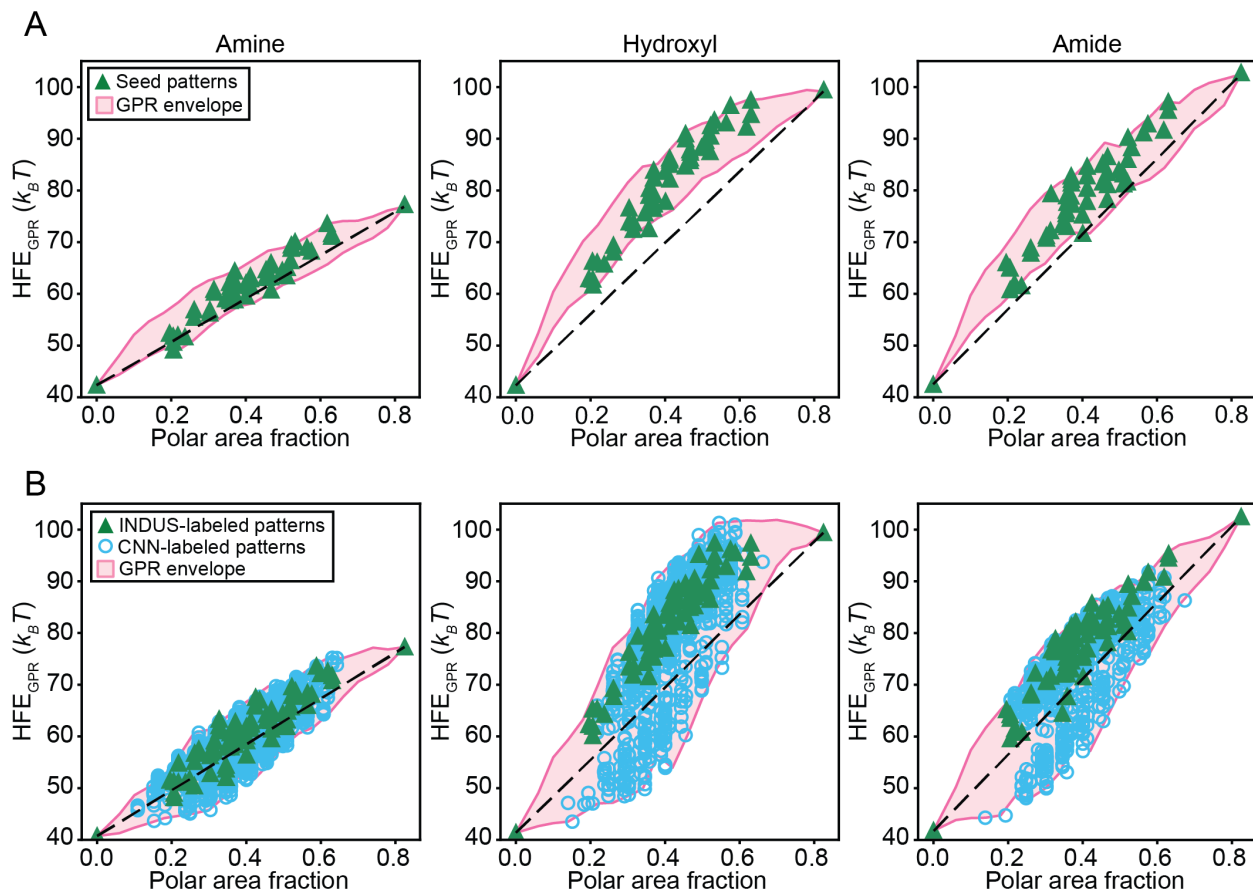
<b>Fast Loop</b>	
<b>1</b>	Maximize the acquisition function to select next pattern to sample
<b>2</b>	Perform unbiased MD simulation and predict HFE using a CNN
<b>3</b>	Append the CNN-predicted HFE to the training data with proper replicate error
<b>4</b>	Use GPR model to map patterns to HFE labels

---

<b>Slow Loop</b>	
<b>1</b>	Choose pattern with max deviation from expected additive behavior from last 50 fast loop iterations
<b>2</b>	Perform INDUS simulation of chosen pattern to calculate HFE
<b>3</b>	Append INDUS-calculated HFE to training data with proper replicate error

fraction; predicted HFEs for all possible SAM patterns lie within the envelopes. The linear interpolation line depicts additive behavior. The differences in HFEs for patterns with equivalent polar area fractions indicate the impact of patterning. All HFE envelopes monotonically increase with polar area fraction as observed previously.<sup>32, 35</sup> However, there are substantial differences between polar groups; HFEs for the amine SAMs (*i.e.*, SAMs containing polar ligands with amine end groups) are clustered near the linear interpolation line, hydroxyl SAMs display large positive deviations from additive behavior, and amide SAMs display positive deviations less than those of the hydroxyl SAMs. Figure 3B shows that the envelopes of GPR-predicted HFEs increase in area after convergence of active learning, indicating that active learning amplifies HFE variations associated with patterning.<sup>46</sup> To quantify these changes, we characterize the envelope of INDUS-calculated HFEs before (*i.e.*, for only the seed patterns) and after active learning (*i.e.*, including patterns selected for the slow loop). Only INDUS-calculated HFEs are considered because of their low error and because patterns were selected for the slow loop to maximize deviations from additive behavior. We define the average envelope width as the difference between the maximum and minimum INDUS-calculated HFEs averaged over all polar area fractions and the maximum envelope width as the largest difference in maximum and minimum INDUS-calculated HFEs for any polar area fraction.






**Figure 3: Hydration free energy (HFE) envelopes generated using Gaussian Process Regression (GPR) models.** Each envelope is plotted using the highest and lowest GPR-predicted HFEs for each polar area fraction (binned in intervals of 0.04); GPR-predicted HFEs for all 65,536 possible patterns lie inside each envelopes. The black dashed linear interpolation lines indicate HFEs corresponding to additive behavior (a weighted average of pure polar and pure nonpolar HFEs based on the polar area fraction). **(A)** Envelopes after training only on the set of seed patterns. Green points denote GPR predictions for the set of seed patterns, which have low uncertainty because they are labeled with INDUS. **(B)** Envelopes after active learning. Green points denote GPR predictions for INDUS-labeled patterns (including seed patterns). Blue circles denote GPR predictions for CNN-labeled patterns.

Figure 3B shows that amine SAMs still display behavior close to additive behavior after active learning. However, the average envelope width increases from  $\sim 3.3 k_B T$  to  $\sim 5.5 k_B T$ , indicating that active learning successfully identifies patterns with higher deviations from additivity than present in the seed set. Similarly, the maximum envelope width increases from  $\sim 6.3 k_B T$  to  $\sim 9.7 k_B T$  after active learning. Hydroxyl SAMs exhibit a plateau in the calculated HFEs as the polar area fraction reaches  $\sim 0.6$ , a value lower than a pure polar pattern (0.82), in qualitative agreement with past studies.<sup>32, 35</sup> All INDUS-calculated HFEs exhibit positive deviations from

additive behavior, with some tracking the upper edge of the GPR-predicted envelope. None of the INDUS-calculated HFEs lie below the linear interpolation line, where the GPR model predicts negative deviations from additive behavior. This discrepancy is an artifact of the CNN which tends to systematically underpredict values at the lower end of the HFE range. Despite this quantitative inaccuracy, active learning still discovers new patterns with large deviations from additive behavior; the average envelope width for hydroxyl patterns increases from  $\sim 4.7 k_B T$  to  $\sim 7.1 k_B T$  after active learning and the maximum envelope width substantially increases from  $\sim 8.1 k_B T$  to  $\sim 15.7 k_B T$ . The largest positive deviation from additivity for INDUS-calculated HFEs also increases from  $\sim 17.9 k_B T$  to  $\sim 20.1 k_B T$ . Amide SAMs exhibit behavior between that of the amine and hydroxyl SAMs, with most INDUS-calculated displaying positive deviations from additive behavior but without plateauing at intermediate polar area fractions. The average envelope width increases from  $\sim 5.6 k_B T$  to  $\sim 7.2 k_B T$  while the maximum envelope width increases from  $\sim 12.7 k_B T$  to  $\sim 15.4 k_B T$  after active learning. Together, these data indicate that active learning identifies larger deviations from additivity for all three sets of SAMs while confirming that varying patterning at fixed polar area fraction substantially impacts HFEs, even for amine SAMs that only moderately deviate from additive behavior.

### **The trained GPR model identifies patterns with large deviations from additive behavior**

A chief benefit of the GPR model is the ability to label all possible SAM patterns, permitting analysis of the distribution of GPR-predicted HFEs as a function of polar area fraction. We quantified this distribution by binning the number of patterns with respect to the polar area fraction (in increments of 0.04) and GPR-predicted HFE (in increments of 0.5). Each bin was normalized by the largest number of patterns in any bin for the same polar area fraction to obtain the relative probability of obtaining a HFE for any given polar area fraction. Figure 4 plots the distribution of GPR-predicted HFEs for the three end groups as a function of the polar area fraction. For the amine and amide SAMs, most patterns have HFEs clustered near the linear interpolation line. The inset in Figure 4 shows the relative probability of patterns for an amine SAM with a polar area fraction of 0.38, indicating that most patterns also have HFEs clustered near the average HFE. Together, these features indicate that randomly sampled patterns for these SAMs would likely exhibit additive behavior, emphasizing the need for active learning to identify patterns with large deviations from ideality. Conversely, average HFEs for hydroxyl SAMs cluster near the linear



**Figure 4: Distribution of hydration free energies (HFE) predicted by Gaussian Process Regression.** Each plot shows the relative probability of HFEs as a function of polar area fraction (normalized per polar area fraction). Black crosses denote INDUS-calculated HFEs of seed SAMs and red plusses denote INDUS-calculated HFEs of patterns with the largest GPR-predicted HFEs after active learning. The inset of the left panel shows the relative probabilities of HFEs for amine SAMs with a polar area fraction of 0.38 (shown by the gray arrow). While the seed patterns cluster near the most probable HFE (black crosses), after active learning the GPR model correctly identifies low-probability patterns with the largest GPR-predicted HFEs (red plusses).

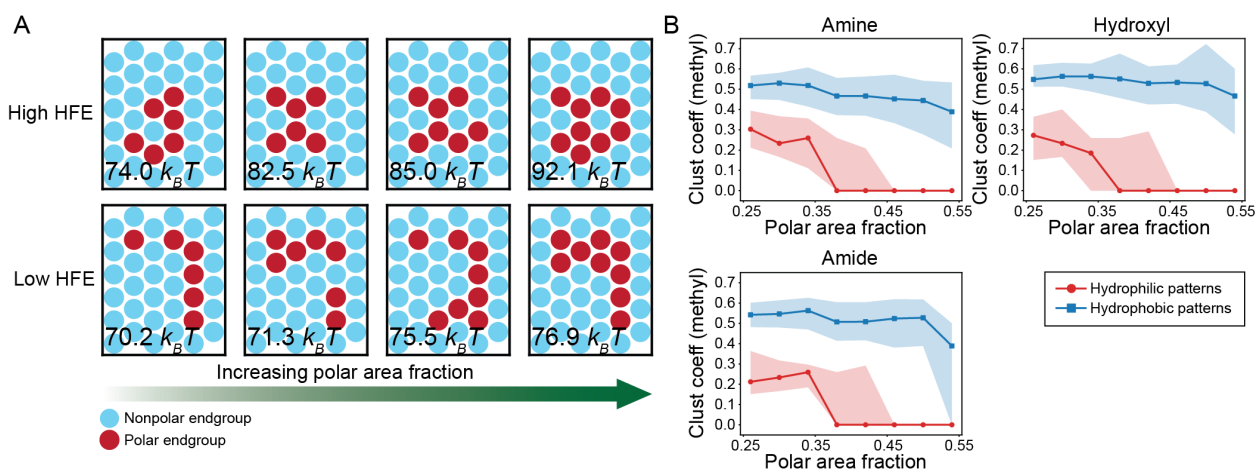
interpolation line at small polar area fractions but deviate from this line at larger polar area fractions, indicating that most patterns for SAMs with large numbers of hydroxyl groups would exhibit deviations from additive behavior (as noted in prior computational studies).<sup>35</sup>

We next sought to confirm the patterns lying on the extremes of the GPR-predicted HFE distributions with higher accuracy INDUS calculations. We selected four patterns—corresponding to the two largest and two smallest GPR-predicted HFEs—in each polar area fraction bin that contained at least 4000 patterns, resulting in 32 patterns per polar group. GPR-predicted and INDUS-calculated HFEs are highly correlated for these patterns with a Pearson’s  $r$  of  $\sim 0.9$  for amine SAMs and  $\sim 0.8$  for hydroxyl and amide SAMs. SI Figure S8 shows INDUS-calculated HFEs superimposed on the final GPR-predicted HFE envelope for the selected SAMs. Figure 4 shows that patterns predicted by the GPR model to have the highest positive deviations from ideality have INDUS-calculated HFEs (denoted by red plusses) that consistently lie on the upper edge of the envelope in regions where the relative probability of finding patterns is extremely low, with outliers noted only for the amide SAMs. On average, these patterns have INDUS-calculated HFEs that deviate by  $\sim 3.0 k_B T$  more from additivity than the set of seed patterns (shown as black crosses in Figure 4). These results indicate that the GPR model reliably identifies patterns with large deviations from additivity in addition to those patterns chosen for the active learning slow

loop, enabling the study of pattern features that distinguish hydrophobicity utilizing a much larger number of patterns than in previous studies.<sup>12, 30, 35</sup>

### Clustering of nonpolar ligands distinguishes hydrophobic and hydrophilic patterns

Having identified large numbers of patterns with substantial deviations from additive behavior for all three sets of SAMs, we sought to identify features of these patterns that distinguish extremes of hydrophobicity. Figure 5A visualizes patterns with high and low HFEs at constant polar area fraction. Visually, patterns with high HFEs appear to have more densely clustered polar groups, whereas patterns with low HFEs appear to have a higher degree of “connectedness” between nonpolar groups. To quantify these observations, we defined a variety of clustering and spatial autocorrelation metrics that could be computed for patterns. Clustering metrics were computed using a graph-based approach; either polar or nonpolar ligands were defined as vertices and edges were defined between neighboring ligands. We calculated the edge connectivity, node connectivity, and the average clustering coefficient for distinct graphs of polar and nonpolar end groups (a total of 6 metrics for each pattern). Spatial autocorrelation metrics included Moran’s I



**Figure 5: Analysis of ligand clustering for hydrophobic and hydrophilic patterns.** (A) Representative high HFE (hydrophilic) and low HFE (hydrophobic) patterns for different hydroxyl polar area fractions. Each column is at fixed polar area fraction. Values noted in the inset for all patterns are INDUS-calculated HFEs. Patterns include an additional layer of nonpolar ligands for ease of visualization. (B) Median of the average clustering coefficient of nonpolar ligands computed for hydrophilic (blue circles) and hydrophobic (red squares) patterns. Shaded regions encompass all patterns between the first and third quartile of clustering coefficient values; only polar area fractions with >4000 patterns are included. Points with no shaded region have the same median, first quartile, and third quartile values. Hydrophobic patterns have more clustered nonpolar regions on average than hydrophilic patterns (seen as breaks in the polar ligand patterns in panel A).

and adjusted Geary's C. Patterns displaying extremes of hydrophobicity were defined based on their GPR-predicted HFEs: patterns with HFEs at least two standard deviations larger than the mean HFE for any given polar area fraction (binned following the same approach as in Figure 4) were defined as hydrophilic for this analysis, whereas patterns with HFEs at least two standard deviations smaller than the mean HFE were defined as hydrophobic. We only included polar area fraction bins with more than 4000 patterns to ensure statistical significance. We then used a binary support vector machine (SVM) classifier with an  $l_1$  weight penalty to select the metrics that can accurately classify a pattern as hydrophobic or hydrophilic. More details on the calculation of these metrics and the SVM classifier can be found in the Supplementary Information.

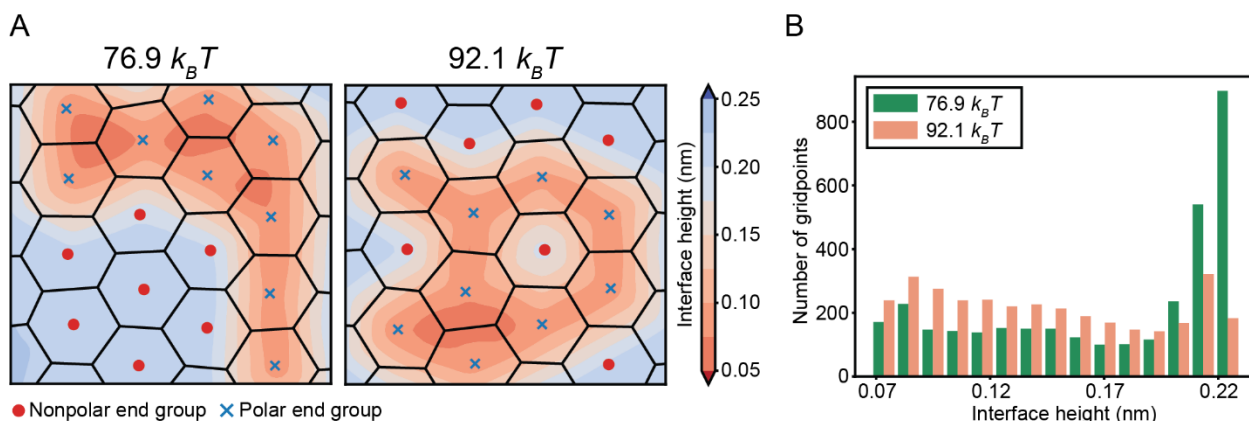
The SVM classifier identified the average clustering coefficient of nonpolar ligands as the most important metric for distinguishing hydrophobic and hydrophilic patterns. Figure 5B shows that there is a distinct difference between the average clustering coefficient of nonpolar ligands computed for hydrophobic and hydrophilic patterns for all end groups and polar area fractions, with the difference most pronounced as the polar area fraction reaches an intermediate value (consistent with the widest HFE envelopes in Figure 3). The impact of clustered nonpolar group on hydrophobicity has been demonstrated in previous computational studies of patterned chemically heterogeneous surfaces<sup>12, 32, 35, 36, 42</sup> and in experimental studies of  $\beta$ -peptides with extended nonpolar domains,<sup>11, 26</sup> but to our knowledge has not been demonstrated quantitatively for a large number of patterns and different polar end groups. The other clustering and spatial autocorrelation metrics do not distinguish between patterns with large variations in HFEs, emphasizing the importance of clustered nonpolar domains over other possible pattern features. This finding underscores the importance of using a quantitative, data-driven approach with enough sampling to find measurable differences in factors affecting hydrophobicity.

### **Clustering of nonpolar ligands impacts the pinning of water molecules**

To explain why the clustering coefficient of nonpolar ligands distinguishes hydrophobic and hydrophilic patterns, we hypothesized that this metric reports on the pinning of the SAM-water interface. Pinning refers to the strong attraction of water molecules to hydrophilic regions of the surface, which reduces water density fluctuations and decreases the likelihood of dewetting (thus contributing to a smaller value of  $P_v(0)$  and larger HFE).<sup>12, 33</sup> Because water molecules are highly correlated due to strong water-water interactions, pinning impacts water density fluctuations across

multiple molecular lengths and is influenced by the patterning of chemically heterogeneous surfaces. Pinning can be quantified by using MD simulations to analyze spatial variations in the average height of the SAM-water interface, which tends to decrease near hydrophilic regions due to pinning.<sup>12, 42</sup> In our recent study of surfaces with hand-selected patterns, we found that SAMs with uniformly distributed polar and nonpolar groups more tightly pinned water molecules (leading to lower interfacial heights) and appeared more hydrophilic than SAMs with large nonpolar domains.<sup>42</sup> We thus sought to investigate if the clustering coefficient of nonpolar ligands similarly captures a relationship between patterning and pinning for SAMs with patterns that do not fit these hand-selected extremes. We characterized pinning by performing MD simulations of representative SAMs and calculating the SAM-water interface height, which we defined as the difference between the height of a constant water density isosurface (described in the Supplementary Information) and the height of the center of mass of ligand end groups.

Figure 6 shows interface height contours for a representative pair of hydroxyl SAMs with the same polar area fraction but with INDUS-calculated HFEs differing by  $\sim 15 k_B T$ . The average clustering coefficient of nonpolar ligands is 0.63 for the more hydrophobic pattern (HFE of  $76.9 k_B T$ ) and 0 for the more hydrophilic pattern (HFE of  $92.1 k_B T$ ). A larger region of loose pinning



**Figure 6: Differences in interface height for a representative pair of patterns. (A)** Contours of SAM-water interface height for hydroxyl SAMs with the same polar area fraction (0.47) but HFEs differing by  $\sim 15 k_B T$ . A larger region of tight pinning (denoted by the contours in shades of red) is apparent for the hydrophilic pattern ( $92.1 k_B T$ ) compared to the hydrophobic pattern ( $76.9 k_B T$ ). **(B)** Histogram of SAM-water interface heights. The area covered by the cavity is discretized into a grid and the average SAM-water interface height is calculated at each grid point. The hydrophobic pattern has a larger proportion of loosely pinned areas (height  $> 0.21$  nm) whereas the hydrophilic pattern has a larger proportion of tightly pinned areas (height  $< 0.15$  nm).

(*i.e.*, a region with a high interface height) is apparent for the hydrophobic pattern. In the hydrophilic pattern, the interface height is reduced near the nonpolar ligand that is bordered by polar ligands, reflecting the perturbation of water structure due these ligands. Figure 6B quantifies these differences by plotting the distribution of interface heights for both patterns. In agreement with prior studies of hand-selected patterns, there is a greater prevalence of larger areas of loose pinning (interface height  $\sim 0.21$  nm) for the hydrophobic pattern and tight pinning (interface height  $< 0.15$  nm) the hydrophilic pattern. Similar distributions were identified for other patterns (examples shown in SI Figure S11), pointing to a general relationship between pinning and the clustering of nonpolar ligands. These findings illustrate the non-obvious changes to interfacial water structure that can be related to patterns identified by the GPR model without being selected by hand.

## Discussion

We investigated nonadditive contributions to the hydrophobicity of chemically heterogeneous surfaces by training a GPR model that relates patterning to HFEs for SAMs with three different polar end groups (amine, hydroxyl, and amide). To efficiently train the GPR model, we developed a dual-loop active learning scheme that leverages the strengths of two simulation methods to label HFEs: a fast method with lower accuracy (a pre-trained CNN) that is used to explore pattern space and a slow method with high accuracy (INDUS) that is used to identify patterns with large deviations from additive behavior. The active learning scheme enables the GPR model to map 65536 patterns to HFEs by training on simulation data for only  $\sim 2\%$  of the patterns (labeled with 82 INDUS and 1448 unbiased simulations). GPR-predicted HFEs reveal substantial differences in the impact of patterning and composition on hydrophobicity for SAMs with different polar groups, emphasizing the need for inclusion of diverse polar groups in studies of chemically heterogeneous surfaces. The trained GPR model also consistently identifies patterns with HFEs that substantially deviate from additive behavior. Analysis of pattern features further shows that a metric quantifying the clustering of nonpolar ligands distinguishes more hydrophobic from more hydrophilic patterns at constant polar area fraction, which we attribute to differences in the local pinning of the SAM-water interface. To our knowledge, this is the first analysis of a large number (thousands) of patterns to quantitatively demonstrate the effect of clustering on the hydrophobicity of diverse chemically heterogeneous SAMs, enabled by the new active learning approach.

The success of the dual-loop active learning scheme suggests several possibilities for future investigation. To correct the tendency of the GPR model to underpredict the lower bounds of possible HFEs, future work will investigate machine learning models that exploit more informative data representations (*e.g.*, hydrogen bond graphs) or alternative deep learning architectures (*e.g.*, recurrent neural networks) to improve the accuracy of HFE predictions during the fast loop. We also envision extending the active learning scheme to train a single GPR model that maps both polar group chemistry (encoded as a molecular fingerprint, for example) and patterning to HFEs, permitting efficient investigation of a broader range of polar groups. Finally, the dual-loop active learning scheme could be modified to include experimental measurements during the slow loop, thereby guiding the design of experiments for the targeted exploration of systems expected to exhibit deviations from additive behavior.

## Methods

### Molecular dynamics simulations

Each SAM was constructed by positioning 64 ligands in the  $x$ - $y$  plane with a grafting density of  $21.6 \text{ \AA}^2/\text{ligand}$  to mimic self-assembly onto a gold (111) lattice.<sup>47</sup> Gold atoms were not modeled because prior work showed that removing the gold substrate leads to SAM properties in better agreement with experiments without influencing water structure.<sup>31</sup> Ligands were oriented with the end groups pointing in the positive  $z$ -direction. A 5-nm thick water layer was placed above the SAM and in contact with the ligand end groups. A 3-nm thick buffering vacuum layer was added above the water layer. Periodic boundary conditions were applied in all dimensions. Ligands were modeled using the CHARMM36 force field<sup>48</sup> with the TIP4P/2005 water model.<sup>49</sup> Electrostatic interactions were calculated using the smooth Particle Mesh Ewald algorithm<sup>50</sup> with short-range Coulomb, van der Waals, and neighbor list cutoffs set to 1.2 nm. Because gold-sulfur bonds were omitted in the simulations, harmonic restraints with a spring constant of  $50,000 \text{ kJ/mol/nm}^2$  were applied to the sulfur atoms to maintain SAM structure. All MD simulations were performed in the  $NVT$  ensemble using *Gromacs 2016* with a 2-fs timestep. The temperature was maintained at 300 K using a velocity rescaling thermostat with a temperature-coupling time of 0.1 ps.<sup>51</sup>



### **Seed patterns and hydration free energy calculations**

The set of seed patterns included patterns generated from 50 random binary matrices and two patterns corresponding to all polar and all nonpolar groups. The same seed patterns were used for all polar end groups. Seed patterns and patterns chosen for the slow loop during active learning were labeled with HFEs using INDUS. INDUS samples  $P_\nu(0)$  (Equation 1) by biasing the number of water molecules in the cavity  $\nu$  because  $P_\nu(0)$  is too small to be sampled during unbiased simulations if the cavity is large.<sup>27, 28</sup> We performed INDUS using *GROMACS 2016.6*<sup>52</sup> patched with the PLUMED 2.5.1 plugin.<sup>53</sup> Each SAM was equilibrated for 5 ns, INDUS was performed using 13 independent, 5-ns simulation windows, and the unbiased value of  $P_\nu(0)$  was obtained using the Weighted Histogram Analysis Method.<sup>54</sup> Details are included in the Supplementary Information. The error associated with HFEs calculated using INDUS is  $\sim 2 k_B T$ .<sup>32, 42</sup>

Patterns chosen for the fast loop were labeled using a CNN developed in our prior work.<sup>32</sup> The CNN was previously trained on a set of homogeneous SAMs<sup>36</sup> containing amine, amide, and hydroxyl ligands with scaled partial charges. Input to the CNN consists of a short unbiased MD simulation ( $\sim 1$  ns) converted into a sequence of atomic density histograms. The CNN predicts HFEs that are highly correlated with INDUS-calculated HFEs and accurate up to a multiplicative constant. The error associated with HFEs predicted by the CNN is  $\sim 6 k_B T$ . CNN details are included in the Supplementary Information.

### **Gaussian Process Regression to map patterns to HFE labels**

GPR is a non-parametric Bayesian model that estimates the mean and uncertainty of unlabeled test data based on knowledge of labeled training data under the assumption that the joint distribution of training and test data is Gaussian.<sup>55</sup> In our approach, input data are patterns encoded as  $4 \times 4$  binary matrices and labels are HFEs. We define elements of the symmetric covariance matrix  $\mathbf{K}$  using a radial basis kernel function  $k(\mathbf{x}_i, \mathbf{x}_j)$ :<sup>56</sup>

$$k(\mathbf{x}_i, \mathbf{x}_j) = \exp\left(-\frac{1}{2\gamma^2} \|\mathbf{x}_i - \mathbf{x}_j\|^2\right) \quad (2)$$

$\mathbf{x}_i$  and  $\mathbf{x}_j$  denote individual patterns and  $\gamma = \sqrt{2}$  defines the characteristic length scale in pattern space. For any pattern  $\mathbf{x}^*$  not in the training data, estimates for the mean ( $\mu_{f^*}$ ) and uncertainty ( $\sigma_{f^*}$ ) of the HFE for that pattern are:

$$\mu_f^* = \mathbf{k}^T (\mathbf{K} + \sigma_{\text{rep}}^2 \mathbf{I})^{-1} \mathbf{y} \quad (3)$$

$$\sigma_f^* = k(\mathbf{x}^*, \mathbf{x}^*) - \mathbf{k}^T (\mathbf{K} + \sigma_{\text{rep}}^2 \mathbf{I})^{-1} \mathbf{k} \quad (4)$$

$\mathbf{y}$  is a vector of HFE labels (obtained from either INDUS or the CNN) for the  $t$  training patterns,  $\sigma_{\text{rep}}$  is a vector of error values for the training pattern HFEs (equal to  $2 k_B T$  for HFEs from INDUS and  $6 k_B T$  for HFEs from the CNN), and  $\mathbf{k} = [k(\mathbf{x}_1, \mathbf{x}^*), k(\mathbf{x}_2, \mathbf{x}^*), \dots, k(\mathbf{x}_t, \mathbf{x}^*)]^T$  is a vector of kernel distances between each training pattern  $\mathbf{x}_i$  and  $\mathbf{x}^*$ . Additional details are included in the Supplementary Information.

### Acquisition function for pattern selection

Active learning requires an acquisition function,  $u(\mathbf{x}^*)$ , which is maximized to choose a new pattern,  $\mathbf{x}_{t+1}$ , to label and add to the  $t$  training patterns. We used the ‘‘Expectation of Improvement’’ acquisition function to balance the tradeoff between choosing patterns with high GPR-predicted uncertainties (exploration) and choosing patterns with large deviations from additive behavior (exploitation).<sup>56, 57</sup> To define this tradeoff, we first define the expected additive HFE for a given pattern,  $\text{HFE}_{\text{add}}$ , as the average of INDUS-calculated HFEs for patterns with only polar ( $\text{HFE}_{\text{polar}}$ ) and only nonpolar ( $\text{HFE}_{\text{nonpolar}}$ ) ligands weighted by the mole fraction of polar ligands in the pattern:

$$\text{HFE}_{\text{add}}(\mathbf{x}^*) = \text{HFE}_{\text{polar}} \times \chi_{\mathbf{x}^*} + \text{HFE}_{\text{nonpolar}} \times (1 - \chi_{\mathbf{x}^*}) \quad (5)$$

$\chi_{\mathbf{x}^*}$  is the mole fraction of polar ligands in the pattern. Deviation from the expected additive HFE is quantified by:

$$\mu_{\Delta}(\mathbf{x}^*) = |\text{HFE}_{\text{GPR}}(\mathbf{x}^*) - \text{HFE}_{\text{add}}(\mathbf{x}^*)| \quad (6)$$

$\text{HFE}_{\text{GPR}}(\mathbf{x}^*)$  is the GPR-predicted HFE for a pattern  $\mathbf{x}^*$ . The acquisition function is then:

$$u(\mathbf{x}^*) = \begin{cases} (\mu_{\Delta}(\mathbf{x}^*) - \mu_{\Delta}^+ - \xi) \Phi(Z) + \sigma_f(\mathbf{x}^*) \phi(Z) & \sigma_f(\mathbf{x}^*) \neq 0 \\ 0 & \sigma_f(\mathbf{x}^*) = 0 \end{cases} \quad (7)$$

$$Z = \begin{cases} \frac{\mu_{\Delta}(\mathbf{x}^*) - \mu_{\Delta}^+ - \xi}{\sigma_f(\mathbf{x}^*)} & \sigma_f(\mathbf{x}^*) \neq 0 \\ 0 & \sigma_f(\mathbf{x}^*) = 0 \end{cases} \quad (8)$$

$\mu_{\Delta}^{\dagger}$  is the maximum deviation among the set of training points,  $\phi(Z)$  and  $\Phi(Z)$  are the probability density function and cumulative density function of a standard Gaussian, respectively, and  $\xi$  is a parameter that controls the exploration-exploitation tradeoff. We chose  $\xi = 0.01$  in accordance with past experiments.<sup>57</sup> During the fast loop,  $u(\mathbf{x}^*)$  is computed for all possible SAM patterns and the pattern which maximizes  $u(\mathbf{x}^*)$  is chosen as the next pattern to label.

## Data Availability

All data and code to reproduce this work are available at <https://gitlab.com/atharva-kelkar/dual-loop-active-learning>.

## Acknowledgments

This work was supported by the National Science Foundation under Grant No. 2044997. This work used the Extreme Science and Engineering Discovery Environment (XSEDE), which is supported by National Science Foundation grant number ACI-1548562. R. C. V. also acknowledges support from the 3M Non-Tenured Faculty Award.

## References

1. Friesner, R. A.; Murphy, R. B.; Repasky, M. P.; Frye, L. L.; Greenwood, J. R.; Halgren, T. A.; Sanschagrin, P. C.; Mainz, D. T., Extra precision glide: docking and scoring incorporating a model of hydrophobic enclosure for protein-ligand complexes. *J Med Chem* **2006**, *49* (21), 6177-96.
2. Nicholls, A.; Sharp, K. A.; Honig, B., Protein folding and association: insights from the interfacial and thermodynamic properties of hydrocarbons. *Proteins* **1991**, *11* (4), 281-96.
3. Snyder, P. W.; Mecinovic, J.; Moustakas, D. T.; Thomas, S. W., 3rd; Harder, M.; Mack, E. T.; Lockett, M. R.; Heroux, A.; Sherman, W.; Whitesides, G. M., Mechanism of the hydrophobic effect in the biomolecular recognition of arylsulfonamides by carbonic anhydrase. *Proc Natl Acad Sci U S A* **2011**, *108* (44), 17889-94.
4. Dill, K. A., Dominant forces in protein folding. *Biochemistry* **1990**, *29* (31), 7133-55.
5. Dobson, C. M., Protein folding and misfolding. *Nature* **2003**, *426* (6968), 884-90.
6. Van Lehn, R. C.; Alexander-Katz, A., Ligand-mediated short-range attraction drives aggregation of charged monolayer-protected gold nanoparticles. *Langmuir* **2013**, *29* (28), 8788-98.
7. Nakanishi, K.; Sakiyama, T.; Imamura, K., On the adsorption of proteins on solid surfaces, a common but very complicated phenomenon. *J Biosci Bioeng* **2001**, *91* (3), 233-44.
8. Van Lehn, R. C.; Atukorale, P. U.; Carney, R. P.; Yang, Y. S.; Stellacci, F.; Irvine, D. J.; Alexander-Katz, A., Effect of particle diameter and surface composition on the spontaneous

- fusion of monolayer-protected gold nanoparticles with lipid bilayers. *Nano Lett* **2013**, *13* (9), 4060-7.
9. Van Lehn, R. C.; Alexander-Katz, A., Pathway for insertion of amphiphilic nanoparticles into defect-free lipid bilayers from atomistic molecular dynamics simulations. *Soft Matter* **2015**, *11* (16), 3165-3175.
  10. Chandler, D., Interfaces and the driving force of hydrophobic assembly. *Nature* **2005**, *437* (7059), 640-7.
  11. Wang, C.; Ma, C. D.; Yeon, H.; Wang, X.; Gellman, S. H.; Abbott, N. L., Nonadditive Interactions Mediated by Water at Chemically Heterogeneous Surfaces: Nonionic Polar Groups and Hydrophobic Interactions. *J Am Chem Soc* **2017**, *139* (51), 18536-18544.
  12. Xi, E.; Venkateshwaran, V.; Li, L.; Rego, N.; Patel, A. J.; Garde, S., Hydrophobicity of proteins and nanostructured solutes is governed by topographical and chemical context. *Proc Natl Acad Sci U S A* **2017**, *114* (51), 13345-13350.
  13. Cassie, A. B. D.; Baxter, S., Wettability of porous surfaces. *T Faraday Soc* **1944**, *40*, 0546-0550.
  14. Israelachvili, J. N.; Gee, M. L., Contact Angles on Chemically Heterogeneous Surfaces. *Langmuir* **1989**, *5* (1), 288-289.
  15. Reynolds, J. A.; Gilbert, D. B.; Tanford, C., Empirical Correlation between Hydrophobic Free-Energy and Aqueous Cavity Surface-Area. *Proceedings of the National Academy of Sciences of the United States of America* **1974**, *71* (8), 2925-2927.
  16. Abraham, M. H.; Whiting, G. S.; Fuchs, R.; Chambers, E. J., Thermodynamics of Solute Transfer from Water to Hexadecane. *Journal of the Chemical Society-Perkin Transactions 2* **1990**, (2), 291-300.
  17. Stephenson, B. C.; Goldsipe, A.; Beers, K. J.; Blankschtein, D., Quantifying the hydrophobic effect. 1. A computer simulation-molecular-thermodynamic model for the self-assembly of hydrophobic and amphiphilic solutes in aqueous solution. *Journal of Physical Chemistry B* **2007**, *111* (5), 1025-1044.
  18. Jiang, L. G.; Cao, S. Q.; Cheung, P. P. H.; Zheng, X. Y.; Leung, C. W. T.; Peng, Q.; Shuai, Z. G.; Tang, B. Z.; Yao, S. H.; Huang, X. H., Real-time monitoring of hydrophobic aggregation reveals a critical role of cooperativity in hydrophobic effect. *Nature Communications* **2017**, *8*.
  19. Chen, J.; Brooks, C. L., Implicit modeling of nonpolar solvation for simulating protein folding and conformational transitions. *Phys Chem Chem Phys* **2008**, *10* (4), 471-481.
  20. MacCallum, J. L.; Tieleman, D. P., Hydrophobicity scales: a thermodynamic looking glass into lipid-protein interactions. *Trends Biochem Sci* **2011**, *36* (12), 653-662.
  21. Moyano, D. F.; Saha, K.; Prakash, G.; Yan, B.; Kong, H.; Yazdani, M.; Rotello, V. M., Fabrication of Corona-Free Nanoparticles with Tunable Hydrophobicity. *ACS Nano* **2014**, *8* (7), 6748-6755.
  22. Li, X. N.; Robinson, S. M.; Gupta, A.; Saha, K.; Jiang, Z. W.; Moyano, D. F.; Sahar, A.; Riley, M. A.; Rotello, V. M., Functional Gold Nanoparticles as Potent Antimicrobial Agents against Multi-Drug-Resistant Bacteria. *Acs Nano* **2014**, *8* (10), 10682-10686.
  23. Moyano, D. F.; Goldsmith, M.; Solfiell, D. J.; Landesman-Milo, D.; Miranda, O. R.; Peer, D.; Rotello, V. M., Nanoparticle Hydrophobicity Dictates Immune Response. *J Am Chem Soc* **2012**, *134* (9), 3965-3967.

24. Chen, K. M.; Rana, S.; Moyano, D. F.; Xu, Y. S.; Guo, X. H.; Rotello, V. M., Optimizing the selective recognition of protein isoforms through tuning of nanoparticle hydrophobicity. *Nanoscale* **2014**, *6* (12), 6492-6495.
25. Woodward, J. T.; Gwin, H.; Schwartz, D. K., Contact angles on surfaces with mesoscopic chemical heterogeneity. *Langmuir* **2000**, *16* (6), 2957-2961.
26. Ma, C. D.; Wang, C.; Acevedo-Velez, C.; Gellman, S. H.; Abbott, N. L., Modulation of hydrophobic interactions by proximally immobilized ions. *Nature* **2015**, *517* (7534), 347-50.
27. Patel, A. J.; Varilly, P.; Chandler, D., Fluctuations of water near extended hydrophobic and hydrophilic surfaces. *J Phys Chem B* **2010**, *114* (4), 1632-7.
28. Patel, A. J.; Varilly, P.; Chandler, D.; Garde, S., Quantifying density fluctuations in volumes of all shapes and sizes using indirect umbrella sampling. *J Stat Phys* **2011**, *145* (2), 265-275.
29. Patel, A. J.; Varilly, P.; Jamadagni, S. N.; Hagan, M. F.; Chandler, D.; Garde, S., Sitting at the edge: how biomolecules use hydrophobicity to tune their interactions and function. *J Phys Chem B* **2012**, *116* (8), 2498-503.
30. Wang, J.; Bratko, D.; Luzar, A., Probing surface tension additivity on chemically heterogeneous surfaces by a molecular approach. *Proc Natl Acad Sci U S A* **2011**, *108* (16), 6374-9.
31. Dallin, B. C.; Yeon, H.; Ostwalt, A. R.; Abbott, N. L.; Van Lehn, R. C., Molecular Order Affects Interfacial Water Structure and Temperature-Dependent Hydrophobic Interactions between Nonpolar Self-Assembled Monolayers. *Langmuir* **2019**, *35* (6), 2078-2088.
32. Kelkar, A. S.; Dallin, B. C.; Van Lehn, R. C., Predicting Hydrophobicity by Learning Spatiotemporal Features of Interfacial Water Structure: Combining Molecular Dynamics Simulations with Convolutional Neural Networks. *J Phys Chem B* **2020**, *124* (41), 9103-9114.
33. Dallin, B. C.; Van Lehn, R. C., Spatially Heterogeneous Water Properties at Disordered Surfaces Decrease the Hydrophobicity of Nonpolar Self-Assembled Monolayers. *J Phys Chem Lett* **2019**, *10* (14), 3991-3997.
34. Wang, J.; Bratko, D.; Luzar, A., Probing surface tension additivity on chemically heterogeneous surfaces by a molecular approach. *Proc Natl Acad Sci U S A* **2011**, *108* (16), 6374-6379.
35. Monroe, J. I.; Shell, M. S., Computational discovery of chemically patterned surfaces that effect unique hydration water dynamics. *Proc Natl Acad Sci U S A* **2018**, *115* (32), 8093-8098.
36. Kanduc, M.; Schlaich, A.; Schneck, E.; Netz, R. R., Water-Mediated Interactions between Hydrophilic and Hydrophobic Surfaces. *Langmuir* **2016**, *32* (35), 8767-82.
37. Giovambattista, N.; Debenedetti, P. G.; Rossky, P. J., Hydration Behavior under Confinement by Nanoscale Surfaces with Patterned Hydrophobicity and Hydrophilicity. *The Journal of Physical Chemistry C* **2007**, *111* (3), 1323-1332.
38. Settles, B., Active learning literature survey. **2009**.
39. Shmilovich, K.; Mansbach, R. A.; Sidky, H.; Dunne, O. E.; Panda, S. S.; Tovar, J. D.; Ferguson, A. L., Discovery of Self-Assembling pi-Conjugated Peptides by Active Learning-Directed Coarse-Grained Molecular Simulation. *J Phys Chem B* **2020**, *124* (19), 3873-3891.
40. Kusne, A. G.; Yu, H.; Wu, C.; Zhang, H.; Hattrick-Simpers, J.; DeCost, B.; Sarker, S.; Oses, C.; Toher, C.; Curtarolo, S.; Davydov, A. V.; Agarwal, R.; Bendersky, L. A.; Li, M.; Mehta, A.; Takeuchi, I., On-the-fly closed-loop materials discovery via Bayesian active learning. *Nat Commun* **2020**, *11* (1), 5966.

41. Jablonka, K. M.; Jothiappan, G. M.; Wang, S.; Smit, B.; Yoo, B., Bias free multiobjective active learning for materials design and discovery. *Nature Communications* **2021**, *12* (1), 2312.
42. Dallin, B. C.; Kelkar, A. S.; Van Lehn, R. C., Generalizable Features of Interfacial Water Structure Predict the Hydrophobicity of Chemically Heterogeneous Surfaces. **2021**.
43. Godawat, R.; Jamadagni, S. N.; Garde, S., Characterizing hydrophobicity of interfaces by using cavity formation, solute binding, and water correlations. *Proc Natl Acad Sci U S A* **2009**, *106* (36), 15119-24.
44. Chew, A. K.; Dallin, B. C.; Van Lehn, R. C., The Interplay of Ligand Properties and Core Size Dictates the Hydrophobicity of Monolayer-Protected Gold Nanoparticles. *ACS Nano* **2021**, *15* (3), 4534-4545.
45. Willard, A. P.; Chandler, D., Instantaneous liquid interfaces. *J Phys Chem B* **2010**, *114* (5), 1954-8.
46. Bloodgood, M.; Vijay-Shanker, K., A method for stopping active learning based on stabilizing predictions and the need for user-adjustable stopping. *ArXiv preprint* **2014**.
47. Love, J. C.; Estroff, L. A.; Kriebel, J. K.; Nuzzo, R. G.; Whitesides, G. M., Self-assembled monolayers of thiolates on metals as a form of nanotechnology. *Chem Rev* **2005**, *105* (4), 1103-69.
48. Vanommeslaeghe, K.; Hatcher, E.; Acharya, C.; Kundu, S.; Zhong, S.; Shim, J.; Darian, E.; Guvench, O.; Lopes, P.; Vorobyov, I.; Mackerell, A. D., Jr., CHARMM general force field: A force field for drug-like molecules compatible with the CHARMM all-atom additive biological force fields. *J Comput Chem* **2010**, *31* (4), 671-90.
49. Abascal, J. L.; Vega, C., A general purpose model for the condensed phases of water: TIP4P/2005. *J Chem Phys* **2005**, *123* (23), 234505.
50. Essmann, U.; Perera, L.; Berkowitz, M. L.; Darden, T.; Lee, H.; Pedersen, L. G., A Smooth Particle Mesh Ewald Method. *Journal of Chemical Physics* **1995**, *103* (19), 8577-8593.
51. Bussi, G.; Donadio, D.; Parrinello, M., Canonical sampling through velocity rescaling. *The Journal of chemical physics* **2007**, *126* (1), 014101.
52. Abraham, M. J.; Murtola, T.; Schulz, R.; Páll, S.; Smith, J. C.; Hess, B.; Lindahl, E. J. S., GROMACS: High performance molecular simulations through multi-level parallelism from laptops to supercomputers. **2015**, *1*, 19-25.
53. Tribello, G. A.; Bonomi, M.; Branduardi, D.; Camilloni, C.; Bussi, G., PLUMED 2: New feathers for an old bird. *Computer Physics Communications* **2014**, *185* (2), 604-613.
54. Kumar, S.; Rosenberg, J. M.; Bouzida, D.; Swendsen, R. H.; Kollman, P. A., The weighted histogram analysis method for free-energy calculations on biomolecules. I. The method. *Journal of computational chemistry* **1992**, *13* (8), 1011-1021.
55. Williams, C. K.; Rasmussen, C. E., Gaussian processes for regression. **1996**.
56. Brochu, E.; Cora, V. M.; De Freitas, N., A tutorial on Bayesian optimization of expensive cost functions, with application to active user modeling and hierarchical reinforcement learning. *arXiv preprint* **2010**.
57. Lizotte, D. J., *Practical bayesian optimization*. University of Alberta: 2008.