



Learning the relationship between nanoscale chemical patterning and hydrophobicity

Nicholas B. Rego^a , Andrew L. Ferguson^{b,1} , and Amish J. Patel^{a,1}

Edited by B. M. Pettitt, University of Texas Medical Branch, Galveston, TX; received January 5, 2022; accepted October 19, 2022 by Editorial Board Member Peter J. Rossky

The hydrophobicity of proteins and similar surfaces, which display chemical heterogeneity at the nanoscale, drives countless aqueous interactions and assemblies. However, predicting how surface chemical patterning influences hydrophobicity remains a challenge. Here, we address this challenge by using molecular simulations and machine learning to characterize and model the hydrophobicity of a diverse library of patterned surfaces, spanning a wide range of sizes, shapes, and chemical compositions. We find that simple models, based only on polar content, are inaccurate, whereas complex neural network models are accurate but challenging to interpret. However, by systematically incorporating chemical correlations between surface groups into our models, we are able to construct a series of minimal models of hydrophobicity, which are both accurate and interpretable. Our models highlight that the number of proximal polar groups is a key determinant of hydrophobicity and that polar neighbors enhance hydrophobicity. Although our minimal models are trained on particular patch size and shape, their interpretability enables us to generalize them to rectangular patches of all shapes and sizes. We also demonstrate how our models can be used to predict hot-spot locations with the largest marginal contributions to hydrophobicity and to design chemical patterns that have a fixed polar content but vary widely in their hydrophobicity. Our data-driven models and the principles they furnish for modulating hydrophobicity could facilitate the design of novel materials and engineered proteins with stronger interactions or enhanced solubilities.

hydrophobicity | chemical correlations | self-assembled monolayer | machine learning | interpretability

Hydrophobicity—the aversion of nonpolar solutes to water—drives diverse molecular assemblies, ranging from micelle formation and protein folding to supramolecular chemistry and biomolecular interactions (1–7). The hydrophobicity of homogeneous surfaces is determined primarily by their polarity (8–10). In contrast, the hydrophobicity of proteins and similar surfaces, which display chemical heterogeneity at the nanoscale, depends not just on their polar content but also on the arrangement of their polar and nonpolar groups (11–17). Consequently, additive approaches for characterizing hydrophobicity such as hydrophathy scales (18, 19), which sum over the contributions from individual chemical groups, fail to accurately quantify the hydrophobicity of heterogeneous surfaces (20–24).

Recent theoretical developments (25–27), coupled with advances in specialized molecular simulation methods (28, 29), have made it possible to rigorously characterize the hydrophobicity of surfaces with nanoscale chemical patterns (30, 31). In particular, a number of studies have shown that the ease with which water can be displaced from the vicinity of a surface to form interfacial cavities is a robust molecular measure of its hydrophobicity (32–34). Moreover, these studies have highlighted that the response of interfacial water structure to surface chemical patterning is complex and collective (15, 35), obfuscating the relationship between patterning and hydrophobicity and making it challenging to develop predictive models for the hydrophobicity of heterogeneous surfaces.

Machine learning (ML) approaches (36–38) offer the prospect of exposing the emergent many-body determinants of hydrophobicity in a data-driven fashion (16, 39). Indeed, such approaches have recently been applied to uncover patterned surfaces with optimal interfacial properties, such as water diffusivity or affinity for small hydrophobic solutes (40, 41). However, data-driven models are typically subject to the accuracy interpretability trade-off—sophisticated “black box” models capable of accurately learning high-dimensional functional mappings can resist interpretation, whereas simple “glass box” models amenable to intuitive understanding can be insufficiently expressive, rendering them inaccurate. Models that are both accurate and interpretable are desirable because they can justify their predictions, divulge the limits of their applicability, and are amenable to

Significance

Diverse biomolecular recognition and self-assembly processes are driven by hydrophobic interactions between surfaces that display nanoscale chemical heterogeneity. However, the relationship between chemical patterning and hydrophobicity is nontrivial and cannot be captured by additive approaches, such as hydrophathy scales. Here, we combine molecular simulations and machine learning to learn this relationship and develop predictive models of hydrophobicity that are accurate, interpretable, and generalizable. The learned models unveil new insights into the chemical determinants of hydrophobicity and highlight the importance of accounting for chemical correlations between surface groups in determining hydrophobicity. Our models could spur the rational design of soft materials and biomolecules with optimal hydrophobicity/hydrophilicity.

Author contributions: N.B.R., A.L.F., and A.J.P. designed research; N.B.R. performed research; N.B.R. analyzed data; and N.B.R., A.L.F., and A.J.P. wrote the paper.

The authors declare no competing interest. A.L.F. is a co-founder and consultant of Evozyne, Inc. and a co-author of US Provisional Patents 62/853,919 and 62/900,420 and International Patent Applications PCT/US2020/035206 and PCT/US20/50466.

This article is a PNAS Direct Submission. B.M.P. is a guest editor invited by the Editorial Board.

Copyright © 2022 the Author(s). Published by PNAS. This article is distributed under [Creative Commons Attribution-NonCommercial-NoDerivatives License 4.0 \(CC BY-NC-ND\)](https://creativecommons.org/licenses/by-nc-nd/4.0/).

¹To whom correspondence should be addressed. Email: andrewferguson@uchicago.edu or amish.patel@seas.upenn.edu.

This article contains supporting information online at <https://www.pnas.org/lookup/suppl/doi:10.1073/pnas.2200018119/-/DCSupplemental>.

Published November 21, 2022.

generalization, which prompts the question: Is it possible to learn models relating chemical patterning and hydrophobicity that are at once accurate, interpretable, and generalizable, and if such models exist, what insights can they provide into the chemical determinants of hydrophobicity?

To address these questions, here we use specialized molecular simulations (28, 29) to characterize the hydrophobicity of numerous patterned surface patches, spanning a range of sizes, shapes, and chemical compositions and employ deep learning techniques (36, 38) to model the relationship between chemical patterning and hydrophobicity. We demonstrate that simple models, based solely on polar content, are poor predictors of hydrophobicity, whereas neural network models are highly accurate predictors. Although the latter fail to provide interpretable principles and actionable design precepts, they nevertheless highlight the importance of local chemical correlations in determining hydrophobicity.

To achieve a seamless trade-off between model accuracy and interpretability, we incorporate chemical correlations between surface groups and construct a series of increasingly complex models of hydrophobicity. Our models emulate additive models in their simplicity but rival the neural network models in their predictive performance. Importantly, our minimal models reveal simple but nonintuitive principles that advance our understanding of the chemical determinants of hydrophobicity. In particular, we find that in addition to the number of polar groups, the number of polar neighbors in a patch also influences its hydrophobicity, and that for patches with the same polar content, clustering the polar groups enhances hydrophobicity. Uncovering the predominant factors that govern the hydrophobicity of patterned patches also allows us to generalize our minimal models, which were trained on 3-nm \times 3-nm patches, to rectangular patches of all shapes and sizes.

Leveraging the efficiency of our models, we further find that patch hydrophobicity is most susceptible to a polar mutation introduced at the center of a nonpolar cluster, whereas patch hydrophilicity is particularly susceptible to nonpolar mutations at isolated polar groups or along the periphery of a polar cluster. We also find that for patches with a given chemical composition, the most hydrophilic patches display significant dispersion of polar and nonpolar groups, whereas the most hydrophobic patches feature a single nonpolar cluster. We hope that the accuracy and interpretability of our data-driven models, as well as the principles they furnish for optimally modulating hydrophobicity, will facilitate the rational design of novel materials with tailored hydrophobicity and engineered proteins with stronger interactions or enhanced solubilities (35, 42–44).

Hydrophobicity of Chemically Patterned Surfaces

To interrogate the relationship between chemical patterns displayed by heterogeneous surfaces and their hydrophobicity, we employ self-assembled monolayer (SAM) surfaces, which are versatile soft material templates with diverse applications ranging from fouling resistance to molecular electronics (45, 46). In particular, we study binary SAM surfaces with end groups that are either polar (hydroxyl) or nonpolar (methyl) (Fig. 1). Initially, we focus on a 6 \times 6 square patch (roughly 3 nm \times 3 nm), which contains 36 end groups, and is embedded in a background of polar end groups, as shown in Fig. 1*A*. Given that simulating all the $2^{36} \approx 7 \times 10^{10}$ possible patches is computationally infeasible, we construct a reasonably large training library of $N = 884$ patches (Fig. 1*B*). To promote the diversity of patterning motifs in our library, we use the Wang–Landau algorithm (47) to sample patches

spanning a wide range of polar contents and degrees of clustering. We quantify patch hydrophobicity/hydrophilicity, f , by displacing water molecules from an interfacial probe volume, v , and estimating the free energetic cost, ΔG_{cav} , of creating a cavity adjacent to the patch (Fig. 1*C*); the greater the cost, the more hydrophilic the patch is (33). We use molecular simulations and the indirect umbrella sampling (INDUS) technique (29) to characterize the hydrophobicity, $f \equiv \Delta G_{\text{cav}}$, of each of the 884 patches in our library. A detailed description of how we curate our library of SAM patches and characterize their hydrophobicity is provided in *SI Appendix*, Figs. S1–S4.

Using estimates of f for every patch in our library, we then seek to construct a model, \hat{f} , that is capable of predicting the hydrophobicity of a patch from its chemical pattern. Cassie's law, which is based on macroscopic interfacial thermodynamics, suggests that the greater the polar content of a surface, the more hydrophilic it should be (48). Accordingly, we plot f for every patch in our library against the number of polar end groups, n_o , in that patch (Fig. 1*D*, black dots), and find that patches with more polar groups (higher n_o) indeed tend to be more hydrophilic (higher f). In fact, the general trend of f increasing with n_o is captured reasonably well by a linear fit, $\hat{f}_{\text{M1}} = 110.7 + 4.0n_o$ (model M1, blue line), where \hat{f}_{M1} is expressed in units of thermal energy, $k_B T$; k_B is Boltzmann's constant and T is temperature. We quantify the accuracy of model M1 and all subsequent models that we consider by estimating the root mean square error (RMSE), ϵ , obtained using fivefold cross-validation. We find model M1 error ($\epsilon_{\text{M1}} = 3.67 k_B T$) to be significantly higher than the underlying uncertainty in our computational estimates of f ($\epsilon_0 = 1.02 k_B T$). Moreover, as shown in *SI Appendix*, Fig. S5, other reasonable functional forms of $\hat{f}(n_o)$ do not perform much better than model M1 (e.g., the error of a quadratic model is $3.42 k_B T$). Thus, the predominant contribution to ϵ_{M1} stems not from our choice of the simple (linear) functional form of $\hat{f}_{\text{M1}}(n_o)$ but from the observed variation in f at any given value of n_o . Importantly, although the magnitude of ϵ_{M1} may seem small when compared to the range spanned by f , significant shortcomings of model M1 are exposed upon closer examination of the variation in f for patches with the same n_o . For example, the two patches shown in Fig. 1*D*, *Bottom, Right*, which have the same chemical composition ($n_o = 16$), differ in their hydrophobicity by $\Delta f = 16.1 k_B T$, corresponding to a difference of roughly 10° in their water droplet contact angles and translating into a seven orders of magnitude difference in their binding affinity for extended nonpolar surfaces (33). These findings highlight the important role that the chemical patterning of a patch plays in determining its emergent hydrophobicity (49–51) and further emphasize the need to go beyond simple additive models for predicting the hydrophobicity of heterogeneous surfaces (17, 34).

Neural Network Models for Predicting Hydrophobicity

To go beyond model M1 and learn a predictive model capable of capturing how the hydrophobicity, f , of a patch depends not just on its polar content, n_o , but on its entire chemical pattern, we encode the pattern using a feature vector \mathbf{x} whose elements are assigned values (–1) and (+1) for polar and nonpolar patch end groups, respectively, and a value of (0) for polar groups exterior to the patch (Fig. 2*A*). To learn the functional mapping, $f(\mathbf{x})$, we then use artificial neural networks (ANNs) (36), which are capable of accepting an arbitrarily high-dimensional input, such as \mathbf{x} , and learning complex, nonlinear functional forms

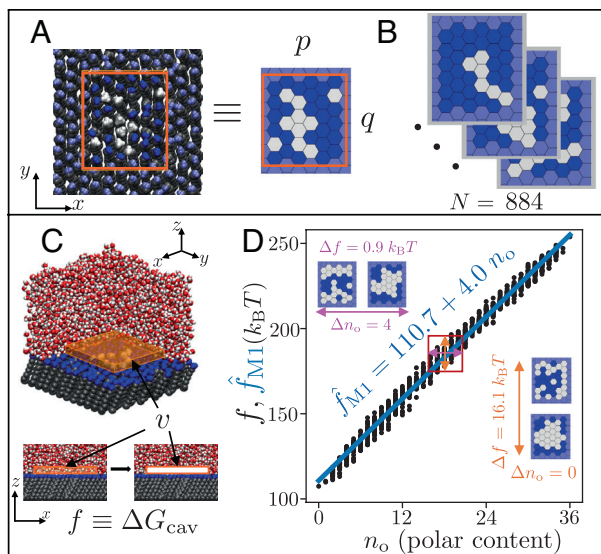


Fig. 1. Characterizing the hydrophobicity/hydrophilicity of heterogeneous surfaces. (A) Simulation snapshot (Left) and schematic (Right) highlighting a chemically patterned patch on a binary SAM surface. The rectangular patch (orange) spans p and q end groups in the x and y directions, respectively, and comprises polar (hydroxyl, blue) and nonpolar groups (methyl, white); groups outside the patch are polar (hydroxyl, light blue). (B) A library of 884 patches with $p = q = 6$ encompassing a range of chemical compositions and patterning motifs is constructed, and the hydrophilicity, f , of each patch is characterized. (C) Patch hydrophilicity, f , is defined using the free energy, ΔG_{cav} , required to empty a roughly 0.3-nm thin, cuboidal probe volume, v (orange) adjacent to the patch; the more hydrophilic a patch, the greater the value of f . (D) The hydrophilicity, f , of every patch in the library is plotted as a function of the number of polar end groups, n_o , in the patch. Although a linear fit (model M1, blue), denoted by \hat{f}_{M1} , captures the general trend of patch hydrophilicity, f , increasing with polar content, n_o , it is incapable of capturing the variation in f for patches with the same n_o . Consequently, model M1 has an error, $\epsilon_{\text{M1}} = 3.7 k_{\text{B}}T$, which is much greater than the inherent uncertainty, $\epsilon_0 = 1.02 k_{\text{B}}T$, in our estimates of f . Representative patches with the same polar content n_o and different hydrophilicity, f (orange), as well as those with different n_o and similar f (magenta) are shown.

(37). We train the ANNs on the residuals, $f(\mathbf{x}) - \hat{f}_{\text{M1}}(n_o(\mathbf{x}))$, of model M1 predictions, thereby adopting a delta learning paradigm (52), which seeks to learn the influence of chemical patterning not contained within model M1. We term this model M1A. Since the hydrophobicity of a patch is invariant under rotation, reflection, and translation, we augment the training dataset with the corresponding isomorphic variants of each patch (38). By exploring a number of fully connected feed-forward architectures, trained in PyTorch (53), we find that an ANN with a single hidden layer and 48 neurons possesses the lowest cross-validation error. A detailed description of our data augmentation procedure as well as ANN training, architecture selection, and hyperparameter optimization is included in *SI Appendix, Figs. S6 and S7*.

The optimal model M1A is able to predict patch hydrophobicity with an error ($\epsilon_{\text{M1A}} = 2.75 k_{\text{B}}T$) that is roughly $1 k_{\text{B}}T$ lower than that of model M1 (Fig. 2D). Moreover, unlike model M1, the accuracy of model M1A is expected to continue to improve as the model is trained with more data (*SI Appendix, Fig. S11*). The substantial improvement in the accuracy of model M1A stems from its ability to capture the variation of f with chemical patterning, \mathbf{x} , for patches with the same polar content, n_o . However, to achieve this improvement, the modestly sized ANN used in model M1A employed 6,433 trainable parameters, making the learned model challenging to interpret. In contrast, model M1 is easy to understand because it employs an intuitive physical descriptor, n_o ; however, this choice also renders model M1

incapable of capturing how hydrophobicity, f , varies with chemical patterning (for fixed n_o), thereby degrading its accuracy and limiting its usefulness. The contrast between models M1 and M1A thus raises the question: Are model accuracy and interpretability mutually exclusive, or is it possible to construct a hydrophobicity model that is both highly accurate and simple enough that it lends itself to physical interpretation?

To obtain simpler models with improved interpretability, but without sacrificing accuracy, we explore the use of CNNs to learn the model M1 residuals, $f(\mathbf{x}) - \hat{f}_{\text{M1}}(n_o(\mathbf{x}))$. We term this model M1C. Traditionally applied to image recognition (38), CNNs include a preprocessing operation to extract relevant high-level image features, which are then input to an ANN. Fig. 2B illustrates this procedure: a filter, \mathbf{w}_i , is rastered over the patch pattern, \mathbf{x} , aggregating local chemical information to produce a convolved pattern, \mathbf{c}_i ; the weights associated with the filter are trained parameters. Multiple filters are applied in parallel to produce a set of convolved patterns, $\{\mathbf{c}_i\}$, which are then coarse-grained under a max-pooling operation to produce a set of lower-dimensional patterns, $\{\mathbf{p}_i\}$. An additional round of convolution and pooling produces a final set of patterns, $\{\mathbf{p}'_i\}$, which serve as inputs to a fully connected ANN that predicts the residuals, $f(\mathbf{x}) - \hat{f}_{\text{M1}}(n_o(\mathbf{x}))$. In general, the convolution and pooling operations enable optimal network performance with far fewer trainable parameters. Moreover, the filters, $\{\mathbf{w}_i\}$, are amenable to interpretability and can inform aspects of local chemical patterning that influence patch hydrophobicity. Once again, the training dataset is augmented with rotational variants of every patch, and models are trained in PyTorch (53) calling the HexagDLY libraries (54); details of the CNN architecture and training are provided in *SI Appendix, Fig. S8*.

We find that model M1C with five filters performs as well as model M1A ($\epsilon_{\text{M1C}} = 2.72 k_{\text{B}}T$) but does so using only 349 trainable parameters (Fig. 2D). The reduced complexity of model M1C suggests that the preprocessing operations (i.e., convolution and pooling) are able to efficiently extract from the patch patterns, \mathbf{x} , local features that are the key determinants of patch hydrophobicity. To shed light on these features, we inspect the five learned filters used in the first convolution layer; the filters are shown in Fig. 2C and are colored according to their trained weights. During the convolution operation, the greatest signal is obtained when nonpolar groups are convolved with positive filter weights (green) and polar groups are convolved with negative weights (purple). Accordingly, patches with clusters of nonpolar groups ought to provide the greatest signal when convolved with filters that feature predominantly positive filter weights, such as the first filter in Fig. 2C. Correspondingly, the presence of negative filter weights at neighboring locations (e.g., as seen in the third filter) suggests that the filter seeks motifs featuring proximal polar groups. Thus, even though model M1C is too complex to fully interpret, an examination of the learned filters in the first convolution layer points to the importance of spatial end-group correlations in determining the hydrophobicity of heterogeneous surfaces.

Incorporating Local Chemical Correlations into Hydrophobicity Models

Our use of neural network models represents a top-down strategy wherein the networks are presented with the entire patch pattern and asked to learn the chemical determinants of hydrophobicity not contained in model M1. As universal function approximators, sufficiently deep neural networks can, in principle, learn end-group correlations of all orders (37). However, to obtain more interpretable

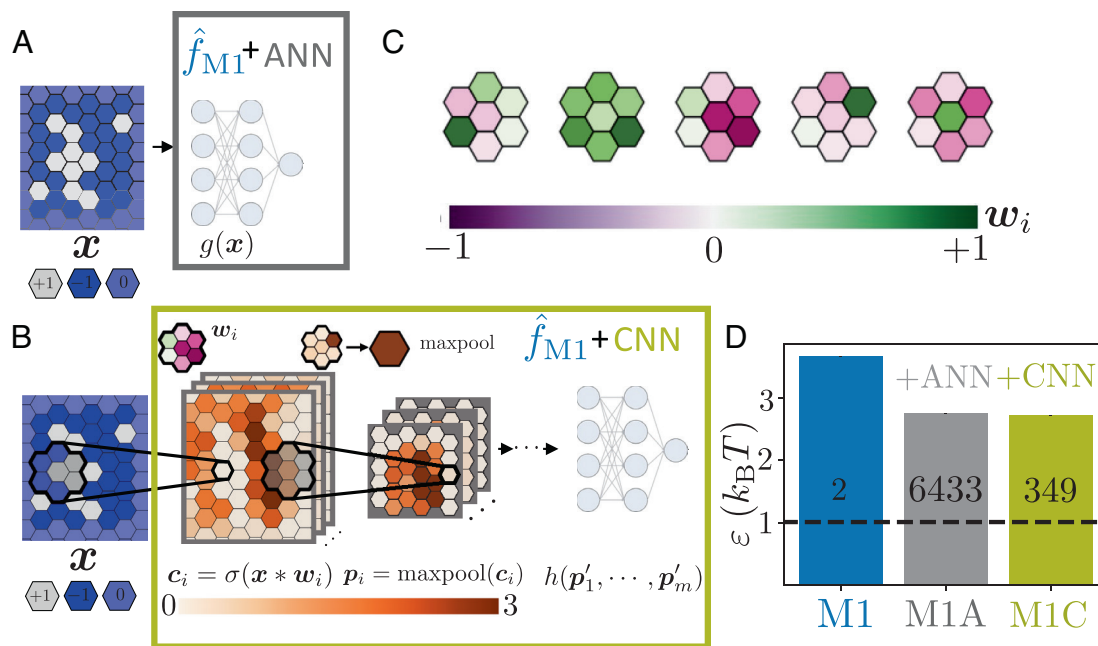


Fig. 2. Using neural networks to accurately predict the hydrophobicity of patterned SAM surfaces. (A) An ANN model M1A is trained to predict patch hydrophobicity, f , as a function of the chemical patterning, \mathbf{x} , by learning over the residuals, $f(\mathbf{x}) - \hat{f}_{M1}(n_o(\mathbf{x}))$, of Model M1. The components of the feature vector, \mathbf{x} , correspond to SAM end groups and assume values of (+1), (-1), or 0, for end groups that are nonpolar, polar, or external to the patch, respectively. The model employs 6,433 trainable parameters, making it challenging to interpret. (B) A convolutional neural network (CNN) model M1C is trained to convolve the patch pattern, \mathbf{x} , with a set of filters, $\{\mathbf{w}_i\}$ ($i = 1, 2, \dots, 5$), that extract information about the local chemical environment into a set of patterns, $\{\mathbf{c}_i\}$, which are then coarse-grained using max-pool operations into reduced patterns, $\{\mathbf{p}_i\}$. The reduced patterns are subjected to additional convolution and pooling before being passed through a fully connected ANN to predict the model M1 residuals, $f(\mathbf{x}) - \hat{f}_{M1}(n_o(\mathbf{x}))$. (C) The CNN architecture offers improved interpretability via the comprehensible nature of the learned filters. The five filters used in the first convolution layer are shown here with positive values colored green and negative ones shown in purple; darker shades represent larger absolute filter weights, whereas lighter shades correspond to smaller weights. The presence of large filter weights (dark shades) at neighboring locations points to the importance of local chemical correlations in determining patch hydrophobicity. (D) The $3.67 k_B T$ error of model M1 is roughly $1 k_B T$ greater than the errors of the neural network models M1A and M1C, highlighting the importance of accounting for the detailed chemical patterning of a patch, \mathbf{x} , in predicting its hydrophobicity. The underlying uncertainty in the computed values of f is $\epsilon_0 = 1.02 k_B T$ (dashed line). The number of trainable parameters in each model is also reported, highlighting the accuracy interpretability trade-off.

models, we now include increasingly higher-order end-group correlations into our baseline model M1 in a complementary bottom-up approach and seek to match the performance of the neural network models. We first incorporate the number of polar-polar nearest neighbors (with at least one neighbor belonging to the patch), n_{oo} , as a descriptor of local chemical patterning to construct model M2, which is linear in n_o and n_{oo} (Fig. 3A). Interestingly, both the number of polar groups adjacent to nonpolar groups, n_{oc} , and the number of neighboring nonpolar groups, n_{cc} , are equivalent descriptor of two-body nearest neighbor correlations. Indeed, due to geometric and stoichiometric constraints, n_{oc} and n_{cc} are uniquely determined by n_o , n_{oo} and the patch dimensions, as shown in *SI Appendix*; thus, models that are isomorphic to model M2 can be constructed using n_o and either n_{oc} or n_{cc} in lieu of n_{oo} .

Importantly, the very simple model M2, with only three learned parameters, is just as accurate ($\epsilon_{M2} = 2.67 k_B T$) as the neural network models M1C and M1A, which employ 349 and 6,433 parameters, respectively (Fig. 3D, solid bars). Even more importantly, the relationship between patch hydrophobicity, f , and chemical patterning, \mathbf{x} , encoded in model M2, $\hat{f}_{M2} = 105.7 + 5.4 n_o - 0.4 n_{oo}$, is highly interpretable and exposes the essential physics underpinning the learned relationship: the greater the polar content, n_o , of a patch, the more hydrophilic it is, and for a given polar content n_o , the more polar neighbors there are, the less hydrophilic the patch is. In other words, the ability of a polar group to enhance patch hydrophilicity is reduced when it is placed adjacent to other polar groups. Indeed, when a patch with proximal polar groups undergoes dewetting, those groups can hydrogen-bond with one another, making the patch

easier to dewet and therefore less hydrophilic; conversely, when a patch with well-separated polar groups undergoes dewetting, those groups can hydrogen-bond with and pin more hydration waters, making it harder to dewet (15). Our finding that polar neighbors lower hydrophilicity (for patches with a fixed polar content) is consistent with the results of both Xi et al. (15) and Kelkar et al. (16), who found that patches with adjoining polar groups were less hydrophilic than those with alternating polar groups.

We now incorporate three-body correlations between end groups into our linear models and seek to ascertain whether explicit accounting of such correlations improves model performance (55). As detailed in *SI Appendix*, seven independent descriptors are needed for this purpose (*SI Appendix*, Figs. S9 and S10); as with model M2, different descriptor combinations can be chosen, giving rise to isomorphic models. We choose to work with the descriptors shown in Fig. 3B (green rectangle), which correspond to the prevalence of the following motifs: polar groups (n_o), polar dimers (n_{oo}), compact polar trimers (n_{ooo}^{compact}), straight polar trimers ($n_{ooo}^{\text{straight}}$), bent polar trimers (n_{ooo}^{bent}), straight mixed trimers ($n_{oco}^{\text{straight}}$), and bent mixed trimers (n_{oco}^{bent}). We note that these motifs can feature polar end groups external to the patch as long as at least one of the end groups in the motif belongs to the patch. We term this model M7. As shown in Fig. 3D, model M7 is slightly more accurate ($\epsilon_{M7} = 2.60 k_B T$) than the neural network models M1C and M1A, which consider the entire patch pattern, \mathbf{x} , as input, and therefore have access to many-body end-group correlations of all orders; this finding suggests that including two- and three-body correlations between end groups is sufficient

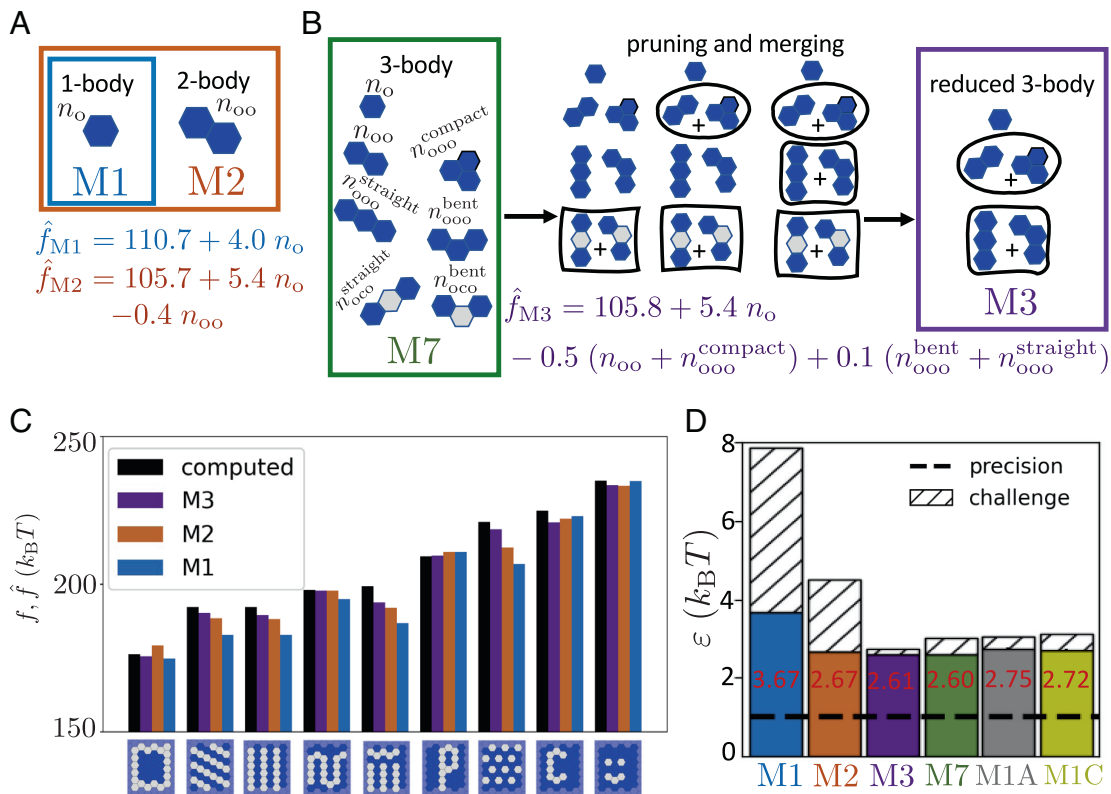


Fig. 3. Bottom-up incorporation of spatial end-group correlations results in models that are both accurate and interpretable. (A) To construct a hydrophobicity model that is cognizant of local chemical patterning, and not just polar content, n_o , we employ the number of polar-polar nearest neighbors, n_{oo} , as an additional descriptor that quantifies the two-body correlations among proximal end groups. Model M2, which is linear in n_o and n_{oo} , highlights that for a given polar content n_o , the presence of neighboring polar groups decreases patch hydrophobicity. (B) To incorporate three-body correlations, we construct model M7, which is linear in the seven descriptors that quantify the prevalence of the motifs shown in the green rectangle: the number of polar groups (n_o) and dimers (n_{oo}); the number of polar trimers that are compact (n_{ooo}^{compact}), straight ($n_{ooo}^{\text{straight}}$), and bent (n_{ooo}^{bent}); and mixed trimers that are straight ($n_{ooo}^{\text{straight}}$) and bent (n_{ooo}^{bent}). To reduce the dimensionality of model M7, we perform a series of ‘pruning’ (removing unimportant descriptors) and ‘merging’ (combining two descriptors) steps, which leads to the three-descriptor model M3 (purple rectangle). (C) To assess the predictive performance of models M2 and M3 when confronted with novel patches, we curate a challenge dataset with nine chemical patterns that contain uncommon motifs. The hydrophobicities, f , of the nine patches in the challenge dataset are compared against the corresponding predictions, \hat{f} , of models M1, M2, and M3. (D) A comparison of the accuracy of the different models, as quantified by ϵ (solid bars), highlights the importance of accounting for end-group correlations. The inclusion of two-body correlations in model M2 renders it 1.00 $k_B T$ more accurate than model M1 and comparable in accuracy to the black-box models M1A and M1C. In contrast, the 0.07 $k_B T$ improvement in accuracy from incorporating three-body correlations is modest, with models M3 and M7 being comparable in accuracy. The errors associated with the challenge dataset (hatched bars) are greater than those with the training dataset (solid bars) for all the models we consider; however, the difference between those errors reduces as higher-order correlations are incorporated, suggesting improved model generalizability. The precision in our measurements of f is $\epsilon_0 = 1.02 k_B T$ (dashed line).

to accurately predict patch hydrophobicity, and incorporating higher-order correlations is unlikely to lead to further improvements in model performance. Indeed, as shown in *SI Appendix*, surrogate linear models, which incorporate two- and three-body end-group correlations, are able to explain the predictions of the neural network models remarkably well.

Although model M7 requires five more parameters than model M2, the relative improvement in its accuracy, $\epsilon_{M2} - \epsilon_{M7} = 0.07 k_B T$, is an order of magnitude smaller than the improvement in model performance achieved upon introducing two-body correlations, $\epsilon_{M1} - \epsilon_{M2} = 1.00 k_B T$ (Fig. 3D). This finding motivates us to ask whether all the terms in model M7, and correspondingly, the motifs in Fig. 3B (green rectangle), are equally important, or whether some may be omitted to arrive at a more interpretable model. To this end, we perform a series of operations, which involve either discarding unimportant descriptors (pruning) or combining two descriptors with similar importance (merging); all possible pruning and merging operations are attempted, and the operation that results in a model with the greatest accuracy is accepted at each stage (Fig. 3B). After four iterations of this procedure, we obtain model M3, which

expresses hydrophobicity as a linear function of three independent descriptors, $\hat{f}_{M3} = 105.8 + 5.4 n_o - 0.49 (n_{oo} + n_{ooo}^{\text{compact}}) + 0.13 n_{ooo}^{\text{straight}} + n_{ooo}^{\text{bent}}$. Model M3 employs four fewer descriptors than model M7 but is just as accurate ($\epsilon_{M3} = 2.61 k_B T$), indicating that it is capable of capturing the salient three-body correlations (Fig. 3D). Moreover, the simplicity of model M3 makes it amenable to interpretation. Like model M2, model M3 predicts that polar groups increase patch hydrophobicity, whereas polar neighbors lower hydrophobicity. Model M3 additionally accounts for three-body correlations through terms that are linear in the lumped descriptors, $(n_{oo} + n_{ooo}^{\text{compact}})$ and $(n_{ooo}^{\text{straight}} + n_{ooo}^{\text{bent}})$; the former highlights that a compact polar trimer lowers patch hydrophobicity as much as a polar dimer, whereas the latter with its small positive coefficient suggests that straight or bent polar trimers lead to a slight increase in hydrophobicity. Thus, model M3 predicts that given a certain number of polar groups and polar dimers, overlaps between polar dimers to form bent or straight trimers lead to small increases in hydrophobicity, whereas a clustering of dimers to form compact trimers results in a decrease in hydrophobicity.

Fig. 3D illustrates that the lowest error across all the models considered, i.e., $\epsilon_{M7} = 2.60 k_B T$, is $1.58 k_B T$ larger than the $1.02 k_B T$ uncertainty in our estimates of f . As shown in *SI Appendix*, Fig. S11, the difference between the error of our best-performing model and the uncertainty in our data decreases as the number of patches included in the training set size is increased. Indeed, although our training dataset contains patches with diverse chemical patterns, the 884 patches in our dataset nevertheless represent a tiny fraction (0.0000013%) of all 2^{36} possible 6×6 patches. To challenge our data-centric models and test how well they perform when presented with novel patches, we designed a set of nine patterns that do not belong to the training dataset and contain uncommon motifs (Fig. 3C). As expected, every model performs worse on this test set than predicted by their cross-validated RMSE estimates; however, model performance on the test dataset improves substantially upon incorporating higher-order correlations (Fig. 3D). In particular, we find that the test set error improves from $7.85 k_B T$ under model M1 to $4.50 k_B T$ under model M2 to $2.73 k_B T$ under model M3. Furthermore, the test set performance of model M3 is superior to that of model M7 (and of models M1A and M1C), suggesting that the latter suffers from overfitting. Thus, model M3 strikes an excellent bias–variance trade-off and is thus simultaneously accurate, generalizable, and interpretable.

Generalizing Minimal Models of Hydrophobicity to Patterned Patches of Different Sizes and Shapes

Encouraged by the accuracy of models M2 and M3 in predicting the hydrophobicity of chemically patterned 6×6 patches, we now seek to generalize these models to all rectangular $p \times q$ patches. We do so by exploiting the interpretability of our models and augmenting it with an understanding of interfacial thermodynamics (33). We describe this analysis for model M2 due to its relative simplicity; the analogous exercise for model M3 yields similar results and is presented in *SI Appendix*, Fig. S14. In model M2 ($\hat{f}_{M2} = \alpha_c + \alpha_o n_o + \alpha_{oo} n_{oo}$), the intercept α_c , corresponds to the hydrophobicity, f_c , of a uniform nonpolar patch (for which $n_o = n_{oo} = 0$), whereas the coefficients α_o and α_{oo} capture the extent to which the presence of polar groups, as quantified by n_o , and the nearest-neighbor chemical patterning, as quantified by n_{oo} , respectively, contribute to patch hydrophobicity, f (Fig. 4A). To construct a generalized version of model M2, we hypothesize that the intercept, α_c (which corresponds to the hydrophobicity of a nonpolar patch, f_c), must depend on patch size and shape, whereas the coefficients, α_o and α_{oo} , should be intrinsic properties of the chosen end-group chemistries and should therefore be independent of patch geometry.

First, we seek to understand how the intercept, α_c , should depend on patch dimensions, p and q , by estimating the hydrophobicity, f_c , of 62 nonpolar patches with sizes that ranged from 10 to 96 end groups and aspect ratios that ranged from 1 (for square patches) to 16 (for a 2×32 patch). We then propose a physically motivated model for f_c , which is a linear function of the patch area, pq , as prescribed by macroscopic interfacial thermodynamics, and is also linear in the patch dimensions, p and q , to account for edge effects that stem from finite patch sizes. As shown in Fig. 4B, the linear model $\hat{f}_c(p, q) = 1.22pq + 5.86p + 4.10q + 3.85$ fits the f_c -data very well with an RMSE of only $1.9 k_B T$. Moreover, this model can also be generalized to nonrectangular patches, as shown in *SI Appendix*, Fig. S12. The intercept, α_c , in the generalized model M2 is then defined as $\alpha_c \equiv \hat{f}_c$.

Next, we test our hypothesis that the coefficients, α_o and α_{oo} , should be independent of patch size and shape. To this end, we generate an expanded dataset that includes 228 square 4×4 patches and 682 rectangular 4×9 patches in addition to the 884 square 6×6 patches in the original dataset (Fig. 4C) and calculate the hydrophobicity of the additional patches. From the hydrophobicity, f , of every patch in the expanded (or original) dataset, we then subtract the hydrophobicity of the corresponding nonpolar patch, f_c , and determine the coefficients, α_o and α_{oo} , by performing a least squares fit to the residuals: $f - f_c = \alpha_o n_o + \alpha_{oo} n_{oo}$. As illustrated in Fig. 4D, the coefficients computed over the expanded dataset, $\alpha_o = (5.07 \pm 0.03) k_B T$ and $\alpha_{oo} = (-0.32 \pm 0.01) k_B T$, are in good agreement with those calculated over the original dataset, $\alpha_o = (5.39 \pm 0.05) k_B T$ and $\alpha_{oo} = (-0.38 \pm 0.02) k_B T$, thereby supporting our hypothesis that α_o and α_{oo} are approximately independent of patch geometry. Collectively, our findings yield the generalized model M2, which can be used to predict the hydrophobicity of heterogeneous rectangular patches of all shapes and sizes (including uniformly polar patches): $\hat{f}_{M2}(p, q, n_o, n_{oo}) = \alpha_c(p, q) + \alpha_o n_o + \alpha_{oo} n_{oo} = [1.22pq + 5.86p + 4.10q + 3.85] + 5.07n_o - 0.32n_{oo}$ (*SI Appendix*, Figs. S13 and S15).

Rational Design of Chemically Patterned Patches

Our learned models M2 and M3 represent accurate, generalizable and interpretable models for predicting patch hydrophobicity, f , as a function of its chemical pattern, \mathbf{x} . We now seek to exploit these models to efficiently navigate through the high-dimensional chemical pattern space and uncover design rules for optimally enhancing patch hydrophobicity or hydrophilicity. In particular, we first interrogate the characteristic features of patch locations known as “hot spots,” which when mutated bring about the largest change in patch hydrophobicity or hydrophilicity (56, 57). We then seek to uncover patches that have the same polar content but vary widely in their hydrophobicity. We present results obtained using model M2 here and include the corresponding results obtained using model M3 in *SI Appendix*, Figs. S16 and S17.

To identify hot-spot end groups, we systematically mutate every end group and estimate the predicted change in hydrophobicity, $\Delta\hat{f}$, using model M2. As an example, consider the patch pattern shown in Fig. 5A. Although nonpolar to polar mutations result in an increase in \hat{f} and polar to nonpolar mutations result in a decrease in \hat{f} , certain end groups have much larger marginal impacts on patch hydrophobicity/hydrophilicity than others (Fig. 5B). Such hot spots are highlighted in Fig. 5A and B (red outline). The largest increase in \hat{f} results from the mutation of a nonpolar group at the center of the nonpolar cluster; this location is optimal because it enables the newly introduced polar group to have the fewest (zero) polar neighbors. In contrast, the largest decrease in \hat{f} results from the mutation of a polar end group that is at the periphery of the polar cluster and has the fewest polar neighbors (including those outside the patch). Thus, polar mutations optimally enhance patch hydrophilicity when introduced at the center of a nonpolar cluster, whereas nonpolar mutations optimally enhance patch hydrophobicity when introduced at the edges of a polar cluster or at isolated polar groups.

In addition to uncovering the characteristic features of hot spots, it is also instructive to consider mutations that bring about the smallest change in hydrophobicity or hydrophilicity (Fig. 5A and B, cyan outline). We find that the smallest increase in \hat{f} occurs on mutating a nonpolar group that is surrounded by polar groups, and conversely, the smallest decrease in \hat{f} occurs on

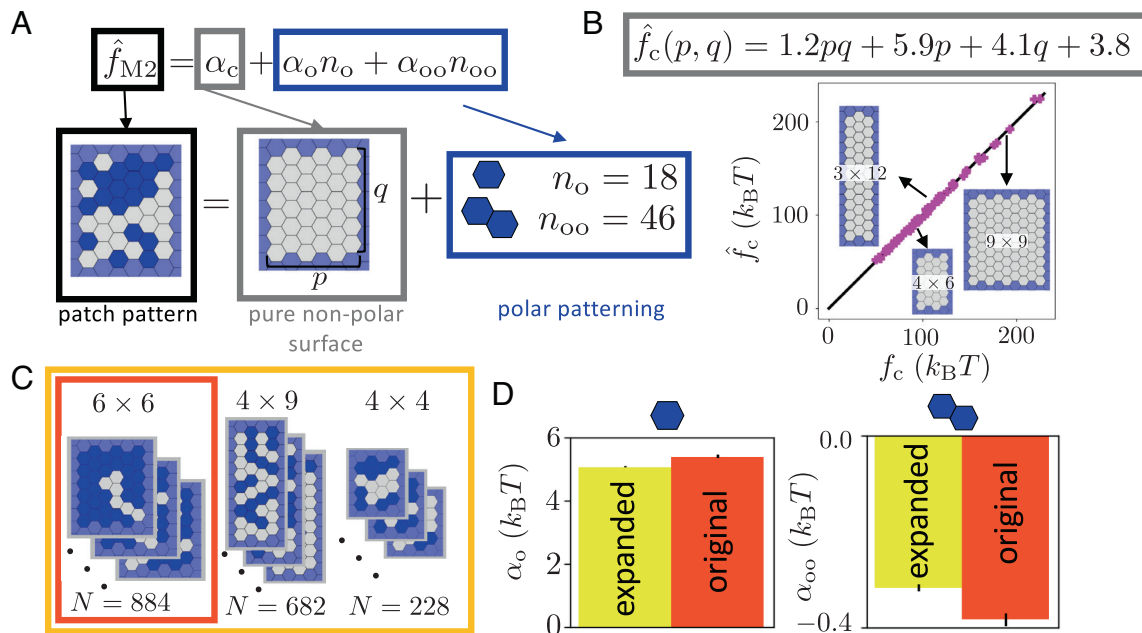


Fig. 4. Physics-based generalization of model M2 to rectangular patches of different shapes and sizes. (A) We first recognize that the intercept, α_c (gray box), in model M2 corresponds to the hydrophobicity of a uniformly nonpolar patch for which $n_o = n_{oo} = 0$. The linear terms in n_o and n_{oo} (blue box) determine how the presence and arrangement of polar groups modulate patch hydrophobicity. (B) The hydrophobicity of a nonpolar patch, f_c , is captured well by a model, \hat{f}_c , that is linear in the area of the patch, pq , as well as its dimensions, p and q . (C) To test the hypothesis that the coefficients, α_o and α_{oo} , are independent of patch size and shape, we generate an expanded dataset (yellow box) containing $N = 228$ patches of a different size (4×4) and $N = 682$ patches with a different shape (4×9) to augment the original dataset (red box) of $N = 884$ patches (6×6). (D) The coefficients, α_o and α_{oo} , obtained by performing a least-squares linear fit to the original (red) and expanded (yellow) datasets are in good agreement, validating the hypothesis that they are insensitive to patch size and shape.

mutating a polar group in a polar cluster. Collectively, our findings highlight that polar clusters are relatively insensitive to mutations, whereas nonpolar clusters are particularly susceptible to mutations. Such an asymmetry between the susceptibility of nonpolar and polar clusters to mutations was also observed by Acharya et al. (13), who found that introducing a polar group at the center of the nonpolar patch substantially suppressed interfacial water density fluctuations, whereas the introduction of a nonpolar group at the center of a polar patch had a much smaller effect on fluctuations. The sensitivity of nonpolar clusters to mutations is also consistent with the findings of Patel et al. (30), who showed that water near extended nonpolar patches sits at the edge of a dewetting transition and is therefore particularly susceptible to perturbations.

By building on our ability to identify hot spots, we now seek to solve the constrained optimization problem of identifying 6×6 patches that maximize hydrophobicity or hydrophilicity under a fixed budget of polar end groups, n_o . Because exhaustive optimization over all 2^{36} possible patches is infeasible, we adopt an iterative greedy strategy to determine approximate solutions to the optimization problem for every value of n_o from 0 to 36. To design maximally hydrophilic patterns, we start with a nonpolar patch ($n_o = 0$) and introduce a polar mutation at a hot-spot location that results in the largest increase in \hat{f} ; if multiple equally favorable hot spots exist, we choose one at random. The result of this operation is a maximally hydrophilic patch containing one polar end group ($n_o = 1$); we then identify the hot-spot location on this patch and mutate it to obtain the most hydrophilic patch with $n_o = 2$. Repeating this

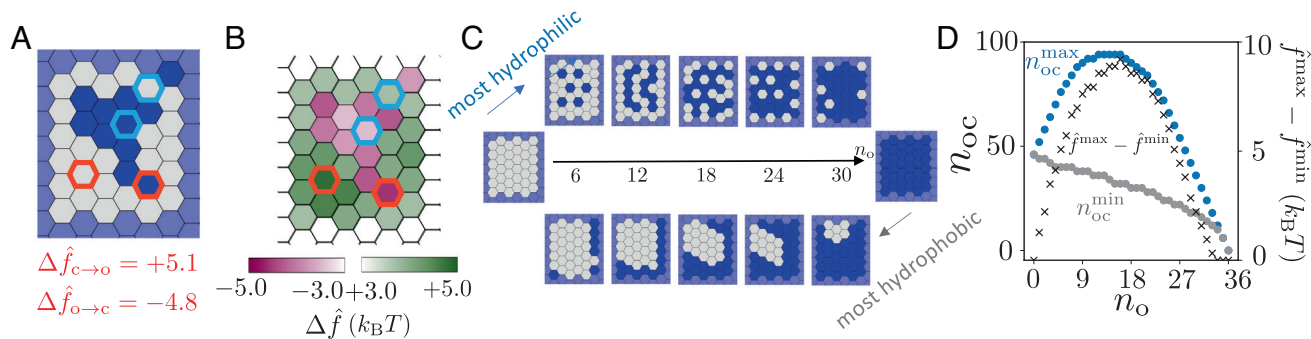


Fig. 5. Rational design of heterogeneous SAM patches with maximal hydrophobicity/hydrophilicity. (A) For the patch pattern shown, we perform an exhaustive screening of all point mutations and interrogate their influence on patch hydrophobicity using model M2. (B) An inspection of the changes in patch hydrophobicity, $\Delta \hat{f}$, in response to point mutations, enables identification of hot-spot locations (red outline). End groups mutations that result in the smallest absolute change in patch hydrophobicity are also highlighted (cyan outline). (C) We employ an iterative greedy design protocol to discover maximally hydrophobic or hydrophilic patches with a particular polar content, n_o . The most hydrophilic patches (Top) tend to feature well-dispersed polar and nonpolar end groups, whereas the most hydrophobic patches (Bottom) feature a single nonpolar cluster. (D) To quantify the relative dispersion of the functional patches, we plot the number of polar-nonpolar neighbors, n_{oc} , for the most hydrophilic patches (n_{oc}^{\max}) and the most hydrophobic patches (n_{oc}^{\min}) as a function of their polar content, n_o . Moreover, the difference between n_{oc}^{\max} and n_{oc}^{\min} (for any n_o) informs the corresponding range of attainable hydrophobicities, $\hat{f}^{\max} - \hat{f}^{\min}$.

procedure for $n_o = 3, 4, \dots, 36$, we obtain a maximally hydrophilic patch at each value of n_o . The design of maximally hydrophobic patches proceeds in an analogous fashion, commencing from a polar patch ($n_o = 36$) and iteratively introducing nonpolar mutations that result in the largest decreases in \hat{f} . We repeat this greedy optimization protocol 100 times and identify the most hydrophilic and hydrophobic patches for every value of n_o ; the optimal patches are illustrated in Fig. 5C for select values of n_o . We find that the most hydrophilic patches feature highly dispersed patterns with polar and nonpolar groups occupying adjacent sites (to minimize polar neighbors), whereas the most hydrophobic patches cluster their nonpolar groups (to maximize polar neighbors).

To characterize the dispersion of the end groups in the optimal patches, we plot the number of polar–nonpolar neighbors, n_{oc} , for the most hydrophilic (n_{oc}^{\max}) and hydrophobic (n_{oc}^{\min}) patches identified at every value of n_o (Fig. 5D). We also show the range of hydrophobicities, $\hat{f}^{\max} - \hat{f}^{\min}$, attainable for a given polar content, n_o , which is proportional to the difference, $n_{oc}^{\max} - n_{oc}^{\min}$, emphasizes the importance of chemical patterning in determining hydrophobicity. We find this range to be maximal for $n_o = 16$, where the most hydrophilic patch has $n_{oc}^{\max} = 94$ and the most hydrophobic patch has $n_{oc}^{\min} = 32$, resulting in a hydrophobicity range of roughly $10.9 k_B T$. Our findings that the most hydrophilic patches feature dispersed chemical patterns and that the most hydrophobic patches display clustering of like groups are in good agreement with the results of Monroe and Shell, who used a genetic algorithm to identify heterogeneous patches with the lowest (or highest) interfacial water diffusivities and found that the optimal patches feature well-mixed (or separated) patterns of polar and nonpolar groups (40); the authors also found that dispersed patches bind small hydrophobic solutes weakly relative to well-separated patches (41).

Conclusions and Outlook

In this work, we seek to understand the relationship between the nanoscale chemical patterns displayed by heterogeneous SAM surfaces and their hydrophobicity using molecular simulations, enhanced sampling techniques, and ML. To this end, we characterize the hydrophobicity of a diverse library of $\mathcal{O}(10^3)$ SAM patches that span a range of sizes, shapes, and chemical compositions. Although the hydrophilicity of a patch generally increases with its polar content, we find that patches with the same polar content can nevertheless display substantial variation in hydrophobicity and that patches with different polar contents can have similar hydrophobicity. Thus, simple additive models of hydrophobicity, which account only for the polar content of a patch, are unable to accurately quantify its hydrophobicity. In contrast, neural network models, which consider the entire chemical pattern of a patch, are able to capture the variation in the hydrophobicity of patches with the same polar content; however, such models are challenging to interpret from a physical standpoint.

To obtain hydrophobicity models that are both accurate and interpretable, we incorporate local spatial correlations between chemical end groups into our models. By constructing a series of models that explicitly account for two-body and higher-order correlations among end groups, we are able to obtain two- and three-descriptor minimal models M2 and M3, respectively, which are just as accurate as the neural network models in capturing the relationship between chemical patterning and hydrophobicity. Models M2 and M3 incorporate not just the polar content of patch but also the number of polar–polar nearest neighbors; model M3 additionally accounts for the number of polar trimers. Importantly, our minimal models are amenable to interpretation and elucidate that for patches with the same polar content, clustering of polar

groups results in enhanced patch hydrophobicity. Model M3 further predicts that bent or straight trimers confer a small increase in hydrophilicity, whereas compact trimers lead to a decrease in hydrophilicity. The interpretability of models M2 and M3 also enables their generalization to rectangular patches of all sizes and shapes. The accuracy of models M2 and M3 emphasizes the importance of chemical correlations between surface groups in determining the hydrophobicity of heterogeneous surfaces and exposes the inherent limitations of additive approaches, such as hydrophathy scales, which seek to express hydrophobicity as a sum of contributions from individual chemical groups.

An interesting consequence of the importance of correlations is that different nonpolar (or polar) groups do not contribute equally to the hydrophobicity of a patch. By leveraging the efficiency of models M2 and M3, we identify hot-spot locations that contribute disproportionately to the overall patch hydrophobicity or hydrophilicity. In particular, we find that the center of a nonpolar cluster is an optimal location for introducing a polar mutation to increase patch hydrophilicity, whereas an isolated polar group or the edge of a polar cluster is an optimal site for introducing a nonpolar mutation to enhance patch hydrophobicity. We also find that the hydrophobicity of nonpolar clusters is susceptible to mutations, whereas that of polar clusters is robust against mutations. These observations may find applications in protein engineering. For example, the identification and mutation of nonpolar hot-spot residues to polar or charged residues could be used to optimally enhance the hydrophilicity of proteins and thereby their aqueous solubility (58, 59). Similarly, the mutation of polar hot spots to nonpolar residues, which optimally enhances protein hydrophobicity, could facilitate tighter binding between a protein and its binding partners (35, 60). By exploiting our ability to identify hot spots, we also perform iterative greedy design of patches, which for a given polar content are maximally hydrophobic or hydrophilic. We find that the most hydrophilic patches display well-dispersed patterns with alternating polar and nonpolar groups, whereas the most hydrophobic patches are distinguished by a single, contiguous nonpolar cluster. These findings may have implications for the rational design of soft materials (61), such as superhydrophilic surface coatings that can resist fouling by proteins (42). Our findings may also facilitate the design of supramolecular hosts that seek to mimic proteins and must similarly strike a delicate balance between being soluble in water (hydrophilic) and being able to bind to their guests (hydrophobic) (62, 63).

Protein surfaces have evolved to use chemical patterning, and not just polar content, to modulate their hydrophobicity and tune their interactions (30, 64). Although protein surfaces are substantially more complex than the binary SAM surfaces studied here (65), employing a much larger palette of chemical moieties and displaying nanoscale roughness, we hope that our work will provide a template for developing data-driven models of protein hydrophobicity that are both accurate and interpretable. Along with parallel advances in the development of methods for accurately and efficiently characterizing protein hydrophobicity (35, 66), our work thus offers the promise of being able to not only predict but also understand how the hydrophobicity of a protein depends on the nanoscale chemical and topographical patterns it displays. Our approach should also be useful in informing the hydrophobicity of other heterogeneous solutes, such as patchy nanoparticles, dendrimers, or supramolecular hosts (67–70). Finally, because our characterization of hydrophobicity accounts for the collective solvent response to the chemical patterns displayed by heterogeneous surfaces, our approach can also be generalized to investigate how the presence of cosolutes (e.g., salt) or cosolvents (e.g., ethanol) modulates the relationship between patterning and hydrophobicity (71–73).

Data, Materials, and Software Availability. All study data are included in the article and/or *SI Appendix*.

ACKNOWLEDGMENTS. N.B.R. was supported by the National Science Foundation (NSF CBET-1652646, DMR-1844514). A.L.F. gratefully acknowledges financial support from the NSF (DMR-1844505). A.J.P. gratefully acknowledges financial support from the NSF (CBET-1652646, DMR-1720530), the Alfred P.

Sloan Research Foundation (FG-2017-9406), and the Camille & Henry Dreyfus Foundation (TC-19-033).

Author affiliations: ^aDepartment of Chemical and Biomolecular Engineering, University of Pennsylvania, Philadelphia, PA 19104; and ^bPritzker School of Molecular Engineering, University of Chicago, Chicago, IL 60637

1. C. Tanford, *The Hydrophobic Effect: Formation of Micelles and Biological Membranes* (Wiley, New York, 1973).
2. D. Chandler, Interfaces and the driving force of hydrophobic assembly. *Nature* **437**, 640–647 (2005).
3. Y. Levy, J. N. Onuchic, Water mediation in protein folding and molecular recognition. *Annu. Rev. Biophys. Biomol. Struct.* **35**, 389–415 (2006).
4. H. S. Ashbaugh, L. R. Pratt, Colloquium: Scaled particle theory and the length scales of hydrophobicity. *Rev. Mod. Phys.* **78**, 159 (2006).
5. D. Thirumalai, E. P. O'Brien, G. Morrison, C. Hyeon, Theoretical perspectives on protein folding. *Annu. Rev. Biophys.* **39**, 159–183 (2010).
6. E. Rabani, D. R. Reichman, P. L. Geissler, L. E. Brus, Drying-mediated self-assembly of nanoparticles. *Nature* **426**, 271–274 (2003).
7. J. W. Barnett *et al.*, Spontaneous drying of non-polar deep-cavity cavitation pockets in aqueous solution. *Nat. Chem.* **12**, 589–594 (2020).
8. S. Granick, S. C. Bae, A curious antipathy for water. *Science* **322**, 1477–1478 (2008).
9. R. Evans, M. C. Stewart, N. B. Wilding, A unified description of hydrophilic and superhydrophobic surfaces in terms of the wetting and drying transitions of liquids. *Proc. Natl. Acad. Sci. U.S.A.* **116**, 23901–23908 (2019).
10. H. Jiang, A. J. Patel, Recent advances in estimating contact angles using molecular simulations and enhanced sampling methods. *Curr. Opin. Chem. Engng.* **23**, 130–137 (2019).
11. N. Giovambattista, P. G. Debenedetti, P. J. Rossky, Hydration behavior under confinement by nanoscale surfaces with patterned hydrophobicity and hydrophilicity. *J. Phys. Chem. C* **111**, 1323–1332 (2007).
12. N. Giovambattista, C. F. Lopez, P. J. Rossky, P. G. Debenedetti, Hydrophobicity of protein surfaces: Separating geometry from chemistry. *Proc. Natl. Acad. Sci. U.S.A.* **105**, 2274–2279 (2008).
13. H. Acharya, S. Vembanur, S. N. Jamadagni, S. Garde, Mapping hydrophobicity at the nanoscale: Applications to heterogeneous surfaces and proteins. *Faraday Discuss.* **146**, 353–365 (2010).
14. R. C. Harris, B. M. Pettitt, Effects of geometry and chemistry on hydrophobic solvation. *Proc. Natl. Acad. Sci. U.S.A.* **111**, 14681–14686 (2014).
15. E. Xi *et al.*, Hydrophobicity of proteins and nanostructured solutes is governed by topographical and chemical context. *Proc. Natl. Acad. Sci. U.S.A.* **114**, 13345–13350 (2017).
16. A. S. Kelkar, B. C. Dallin, R. C. Van Lehn, Predicting hydrophobicity by learning spatiotemporal features of interfacial water structure: Combining molecular dynamics simulations with convolutional neural networks. *J. Phys. Chem. B* **124**, 9103–9114 (2020).
17. N. B. Rego, A. J. Patel, Understanding hydrophobic effects: Insights from water density fluctuations. *Annu. Rev. Cond. Matt. Phys.* **13**, 303–324 (2022).
18. J. Kyte, R. F. Doolittle, A simple method for displaying the hydropathic character of a protein. *J. Mol. Biol.* **157**, 105–132 (1982).
19. J. L. Cornette *et al.*, Hydrophobicity scales and computational techniques for detecting amphipathic structures in proteins. *J. Mol. Biol.* **195**, 659–685 (1987).
20. A. E. Kister, J. C. Phillips, A stringent test for hydrophobicity scales: Two proteins with 88% sequence identity but different structure and function. *Proc. Natl. Acad. Sci. U.S.A.* **105**, 9233–9237 (2008).
21. J. Wang, D. Bratko, A. Luzar, Probing surface tension additivity on chemically heterogeneous surfaces by a molecular approach. *Proc. Natl. Acad. Sci. U.S.A.* **108**, 6374–6379 (2011).
22. D. S. Tomar, V. Weber, B. M. Pettitt, D. Asthagiri, Conditional solvation thermodynamics of isoleucine in model peptides and the limitations of the group-transfer model. *J. Phys. Chem. B* **118**, 4080–4087 (2014).
23. C. D. Ma, C. Wang, C. Acevedo-Vélez, S. H. Gellman, N. L. Abbott, Modulation of hydrophobic interactions by proximally immobilized ions. *Nature* **517**, 347–350 (2015).
24. D. S. Tomar, M. E. Paulaitis, L. R. Pratt, D. N. Asthagiri, Hydrophilic interactions dominate the inverse temperature dependence of polypeptide hydration free energies attributed to hydrophobicity. *J. Phys. Chem. Lett.* **11**, 9965–9970 (2020).
25. K. Lum, D. Chandler, J. D. Weeks, Hydrophobicity at small and large length scales. *J. Phys. Chem. B* **103**, 4570–4577 (1999).
26. J. D. Weeks, Connecting local structure to interface formation: A molecular scale van der Waals theory of nonuniform liquids. *Annu. Rev. Phys. Chem.* **53**, 533–562 (2002).
27. S. Vaikuntanathan, G. Rotskoff, A. Hudson, P. L. Geissler, Necessity of capillary modes in a minimal model of nanoscale hydrophobic solvation. *Proc. Natl. Acad. Sci. U.S.A.* **113**, E2224–E2230 (2016).
28. A. J. Patel, P. Varilly, D. Chandler, Fluctuations of water near extended hydrophobic and hydrophilic surfaces. *J. Phys. Chem. B* **114**, 1632–1637 (2010).
29. A. J. Patel, P. Varilly, D. Chandler, S. Garde, Quantifying density fluctuations in volumes of all shapes and sizes using indirect umbrella sampling. *J. Stat. Phys.* **145**, 265–275 (2011).
30. A. J. Patel *et al.*, Sitting at the edge: How biomolecules use hydrophobicity to tune their interactions and function. *J. Phys. Chem. B* **116**, 2498–2503 (2012).
31. A. J. Patel, S. Garde, Efficient method to characterize the context-dependent hydrophobicity of proteins. *J. Phys. Chem. B* **118**, 1564–1573 (2014).
32. R. Godawat, S. N. Jamadagni, S. Garde, Characterizing hydrophobicity of interfaces by using cavity formation, solute binding, and water correlations. *Proc. Natl. Acad. Sci. U.S.A.* **106**, 15119–15124 (2009).
33. A. J. Patel *et al.*, Extended surfaces modulate hydrophobic interactions of neighboring solutes. *Proc. Natl. Acad. Sci. U.S.A.* **108**, 17678–17683 (2011).
34. S. N. Jamadagni, R. Godawat, S. Garde, Hydrophobicity of proteins and interfaces: Insights from density fluctuations. *Ann. Rev. Chem. Biomol. Eng.* **2**, 147–171 (2011).
35. N. B. Rego, E. Xi, A. J. Patel, Identifying hydrophobic protein patches to inform protein interaction interfaces. *Proc. Natl. Acad. Sci. U.S.A.* **118**, e2018234118 (2021).
36. M. Hassoun, *Fundamentals of Artificial Neural Networks* (MIT Press, 1995).
37. K. Hornik, M. Stinchcombe, H. White, Multilayer feedforward networks are universal approximators. *Neural Netw.* **2**, 359–366 (1989).
38. A. Krizhevsky, I. Sutskever, G. E. Hinton, Imagenet classification with deep convolutional neural networks. *Adv. Neural Inf. Process. Syst.* **25**, 1097–1105 (2012).
39. S. Shin, A. P. Willard, Characterizing hydration properties based on the orientational structure of interfacial water molecules. *J. Chem. Theory Comput.* **14**, 461–465 (2018).
40. J. I. Monroe, M. S. Shell, Computational discovery of chemically patterned surfaces that effect unique hydration water dynamics. *Proc. Natl. Acad. Sci. U.S.A.* **115**, 8093–8098 (2018).
41. J. I. Monroe *et al.*, Affinity of small-molecule solutes to hydrophobic, hydrophilic, and chemically patterned interfaces in aqueous solution. *Proc. Natl. Acad. Sci. U.S.A.* **118**, e2020205118 (2021).
42. S. Chen, Z. Cao, S. Jiang, Ultra-low fouling peptide surfaces derived from natural amino acids. *Biomaterials* **30**, 5892–5896 (2009).
43. N. Chennamsetty, V. Vovnov, V. Kayser, B. Helk, B. L. Trout, Design of therapeutic proteins with enhanced stability. *Proc. Natl. Acad. Sci. U.S.A.* **106**, 11937–11942 (2009).
44. C. C. Lee, J. M. Perchiacca, P. M. Tessier, Toward aggregation-resistant antibodies by design. *Trends Biotechnol.* **31**, 612–620 (2013).
45. M. Mrksich, G. M. Whitesides, Using self-assembled monolayers to understand the interactions of man-made surfaces with proteins and cells. *Annu. Rev. Biophys. Biomol. Struct.* **25**, 55–78 (1996).
46. F. R. F. Fan *et al.*, Determination of the molecular electrical properties of self-assembled monolayers of compounds of interest in molecular electronics. *J. Am. Chem. Soc.* **123**, 2454–2455 (2001).
47. F. Wang, D. P. Landau, Efficient, multiple-range random walk algorithm to calculate the density of states. *Phys. Rev. Lett.* **86**, 2050–2053 (2001).
48. A. B. D. Cassie, Contact angles. *Discuss. Faraday Soc.* **3**, 11–16 (1948).
49. P. J. Rossky, Exploring nanoscale hydrophobic hydration. *Faraday Discuss.* **146**, 13–18 (2010).
50. N. Giovambattista, P. Rossky, P. Debenedetti, Computational studies of pressure, temperature, and surface effects on the structure and thermodynamics of confined water. *Annu. Rev. Phys. Chem.* **63**, 179–200 (2012).
51. J. Monroe *et al.*, Water structure and properties at hydrophilic and hydrophobic surfaces. *Annu. Rev. Chem. Biomol. Eng.* **11**, 523–557 (2020).
52. A. Nandi, C. Qu, P. L. Houston, R. Conte, J. M. Bowman, δ -machine learning for potential energy surfaces: A PIP approach to bring a DFT-based PES to CCSD(T) level of theory. *J. Chem. Phys.* **154**, 051102 (2021).
53. A. Paszke, "Pytorch: An imperative style, high-performance deep learning library" in *Advances in Neural Information Processing Systems*, H. Wallach *et al.*, Eds. (Curran Associates, Inc., 2019), vol. **32**, pp. 8024–8035.
54. C. Steppa, T. L. Holch, HexagDLY—processing hexagonally sampled data with CNNs in PyTorch. *SoftwareX* **9**, 193–198 (2019).
55. S. Shin, A. P. Willard, Three-body hydrogen bond defects contribute significantly to the dielectric properties of the liquid water–vapor interface. *J. Phys. Chem. Lett.* **9**, 1649–1654 (2018).
56. T. Kortemme, D. Baker, A simple physical model for binding energy hot spots in protein–protein complexes. *Proc. Natl. Acad. Sci. U.S.A.* **99**, 14116–14121 (2002).
57. A. Ganesan *et al.*, Structural hot spots for the solubility of globular proteins. *Nat. Commun.* **7**, 1–15 (2016).
58. R. C. Remsing, E. Xi, A. J. Patel, Protein hydration thermodynamics: The influence of flexibility and salt on hydrophobin II hydration. *J. Phys. Chem. B* **122**, 3635–3646 (2018).
59. B. A. Rogers *et al.*, A stepwise mechanism for aqueous two-phase system formation in concentrated antibody solutions. *Proc. Natl. Acad. Sci. U.S.A.* **116**, 15784–15791 (2019).
60. T. M. Przybycien, Protein-protein interactions as a means of purification. *Curr. Opin. Biotechnol.* **9**, 164–170 (1998).
61. Y. X. Shen *et al.*, Achieving high permeability and enhanced selectivity for angstrom-scale separations using artificial water channel membranes. *Nat. Commun.* **9**, 2294 (2018).
62. M. B. Hillyer, B. C. Gibb, Molecular shape and the hydrophobic effect. *Annu. Rev. Phys. Chem.* **67**, 307–329 (2016).
63. S. Garde, Sticky when dry. *Nat. Chem.* **12**, 587–588 (2020).
64. N. B. Rego, E. Xi, A. J. Patel, Protein hydration waters are susceptible to unfavorable perturbations. *J. Am. Chem. Soc.* **141**, 2080–2086 (2019).
65. P. Gainza *et al.*, Deciphering interaction fingerprints from protein molecular surfaces using geometric deep learning. *Nat. Methods* **17**, 184–192 (2020).
66. E. Xi, R. C. Remsing, A. J. Patel, Sparse sampling of water density fluctuations in interfacial environments. *J. Chem. Theory Comput.* **12**, 706–713 (2016).
67. I. C. Pons-Siepmann, S. C. Glotzer, Design of patchy particles using quaternary self-assembled monolayers. *ACS Nano* **6**, 3919–3924 (2012).
68. Q. Xiao *et al.*, Janus dendrimersomes coassembled from fluorinated, hydrogenated, and hybrid janus dendrimers as models for cell fusion and fission. *Proc. Natl. Acad. Sci. U.S.A.* **114**, E7045–E7053 (2017).
69. Z. Luo *et al.*, Determination and evaluation of the nonadditivity in wetting of molecularly heterogeneous surfaces. *Proc. Natl. Acad. Sci. U.S.A.* **116**, 25516–25523 (2019).
70. C. B. Rodell, J. E. Mealy, J. A. Burdick, Supramolecular guest–host interactions for the preparation of biomedical materials. *Bioconjug. Chem.* **26**, 2279–2289 (2015).
71. M. Lund, L. Vrbka, P. Jungwirth, Specific ion binding to nonpolar surface patches of proteins. *J. Am. Chem. Soc.* **130**, 11582–11583 (2008).
72. S. Bhardwaj, N. F. van der Vegt, Does preferential adsorption drive consolvency? *Macromolecules* **52**, 4131–4138 (2019).
73. Q. Gu, Z. Li, J. L. Coffman, T. M. Przybycien, A. L. Zydney, High throughput solubility and redissolution screening for antibody purification via combined peg and zinc chloride precipitation. *Biotechnol. Prog.* **36**, e3041 (2020).