Learning the Relationship between Nanoscale Chemical Patterning and Hydrophobicity

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

^a Department of Chemical and Biomolecular Engineering, University of Pennsylvania, Philadelphia, PA 19104.; ^b Pritzker School of Molecular Engineering, University of Chicago, Chicago, IL 60637.

This manuscript was compiled on September 20, 2022

11

12

13

14

15

21

24

11

13

14

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 a 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

ydrophobicity – the aversion of non-polar 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 non-polar groups (11–17). Consequently, additive approaches for characterizing hydrophobicity, such as hydropathy 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.

26

27

28

29

30

31

33

35

37

38

41

42

43

44

45

46

47

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 accuracyinterpretability trade-off – sophisticated "black box" models capable of accurately learning high-dimensional functional mappings can resist interpretation, whereas simple "white 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 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 molec-

Significance Statement

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 non-trivial, and cannot be captured by additive approaches, such as hydropathy 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.

Please provide details of author contributions here

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.

¹To whom correspondence should be addressed. E-mail: andrewferguson@uchicago.edu, amish.patel@seas.upenn.edu

ular 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.

49

50

51 52

55

56

57

58

59

60

61

62

63

64

65

66

67

69

70

71

72

73

74

75

76

77

78

79

81

82

83

84

85

87

88

89

90

92

93

94

95

96

97

98

100

101

102

103

104

105

108

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 non-intuitive 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 non-polar cluster, whereas patch hydrophilicity is particularly susceptible to non-polar 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 non-polar groups, whereas the most hydrophobic patches feature a single non-polar 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 endgroups that are either polar (hydroxyl) or non-polar (methyl) (Figure 1). Initially, we focus on a 6×6 square patch (roughly $3 \text{ nm} \times 3 \text{ nm}$), which contains 36 end-groups, and is embedded in a background of polar end-groups, as shown in Figure 1A. 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 (Figure 1B). 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

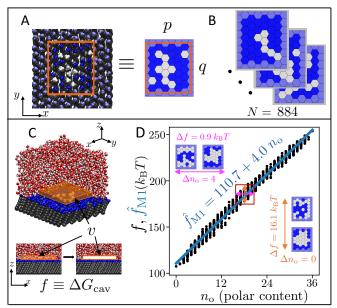


Fig. 1. Characterizing the hydrophobicity/hydrophilicity of heterogeneous surfaces. (A) Simulation snapshot (left) and schematic (right) highlighting a chemically patterned patch on a binary self-assembled monolayer (SAM) surface. The rectangular patch (orange) spans p and q end-groups in the x and y directions, respectively, and is comprised of polar (hydroxyl, blue) and non-polar 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, $\Delta G_{\rm cay}$, required to empty a roughly 0.3 nm thin, cuboidal probe volume, \boldsymbol{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_0 , in the patch. Although a linear fit (model M1, blue), denoted by $\hat{f}_{\rm M1}$, captures the general trend of patch hydrophilicity, f, increasing with polar content, n_0 , it is incapable of capturing the variation in f for patches with the same $n_{
m o}$. Consequently, model M1 has an error, $\varepsilon_{\rm M1}=3.7~k_{\rm B}T$, which is much greater than the inherent uncertainty, $\varepsilon_0=1.02~k_{\rm B}T$, in our estimates of f. Representative patches with the same polar content $n_{\scriptscriptstyle \mathrm{O}}$ and different hydrophilicity, f (orange), as well as those with different $n_{\rm o}$ and similar f (magenta) are shown.

water molecules from an interfacial probe volume, v, and estimating the free energetic cost, $\Delta G_{\rm cav}$, of creating a cavity adjacent to the patch (Figure 1C); the greater the cost, the more hydrophilic the patch (33). We use molecular simulations and the Indirect Umbrella Sampling (INDUS) technique (29) to characterize the hydrophobicity, $f \equiv \Delta G_{\rm cav}$, 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 the SI (Figures S1-S4).

Using estimates of f for every patch in our library, we then seek to construct a model, \tilde{f} , that is capable of predicting the hydrophobicity of a patch from its chemical pattern. The venerable 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 (Figure 1D, black dots), and find that patches with more polar groups (higher n_0) indeed tend to be more hydrophilic (higher f). In fact, the general trend of f increasing with n_0 is captured reasonably well by a linear fit, $\hat{f}_{M1} = 110.7 + 4.0n_o$ (model M1, blue line), where f_{M1} is expressed in units of thermal energy, $k_{\text{B}}T$; 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),

109

111

112

113

114

115

116

117

118

119

120

121

122

125

126

127

128

 ε , obtained using five-fold cross-validation. We find model M1 error ($\varepsilon_{\rm M1} = 3.67 \ k_{\rm B}T$) to be significantly higher than the underlying uncertainty in our computational estimates of f $(\varepsilon_0 = 1.02 \ k_B T)$. Moreover, as shown in the SI (Figure S5), other reasonable functional forms of $f(n_0)$ do not perform much better than model M1 (e.g., the error of a quadratic model is 3.42 $k_{\rm B}T$). Thus, the predominant contribution to $\varepsilon_{\rm M1}$ stems not from our choice of the simple (linear) functional form of $\hat{f}_{M1}(n_0)$, but from the observed variation in f at any given value of n_0 . Importantly, although the magnitude of $\varepsilon_{\rm M1}$ 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_0 . For example, the two patches shown in Figure 1D (bottom, right), which have the same chemical composition $(n_0 = 16)$, 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 7 orders of magnitude difference in their binding affinity for extended non-polar 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).

134

135

136

137

139

140

141

142

143

144

145

147

148

149

150

151

152

153

154

155

156

157

158

159

160

161

162

163

164

165

166

167

168

169

170

171

172

173

174

175

176

177

179

180

181

182

183

184

185

186

187

188

189

190

193

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, $n_{\rm o}$, but on its entire chemical pattern, we encode the pattern using a feature vector \boldsymbol{x} whose elements are assigned values (-1) and (+1) for polar and non-polar patch end-groups, respectively, and a value of (0) for polar groups exterior to the patch (Figure 2A). To learn the functional mapping, f(x), we then use artificial neural networks (ANNs) (36), which are capable of accepting an arbitrarily high-dimensional input, such as x, and learning complex, non-linear functional forms (37). We train the ANNs on the residuals, $f(\mathbf{x}) - \hat{f}_{M1}(n_o(\mathbf{x}))$, 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 the SI (Figures S6, S7).

The optimal model M1A is able to predict patch hydrophobicity with an error ($\varepsilon_{\rm M1A}=2.75~k_{\rm B}T$) that is roughly 1 $k_{\rm B}T$ lower than that of model M1 (Figure 2D). Moreover, unlike model M1, the accuracy of model M1A is expected to continue to improve as the model is trained with more data (Figure S11). The substantial improvement in the accuracy of model M1A stems from its ability to capture the variation of f with chemical patterning, x, for patches with the same polar content, $n_{\rm o}$. However, to achieve this improvement, the modestly-sized ANN used in model M1A employed 6,433 trainable parame-

ters, making the learned model challenging to interpret. In contrast, model M1 is easy to understand because it employs an intuitive physical descriptor, $n_{\rm o}$; however, this choice also renders model M1 incapable of capturing how hydrophobicity, f, varies with chemical patterning (for fixed $n_{\rm o}$), 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?

194

195

196

197

200

201

202

203

204

205

206

208

209

210

211

212

213

215

216

217

218

219

220

221

223

224

225

226

227

228

229

230

231

232

233

234

235

236

237

239

240

241

242

243

244

245

247

248

249

250

251

254

To obtain simpler models with improved interpretability, but without sacrificing on accuracy, we explore the use of convolutional neural networks (CNNs) to learn the model M1 residuals, $f(\mathbf{x}) - \hat{f}_{M1}(n_o(\mathbf{x}))$. We term this model M1C. Traditionally applied to image recognition (38), CNNs include a pre-processing operation to extract relevant high-level image features, which are then input to an ANN. Figure 2B illustrates this procedure: a filter, w_i , is rastered over the patch pattern, \boldsymbol{x} , aggregating local chemical information to produce a convolved pattern, c_i ; the weights associated with the filter are trained parameters. Multiple filters are applied in parallel to produce a set of convolved patterns, $\{c_i\}$, which are then coarse-grained under a max-pooling operation to produce a set of lower-dimensional patterns, $\{p_i\}$. An additional round of convolution and pooling produces a final set of patterns, $\{p_i'\}$, which serve as inputs to a fully-connected ANN that predicts the residuals, $f(\mathbf{x}) - \hat{f}_{M1}(n_o(\mathbf{x}))$. In general, the convolution and pooling operations enable optimal network performance with far fewer trainable parameters. Moreover, the filters, $\{w_i\}$, 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 the SI (Figure S8).

We find that model M1C with five filters performs as well as model M1A ($\varepsilon_{\rm M1C} = 2.72~k_{\rm B}T$) but does so using only 349 trainable parameters (Figure 2D). The reduced complexity of model M1C suggests that the pre-processing operations (i.e., convolution and pooling) are able to efficiently extract from the patch patterns, 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 Figure 2C and are colored according to their trained weights. During the convolution operation, the greatest signal is obtained when non-polar groups are convolved with positive filter weights (green) and polar groups are convolved with negative weights (purple). Accordingly, patches with clusters of non-polar groups ought to provide the greatest signal when convolved with filters that feature predominantly positive filter weights, such as the first filter in Figure 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.

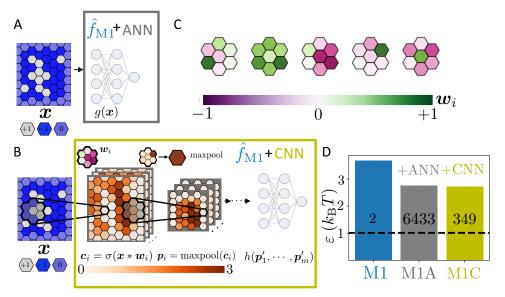


Fig. 2. Using neural networks to accurately predict the hydrophobicity of patterned SAM surfaces. (A) An artificial neural network (ANN) model M1A is trained to predict patch hydrophobicity, f, as a function of the chemical patterning, x, by learning over the residuals, $f(x) - \hat{f}_{M1}(n_o(x))$, of Model M1. The components of the feature vector, x, correspond to SAM end-groups, and assume values of (+1), (-1) or 0, for end groups that are non-polar, 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, x, with a set of filters, $\{w_i\}$ ($i=1,2,\ldots 5$), that extract information about the local chemical environment into a set of patterns, $\{c_i\}$, which are then coarse-grained using mx-pool operations into reduced patterns, $\{p_i\}$. The reduced patterns are subjected to additional convolution and pooling before being passed through a fully-connected ANN to predict the model M1 residuals, $f(x) - \hat{f}_{M1}(n_o(x))$. (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 M1 and M1C, highlighting the importance of accounting for the detailed chemical patterning of a patch, x, in predicting its hydrophobicity. The underlying uncertainty in the computed values of f is $\varepsilon_0 = 1.02 k_B T$ (dashed line). The number of trainable parameters in each model

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 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} (Figure 3A). Interestingly, the number of polar groups adjacent to non-polar groups, $n_{\rm oc}$, or the number of neighboring non-polar groups, n_{cc} , are equivalent descriptors of two-body nearest neighbor correlations. Indeed, due to geometric and stoichiometric constraints, $n_{\rm oc}$ and n_{cc} are uniquely determined by n_{cc} , n_{cc} and the patch dimensions, as shown in the SI; thus, models that are isomorphic to model M2 can be constructed using n_0 and either n_{oc} or $n_{\rm cc}$ in lieu of $n_{\rm oo}$.

Importantly, the very simple model M2, with only three learned parameters, is just as accurate ($\varepsilon_{\rm M2} = 2.67~k_{\rm B}T$) as the neural network models M1C and M1A, which employ 349 and 6,433 parameters, respectively (Figure 3D, solid bars). Even more importantly, the relationship between patch hydrophobicity, f, and chemical patterning, x, encoded in model M2,

 $\hat{f}_{\rm M2} = 105.7 + 5.4 \ n_{\rm o} - 0.4 \ n_{\rm oo}$, 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. 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 endgroups into our linear models, and seek to ascertain whether explicit accounting of such correlations improves model performance (55). As detailed in the SI, seven independent descriptors are needed for this purpose (Figures S9, 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 Figure 3B (green rectangle), which correspond to the prevalence of the following motifs: polar groups (n_o) , polar dimers $(n_{oo}^{\text{straight}})$, compact polar trimers $(n_{ooo}^{\text{compact}})$, straight polar trimers $(n_{ooo}^{\text{straight}})$, bent polar trimers (n_{ooo}^{bent}) , straight mixed trimers $(n_{ooo}^{\text{straight}})$ and bent mixed trimers (n_{ooo}^{bent}) . We note that these motifs can feature polar end287

289

290

291

292

293

296

297

298

299

300

301

302

303

304

305

307

308

309

310

311

312

255

256

257

258

259

260

261

262

263

266

267

268

269

270

271

272

273

274

275

276

277

278

279

281

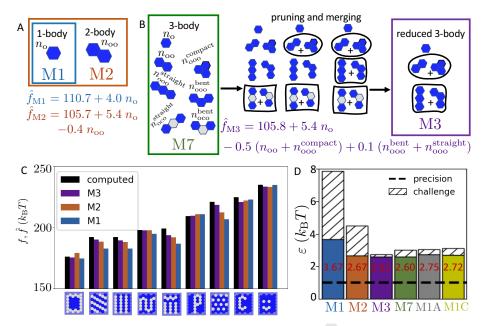


Fig. 3. Bottom-up incorporation of spatial end-group correlations results in models that are both accurate and interpretable. (A) To construct a hydrophobicity model that is cognizant of local chemical patterning, and not just polar content, $n_{\rm o}$, we employ the number of polar-polar nearest neighbors, $n_{\rm oo}$, as an additional descriptor that quantifies the two-body correlations amongst proximal end-groups. Model M2, which is linear in $n_{\rm o}$ and $n_{\rm oo}$, highlights that for a given polar content $n_{\rm oo}$, the presence of neighboring polar groups decreases patch hydrophilicity. (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 $(n_{\rm oo})$ and dimers $(n_{\rm oo})$; the number of polar trimers that are compact $(n_{\rm ooo}^{\rm compact})$, straight $(n_{\rm ooo}^{\rm straight})$ and bent $(n_{\rm ooo}^{\rm bent})$; and mixed trimers that are straight $(n_{\rm ooo}^{\rm straight})$ and bent $(n_{\rm ooo}^{\rm bent})$. 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 $1.00~k_{\rm B}T$ more accurate than model M1, and comparable in accuracy to the black-box models M1A and M1C. In contrast, the $0.07~k_{\rm B}T$ improvement in accuracy from i

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 Figure 3D, model M7 is slightly more accurate ($\varepsilon_{\rm M7}=2.60~k_{\rm B}T$) than the neural network models M1C and M1A, which consider the entire patch pattern, 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 to accurately predict patch hydrophobicity, and incorporating higher-order correlations is unlikely to lead to further improvements in model performance. Indeed, as shown in the SI, 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, $\epsilon_{\rm M2} - \epsilon_{\rm M7} = 0.07~k_{\rm B}T$, is an order of magnitude smaller than the improvement in model performance achieved upon introducing two-body correlations, $\epsilon_{\rm M1} - \epsilon_{\rm M2} = 1.00~k_{\rm B}T$ (Figure 3D). This finding motivates us to ask whether all the terms in model M7, and correspondingly, the motifs in Figure 3B (green rectangle), are equally important, or if 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 (Figure 3B). After four iterations of this procedure, we obtain model M3, which expresses hydrophobicity as a linear function of three independent descriptors, $\hat{f}_{M3} =$ $105.8 + 5.4 \ n_{\rm o} - 0.49 \ (n_{\rm oo} + n_{\rm ooo}^{\rm compact}) + 0.13 \ (n_{\rm ooo}^{\rm straight} + n_{\rm ooo}^{\rm bent}).$ Model M3 employs four fewer descriptors than model M7, but is just as accurate ($\varepsilon_{\rm M3}=2.61~k_{\rm B}T$), indicating that it is capable of capturing the salient three-body correlations (Figure 3D). Moreover, the simplicity of model M3 makes it amenable to interpretation. Like model M2, model M3 predicts that polar groups increase patch hydrophilicity, whereas polar neighbors lower hydrophilicity. Model M3 additionally accounts for three-body correlations through terms that are linear in the lumped descriptors, $(n_{oo} + n_{ooo}^{compact})$ and $(n_{\text{ooo}}^{\text{straight}} + n_{\text{ooo}}^{\text{bent}})$; the former highlights that a compact polar trimer lowers patch hydrophilicity 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 hydrophilicity. 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 hydrophilicity, whereas a clustering of dimers to form compact trimers results in a decrease in hydrophilicity.

Figure 3D illustrates that the lowest error across all the models considered, i.e., $\varepsilon_{\rm M7}=2.60~k_{\rm B}T$, is 1.58 $k_{\rm B}T$ larger than the 1.02 $k_{\rm B}T$ uncertainty in our estimates of f. As shown in the SI (Figure S11), the difference between the error of our

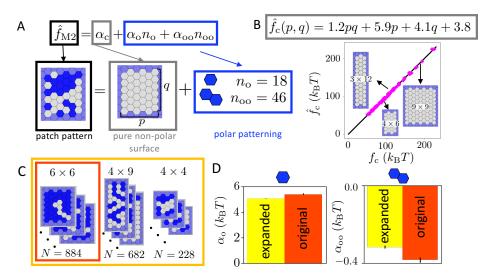


Fig. 4. Physics-based generalization of model M2 to rectangular patches of different shapes and sizes. (A) We first recognize that the intercept, α_c (gray box), in model M2 corresponds to the hydrophobicity of a uniformly non-polar patch for which $n_o=n_{oo}=0$. The linear terms in n_o and n_{oo} (blue box) determine how the presence and arrangement of polar groups modulates patch hydrophobicity. (B) The hydrophobicity of a non-polar patch, f_c , is captured well by a model, 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.

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 (Figure 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 (Figure 3D). In particular, we find that the test set error improves from 7.85 $k_{\rm B}T$ under model M1 to 4.50 $k_{\rm B}T$ under model M2 to $2.73 k_BT$ 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 tradeoff, 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 the SI (Figure S14). In model M2 ($\hat{f}_{\rm M2} = \alpha_{\rm c} + \alpha_{\rm o} n_{\rm o} + \alpha_{\rm oo} n_{\rm oo}$), the intercept, $\alpha_{\rm c}$, corresponds to the hydrophobicity, $f_{\rm c}$, of a uniform non-polar patch (for which $n_{\rm o} = n_{\rm oo} = 0$), whereas the coefficients $\alpha_{\rm o}$ and $\alpha_{\rm oo}$ capture the extent to which the presence of

polar groups, as quantified by $n_{\rm o}$, and the nearest-neighbor chemical patterning, as quantified by $n_{\rm oo}$, respectively contribute to patch hydrophobicity, f (Figure 4A). To construct a generalized version of model M2, we hypothesize that the intercept, $\alpha_{\rm c}$ (which corresponds to the hydrophobicity of a non-polar patch, $f_{\rm c}$), must depend on patch size and shape, whereas the coefficients, $\alpha_{\rm o}$ and $\alpha_{\rm oo}$, 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 non-polar 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 Figure 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 $1.9~k_{\rm B}T$. Moreover, this model can also be generalized to non-rectangular patches, as shown in the SI (Figure 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 (Figure 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 non-polar 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 Figure 4D, the coefficients computed over the

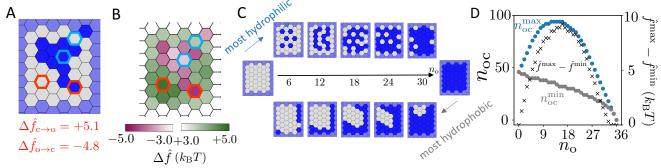


Fig. 5. Rational design of heterogeneous SAM patches with maximal hydrophobicity/hydrophilicity. (A) For the patch pattern shown, we perform an exhaustive screening of all point mutations, and interrogate their influence on patch hydrophobicity using model M2. (B) An inspection of the changes in patch hydrophilicity, $\Delta \hat{f}$, 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 non-polar end-groups, whereas the most hydrophobic patches (bottom) feature a single non-polar cluster. (D) To quantify the relative dispersion of the optimal patches, we plot the number of polar – non-polar neighbors, n_{oc} , for 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, $\hat{f}^{max} - \hat{f}^{min}$.

expanded dataset, $\alpha_{\rm o}=(5.07\pm0.03)~k_{\rm B}T$ and $\alpha_{\rm oo}=(-0.32\pm0.01)~k_{\rm B}T$, are in good agreement with those calculated over the original dataset, $\alpha_{\rm o}=(5.39\pm0.05)~k_{\rm B}T$ and $\alpha_{\rm oo}=(-0.38\pm0.02)~k_{\rm B}T$, thereby validating our hypothesis that $\alpha_{\rm o}$ and $\alpha_{\rm oo}$ are 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): $\hat{f}_{\rm M2}(p,q,n_{\rm o},n_{\rm oo})=\alpha_{\rm c}(p,q)+\alpha_{\rm o}n_{\rm o}+\alpha_{\rm oo}n_{\rm oo}=[1.22pq+5.86p+4.10q+3.85]+5.07n_{\rm o}-0.32n_{\rm oo}$ (Figures 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, 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 the SI (Figures 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 Figure 5A. Although non-polar to polar mutations result in an increase in \hat{f} , and polar to non-polar mutations result in a decrease in \hat{f} , certain end-groups have much larger marginal impacts on patch hydrophobicity/hydrophilicity than others (Figure 5B). Such hot-spots are highlighted in Figures 5A,B (red outline). The largest increase in \hat{f} results from the mutation of a non-polar group at the center of the non-polar 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 non-polar cluster, whereas non-polar 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 hotspots, it is also instructive to consider mutations that bring about the smallest change in hydrophobicity or hydrophilicity (Figures 5A,B, cyan outline). We find that the smallest increase in \hat{f} occurs on mutating a non-polar group that is surrounded by polar groups, and conversely, the smallest decrease in \hat{f} occurs on mutating a polar group in a polar cluster. Collectively, our findings highlight that polar clusters are relatively insensitive to mutations, whereas non-polar clusters are particularly susceptible to mutations. Such an asymmetry between the susceptibility of non-polar 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 non-polar patch substantially suppressed interfacial water density fluctuations, whereas the introduction of a non-polar group at the center of a polar patch had a much smaller effect on fluctuations. The sensitivity of non-polar clusters to mutations is also consistent with the findings of Patel et al. (30), who showed that water near extended non-polar 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 non-polar 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 procedure for n_o =3, 4, ...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 non-polar 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 Figure 5C for select values of n_o . We find that the most hydrophilic patches feature highly dispersed patterns with polar and non-polar groups occupying adjacent sites (to minimize polar neighbors), whereas the most hydrophobic patches cluster their non-polar groups (to maximize polar neighbors).

To characterize the dispersion of the end-groups in the optimal patches, we plot the number of polar – non-polar neighbors, $n_{\rm oc}$, for the most hydrophilic $(n_{\rm oc}^{\rm max})$ and hydrophobic $(n_{\text{oc}}^{\text{min}})$ patches identified at every value of n_{o} (Figure 5D). We also show the range of hydrophobicities, $\hat{f}^{\text{max}} - \hat{f}^{\text{min}}$, attainable for a given polar content, n_0 , which is proportional to the difference, $n_{\text{oc}}^{\text{max}} - n_{\text{oc}}^{\text{min}}$, and emphasizes the importance of chemical patterning in determining hydrophobicity. We find this range to be maximal for $n_0 = 16$, where the most hydrophilic patch has $n_{\rm oc}^{\rm max} = 94$ and the most hydrophobic patch has $n_{\rm oc}^{\rm min} = 32$, resulting in a hydrophobicity range of roughly 10.9 $k_{\rm B}T$. 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 Shell and co-workers, 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 non-polar groups (40); the authors also found that dispersed patches bind small hydrophobic solutes weakly relative to well-separated patches (41).

Conclusions and Outlook

524

525

526

527

530

531

532

533

534

537

538

539

540

541

544

545

546

547

552

553

554

555

558

561

562

563

564

565

570

571

572

573

574

576

577

578

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 machine learning. 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 amongst 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 hydropathy scales, which seek to express hydrophobicity as a sum of contributions from individual chemical groups.

585

586

587

590

591

592

593

594

597

598

599

600

601

602

603

605

606

607

608

609

610

611

613

614

615

616

617

618

620

621

622

623

624

629

630

631

632

633

635

636

637

638

639

An interesting consequence of the importance of correlations is that different non-polar (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 non-polar 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 are optimal sites for introducing a non-polar mutation to enhance patch hydrophobicity. We also find that the hydrophobicity of non-polar 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 non-polar 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 non-polar 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 non-polar groups, whereas the most hydrophobic patches are distinguished by a single, contiguous non-polar cluster. These findings may have implications for the rational design of soft materials (61), such as super-hydrophilic surface coatings that can resist fouling by proteins (42). Our findings may also facilitate the design of supramolecular hosts, who like the proteins they seek to mimic, must 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 co-solutes (e.g., salt) or co-solvents (e.g., ethanol) modulates the relationship between patterning and hydrophobicity (71–73).

645

646

647

648

650

651

652

653

654

655

658

659

660

662

663

664

665

667

668

671

672

673

674

675

676

677

678

679

680

681

682

683

684

685

686

687

688

689

690

691

692

693

694

695

696

697

698

699

700

701

707

708

709

710

711

712

713

719

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).

- Tanford C (1973) The hydrophobic effect: Formation of micelles and biological membranes. (Wiley, New York).
- Chandler D (2005) Interfaces and the driving force of hydrophobic assembly. Nature 437:640–670
 647.
 - Levy Y, Onuchic JN (2006) Water mediation in protein folding and molecular recognition. Annu. Rev. Biophys. Biomol. Struct. 35:389–415.
 - Ashbaugh HS, Pratt LR (2006) Colloquium: Scaled particle theory and the length scales of hydrophobicity. Rev. Mod. Phys. 78(1):159.
 - Thirumalai D, O'Brien EP, Morrison G, Hyeon C (2010) Theoretical Perspectives on Protein Folding. Annu. Rev. Biophys. 39(1):159–183.
 - Rabani E, Reichman DR, Geissler PL, Brus LE (2003) Drying-mediated self-assembly of nanoparticles. Nature 426(6964):271–274.
 - Barnett JW, et al. (2020) Spontaneous drying of non-polar deep-cavity cavitand pockets in aqueous solution. Nat. Chem. 12:589–594.
 - 8. Granick S, Bae SC (2008) A Curious Antipathy for Water. Science 322:1477-1478.
 - Evans R, Stewart MC, Wilding NB (2019) 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(48):23901–23908.
 - Jiang H, Patel AJ (2019) Recent advances in estimating contact angles using molecular simulations and enhanced sampling methods. Curr. Opin. Chem. Engg. 23:130–137.
 - Giovambattista N, Debenedetti PG, Rossky PJ (2007) Hydration Behavior under Confinement by Nanoscale Surfaces with Patterned Hydrophobicity and Hydrophilicity. *J. Phys. Chem. C* 111(3):1323–1332.
 - Giovambattista N, Lopez CF, Rossky PJ, Debenedetti PG (2008) Hydrophobicity of protein surfaces: Separating geometry from chemistry. Proc. Natl. Acad. Sci. U.S.A. 105(7):2274– 2279.
 - Acharya H, Vembanur S, Jamadagni SN, Garde S (2010) Mapping hydrophobicity at the nanoscale: Applications to heterogeneous surfaces and proteins. Faraday Discuss. 146(0):353–365.
 - Harris RC, Pettitt BM (2014) Effects of geometry and chemistry on hydrophobic solvation Proc. Natl. Acad. Sci. U.S.A. 111(41):14681–14686.
 - Xi E, et al. (2017) Hydrophobicity of proteins and nanostructured solutes is governed by topographical and chemical context. Proc. Natl. Acad. Sci. U.S.A. 114(51):13345–13350.
 - Kelkar AS, Dallin BC, Van Lehn RC (2020) Predicting Hydrophobicity by Learning Spatiotemporal Features of Interfacial Water Structure: Combining Molecular Dynamics Simulations with Convolutional Neural Networks. J. Phys. Chem. B 124(41):9103–9114.
- with Convolutional Neural Networks. J. Phys. Chem. B 124(41):9103–9114.
 17. Rego NB, Patel AJ (2022) Understanding hydrophobic effects: Insights from water density fluctuations. Annu. Rev. Cond. Matt. Phys. 13(1):303 324.
- Kyte J, Doolittle RF (1982) A simple method for displaying the hydropathic character of a protein. J Mol. Biol. 157(1):105–132.
 - Cornette JL, et al. (1987) Hydrophobicity scales and computational techniques for detecting amphipathic structures in proteins. J Mol. Biol. 195(3):659–685.
 - Kister AE, Phillips JC (2008) 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(27):9233–9237.
 - Wang J, Bratko D, Luzar A (2011) Probing surface tension additivity on chemically heterogeneous surfaces by a molecular approach. Proc. Natl. Acad. Sci. U.S.A. 108(16):6374–6379.
- 714
 22. Tomar DS, Weber V, Pettitt BM, Asthagiri D (2014) Conditional solvation thermodynamics of isoleucine in model peptides and the limitations of the group-transfer model. *J. Phys. Chem.* 716
 8 118(15):4080–4087.
- 23. Ma CD, Wang C, Acevedo-Vélez C, Gellman SH, Abbott NL (2015) Modulation of hydrophobic
 interactions by proximally immobilized ions. *Nature* 517(7534):347–350.
 - 24. Tomar DS, Paulaitis ME, Pratt LR, Asthagiri DN (2020) Hydrophilic interactions dominate

the inverse temperature dependence of polypeptide hydration free energies attributed to hydrophobicity. *J. Phys. Chem. Lett.* 11(22):9965–9970.

720

721

722

723

724

725

726

727

728

729

730

731

732

733

734

735

736

737

738

739

740

741

742

743

744

745

746

747

748

749

750

751

752

753

754

755

756

757

758

759

760

761

762

763

764

765

766

767

768

769

770

771

772

773

774

775

776

777

778

779

780

781

782

783

784

785

786

787

788

789

790

791

792

793

794

795

796

797

798

799

800

801

- Lum K, Chandler D, Weeks JD (1999) Hydrophobicity at Small and Large Length Scales. J Phys. Chem. B 103(22):4570–4577.
- Weeks JD (2002) Connecting local structure to interface formation: A molecular scale van der Waals theory of nonuniform liquids. Annu. Rev. Phys. Chem. 53(1):533–562.
- Vaikuntanathan S, Rotskoff G, Hudson A, Geissler PL (2016) Necessity of capillary modes in a minimal model of nanoscale hydrophobic solvation. PNAS 113(16):E2224–E2230.
- Patel AJ, Varilly P, Chandler D (2010) Fluctuations of water near extended hydrophobic and hydrophilic surfaces. J Phys. Chem. B 114(4):1632–1637.
- Patel AJ, Varilly P, Chandler D, Garde S (2011) Quantifying density fluctuations in volumes of all shapes and sizes using indirect umbrella sampling. J. Stat. Phys. 145(2):265–275.
- Patel AJ, et al. (2012) Sitting at the edge: How biomolecules use hydrophobicity to tune their interactions and function. J. Phys. Chem. B 116(8):2498–2503.
- Patel AJ, Garde S (2014) Efficient method to characterize the context-dependent hydrophobicity of proteins. J. Phys. Chem. B 118(6):1564–1573.
- Godawat R, Jamadagni SN, Garde S (2009) Characterizing hydrophobicity of interfaces by using cavity formation, solute binding, and water correlations. *Proc. Natl. Acad. Sci. U.S.A.* 106(36):15119–15124.
- Patel AJ, et al. (2011) Extended surfaces modulate hydrophobic interactions of neighboring solutes. Proc. Natl. Acad. Sci. U.S.A. 108(43):17678–17683.
- Jamadagni SN, Godawat R, Garde S (2011) Hydrophobicity of proteins and interfaces: Insights from density fluctuations. Ann. Rev. Chem. Biomol. Eng. 2(1):147–171.
- Rego NB, Xi E, Patel AJ (2021) Identifying hydrophobic protein patches to inform protein interaction interfaces. Proc. Natl. Acad. Sci. U.S.A. 118(6):e2018234118.
- 36. Hassoun M (1995) Fundamentals of artificial neural networks. (MIT Press).
- Hornik K, Stinchcombe M, White H (1989) Multilayer feedforward networks are universal approximators. Neural netw. 2(5):359–366.
- Krizhevsky A, Sutskever I, Hinton GE (2012) Imagenet classification with deep convolutional neural networks. Adv. Neural Inf. Process. Syst. 25:1097–1105.
- Shin S, Willard AP (2018) Characterizing hydration properties based on the orientational structure of interfacial water molecules. J Chem. Theory Comput. 14(2):461–465.
- Monroe JI, Shell MS (2018) Computational discovery of chemically patterned surfaces that effect unique hydration water dynamics. Proc. Natl. Acad. Sci. U.S.A. 115(32):8093–8098.
- Monroe JI, et al. (2021) Affinity of small-molecule solutes to hydrophobic, hydrophilic, and chemically patterned interfaces in aqueous solution. *Proc. Natl. Acad. Sci. U.S.A.* 118(1):e2020205118.
- Chen S, Cao Z, Jiang S (2009) Ultra-low fouling peptide surfaces derived from natural amino acids. Biomaterials 30(29):5892–5896.
- 43. Chennamsetty N, Voynov V, Kayser V, Helk B, Trout BL (2009) Design of therapeutic proteins with enhanced stability. *Proc. Natl. Acad. Sci. U.S.A.* 106(29):11937–11942.
- Lee CC, Perchiacca JM, Tessier PM (2013) Toward aggregation-resistant antibodies by design. Trends Biotechnol. 31(11):612 – 620.
- Mrksich M, Whitesides GM (1996) Using self-assembled monolayers to understand the interactions of man-made surfaces with proteins and cells. *Annu. Rev. Biophys. Biomol. Struct.* 25(1):55–78.
- Fan FRF, et al. (2001) Determination of the molecular electrical properties of self-assembled monolayers of compounds of interest in molecular electronics. J. Am. Chem. Soc. 123(10):2454–2455.
- Wang F, Landau DP (2001) Efficient, multiple-range random walk algorithm to calculate the density of states. Phys. Rev. Lett. 86(10):2050–2053.
- 48. Cassie ABD (1948) Contact angles. Discuss. Faraday Soc. 3:11-16.
- Rossky PJ (2010) Exploring nanoscale hydrophobic hydration. Faraday Discuss. 146:13–18.
 Giovambattista N, Rossky P, Debenedetti P (2012) Computational studies of pressure, tem-
- perature, and surface effects on the structure and thermodynamics of confined water. *Annu. Rev. Phys. Chem.* 63(1):179–200.
- Monroe J, et al. (2020) Water structure and properties at hydrophilic and hydrophobic surfaces. Annu. Rev. Chem. Biomol. Eng. 11(1):523–557.
- 52. Nandi A, Qu C, Houston PL, Conte R, Bowman JM (2021) 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(5):051102.
- Paszke A, et al. (2019) Pytorch: An imperative style, high-performance deep learning library in Advances in Neural Information Processing Systems 32, eds. Wallach H, et al. (Curran Associates, Inc.), pp. 8024–8035.
- Steppa C, Holch TL (2019) HexagDLy—Processing hexagonally sampled data with CNNs in PyTorch. SoftwareX 9:193–198.
- Shin S, Willard AP (2018) Three-body hydrogen bond defects contribute significantly to the dielectric properties of the liquid water–vapor interface. J Phys. Chem. Lett. 9(7):1649–1654.
- Kortemme T, Baker D (2002) A simple physical model for binding energy hot spots in protein– protein complexes. *Proc. Natl. Acad. Sci. U.S.A.* 99(22):14116–14121.
- Ganesan A, et al. (2016) Structural hot spots for the solubility of globular proteins. Nat. Commun. 7(1):1–15.
- Remsing RC, Xi E, Patel AJ (2018) Protein hydration thermodynamics: The influence of flexibility and salt on hydrophobin II hydration. J. Phys. Chem. B 122(13):3635–3646.
- Rogers BA, et al. (2019) A stepwise mechanism for aqueous two-phase system formation in concentrated antibody solutions. *Proc. Natl. Acad. Sci. U.S.A.* 116(32):15784–15791.
- Przybycien TM (1998) Protein-protein interactions as a means of purification. Curr. Opin. Biotechnol. 9(2):164–170.
- Shen Yx, et al. (2018) Achieving high permeability and enhanced selectivity for angstromscale separations using artificial water channel membranes. Nat. Commun. 9(1):2294.
- Hillyer MB, Gibb BC (2016) Molecular shape and the hydrophobic effect. Annu. Rev. Phys. Chem. 67:307–329.
- 63. Garde S (2020) Sticky when dry. Nat. Chem. 12(7):587-588.
- 64. Rego NB, Xi E, Patel AJ (2019) Protein Hydration Waters Are Susceptible to Unfavorable

Perturbations. J. Am. Chem. Soc. 141(5):2080-2086.

804

805

806

814

- 65. Gainza P, et al. (2020) Deciphering interaction fingerprints from protein molecular surfaces using geometric deep learning. Nat. Methods 17(2):184-192.
- 66. Xi E, Remsing RC, Patel AJ (2016) Sparse sampling of water density fluctuations in interfacial 807 environments. J. Chem. Theory Comput. 12(2):706-713. 808
- 67. Pons-Siepermann IC, Glotzer SC (2012) Design of patchy particles using quaternary self-809 assembled monolayers. ACS Nano 6(5):3919-3924. 810
- 68. Xiao Q, et al. (2017) Janus dendrimersomes coassembled from fluorinated, hydrogenated, 811 and hybrid janus dendrimers as models for cell fusion and fission. Proc. Natl. Acad. Sci. 812 U.S.A. 114(34):E7045-E7053. 813
 - 69. Luo Z, et al. (2019) Determination and evaluation of the nonadditivity in wetting of molecularly heterogeneous surfaces. Proc. Natl. Acad. Sci. U.S.A. 116(51):25516-25523.
- 70. Rodell CB, Mealy JE, Burdick JA (2015) Supramolecular guest-host interactions for the 816 817 preparation of biomedical materials. *Bioconjug. Chem.* 26(12):2279–2289.
- 71. Lund M, Vrbka L, Jungwirth P (2008) Specific ion binding to nonpolar surface patches of 818 proteins. J Am. Chem. Soc. 130(35):11582-11583. 819
- 820 72. Bharadwaj S, van der Vegt NF (2019) Does preferential adsorption drive cononsolvency? 821 Macromolecules 52(11):4131-4138.
- 73. Gu Q, Li Z, Coffman JL, Przybycien TM, Zydney AL (2020) High throughput solubility and 822 823 redissolution screening for antibody purification via combined peg and zinc chloride precipi-824 tation. Biotechnol. Prog. 36(6):e3041.





Supplementary Information for

- Learning the Relationship between Nanoscale Chemical Patterning and Hydrophobicity
- 4 Nicholas B. Rego, Andrew L. Ferguson, and Amish J. Patel
- 5 Andrew Ferguson.
- 6 E-mail: andrewferguson@uchicago.edu
- 7 Amish Patel.
- 8 E-mail: amish.patel@seas.upenn.edu

9 This PDF file includes:

- Supplementary text
- Figs. S1 to S17
- SI References

Supporting Information Text

15

18

19

20

21

22

25

26

27

28

29

30

31

33

36

37

38

39

40

41

43

44

47

48

49

50

51

52

53

54

55

56

57

58

59

61

62

1. Generating Libraries of Diverse Chemical Patterns

Given that binary self-assembled monolayer (SAM) patches span a large and high-dimensional chemical pattern space, with 2^{36} patterns possible for the 6×6 patches, here we seek to sample a small but diverse subset of such patches. If the patches are chosen randomly, certain polar contents and patterning motifs are much more likely to appear than others; in particular, randomly generated patterns are likely to favor roughly equal numbers of polar and non-polar end-groups with the groups being well-dispersed rather than clustered. To generate a library of diverse patches that span a wide range of polar contents and degrees of dispersion/clustering, we quantify the former using the number of polar groups, n_0 , and the latter using the degree of polar clustering, do, which is defined as the root mean squared deviation of the polar groups from their center of mass (Figure S1A). For every value of n_0 , we then estimated the number of patterns, $\Omega_{n_0}(d_0)$, as a function of d_0 either by exhaustive sampling (for n_0 -values with fewer than 60,000 patterns) or using the Wang-Landau (1) algorithm with a tolerance of 10^{-10} and a maximum of 60,000 iterations per round. As illustrated in Figure S1B for $n_o = 18$, $\Omega_{n_o}(d_o)$ is sharply peaked, highlighting that randomly-generated patterns would span a narrow range of d_o -values. To engender pattern diversity, we instead sample uniformly across all possible values of n_0 and d_0 and obtain 442 distinct patches. To obtain patterns that span a wide range, not only in their d_0 -values, but also in their degrees of non-polar clustering, we swap the identities of polar and non-polar end-groups in each of our 442 patches to generate another 442 'inverted' patches (Figure S1C) and arrive at a total of N=884 patches for our training library of 6×6 patches. This procedure was repeated to additionally generate diverse 4×4 patches (N = 228) and 4×9 patches (N = 682) for inclusion in the expanded dataset (Figure 4C).

2. Molecular Simulations of Self-Assembled Monolayer (SAM) Surfaces

Simulation set-up. Self-assembled monolayer (SAM) surfaces composed of hexagonally-packed alkyl chains were prepared following references (2, 3). Each chain contains a sulfur atom, 10 united-atom methylene groups, and either a methyl (non-polar) or hydroxyl (polar) end-group that is exposed to solvent. Molecular visualizations of the methyl- and hydroxyl-terminated chains are shown in Figure S2A. The majority of our calculations were performed using a SAM surface, consisting of $12 \times 12 = 144$ alkyl chains, with chemically patterned patches embedded in a polar background (Figure 1A). An elongated SAM surface, composed of $6 \times 36 = 216$ alkyl chains, was used to estimate the hydrophobicity of the 2×12 , 2×18 , 2×32 and 3×12 uniform non-polar patches (Figure 4B). Moreover, a larger SAM surface, composed of 16×16 alkyl chains, was used to study the 12×8 non-polar patch, the non-rectangular non-polar patches (Figure S12), and the 12×12 patterned patches (Figure S15). SAM surfaces were solvated using a roughly 4 nm thick slab of water and placed normal to the z-dimension in a $5.2 \times 6.0 \times 8.5$ nm³ simulation box, as shown in Figures 1C and S2B; for the elongated and larger SAM surfaces, simulation boxes of dimensions, $2.6 \times 18 \times 8.5$ nm³ and $6.93 \times 8.0 \times 8.5$ nm³, respectively, were used. Following reference (4), a repulsive wall of spherical atoms was positioned at z = 8.2 nm to induce a buffering vapor layer, which serves to accommodate any water molecules displaced from the vicinity of the SAM surfaces during the Indirect Umbrella Sampling (INDUS) calculations.

Simulation parameters. All simulations were performed using GROMACS version 4.5.3 (5), suitably modified to incorporate the biasing potentials used in the INDUS calculations. The equations of motion were integrated using the leap-frog integrator (6) with a time step of 2 fs and periodic boundary conditions were used in all dimensions. All systems were simulated in the canonical ensemble with temperature maintained at T = 300 K using the stochastic velocity-rescale thermostat (7) and a coupling constant of $\tau = 0.5$ ps. Before performing biased (INDUS) simulations, every patterned SAM system was energy-minimized using the steepest-descent algorithm and equilibrated for 100 ps; biased simulations were run for 500 ps. Water molecules were represented explicitly using the three-point SPC/E model (8). The united-atom methylene groups were modeled using Model UAc from reference (9). The Lennard-Jones parameters and bonded potentials for the methyl and hydroxyl end-groups were modeled using the Amber99SB force field (10), whereas their partial charges were taken from the OPLS-AA force field (11). Lennard Jones parameters were combined using the Lorentz-Bertholet mixing rules. Van der Waals and short-range electrostatic interactions were truncated at 1 nm, and long-range electrostatics were treated using the Particle Mesh Ewald method (12). All bonds to hydrogen atoms were constrained; SAM end-group bonds were constrained using the LINCS algorithm (13), whereas the bonds in water were constrained using the SETTLE algorithm (14). The sulfur atoms were restrained to lie on a 0.5 nm hexagonal lattice using harmonic restraints with a spring constant of 40,000 kJ/mol/nm². The third methylene groups from the end of the alkyl chain were also position-restrained with the same spring constant (4) to hinder the slow circular motion of the alkyl chains.

3. Characterizing Hydrophobicity using Indirect Umbrella Sampling (INDUS)

The hydrophobicity, f, of a patterned patch was quantified using the free energetic cost, ΔG_{cav} , of displacing all waters from a cuboidal probe volume, v, placed adjacent to the patch (Figure 1B).

Probe volume dimensions. The dimensions, l_x and l_y , of the cuboidal probe volume, v, in the x and y directions, respectively, were chosen to span all end-groups within the patch, as illustrated in Figure S3A. In particular, these dimensions were determined by running a 5 ns equilibrium simulation of a purely polar patch, and padding the average x and y coordinates of the top-right and bottom-left hydroxyl end-groups in the patch with a buffer distance of $\delta = 0.275$ nm; see Figure S3B) The dimensions, l_x and l_y , of v are related to the patch dimensions, p and q, as follows:

$$l_x(p) = (2\delta) + r_x(p-1)$$
, and $l_y(q) = (2\delta) + r_y(q-0.5)$, [1]

where $r_y = 0.5$ nm is the lattice spacing in the y direction and $r_x = \frac{\sqrt{3}}{2} r_y$ is the spacing between end-groups in the x direction. The width, l_z , of v was chosen to include waters within the first hydration shell of the SAM patches. The lower bound of v 70 was chosen to be $z_{\min} = 1.7$, but to ensure that patches of the same size/shape have the same average numbers of waters in v, the upper bound, z_{max} , of v was varied across different patches. An inspection of the normalized density of water oxygens, 72 g(z), adjacent to the uniform 6×6 patches (Figure S3 C-F) makes it clear that a choice of $z_{\rm max} = 2.175$ nm for the polar patch and $z_{\text{max}} = 2.3$ nm for the non-polar patch results in an average of 100 waters in v for both patches. For patterned patches, $z_{\rm max}$ was chosen to be between 2.175 nm and 2.3 nm such that the corresponding v contained roughly 100 waters.

Estimating ΔG_{cav} using INDUS. We use the Indirect Umbrella Sampling (INDUS) method (15, 16) to calculate the probability, $P_v(N)$, of observing N water molecules in a probe volume v, placed adjacent to a patch pattern of interest, and estimate 77 $\Delta G_{\rm cav}$ through the relationship, $\Delta G_{\rm cav} = -k_{\rm B}T \ln P_v(0)$, where $k_{\rm B}T$ is the thermal energy and $k_{\rm B}$ is Boltzmann's constant. In particular, we indirectly sampled the number of waters in v by sampling the coarse-grained number, N_v , of waters in v, defined using a coarse-graining length of 0.01 nm and a cutoff of 0.02 nm (16). To systematically sample \tilde{N}_v , we used a series of harmonic biasing potentials, $\mathcal{U}_{\kappa,N^*}(\tilde{N}_v) = \frac{\kappa}{2}(\tilde{N}_v - N^*)^2$, where the values of κ and N^* were chosen following reference (17). 81 The free energetics, $-k_{\rm B}T \ln P_v(N)$, of water density fluctuations in v were computed following reference (16), and are shown 82 in Figure S4 for three 6×6 patches along with the corresponding values of $\Delta G_{\rm cav}$.

4. Training and Evaluating Linear Models

71

101

104

105

106

108

109

110

111

112

113

116

All linear models (M1, M2, M3, and M7) were trained using ordinary least squares regression using the scikit-learn Python package (18). The performance of each model was evaluated by five-fold randomized cross-validation (CV), and quantified using the root-mean squared error (RMSE), ε , over the five CV rounds. Five-fold CV was repeated 1000 times to obtain error 87 bars for model errors ε and for the parameters of models M2 and M3 (Figure 4D and Figure S14B).

Performance of one-feature models. To asses the usefulness of polar content, n_0 , as a predictor of hydrophobicity, we compared 89 the linear model M1 (Figure S5A, $\varepsilon = 3.67 k_B T$) to non-linear functional forms of n_0 , such as quadratic (Figure S5B, $\varepsilon = 3.42~k_{\rm B}T$) or higher order polynomials. We also considered a 'conditional average' model (Figure S5C, $\varepsilon = 3.53~k_{\rm B}T$), where the hydrophobicity of a patch with n_0 polar groups is predicted to be the average hydrophobicity, \bar{f}_{n_0} , of all patches (in the training set) with n_0 polar groups. The non-linear models perform only marginally better than model M1 (Figure S5D), suggesting that the primary limitation on the accuracy of these one-feature models is their inability to account for the variation in the hydrophobicity of patches with the same polar content, n_0 .

5. Artificial and Convolutional Neural Network Models

Artificial (ANN) and convolutional (CNN) neural networks were trained on the N=884 dataset of 6×6 patch patterns using 97 the PyTorch (19) Python library. All ANN and CNN models were trained on model M1 residuals, $f(x) - \hat{f}_{M1}(n_0(x))$.

Data augmentation and featurization of patch patterns. To account for the invariance of patch hydrophobicity under translation, rotation, and reflection, we augmented our N=884 training set of 6×6 patches with isomorphic variants, as shown in 100 Figure S6. Every patch in the augmented dataset (N = 74, 256) was embedded in a 12×11 SAM surface with a background of polar groups, and was represented using a 132-component feature vector, x. The components of x correspond to end-group 102 positions, and are assigned a value of (+1) for non-polar patch groups, (-1) for polar patch groups and (0) for polar groups 103 outside the patch.

Training neural network models. The ANN and CNN models were trained by minimizing the mean squared error (MSE) using stochastic gradient descent with the Adam (20) optimization algorithm included in the PyTorch package. For the Adam optimizer, the learning rate was set to 1×10^{-3} , β values were chosen to be 0.9 and 0.999, and $\epsilon = 1 \times 10^{-8}$. All models were trained with mini-batch gradient descent using a batch size of 200 samples. Early-stopping was used to terminate training if model MSE did not improve for 10 consecutive epochs. Rectified linear unit (ReLU) activation functions were employed in all nodes. CNNs were constructed using the Hexagdly (21) PyTorch plugin to enable convolutions over the hexagonally-packed lattice of end-groups. Convolutional filters used a kernel size of 1 with a stride length of 1 and max pooling filters employed a kernel size of 1 with a stride length of 2; the kernel size employed corresponds to the nearest-neighbor kernel shown in Figure 2B. A padding length of 1 was used for all filters. All CNN architectures employed two convolutional layers applied in sequence. In each layer, the same number of convolutional filters was applied, in parallel, to each patch pattern; the convolved patterns were passed through a ReLU activation filter and then through a max-pooling operation to reduce the dimensionality of the patterns. The output of the second convolutional layer was passed to a fully connected ANN.

Evaluating ANN and CNN model performance. As with the linear models, the performance of the ANN and CNN models was evaluated using 5-fold CV. To determine optimal network architectures for models M1A and M1C, we trained models with different hyper-parameter combinations, and evaluated their performance as an average over 5 independent CV rounds.

Determining the optimal ANN architecture. The optimal number of hidden layers and nodes per hidden layer in the ANN model were determined by training a series of models with different architectures. The performance, ε , of different ANN architectures is shown in Figure S7. The best performing ANN architecture had 1 hidden layer with 48 neurons and an output layer with one neuron, employing a total of 6,433 trained parameters. This architecture was chosen for model M1A, and was trained over 4 independent CV rounds to evaluate its performance ($\varepsilon = 2.75 k_{\rm B}T$).

Determining the optimal CNN architecture. The number of convolutional filters, hidden ANN layers, and nodes per hidden layer of the CNN model were similarly determined by training a variety of CNN models with different hyper-parameter combinations (Figure S8). We chose a near-optimal CNN model that had 5 convolutional filters, 1 hidden layer with 4 nodes and an output layer with one neuron, employing a total of 349 trained parameters. Although more complicated CNN architectures performed slightly better, we selected this architecture for model M1C because it corresponds to a very high performing model with the fewest parameters. The error of model M1C, estimated over 4 independent CV rounds, was $\varepsilon = 2.72~k_{\rm B}T$.

6. Two-body Feature Constraints and Isomorphic Forms of Model M2

To incorporate two-body or nearest-neighbor end-group correlations into our models, we must consider the number of polar $(n_{\rm oo})$, non-polar $(n_{\rm cc})$ and mixed $(n_{\rm oc})$ neighbors. Although the number of neighbors of each type varies with the patch pattern, x, the total number of nearest neighbors, $n_{\rm xx} \equiv n_{\rm oo} + n_{\rm oc} + n_{\rm cc}$ depends only on patch dimensions, p and q, for a rectangular patch:

$$n_{\rm oo} + n_{\rm oc} + n_{\rm cc} = n_{\rm xx}(p, q) = 3pq + 2(p + q) - 1.$$
 [2]

Moreover, stoichiometric and geometric considerations dictate that the number of polar and mixed neighbors must be related to the total number of polar end-groups through:

$$2n_{oo} + n_{oc} = 6n_o + [2n_{xx}(p,q) - 6pq],$$
 [3]

where the last term on the right hand side depends only on the patch dimensions. Thus, once the patch dimensions and n_o are specified, the nearest-neighbor numbers, n_{oo} , n_{oc} and n_{cc} are related to one another through equations 2 and 3 with only one of the three variables being independent. By choosing n_{oo} to be the independent variable (along with n_o), we obtain the linear model M2 for 6×6 patches as:

$$\hat{f}_{M2}(n_o, n_{oo}) = 105.68 + 5.39 \ n_o - 0.38 \ n_{oo}$$
 [4]

By equivalently choosing $n_{\rm oc}$ and $n_{\rm o}$ as the independent variables, we obtain the following isomorphic model M2:

$$\hat{f}_{M2}(n_o, n_{oc}) = 96.86 + 4.24 \, n_o + 0.19 \, n_{oc}$$
 [5]

This form of model M2 emphasizes that placing polar and non-polar groups next to one another increases patch hydrophilicity.

We can similarly choose n_{cc} and the number of non-polar groups, $n_c = pq - n_o$ as the independent variables to obtain:

$$\hat{f}_{M2}(n_c, n_{cc}) = 249.57 - 3.09 \, n_c - 0.38 \, n_{cc} \tag{6}$$

We emphasize that Equations 4, 5 and 6 contain the same information and that other isomorphic models can also be constructed by choosing one feature each from the sets, $\{n_{\rm o}, n_{\rm c}\}$ and $\{n_{\rm oo}, n_{\rm oc}, n_{\rm cc}\}$ with different feature choices providing different perspectives on the chemical determinants of hydrophobicity. For instance, the intercept in equation 6 corresponds to the hydrophobicity of a uniform polar patch (for which $n_{\rm c} = n_{\rm cc} = 0$), whereas the coefficients capture the extent to which the presence of non-polar groups, as quantified by $n_{\rm c}$, and their non-polar patterning, as quantified by $n_{\rm cc}$, contribute to patch hydrophobicity. This form of model M2 clarifies that the addition of a non-polar end-group to a purely polar patch increases its hydrophobicity (i.e., decreases \hat{f} by 3.09 $k_{\rm B}T$), its effect is smaller than the addition of a polar group to a purely non-polar patch (which increases \hat{f} by 5.39 $k_{\rm B}T$ according to equation 4).

7. Three-body Feature Constraints

There are 16 possible three-body features (Figure S9) that span different chemistries ('ooo', 'ccc', 'ooc', 'cco', 'cco'

$$n_{\text{ooo}}^{\text{compact}} + n_{\text{ccc}}^{\text{compact}} + n_{\text{ooc}}^{\text{compact}} + n_{\text{occ}}^{\text{compact}} = n_{\text{xxx}}^{\text{compact}}(p,q) = 2pq + 2(p+q)$$

$$n_{\text{ooo}}^{\text{bent}} + n_{\text{ccc}}^{\text{bent}} + n_{\text{ooc}}^{\text{bent}} + n_{\text{occ}}^{\text{bent}} + n_{\text{occ}}^{\text{bent}} + n_{\text{ccc}}^{\text{bent}} = n_{\text{xxx}}^{\text{bent}}(p,q) = 6pq + 7p + 8q - 2$$

$$n_{\text{ooo}}^{\text{straight}} + n_{\text{ccc}}^{\text{straight}} + n_{\text{occ}}^{\text{straight}} + n_{\text{occ}}^{\text{straight}} + n_{\text{ccc}}^{\text{straight}} = n_{\text{xxx}}^{\text{straight}}(p,q) = 3pq + 4(p+q) - 2$$

We chose to eliminate the purely non-polar terms $n_{\text{ccc}}^{\text{compact}}$, $n_{\text{ccc}}^{\text{bent}}$, and $n_{\text{ccc}}^{\text{straight}}$. Furthermore, by incorporating polar content, n_{o} , and the total number of polar-polar neighbors, n_{oo} , we can construct 8 additional constraint equations, enumerated in Figure S10:

$$\begin{split} n_{\text{ooc}}^{\text{compact}} + 3n_{\text{ooo}}^{\text{compact}} - 2n_{\text{oo}} &= 2(p+q) + 2 \\ n_{\text{ooc}}^{\text{bent}} + 2n_{\text{ooo}}^{\text{bent}} - 4n_{\text{oo}} &= 6p + 8q \\ n_{\text{ooc}}^{\text{straight}} + 2n_{\text{ooo}}^{\text{straight}} - 2n_{\text{oo}} &= 4(p+q) - 2 \\ n_{\text{occ}}^{\text{compact}} - 3n_{\text{ooo}}^{\text{compact}} + 4n_{\text{oo}} - 6n_{\text{o}} &= 2(p+q) - 4 \\ n_{\text{occ}}^{\text{bent}} + 2n_{\text{oco}}^{\text{bent}} + 4n_{\text{oo}} - 12n_{\text{o}} &= 8(p+q) - 4 \\ n_{\text{occ}}^{\text{straight}} + 2n_{\text{oco}}^{\text{straight}} + 2n_{\text{oo}} - 6n_{\text{o}} &= 4(p+q) - 2 \\ n_{\text{coc}}^{\text{bent}} - n_{\text{ooo}}^{\text{bent}} + 4n_{\text{oo}} - 6n_{\text{o}} &= p - 2 \\ n_{\text{coc}}^{\text{straight}} - n_{\text{ooo}}^{\text{straight}} + 2n_{\text{oo}} - 3n_{\text{o}} &= 0 \end{split}$$

These constraints result in the set of 7 linearly independent features, $\{n_{\rm o}, n_{\rm oo}, n_{\rm ooo}^{\rm compact}, n_{\rm ooo}^{\rm bent}, n_{\rm ooo}^{\rm straight}, n_{\rm oco}^{\rm bent}, n_{\rm oco}^{\rm straight}\}$, which were used to construct model M7.

8. Learning Curves: Model Performance vs Size of the Training Dataset

161

162

163

166

167

168

169

170

171

174

175

176

177

178

180

181

182

183

184

185

186

187

188

189

191

192

All our models for predicting the hydrophobicity of 6×6 patterned SAM patches were trained using a library comprising N=884 patch patterns. To assess how model performance depends on the size of the training dataset, we randomly select M=177 (roughly one-fifth of N) patches from the library and designate them to be the test-set. This procedure is repeated a total of 5 times using a fifth of the dataset as the test set each time. Each model is then trained on increasingly larger subsets of the remaining N-M=767 patches, and model error, ε , is reported using the RMSE calculated on the five different test sets. As the selection of the testing set is randomized, this procedure is repeated 4 times to estimate uncertainties in ε for models M1A and M1C, and 1000 times to estimate the uncertainties in ε for the linear models. The learning curves, which plot model performance, ε , as a function of training set size, are shown for both the linear (Figure S11A) and neural network (Figure S11B) models. The performance of the linear models appears to saturate as the size of the training set is increased, whereas the performance of the neural network models continues to improve (for $< 10^3$ patches), suggesting that increasing the size of the training set is unlikely to significantly improve the performance of models M1 to M7, whereas larger libraries are likely to improve the performance of models M1A and M1C. In Figure S11B, we also plot the performance of the best-performing model as a function of training set size (multi-colored line); model M2 (orange segment) initially out-performs all other models for small libraries (< 30 patches), whereas the model M3 performs optimally when modest training sets (< 300 patches) are used, and model M7 performs best when trained on the largest datasets. These results suggest that as the size of the training dataset is increased, the neural network models will continue to display improved performance by learning higher order correlations in the data; in contrast, the linear models will need to employ additional descriptors that capture increasingly higher order correlations to optimally utilize larger datasets.

9. Explaining Models M1A and M1C using Surrogate Models

To shed light into why our neural network (NN) models perform well as they do, we draw inspiration from the Local Interpretable Model-agnostic Explanations (LIME) technique (22), proposed by Ribeiro and coworkers to explain predictions made by black-box models. LIME relies on training interpretable 'surrogate' models to explain the predictions of more complex, un-interpretable models. Based on the excellent performance of our linear models (M2, M3 and M7), we propose using them as interpretable surrogates for explaining our NN models (M1A and M1C). To this end, we first consider a substantially expanded dataset with 88,002 patches, which contains not only the 884 patches included in the original training set, but also an additional 87,118 patterns generated using Wang-Landau sampling. We then used a NN model M1Y (where $Y \in \{A,C\}$) to predict the hydrophobicity of these 88,002 patches, and using these predictions, we trained a surrogate model MX (where $X \in \{2,3,7\}$). We found that all surrogate models MX | M1Y explain the underlying NN model remarkably well with errors of roughly 1 k_BT , and that the coefficients of the surrogate models are in good agreement with those of linear models MX obtained by training with the simulation dataset of 884 patches. The surrogate models and their performances are as follows:

Name	Model	$\varepsilon (k_{\rm B}T)$
$\hat{f}_{ ext{M2}}$	$105.7 + 5.39 n_{\rm o} - 0.384 n_{\rm oo}$	_
$\hat{f}_{\mathrm{M2 M1A}}$	$105.59 + 5.313 \ n_{\rm o} - 0.3545 \ n_{\rm oo}$	0.99
$\hat{f}_{\mathrm{M2 M1C}}$	$106.07 + 5.319 \ n_{\rm o} - 0.3637 \ n_{\rm oo}$	1.00
$\hat{f}_{ m M3}$	$105.8 + 5.40 n_{\rm o} - 0.49 (n_{\rm oo} + n_{\rm ooo}^{\rm compact}) + 0.13 (n_{\rm ooo}^{\rm bent} + n_{\rm ooo}^{\rm straight})$	_
$\hat{f}_{\mathrm{M3 M1A}}$	$106.00 + 5.154 n_{\rm o} - 0.292 (n_{\rm oo} + n_{\rm ooo}^{\rm compact}) + 0.0547 (n_{\rm ooo}^{\rm bent} + n_{\rm ooo}^{\rm straight})$	1.02
$\hat{f}_{ m M3 M1C}$	$106.42 + 5.191 n_{\rm o} - 0.331 (n_{\rm oo} + n_{\rm ooo}^{\rm compact}) + 0.0693 (n_{\rm ooo}^{\rm bent} + n_{\rm ooo}^{\rm straight})$	0.99
$\hat{f}_{ ext{M7}}$	$105.7 + 5.89 n_{\rm o} - 0.64 n_{\rm oo} - 0.51 n_{\rm oo}^{\rm compact} + 0.16 n_{\rm ooo}^{\rm bent} + 0.11 n_{\rm ooo}^{\rm straight} - 0.05 n_{\rm oco}^{\rm bent} - 0.04 n_{\rm oco}^{\rm straight}$	_
$\hat{f}_{\mathrm{M7 M1A}}$	$105.08 + 5.703 \ n_{\rm o} - 0.658 \ n_{\rm oo} - 0.013 \ n_{\rm oo}^{\rm compact} + 0.045 \ n_{\rm oo}^{\rm best} + 0.078 \ n_{\rm ooo}^{\rm straight} - 0.023 \ n_{\rm oco}^{\rm straight} - 0.032 \ n_{\rm oco}^{\rm straight}$	0.93
$\hat{f}_{\mathrm{M7 M1C}}$	$105.91 + 5.695 \ n_{\rm o} - 0.582 \ n_{\rm oo} - 0.202 \ n_{\rm ooo}^{\rm compact} + 0.103 \ n_{\rm ooo}^{\rm bent} + 0.025 \ n_{\rm ooo}^{\rm straight} - 0.044 \ n_{\rm oco}^{\rm bent} + 0.014 \ n_{\rm oco}^{\rm straight}$	0.93

These findings suggest that the trained NN models, which operate on the entire chemical pattern of the patch, x, extract information pertaining to local chemical correlations between patch end-groups that is encoded in the descriptors (e.g., n_{oo}) used to construct the linear models (M2, M3 and M7).

10. Uniform Non-Polar Patches of Various Shapes and Sizes

193

194

195

196

197

199

200

201

To construct a physically-motivated model for the hydrophobicity of a purely non-polar patch, f_c , as a function of its size and shape, we note that cavity formation is expected to be driven by interfacial physics, and hence its free energetic cost should be proportional to the cross-sectional area, $l_x l_y$, of each cuboidal probe volume v. To account for possible edge-effects due to the finite size of the patches considered here, we additionally included the perimeter of the v in the x and y dimensions, ($2l_x$ and $2l_y$, respectively, for a rectangular $p \times q$ patch) to arrive at the following linear model:

$$\hat{f}_{c}(l_{x}, l_{y}) = c_{1}(l_{x}l_{y}) + c_{2}(l_{x}) + c_{3}(l_{y})$$
[7]

We trained this model on a dataset of 62 purely non-polar patches of different sizes and shapes, resulting in the following learned coefficients:

$$c_1 = 5.63 k_B T/nm^2$$

 $c_2 = 11.85 k_B T/nm$
 $c_3 = 7.54 k_B T/nm$

To express \hat{f}_c as a function of patch dimensions p and q, we substituted the relationships, $l_x(p)$ and $l_y(q)$ (Equation 1), into Equation 7 to obtain:

$$\hat{f}_c(p,q) = \hat{f}_c(l_x(p), l_y(q))$$

$$= [c_1 r_x r_y] pq + [(2\delta c_1 - 0.5 r_y c_1 + c_2) r_x] p + [(2\delta c_1 - r_x c_1 + c_3) r_y] q$$

$$+ [c_1 (2\delta - r_x) (2\delta - 0.5 r_y) + c_2 (2\delta - r_x) + c_3 (2\delta - 0.5 r_y)]$$

The coefficients and intercept of $\hat{f}_c(p,q)$ were then determined from the trained coefficients c_1, c_2 , and c_3 and the known values of δ , r_x and r_y to obtain:

$$\hat{f}_c(p,q) = 1.22pq + 5.86p + 4.10q + 3.85$$

This model is able to capture the hydrophobicity of rectangular non-polar patches with an error of $\varepsilon = 1.87~k_{\rm B}T$.

Extending \hat{f}_c to non-rectangular patches. To generalize \hat{f}_c to non-rectangular patches, we recast the model in terms of the total number of patch end-groups, $n_{\rm tot}$, and the number of end-groups along the perimeter of patch in the x and y dimensions, n_x and n_y , respectively (Figure S12A). Because $n_{\rm tot} = pq$, $n_x = 2p$ and $n_y = 2q$ for a rectangular $p \times q$ patch, our generalized model for non-polar patch hydrophobicity becomes:

$$\hat{f}_{c}(n_{\text{tot}}, n_x, n_y) = 1.22n_{\text{tot}} + 2.93n_x + 2.05n_y + 3.85$$

To test this model, we estimated the hydrophobicity, f_c , of the 7 non-rectangular non-polar patches shown in Figure S12B, and found that the generalized model $\hat{f}_c(n_{\rm tot}, n_x, n_y)$ is able to accurately predict f_c with an error of $\varepsilon = 4.75~k_{\rm B}T$.

11. Uniform Polar Patches of Various Shapes and Sizes

We also construct a physically-motivated model for the hydrophobicity of a purely polar rectangular patch, f_o , as a function of its dimensions, p and q. To do so, we estimate f_o for 62 patches of different sizes and shapes, and use this dataset to train a model $\hat{f}_o(p,q)$ that is linear in pq, p and q, resulting in:

$$\hat{f}_{o}(p,q) = 5.81pq + 5.12p + 1.79q + 1.51$$

As shown in Figure S13, this model is able to predict f_0 quite well with an error of $\varepsilon = 4.01 k_B T$.

We note that the generalized model M2: $\hat{f}_{\text{M2}}(p,q,n_{\text{o}},n_{\text{oo}}) = \alpha_{\text{c}}(p,q) + \alpha_{\text{o}}n_{\text{o}} + \alpha_{\text{oo}}n_{\text{oo}}$, discussed in the main text, can also be used to predict the hydrophobicity, $f_{\text{o}}(p,q)$, of polar rectangular patches of all shapes and sizes. In particular, focusing on the leading order terms (i.e., those proportional to pq), and noting that for a purely polar patch, $n_{\text{o}} = pq$ and n_{oo} grows as 3pq, we obtain the constraint: $5.81 = 1.22 + \alpha_{\text{o}} + 3\alpha_{\text{oo}}$, where 5.81 and 1.22 are the coefficients of pq in the models for $\hat{f}_{\text{o}}(p,q)$ and $\hat{f}_{\text{c}}(p,q)$, respectively. Plugging the values of $\alpha_{\text{o}} = 5.07$ and $\alpha_{\text{oo}} = -0.32$, obtained for model M2, into the right hand side of the constraint equation, we obtain 5.33, which is within 10% of 5.81 (left hand side of the constraint).

208

12. Generalizing Model M3 to Rectangular Patches of all Shapes and Sizes

To generalize model M3 $[\hat{f}_{\text{M3}} = \alpha_{\text{c}} + \alpha_{\text{o}} n_{\text{o}} + \alpha_{\text{oo}} (n_{\text{oo}} + n_{\text{ooo}}^{\text{compact}}) + \alpha_{\text{ooo}} (n_{\text{ooo}}^{\text{straight}} + n_{\text{ooo}}^{\text{bent}})]$ to rectangular patches of all shapes 215 and sizes, we first note that the intercept, α_c , in model M3 corresponds to the hydrophobicity, f_c , of a uniform non-polar 216 patch for which $n_{\rm o} = n_{\rm oo} = n_{\rm ooo}^{\rm compact} = n_{\rm ooo}^{\rm straight} = n_{\rm ooo}^{\rm bent} = 0$. Thus, we expect $\alpha_{\rm c}$ to depend on patch dimensions according 217 to: $\alpha_c(p,q) = \hat{f}_c(p,q) = 1.22pq + 5.86p + 4.10q + 3.85$. To assess whether the model M3 coefficients, α_o , α_{oo} and α_{ooo} , are 218 independent of patch size and shape, we compare their values obtained by training the original 6×6 and the expanded 219 datasets (Figure S14). We find good qualitative agreement across all the trained coefficients, and reasonable quantitative agreement, especially for the coefficients that are larger in magnitude (and therefore more important). Our generalized 221 model M3, trained on the expanded dataset, which can be used to predict the hydrophobicity of heterogeneous rectangular 222 patches of all shapes and size, and is thus given by: $\hat{f}_{M3} = \alpha_{\rm c}(p,q) + \alpha_{\rm o}n_{\rm o} + \alpha_{\rm oo}(n_{\rm oo} + n_{\rm ooo}^{\rm compact}) + \alpha_{\rm ooo}(n_{\rm ooo}^{\rm straight} + n_{\rm ooo}^{\rm bent}) = 1.22pq + 5.86p + 4.10q + 3.85 + 4.98n_{\rm o} - 0.28(n_{\rm oo} + n_{\rm ooo}^{\rm compact}) + 0.06(n_{\rm ooo}^{\rm straight} + n_{\rm ooo}^{\rm bent}).$ 223 224

13. Patterned Patches of Different Sizes

214

225

226

227

231

232

233

234

235

236

237

To interrogate the influence of chemical patterning on hydrophobicity for patches of different sizes, we studied 6×6 , 8×8 , 10×10 and 12×12 patterned patches. For each patch size, we estimated the hydrophobicity, f, of two judiciously chosen patches that had the same number of polar groups but different chemical patterns; both patches had an equal number of polar and non-polar groups, but their non-polar end-groups were either clustered or dispersed, as shown in Figure S15. We chose dispersed and clustered patches with the same n_o because model M1 would predict that these patches should have the same hydrophobicity. In contrast, we found that differences in chemical patterning gave rise to differences in hydrophobicity for all patch sizes, and that these differences increased with patch size (Figure S15). Moreover, our generalized models M2 and M3 correctly predict that patches with well-dispersed end-groups are more hydrophilic than those with clustered end-groups, and they capture the differences in hydrophobicity with chemical patterning reasonably well across the different patch sizes.

14. Rational Design of Patterned Patches using Model M3

We used model M3 to perform the hot-spot analysis and the greedy design described in the main text, and found that the locations of the hot-spot end-groups as well as the characteristics of the maximally hydrophobic or hydrophilic patches (Figure S16) are qualitatively similar to those discovered by model M2 (Figure 5). Additionally, we compare the hydrophobicity of the maximally hydrophilic (\hat{f}^{max}) and the maximally hydrophobic (\hat{f}^{min}) patches with a certain polar content, n_{o} , obtained using models M2 and M3 (Figure S17A). In both cases, the maximum range of attainable hydrophobicities, i.e., $\hat{f}^{\text{max}} - \hat{f}^{\text{min}}$, is realized at roughly 50% polar coverage (Figure S17B).

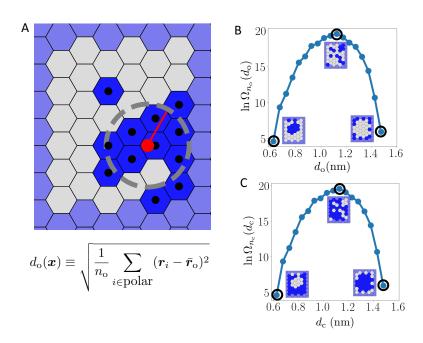


Fig. S1. Using polar end-group clustering to construct a library of diverse patterns. (A) The degree of polar clustering, $d_{\rm o}$, of a patch pattern, x, is defined as the root mean squared deviation of the positions, r_i (black dots), of all $n_{\rm o}$ polar end-groups (hydroxyl, blue) in the patch around their centroid, $\bar{r}_{\rm o}$ (red dot). The radius (red) of the illustrated circle (gray) corresponds to $d_{\rm o}$ of the pattern shown here. Non-polar (methyl) end-groups are shown in white and polar groups outside of the patch are shown in light blue. (B) The number of 6×6 patches with $n_{\rm o}$ polar groups, $\Omega_{n_{\rm o}}(d_{\rm o})$, shown here for $n_{\rm o}=12$, is peaked at $d_{\rm o}\approx 1.1~{\rm nm}$, highlighting that highly clustered (lower left) or dispersed (lower right) patch patterns are improbable. (C) To obtain patterns spanning a wide range of $d_{\rm c}$ (degree of non-polar clustering, defined analogously to $d_{\rm o}$) values, we swap the identities of polar (hydroxyl, blue) and non-polar (methyl, white) end-groups to generate 'inverted' patches. The number of 6×6 patches with $n_{\rm c}$ non-polar groups, $\Omega_{n_{\rm c}}(d_{\rm c})$, is shown here for $n_{\rm c}=12$.

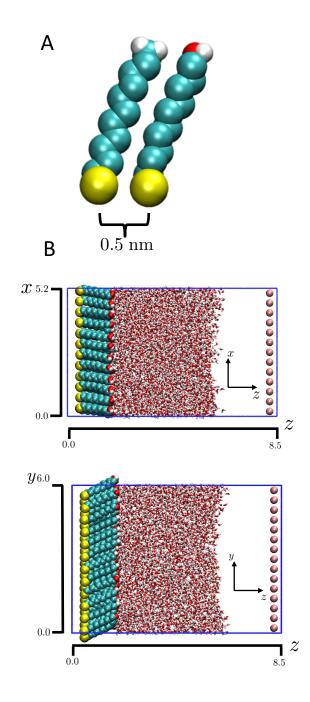


Fig. S2. Simulating self-assembled monolayer (SAM) surfaces. (A) The SAM surfaces are comprised of hexagonally-packed alkyl chains, which contain a sulfur atom (yellow), 10 united-atom methylene groups (teal), and either a non-polar methyl (left, teal/white) or polar hydroxyl (right, red/white) end-group (B) The sulfur atoms are spaced 0.5 nm apart and are arranged on a hexagonal lattice in the x-y plane. The SAM surface (space-fill representation) is solvated using a roughly 4 nm thick slab of water (red/white, licorice representation) A repulsive wall is placed at $z\approx 8.2$ nm (pink spheres) to induce a buffering vapor layer.

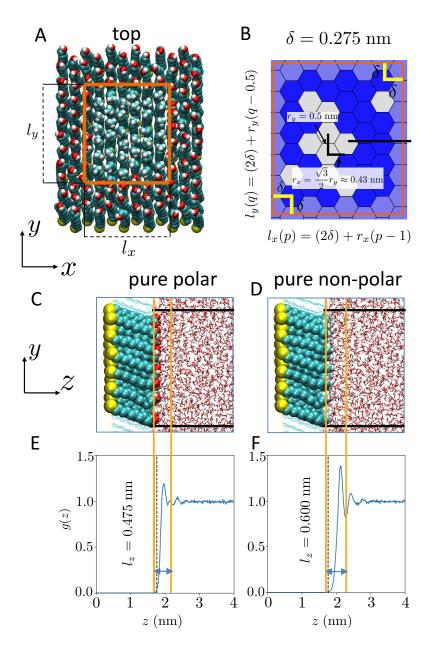


Fig. S3. Defining the cuboidal probe volume, v. (A) A top-down view illustrating a non-polar 6×6 pattern and the dimensions, l_x and l_y , of v (orange) in the x and y directions, respectively. (B) The dimensions, l_x and l_y , are determined by calculating the average positions of the lower-left and upper-right patch end-groups in a patch and padding them with a buffer, δ (yellow lines) to ensure that v encompasses all end-groups in the patch. The relationships between the dimensions of v, i.e., l_x and l_y , and the patch dimensions, p or q are shown, where r_x and r_y are the average end-group spacings in the x and y directions, respectively. (C, D) Simulation snapshots of water near the purely polar and non-polar 6×6 patches, and (E, F) the corresponding normalized water density profiles, g(z). To ensure that probe volumes adjacent to all patches have roughly the same average number of waters, their lower z-bounds were fixed at $z_{\min} = 1.7$ nm, and their upper bounds, z_{\max} , were varied. Choosing $z_{\max} = 2.175$ nm for the polar patch and $z_{\max} = 2.3$ nm for the non-polar patch results in 100 waters in both probe volumes; for patterned surfaces containing a mixture of polar and non-polar end-groups, z_{\max} was chosen to be between these two extremes such that roughly 100 waters were present within the probe volume.

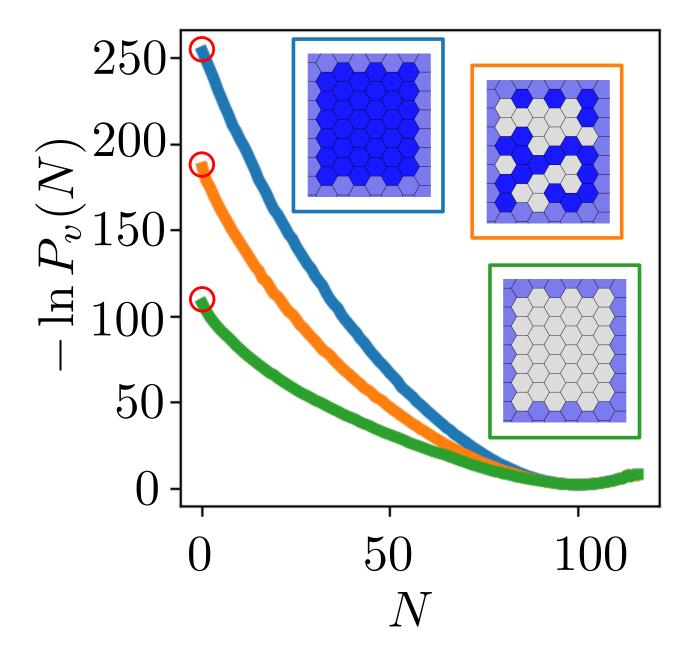
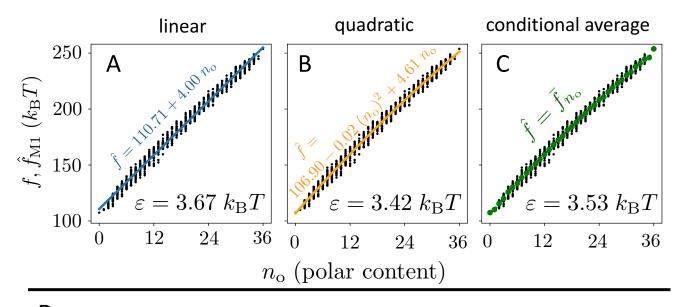


Fig. S4. The free energetics, $-\ln P_v(N)$, of observing N waters in a probe volumes, v, adjacent to polar (blue), non-polar (green) and mixed (orange) 6×6 patches are shown (in units of $k_{\rm B}T$), and the free energy of cavity creation, $-\ln P_v(0)$, which serves to quantify patch hydrophobicity, f, is highlighted (red circles).



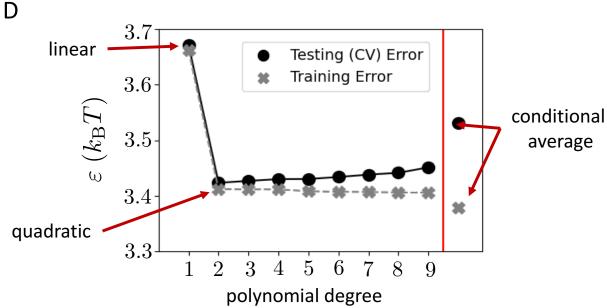


Fig. S5. Non-linear functions of polar content, n_o , do not perform significantly better than model M1, which is linear in n_o . (A-C) The hydrophobicity, f, of every patch in the library is plotted as a function of the polar content, n_o , of the patch, and the data is fit to (A) linear, (B) quadratic and (C) conditional average models. The cross-validated errors, ε , are also shown for each model. Although the models capture the trend of hydrophilicity increasing with polar content, they fail to capture the variation in f for patches with the same n_o ; therefore, the non-linear models perform only marginally better than the linear model M1. (D) The testing (black) and total in-sample training (gray) error of different one-feature (n_o) models are shown. As expected, increasing model complexity (e.g., polynomial degree) results in a decrease in the in-sample training error, but the testing error increases as more complex models overfit the available data. In fact, the conditional average model, which provides the best-possible in-sample training performance, displays as substantial gap between in-sample and testing performance, highlighting the susceptibility of complex models to over-fitting.

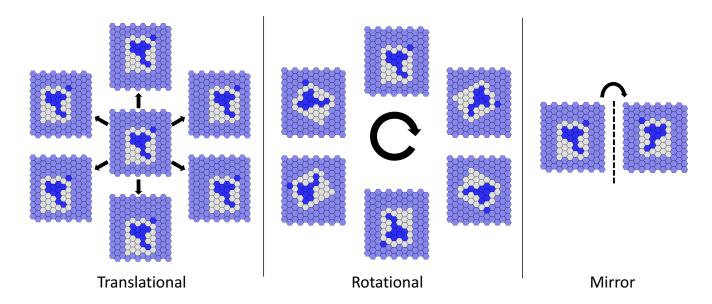


Fig. S6. Augmenting the training library of 6×6 patches with translated, rotated and reflected variants. To account for the invariance of patch hydrophobicity under translation, rotation and reflection, we augmented the training dataset with 83 isomorphic variants for every patch $(7 \times 6 \times 2)$ for translational, rotational, and reflection factors, respectively); select variants are shown for a representative patch pattern. To accommodate the variants, every patch was embedded in a 12×11 SAM surface with a background of polar groups; polar end-groups outside the patch are shown in light blue.

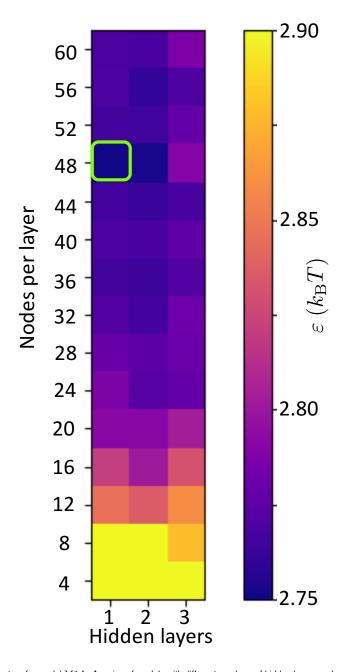


Fig. S7. Determining the optimal ANN architecture for model $\rm M1A$. A series of models with different numbers of hidden layers and nodes per hidden layer was systematically trained and the variation of model performance, ε , with the choice of hyper-parameters is shown. The model with the best performance ($\varepsilon=2.75~k_{\rm B}T$) contains one hidden layer with 48 nodes and is highlighted in green.

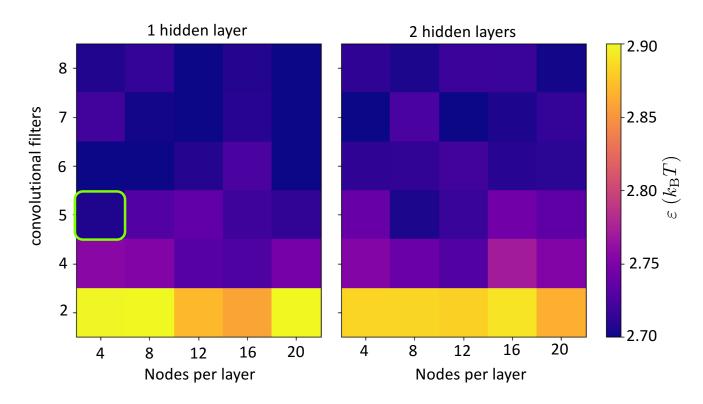


Fig. S8. Determining the optimal CNN architecture for model M1C. A series of models with different numbers of convolutional filters, hidden layers and nodes per hidden layer were systematically trained and the variation of model performance, ε , with the choice of hyper-parameters is shown. The CNN architecture with 5 convolutional filters and 1 hidden ANN layer with 4 hidden nodes (green rectangle) was chosen because it gave rise to the simplest model with high performance.

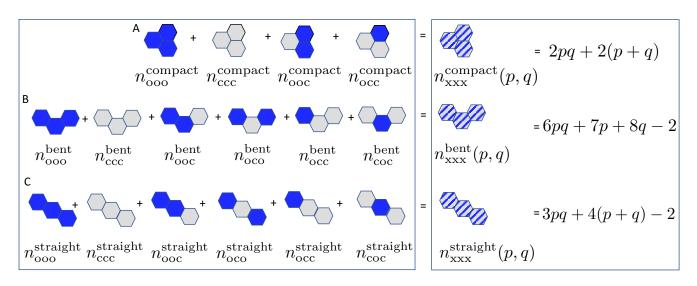


Fig. S9. Enumerating the 16 three-body features (left box). The total number of (A) 'compact', (B) 'bent', or (C) 'straight' three-body features must be constant $(n_{\text{xxx}}^{\text{compact}}, n_{\text{xxx}}^{\text{bent}}, n_{\text{xxx}}^{\text{bent}}, n_{\text{xxx}}^{\text{bent}})$, respectively) for patches with specific dimensions, p and q. Using these constraints, we eliminate the following three features: $n_{\text{ccc}}^{\text{compact}}, n_{\text{ccc}}^{\text{bent}}$ and $n_{\text{ccc}}^{\text{straight}}$.

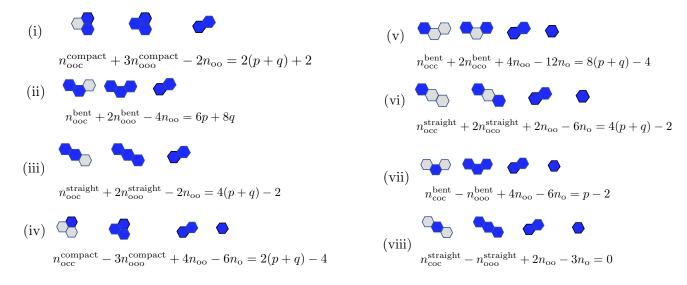


Fig. S10. Stoichiometric and geometric constraints relating three-body features to $n_{\text{\tiny O}}$ and $n_{\text{\tiny OO}}$ as well as the patch dimensions, p and q. The eight constraint equations, listed here, enable elimination of eight features in favor of the remaining features, $n_{\text{\tiny OO}}$, p and q. For each constraint equation, the first feature on the left-hand side is eliminated.

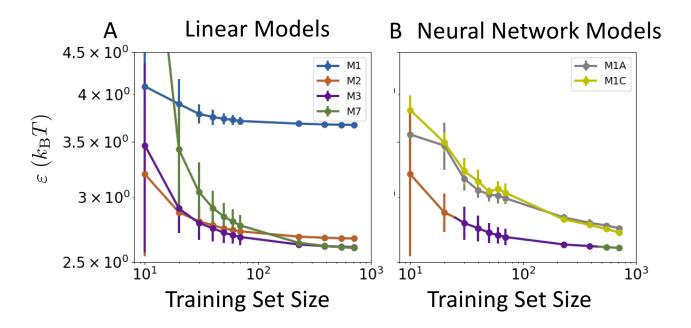


Fig. S11. Dependence of model performance on the size of the training dataset. Learning curves for the (A) linear and (B) neural network models highlight how model error, ε , varies with the size of the training dataset. The performance of the linear models plateaus as the training set size is increased, whereas that of the neural network models continues to improve over the entire range of training set sizes. The orange, purple, and green curve (corresponding to models M2, M3 and M7, respectively) in panel B corresponds to the convex hull of the best-performing models for each training set size.

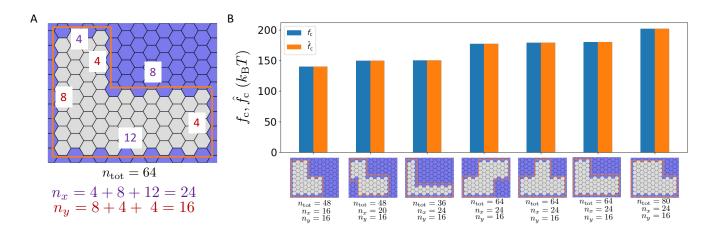


Fig. S12. Predicting the hydrophobicity of uniform non-polar patches with arbitrary (non-rectangular) shapes. (A) To generalize \hat{f}_c to non-rectangular patches, we recast the model in terms of the total number of patch end-groups, $n_{\rm tot}$, and the number of peripheral end-groups, n_x and n_y , in the x and y dimensions, respectively. The values of $n_{\rm tot}$, n_x (purple segments), and n_y (brown segments) are shown for a sample non-rectangular non-polar patch (orange). (B) To test the performance of the generalized model, $\hat{f}_c(n_{\rm tot}, n_x, n_y)$, we estimated the hydrophobicity, f_c , of 7 non-rectangular non-polar patches (gray, outlined in orange), and found that the model accurately predicts f_c ($\varepsilon = 4.75~k_{\rm B}T$); the values of $n_{\rm tot}, n_x$, and n_y for each patch are also shown.

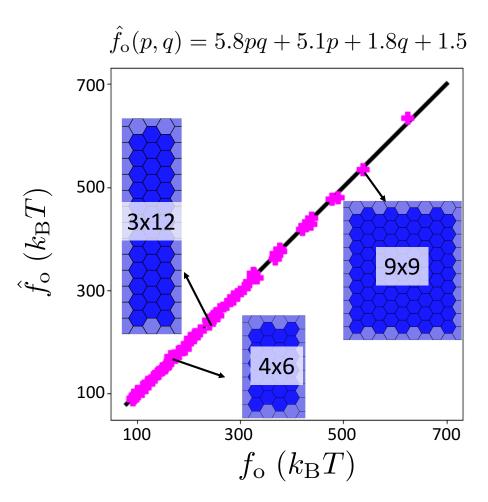


Fig. S13. Predicting the hydrophobicity of uniform polar patches, $f_{\rm o}$, as a function of their dimensions, p and q. We estimate $f_{\rm o}$ for 62 patches of different sizes and shapes, and use this dataset to train a model $\hat{f}_{\rm o}(p,q)$ that is linear in pq, p and q. The resulting model predictions, $\hat{f}_{\rm o}$, plotted against the estimated $f_{\rm o}$ -values, illustrate the performance of the model ($\varepsilon = 4.02~k_{\rm B}T$).

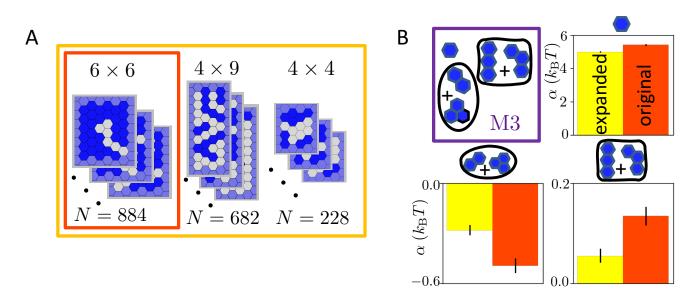


Fig. S14. Generalizing model M3 to rectangular patches of different sizes and shapes. (A) In addition to the N=884 patches (6×6) in the original dataset (red box), the expanded dataset (yellow box) also contains N=682 patches of a different shape (4×9) and N=228 patches of a different size (4×4) . (B) The coefficients obtained by training model M3 on the original and expanded datasets agree reasonably well.

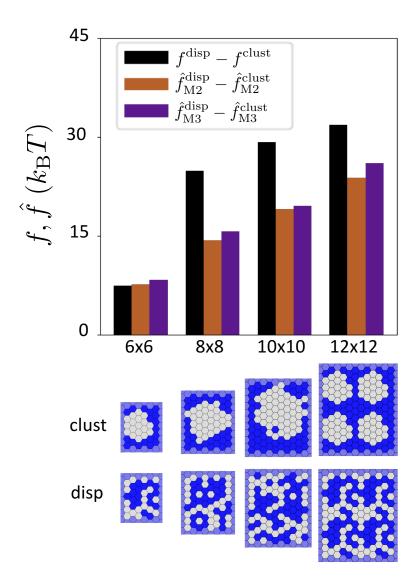


Fig. S15. The influence of chemical patterning on hydrophobicity for different patch sizes. Two patches with different chemical patterns (clustered and dispersed), but the same number of polar groups (equal to half of the total groups), display substantial differences in their hydrophobicity, f, for a range of patch sizes (from 6×6 to 12×12). Because the dispersed and clustered patterns have the same n_o , the additive model M1 would predict that the difference in their hydrophobicities, $f^{\rm disp} - f^{\rm clus}$, should be equal to 0 regardless of patch size. In contrast, the extended models M2 and M3 correctly predict that patches with well-dispersed end-groups are more hydrophilic than those with clustered end-groups, and they capture the differences in hydrophobicity with chemical patterning reasonably well across the different patch sizes.

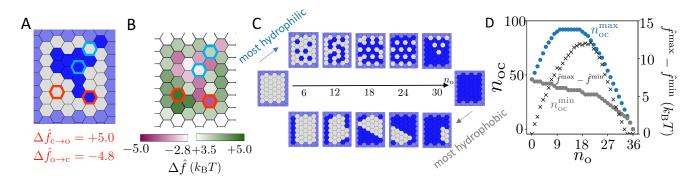


Fig. S16. Using model M3 for the rational design of patterned SAM patches with maximal hydrophobicity/hydrophilicity. (A) For the patch pattern shown, an exhaustive screening of all possible point mutations is performed using model M3. (B) The changes in patch hydrophobicity, $\Delta \hat{f}$, in response to point mutations, enables identification of 'hot-spot' locations (red outline). Mutations that result in the smallest absolute change in patch hydrophobicity are also highlighted (cyan outline). (C) The maximally hydrophilic (top) and hydrophobic (bottom) patches with a particular polar content, $n_{\rm o}$, were uncovered using an iterative, greedy design protocol; the most hydrophilic patches tend to disperse their polar end-groups, whereas the most hydrophobic patches cluster their non-polar end-groups. (D) To quantify the relative dispersion of end-groups in the optimal patches, we plot the number of polar – non-polar neighbors, $n_{\rm oc}$, for the most hydrophilic $(n_{\rm oc}^{\rm max})$, blue) and the most hydrophobic $(n_{\rm oc}^{\rm min})$, gray) patches as a function of polar content, $n_{\rm o}$. We also plot the difference, $\Delta n_{\rm oc} = n_{\rm oc}^{\rm max} - n_{\rm oc}^{\rm min}$.

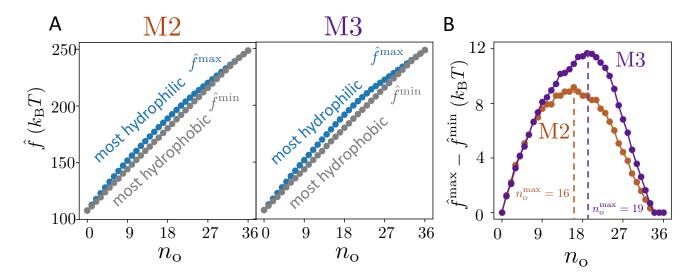


Fig. S17. (A) The hydrophobicities of the most hydrophilic (\hat{f}^{\max}) , blue) and the most hydrophobic (\hat{f}^{\min}) patches, uncovered using the iterative greedy design protocol with either model M2 (left) or model M3 (right), are shown as a function of patch polar content, n_o . (B) The difference, $\hat{f}^{\max} - \hat{f}^{\min}$, corresponds to the range of hydrophobicities attainable by patches with given polar content, n_o , and is shown for the optimal patches obtained using models M2 (orange) and M3 (purple); the models predict that this difference is greatest for patches with roughly 50% polar content.

242 References

243

244

249

250

251

252

253

256

257

258

259

260

261

262

263

264

265

266

267

271

272

- 1. F. Wang and D. P. Landau. Efficient, multiple-range random walk algorithm to calculate the density of states. *Phys. Rev. Lett.*, 86(10):2050–2053, March 2001. ISSN 0031-9007.
- Natalia Shenogina, Rahul Godawat, Pawel Keblinski, and Shekhar Garde. How Wetting and Adhesion Affect Thermal
 Conductance of a Range of Hydrophobic to Hydrophilic Aqueous Interfaces. *Physical Review Letters*, 102(15):156101,
 April 2009.
 - 3. Rahul Godawat, Sumanth N. Jamadagni, and Shekhar Garde. Characterizing hydrophobicity of interfaces by using cavity formation, solute binding, and water correlations. *Proc. Natl. Acad. Sci. U.S.A.*, 106(36):15119–15124, September 2009. ISSN 0027-8424, 1091-6490.
 - 4. Erte Xi, Vasudevan Venkateshwaran, Lijuan Li, Nicholas Rego, Amish J. Patel, and Shekhar Garde. Hydrophobicity of proteins and nanostructured solutes is governed by topographical and chemical context. *Proc. Natl. Acad. Sci. U.S.A.*, 114(51):13345–13350, 2017.
 - Berk Hess, Carsten Kutzner, David van der Spoel, and Erik Lindahl. GROMACS 4: Algorithms for Highly Efficient, Load-Balanced, and Scalable Molecular Simulation. *Journal of Chemical Theory and Computation*, 4(3):435–447, March 2008. ISSN 1549-9618.
 - 6. R. W Hockney, S. P Goel, and J. W Eastwood. Quiet high-resolution computer models of a plasma. *Journal of Computational Physics*, 14(2):148–158, February 1974. ISSN 0021-9991.
 - 7. Giovanni Bussi, Davide Donadio, and Michele Parrinello. Canonical sampling through velocity rescaling. The Journal of Chemical Physics, 126(1):014101, January 2007. ISSN 0021-9606.
 - 8. H. J. C. Berendsen, J. R. Grigera, and T. P. Straatsma. The missing term in effective pair potentials. *The Journal of Physical Chemistry*, 91(24):6269–6271, November 1987. ISSN 0022-3654, 1541-5740.
 - 9. Maurizio Mondello, Gary S. Grest, Edmund B. Webb, and P. Peczak. Dynamics of n-alkanes: Comparison to Rouse model. The Journal of Chemical Physics, 109(2):798–805, June 1998. ISSN 0021-9606.
 - Viktor Hornak, Robert Abel, Asim Okur, Bentley Strockbine, Adrian Roitberg, and Carlos Simmerling. Comparison of multiple Amber force fields and development of improved protein backbone parameters. *Proteins*, 65(3):712–725, November 2006. ISSN 1097-0134.
- William L. Jorgensen, David S. Maxwell, and Julian Tirado-Rives. Development and Testing of the OPLS All-Atom Force
 Field on Conformational Energetics and Properties of Organic Liquids. *Journal of the American Chemical Society*, 118
 (45):11225-11236, November 1996. ISSN 0002-7863.
 - 12. Tom Darden, Darrin York, and Lee Pedersen. Particle mesh Ewald: An N·log(N) method for Ewald sums in large systems. The Journal of Chemical Physics, 98(12):10089–10092, June 1993. ISSN 0021-9606.
- 273 13. Berk Hess, Henk Bekker, Herman J. C. Berendsen, and Johannes G. E. M. Fraaije. LINCS: A linear constraint solver for molecular simulations. *Journal of Computational Chemistry*, 18(12):1463–1472, September 1997. ISSN 1096-987X.
- 275 14. Shuichi Miyamoto and Peter A. Kollman. Settle: An analytical version of the SHAKE and RATTLE algorithm for rigid vater models. *Journal of Computational Chemistry*, 13(8):952–962, October 1992. ISSN 1096-987X.
- 277 15. Amish J. Patel, Patrick Varilly, and David Chandler. Fluctuations of water near extended hydrophobic and hydrophilic 278 surfaces. J Phys. Chem. B, 114(4):1632–1637, February 2010. ISSN 1520-6106.
- 279 16. Amish J. Patel, Patrick Varilly, David Chandler, and Shekhar Garde. Quantifying density fluctuations in volumes of all shapes and sizes using indirect umbrella sampling. *J. Stat. Phys.*, 145(2):265–275, 2011.
- ²⁸¹ 17. Erte Xi, Sean M. Marks, Suruchi Fialoke, and Amish J. Patel. Sparse sampling of water density fluctuations near liquid-vapor coexistence. *Mol. Simul.*, 44(13-14):1124–1135, 2018.
- 18. Fabian Pedregosa, Gaël Varoquaux, Alexandre Gramfort, Vincent Michel, Bertrand Thirion, Olivier Grisel, Mathieu
 Blondel, Peter Prettenhofer, Ron Weiss, Vincent Dubourg, Jake Vanderplas, Alexandre Passos, David Cournapeau,
 Matthieu Brucher, Matthieu Perrot, and Édouard Duchesnay. Scikit-learn: Machine learning in Python. J Mach. Learn
 Res., 12(85):2825–2830, 2011.
- Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming
 Lin, Natalia Gimelshein, Luca Antiga, Alban Desmaison, Andreas Kopf, Edward Yang, Zachary DeVito, Martin Raison,
 Alykhan Tejani, Sasank Chilamkurthy, Benoit Steiner, Lu Fang, Junjie Bai, and Soumith Chintala. Pytorch: An imperative
 style, high-performance deep learning library. In H. Wallach, H. Larochelle, A. Beygelzimer, F. d'Alché Buc, E. Fox, and
 R. Garnett, editors, Advances in Neural Information Processing Systems 32, pages 8024–8035. Curran Associates, Inc.,
 2019.
 - 20. Diederik P. Kingma and Jimmy Ba. Adam: A Method for Stochastic Optimization. arXiv:1412.6980 [cs], January 2017.
- 21. Constantin Steppa and Tim L. Holch. HexagDLy—Processing hexagonally sampled data with CNNs in PyTorch. SoftwareX, 9:193–198, January 2019. ISSN 2352-7110. .
- 22. Marco Tulio Ribeiro, Sameer Singh, and Carlos Guestrin. "Why Should I Trust You?": Explaining the Predictions of
 Any Classifier. In Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data
 Mining, KDD '16, pages 1135–1144, New York, NY, USA, August 2016. Association for Computing Machinery. ISBN 978-1-4503-4232-2. URL https://doi.org/10.1145/2939672.2939778.