



HAL
open science

A new kind of atlas of zeolite building blocks

Benjamin Helfrecht, Rocio Semino, Giovanni Pireddu, Scott Auerbach,
Michele Ceriotti

► **To cite this version:**

Benjamin Helfrecht, Rocio Semino, Giovanni Pireddu, Scott Auerbach, Michele Ceriotti. A new kind of atlas of zeolite building blocks. *Journal of Chemical Physics*, 2019, 151 (15), pp.154112. 10.1063/1.5119751 . hal-02417178

HAL Id: hal-02417178

<https://hal.umontpellier.fr/hal-02417178>

Submitted on 10 Jun 2021

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

A New Kind of Atlas of Zeolite Building Blocks

Benjamin A. Helfrecht,¹ Rocio Semino,^{2,1} Giovanni Pireddu,^{3,1} Scott M. Auerbach,⁴ and Michele Ceriotti¹

¹*Laboratory of Computational Science and Modeling, Institut des Matériaux,
École Polytechnique Fédérale de Lausanne, 1015 Lausanne, Switzerland*

²*Institut Charles Gerhardt Montpellier UMR 5253 CNRS,
Université de Montpellier, Place E. Bataillon, 34095 Montpellier Cedex 05, France*

³*Dipartimento di Chimica e Farmacia, Università degli Studi di Sassari, Via Vienna 2, 01700 Sassari, Italy*

⁴*Department of Chemistry and Department of Chemical Engineering,
University of Massachusetts Amherst, Amherst, MA 01003 USA*

(Dated: July 9, 2019)

We have analysed structural motifs in the Deem database of hypothetical zeolites, to investigate whether the structural diversity found in this database can be well-represented by classical descriptors such as distances, angles, and ring sizes, or whether a more general representation of atomic structure, furnished by the smooth overlap of atomic positions (SOAP) method, is required to capture accurately structure-property relations. We assessed the quality of each descriptor by machine-learning the molar energy and volume for each hypothetical framework in the dataset. We have found that SOAP with a cutoff-length of 6 Å, which goes beyond near-neighbor tetrahedra, best describes the structural diversity in the Deem database by capturing relevant inter-atomic correlations. Kernel principal component analysis shows that SOAP maintains its superior performance even when reducing its dimensionality to those of the classical descriptors, and that the first three kernel principal components capture the main variability in the data set, allowing a 3D point cloud visualization of local environments in the Deem database. This “cloud atlas” of local environments was found to show good correlations with the contribution of a given motif to the density and stability of its parent framework. Local volume and energy maps constructed from the SOAP/machine-learning analyses provide new images of zeolites that reveal smooth variations of local volumes and energies across a given framework, and correlations between local volume and energy in a given framework.

I. INTRODUCTION

Zeolites are nanoporous, crystalline silica-based materials featuring a rich array of structures [1]. They are primarily composed of corner-sharing SiO₄ tetrahedra, and may include heteroatoms such as Al, Ge, and P isomorphically substituted for Si sites. Because of their stable frameworks and versatile structures and compositions, zeolites have enjoyed many applications—especially as catalysts in shape-selective transformations of hydrocarbons, and as molecular sieves for separating mixtures. Databases of real [2] and hypothetical [3–7] zeolites have been screened for new applications such as carbon dioxide storage [8], showing the potential utility of several hypothetical zeolites. Despite this promise, synthesizing new, tailor-made zeolite frameworks has remained a grand challenge in materials science, largely because we lack fundamental understanding of how zeolites crystallize in solution. To shed light on this, recent work has considered assembly [9, 10] of various rings [11] and cages [12] inspired by known or hypothetical zeolites. While such approaches are logical, enumerating zeolite sub-structures through the lens of known rings and cages can leave out many conceivable local silica environments not yet encountered, which may be synthesizable [7] and also important for investigating disordered silica structures leading up to zeolite crystals. In pursuit of a more bias-free understanding of zeolite structures, we implement herein the smooth overlap of atomic positions

(SOAP) [13, 14] descriptor to represent local silica environments in the Deem database of hypothetical zeolite frameworks, and critically compare this descriptor to several well-established descriptors of zeolite structure.

Real and hypothetical zeolite structures are currently understood in terms of framework topologies, which are mathematical connectivities of corner-sharing tetrahedra. These topologies are analyzed, in turn, by identifying rings that form each structure, where an “ n -ring” involves n alternating –Si–O– atomic units. At the time of this writing, there are 234 topologies (of which 223 are fully connected) of fully ordered zeolite materials in the online database of the International Zeolite Association [2], each labeled by a three-letter code. For example, FAU labels a topology with 12-rings exhibited by actual materials such as faujasite, NaX, and NaY, while MFI labels a topology with 10-rings exemplified by zeolites silicalite (all-silica) and ZSM-5. In this database there are roughly 40 topologies that can be synthesized as all-silica zeolites, i.e., polymorphs of α -quartz. Forcefield calculations on many of these all-silica zeolites have found them to be 8.0–20. kJ/mol SiO₂ less stable than α -quartz, with energies that correlate linearly with framework density [15], in good agreement with experimental calorimetry [16].

The Atlas of Zeolite Structures also contains a collection of 60 zeolite building blocks known as composite building units (CBUs) [17], which are motifs such as the double 4-ring (“d4r”) and the double 6-ring (“d6r”). The overwhelming majority of CBUs are cage-like objects centered on empty space. In the present article,

we take a complementary approach to identifying zeolite building blocks by constructing *atom*-centered environments rather than void-centered environments. We pursue such atom-centered environments because we find it simpler logically and mathematically to reconstitute zeolite frameworks and their properties by summing over atoms instead of summing over void spaces. In addition, this construction represents a natural extension of well-established approaches based on fragment decomposition of molecules, and reflects the physical principle that macroscopic properties of materials arise from the combination of local contributions from their atomic or molecular components.

In addition to the database of known all-silica zeolites, there exist databases of hypothetical all-silica zeolites containing millions of candidate framework structures [3–7]. For example, the database constructed by Deem and co-workers includes fully connected, all-silica zeolitic structures with energies within 30 kJ/mol SiO_2 of α -quartz [4], as computed with the shell-model-based Sanders–Leslie–Catlow (SLC) forcefield [18, 19] in the program GULP [20]. While these structures share common local structural features such as Si–O bond lengths near 1.6 Å and O–Si–O angles near 109.5° , they also exhibit substantial diversity. This diversity is often characterized by distributions over descriptors such as Si–O–Si angles [21], Si–Si near-neighbor distances, and ring sizes [3, 4, 22, 23]. In this article, we address the question of whether such descriptors are optimal or even sufficient for describing the structural diversity found in the Deem SLC-PCOD database (which we refer to as simply the “Deem database”), by comparing these conventional zeolite descriptors with a more general geometric representation of network structure focused on atom-centered environments.

To perform this comparative assessment of structural descriptors, we require the following two components: (1) an alternative mathematical description of zeolite structure that is concise yet complete, which incorporates all of the trivial physical symmetries (e.g., translational and rotational symmetries and invariance upon permuting atom numbers) and that is sufficiently general to be applicable to all known structures, as-yet-undiscovered structures, and even disordered structures that may lead to zeolites; (2) an objective way of assessing the performance of a given structural representation in mapping the relevant structural diversity [24].

The recent surge of applications of machine learning (ML) to chemistry and materials science, that includes accurate predictions of molecular properties [14, 25–37], the construction of data-driven interatomic potentials [13, 38, 39], as well as fundamental studies of the mathematical representation of atomic structures, [40–42] provides several suitable mathematical descriptions satisfying criterion (1) above. Previous applications of ML in zeolite science, which include identifying zeolite framework topologies [43] and predicting mechanical properties [44], have relied on zeolite-specific struc-

tural descriptors. Our goal is to establish a more general structural representation, which could be useful for describing both crystal structures and disordered materials leading up to crystals. In particular, we choose the smooth overlap of atomic positions (SOAP) representation [13], which has been used successfully to construct interatomic potentials [45–51], to predict the properties of atomic-scale structures or identify motifs in complex materials such as molecular crystals [52] and polypeptides [53]. SOAP describes local atomic environments in terms of spherically invariant measures of the neighbor distribution around each atom, as described in more detail in the next section, thus providing the atom-centered approach discussed above. In the present article, we investigate the applicability of SOAP for elucidating the structural diversity in the Deem SLC-PCOD database of hypothetical zeolites. This choice is motivated by the very large number of hypothetical, all-silica frameworks that are present in this database (as compared to the 40 frameworks in the IZA database), which allows for ample statistics and thus a proper machine learning study.

To objectively assess the performance of a given structural descriptor, we require that it leads to an effective machine learning model to predict key materials properties [24]. Below we show that the SOAP-based representation includes more information than standard descriptors such as Si–O–Si angles, Si–Si near-neighbor distances, and ring sizes as evidenced by more accurate machine learning predictions of zeolite molar volumes and lattice energies, which are fundamental properties of zeolite frameworks [15]. By comparing the predictions of zeolite volume and energy based on local atomic environments with different nominal sizes, we also discover the degree of locality of each property, i.e., the correlation lengths that are required to determine the overall behavior of that particular property. We also show below new ways to visualize zeolite frameworks in terms of local contributions to energy and volume.

The remainder of this article is organized as follows: in Section II we provide a concise introduction to the SOAP framework, and to the supervised and unsupervised ML algorithms we adopt; in Section III we describe the dataset we used to demonstrate our construction; in Section IV we present and discuss our results, and in Section V we draw our conclusions.

II. METHODS

We start by presenting an overview of the methods we use to represent structural environments in silica-based frameworks—which we define in terms of overlapping motifs centered on each Si atom in the framework—and to investigate quantitatively the relation between such environments and some macroscopic properties of each framework, namely the molar volume and energy.

A. Classical Descriptors

In the context of zeolitic structure analysis, the use of geometric descriptors for classifying and rationalizing structure-property relationships is already well established [1, 2, 54, 55]. The choice of the representation for a given zeolitic structure is often motivated by physico-chemical understanding and intuition regarding which structural features are relevant to the study of certain properties. For example, when investigating catalytic properties of zeolites, one may consider correlating acid-site strengths with Si–O–Al angular distributions. On the other hand, when considering diffusion of guest molecules through zeolite frameworks, one typically uses zeolite ring distributions to rationalize transport properties. We refer to such representations as *classical descriptors* to emphasize the difference between zeolite-specific descriptors and general descriptors such as SOAP.

We focus in particular on three widely used classical descriptors to serve as references for the chemical environment representations: Si–O–Si angles, Si–Si distances, and connected rings. The distance- and angle-based descriptions are local, atom-centered features, and are appropriate for representing properties that could be safely decomposed into additive, local contributions. For instance, bonding interactions in interatomic potentials are often expressed in terms of bond angles and distances [56, 57]. The distance- and angle-based representations characterize the local environment of each Si as a vector of the Si–Si distances or Si–O–Si angles between the central reference Si atom and the four nearest Si atoms, as these structures are mainly composed of Si tetrahedra. In order to make these representations independent of permutations of the atom indices, the vector elements are arranged in descending order.

Ring-based descriptors involve correlations on longer length scales, and consequently may be more suitable for characterizing the topology of a given framework than distance- or angle-based structural representations. Zeolitic frameworks can be described in terms of rings according to various definitions; the ring classifications included in the IZA database are those required to build up a given structure, plus those required to define pores and channels accessible to guest molecules. Such ring systems are readily confirmed by visual inspection of the crystal structure. To apply an automated analysis to a large database of structures, we use two mathematically rigorous definitions in our work, namely, King’s criterion [55, 58] and the shortest path criterion [55, 59–61], as implemented in the R.I.N.G.S. code [55, 62]. In both cases, we translate the list of detected rings in a given zeolite framework into vectors of features x_s associated with local Si environments by counting the number of times the tagged Si atom appears in s -sized rings (since O atoms are also present in the rings, the real ring size is $2s$, here we adopt the frequently used convention of naming a ring counting only Si atoms). For example, the ring vector for a Si atom in silica-sodalite (which are all equivalent by

symmetry) is $[0, 0, 0, 2, 0, 4, 0, \dots, 0]$ indicating that each Si atom is part of two 4-rings and four 6-rings. In general, all ring vectors in the present work are 10-dimensional vectors, ranging from 3-rings to 12-rings. In the present work we focus on the shortest-path definition, which gave marginally better performance than King’s when used as the basis for predicting structure–property relations.

B. The SOAP representation

The SOAP representation can be seen in relation to the symmetrised three-body atom correlation function, defined for an environment \mathcal{X}_j centered around the j -th atom in a structure [41]. One can start by building a smooth atom density centered on atom j , separated in different channels depending on the chemical species α (Si, O in this case) as follows:

$$\psi_{\mathcal{X}_j}^{\alpha}(\mathbf{r}) = \sum_{i \in \mathcal{X}_j^{\alpha}} g(\mathbf{r} - \mathbf{r}_{ij}) f_c(r_{ij}), \quad (1)$$

where Eq. (1) involves a sum of Gaussian functions centered on the interatomic distances \mathbf{r}_{ij} , which are restricted to atoms of type α in the vicinity of atom j using a smooth cutoff function, $f_c(r_{ij})$. For the cutoff function, we chose a section of a cosine function switching from 1 to exactly zero within 0.5\AA of the cutoff distance. The (Gaussian-smoothed) three-body correlation function can be defined as an average over all rotations operators \hat{R} of the density evaluated at two different points

$$\begin{aligned} \langle \alpha r \alpha' r' \omega | \mathcal{X}_j \rangle &= \int d\hat{R} r r' \psi_{\mathcal{X}_j}^{\alpha}(r \hat{R} \hat{\mathbf{e}}_z) \times \\ &\times \psi_{\mathcal{X}_j}^{\alpha'}(r' \hat{R}(\omega \hat{\mathbf{e}}_z + \sqrt{1 - \omega^2} \hat{\mathbf{e}}_x)), \end{aligned} \quad (2)$$

which corresponds to the probability of finding two atoms of type α and α' at distances r and r' from a tagged central atom, in directions forming an angle $\arccos \omega$. In this expression, the density is evaluated at two points, defined in terms of the unit vectors along the z and x axes of a Cartesian reference system. This construction is schematically depicted in Fig. 1, which shows the Gaussian atom density built on the neighbors of a central atom, and the two points at which the density is evaluated. For each value of (r, r', ω) the resulting density correlation is averaged over all the possible orientations of the points in space.

This kind of correlation function is often used in the statistical mechanical treatment of liquids, and is typically computed as an average over an ensemble of configurations. In the case of SOAP, however, each individual atomic environment \mathcal{X}_j is associated with a feature vector $\langle \alpha r \alpha' r' \omega | \mathcal{X}_j \rangle$ that provides a unique [13] description of a given arrangement of atoms. For example, Fig. 2 shows a plot of $\langle \text{Si}r\text{Si}r'\omega | \mathcal{X}_j \rangle$ and $\langle \text{O}r\text{O}r'\omega | \mathcal{X}_j \rangle$ for a Si atom in the framework of sodalite, indicating the correspondence between peaks in the correlation function and Si-Si-Si and O-Si-O atomic motifs.




FIG. 1. Schematic of the SOAP representation applied to a local Si environment in the zeolite CDO. A smooth atom density (here represented only for Si atoms) is built as a superposition of atom-centered Gaussians, and a smooth three-body correlation function (as a function of two distances r and r' , and of the angle between the two directions $\arccos\omega$) is computed by evaluating this density in two points, and averaging the result over all orientations of the (r, r', ω) stencil.

The representation $\langle \alpha r \alpha' r' \omega | \mathcal{X}_j \rangle$ can be expanded on an orthogonal basis built out of radial functions $R_n(r)$ and Legendre polynomials $P_l(\omega)$ [41], leading to the usual form of SOAP [13] as a power spectrum $\langle \alpha n \alpha' n' l | \mathcal{X}_j \rangle$, which is the projection of the three-body density onto radial basis functions n and n' and Legendre polynomial l . The numbers of radial and angular functions are convergence parameters, and the two parameters that are more significant in determining the behavior of the SOAP representation are the cutoff distance of the region included in the environment and the width σ of the Gaussians, which controls how a difference in atomic coordinates between two structures translates in a difference between their respective $\langle \alpha r \alpha' r' \omega | \mathcal{X}_j \rangle$ vectors. Below we consider cutoff distances of 3.5 Å and 6.0 Å, leading to SOAP vectors with approximately 3,000 elements. In the Supporting Information, we list the remaining parameters needed to specify our application of SOAP.

C. Machine learning model and environment-centered properties

To assess the performance of the different representations in quantifying structure–property relations, we built a machine learning model to predict the molar volume and the cohesive energy of each structure based on

the sum of environment feature vectors defined by each representation. The performance of the model, and the rate at which the error on a test dataset decreases as the model is trained on increasing amounts of data (so-called learning curves), offer insights into the completeness and directness of a representation for describing a particular property or set of properties [29, 34, 63]. For the property prediction we use a statistical learning scheme based on Kernel Ridge Regression (KRR), a relatively simple approach that makes it possible to introduce some non-linearity in the regression while having a deterministic learning procedure. Kernel Ridge Regression has been used successfully to build models for several classes of atomic-scale properties [27–34], and a formally equivalent method, Gaussian Process Regression, has been used to create interatomic potentials [45–51].

For each descriptor, we built machine learning models for the molar volume and energy, using a reference data set that we introduce in Section III. To obtain a model that can be evaluated at fixed cost and thus independent on the training set size, we used a sparse KRR model [64–66] employing a Gaussian kernel and based on a set of 2,000 representative environments $M = \{\mathcal{X}_j\}$ selected with Farthest Point Sampling (FPS) [67–69]. [70] For each regression task we optimized the parameters of the kernel and the regression model through a grid search and five-fold cross validation with the aim of minimizing the mean absolute error (MAE) of the regression. Additional details of our approach are discussed together with each example in the Supporting Information.

The sparse KRR model we use approximates the target property y for a zeolite framework \mathcal{A} as a sum over contributions from each environment \mathcal{X} in that framework, i.e.,

$$y(\mathcal{A}) = \sum_{\mathcal{X} \in \mathcal{A}} y(\mathcal{X}) = \sum_{\mathcal{X} \in \mathcal{A}} \sum_{\mathcal{X}_j \in M} x_j k(\mathcal{X}, \mathcal{X}_j), \quad (3)$$

where x_j is a weighting factor and $k(\mathcal{X}, \mathcal{X}_j)$ is the kernel between the local environments \mathcal{X} and \mathcal{X}_j . As a consequence, the regression exercise also permits decomposition of structure-wide properties (like volume or energy) into contributions from the individual atomic environments $y(\mathcal{X})$, which can be used to investigate and visualize the impact of local structural motifs on macroscopic properties of the framework.

D. Dimensionality reduction

While machine learning models based on kernel methods can be quite powerful, they can also be computationally expensive, especially for very large datasets. When working with large amounts of data, as we do here, it is often necessary to sparsify the kernel used in the machine learning model. To this end we use a low rank approximation to the true kernel matrix built by considering the kernel between each environment and a set of




FIG. 2. Contour plots corresponding to the 3-body correlation function for Si–Si (left) and O–O (right) atoms around a central silicon atom. White, blue and orange contours (drawn for $r < r'$) correspond to three isocontours of the density averaged over 1,000 structures from the database. Red isocontours (drawn for $r > r'$) correspond to the density computed for silica sodalite (SOD). The correspondence between some of the peaks and three-atom correlations around a framework node is also indicated. Note that the correlations included in the SOAP features also incorporate some “self-correlation” for the density centered on each atom, which have been eliminated here for clarity of visualization.

representative environments. We selected the representative environments with an iterative strategy based on FPS [67–69], as noted in Section II C, to ensure that the representative environments are structurally diverse. We used approximate kernel matrices constructed in this way to perform KRR and kernel principal component analysis (KPCA) [71]. We applied KPCA to the SOAP-based descriptors to reduce the dimensionality of the descriptor and to assess the information content of the descriptor in the context of predicting the energy and volume of hypothetical zeolite structures. Additional details regarding the kernel approximations for KRR and KPCA are given in the Supporting Information.

III. DATA SELECTION

Our analysis of zeolite structures is based on the DEEM SLC-PCOD database [4], which contains hypothetical zeolite structures that are no more than 30 kJ/mol Si higher in energy than α -quartz. Given that the database contains a few hundred thousand hypothetical zeolites, each zeolite contains several Si-centered environments, and that a single SOAP feature vector can contain thousands of components, computing and analyzing the full SOAP vectors of every structure in the database is computationally intractable at this time. Hence, we reduced the dimensionality of the input space by considering a subset of the database and by selecting only the most diverse SOAP components to use in our ML model, as explained in Section II D. From the approximately 330,000 structures in the database, a subset of 10,000 structures was selected at a fixed stride following the ID number, and the 500 most diverse SOAP vector components were selected via FPS on a random selection of 2,000 structures from the 10,000-structure subset. The Euclidean distance between SOAP vectors was used

as the distance metric for the FPS procedure [69]. The SOAP vectors were then computed for all 10,000 structures in the subset, but only the FPS components were retained. Through this scheme, we considered two SOAP descriptors: one with a radial cutoff of 3.5 Å, and another with a 6.0 Å cutoff, so that the representation incorporates information on the size scale of nearest neighbor Si atoms (3.5 Å) or next-nearest neighbors, possibly including a modestly sized ring (6.0 Å). Each atomic environment comprised a central Si atom and all of the surrounding Si and O atoms within the cutoff radius. Oxygen atoms were not considered as environment centers. Additional details regarding the construction of the SOAP vectors are given in the Supporting Information. The distance-, angle-, and ring-based descriptors were also calculated for the 10,000-structure subset for comparison against SOAP in our machine learning models. Comparisons between the classical and SOAP descriptors are examined in the following section.

In addition, we considered a subset of 1,000 stride-selected structures from the DEEM SLC-PCOD database for our analysis to test whether our results are influenced by the size of our selected subset of structures. We found that the results for the 1,000-structure sample yield similar conclusions to those that can be drawn for the 10,000-structure sample, both in terms of performance when being used as the basis of a machine learning model to predict different geometric features, and in terms of the representation of the data set using unsupervised learning. This robustness with respect to dataset size indicates that our results have likely converged with respect to the structural diversity in the Deem database, and thus, that our results apply to the full Deem database. In Section IV we report our findings based on the 10,000-structure subset. The results for the 1,000-structure subset can be found in the Supporting Information.

IV. RESULTS

A. Comparison of Structural Representations

Many machine learning studies present parity plots to show the quality of learning for a given property for the best model. We prefer to quantify learning by calculating mean average errors (MAEs) of a property estimation, and presenting MAEs as learning curves, i.e., as functions of the number of training points. Building learning curves for the prediction of molar volume and energy enables head-to-head comparisons of the information content in the various structural descriptors. In particular, the value of the MAE for small training set sizes indicates whether the most prominent components of a representation correlate strongly with a given property. The asymptotic behavior in the large training set sizes indicates how complete a representation is: saturation of a learning curve, i.e., a learning curve that flattens to nearly zero slope, indicates that there is not sufficient information to improve the model, even when exposing it to new data; while a substantial asymptotic slope indicates that the model still has sufficient information to improve its learning.

Figures 3(a) and 3(b) show the learning curves for predictions of the molar volume and molar energy, respectively, using the classical and SOAP descriptors. We see in Figures 3(a) and 3(b) that the Si–Si distance (green), Si–O–Si angle (blue), and shortest-path rings (violet) perform poorly at predicting both volume and energy, saturating at MAEs above $1 \text{ \AA}^3/\text{Si atom}$ and 2 kJ/mol Si , respectively. The ring-based descriptor is marginally better at predicting volumes compared to the distance and angle descriptors [see 3(a)], but it is the worst classical descriptor at predicting energies [see 3(b)]. This is likely because lattice energy is sensitive to ring geometries and distortions, thus requiring more information than simply the numbers of atoms in rings.

Using the SOAP descriptor with a 6.0 \AA cutoff (red) results in the best volume predictions for all training set sizes, with MAEs as low as $0.4 \text{ \AA}^3/\text{Si atom}$. This is probably due to the fact that accurate predictions of overall framework densities require information on larger spatial scales. The SOAP descriptor with a 3.5 \AA cutoff (black) performs only slightly better than the classical descriptors in predicting volumes, but yields the best energy predictions off all the descriptors for smaller training set sizes, indicating that relatively local correlations are sufficient for making estimates of lattice energy within 1 kJ/mol Si . This is perhaps not surprising since a substantial contribution to the zeolite lattice energy can be accounted for through nearest-neighbor bond and angle interactions, and thus the 3.5 \AA based model, where only the first neighboring tetrahedral information is included, does a good job at estimating the energy even when the training set has a modest size. However, the learning potential of the 3.5 \AA SOAP descriptor plateaus as the training set size approaches 1,000 structures, while the




FIG. 3. Learning curves of the classical and SOAP descriptors for predictions of (a) volume per Si atom and (b) energy per mol Si.

6.0 \AA SOAP representation continues to improve at larger training set sizes, yielding energy predictions with MAEs as low as 0.4 kJ/mol Si . Overall, Figures 3(a) and 3(b) show that the SOAP descriptor with a 6.0 \AA cutoff does an excellent job of accurately capturing zeolite energy and molar volume, while the classical descriptors do not.

To investigate whether the improved predictive performance and learning ability of the 6.0 \AA SOAP descriptor is a result of the intrinsic quality of the descriptor itself, or merely from the fact that the 6.0 \AA SOAP representation incorporates more information through higher-dimensional vectors, we compare the learning curves of the classical descriptors with those of the SOAP descriptors whose dimensionality has been reduced through kernel principal component analysis (KPCA). Because the distance- and angle-based descriptors are represented as 4-element vectors (because of the tetrahedral coordination around Si) and the ring vectors contain ten elements (because ring sizes in our data set range between 3-rings and 12-rings), we find it instructive to compare the classical descriptors with the first four and ten principal components of the SOAP representations, shown in Figs. 4(a) and 4(b). For the volume prediction (see 4(a)), the SOAP-KPCA 6.0 \AA including ten dimensions (red, dashed) is the best model, followed by the SOAP-KPCA 6.0 \AA including four dimensions (red, solid), which has comparable performance to the ten-dimensional representation for training sets containing fewer than 50 points (structures). In the case of the energy prediction depicted in Figure 4(b), the best model is that based on the SOAP-KPCA 3.5 \AA descriptor including ten dimen-

sions (black, dashed), followed by the four-dimensional descriptor (black, solid). Overall, we note that SOAP-KPCA models equal or surpass the performance of the classical descriptors at predicting energy and volume, even after dimensionality reduction. This suggests that SOAP inherently contains more information about the local structure of a zeolite framework than do the classical descriptors, and this is not a mere result of the flexibility afforded by the higher dimensionality of the feature representation.

For predictions of the molar volume, the SOAP-KPCA descriptors perform better than the classical descriptors regardless of the dimensionality of the representation, though the performance gain of the 3.5 Å SOAP-KPCA descriptor over the classical descriptors is rather small. For predictions of lattice energy per Si atom, the 3.5 Å SOAP-KPCA descriptor performs as well as or better than the classical descriptors at larger training set sizes. The 6.0 Å SOAP-KPCA descriptor, on the other hand, performs worse than the distance- and angle-based descriptors at comparable dimensionality, but better than the rings-based descriptor. The behavior of 6.0 Å SOAP-KPCA can be attributed to the fact that 6.0 Å SOAP considers a larger local environment, and therefore, a much larger amount of information about the surroundings of a single Si atom. This additional information simply becomes noise within the regression model unless more information—in the form of additional principal components—is included. Indeed, in predictions of the energy per Si atom, we find that the 3.5 Å SOAP representation outperforms the 6.0 Å SOAP representation for a given number of principal components unless upwards of 100 components are included. (Learning curves for the SOAP-KPCA representations with a larger range of principal components are provided in the Supporting Information.)

B. Low-Dimensional Maps of the Deem Dataset

The results in Fig. 4 show that the SOAP method provides sufficient information, even when truncated in dimensionality, to outperform classical descriptors in *quantitatively* accounting for the molar volume and energy of a database of zeolites. This finding raises the following question: how many dimensions of SOAP are required to *qualitatively* account for the structural diversity found in the Deem database? To answer this question, we show in Fig. 5(a) the relative variances (σ_{KPCA}^2) of the kernel principal components (KPCs) and their Pearson correlation coefficients with energy ($\rho_{KPCA,E}$) and volume ($\rho_{KPCA,V}$). The KPCs are sorted, by construction, in a way that reflects the intrinsic variability of the dataset, and the first three components in Fig. 5(a) (open circles) capture a substantial portion of the structural diversity as evidenced by the relative variances all exceeding 0.6, while all other relative variances are less than 0.3 and most are much less than 0.1. The Pearson correlation




FIG. 4. Learning curves of classical and SOAP-KPCA descriptors with similar dimensionality for predictions of (a) volume and (b) energy.

coefficients in Fig. 5(a) reveal that KPCs 1 and 2 correlate strongly with volume while KPCs 2 and 3 correlate strongly with energy. These correlations can also be visualized in the contour plots shown in Fig. 5(b)–(g), where the environment energies and volumes are plotted against values of the first three KPCs along with a linear least squares fit to the data. The elongated shapes of the contours indicate the correlations between the environment properties and the SOAP vectors in KPC space, and these correlations are quantified by the Pearson correlation coefficients. Overall, Figs. 5(a)–(g) suggest that a three-dimensional picture of structural diversity using the first three KPCs can reveal the essential features of the Deem zeolite data set.

Consequently, we can use the SOAP-KPCA construction to build a new 3D “atlas” of zeolite building blocks, analogous to the list of 60 composite building units (CBUs) in the IZA database [2, 17], and to the list of “packing units” given by Blatov *et al.* [12]. As emphasized above, virtually all CBUs and packing units are centered on void spaces; in contrast, our zeolite building blocks are atom-centered, allowing such environments to be summed to yield macroscopic properties of overall frameworks such as molar energy and volume. Because of the very large number of local environments in the Deem database, our new atlas has the feel of a point cloud in 3D-KPCA space.

Our SOAP-based “cloud atlas” is presented in Fig. 6, where the SOAP environments are plotted in the space

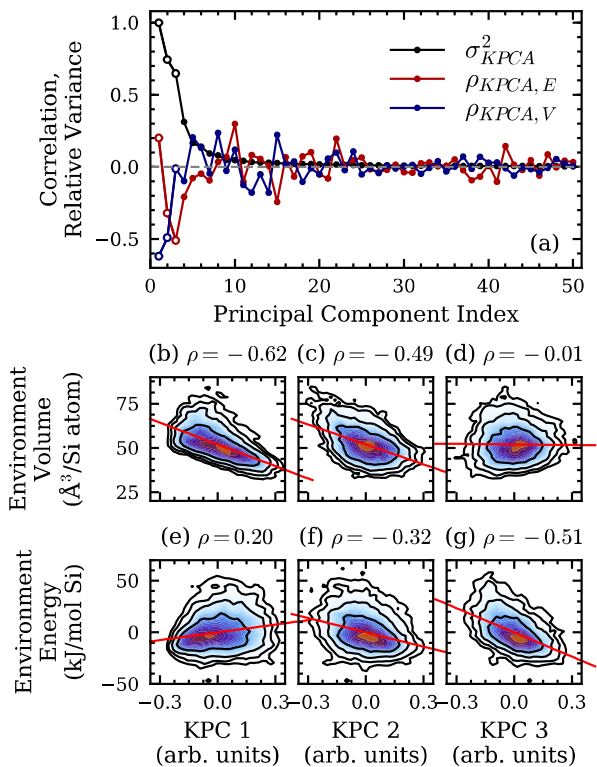


FIG. 5. (a) Pearson correlation coefficients between the first 50 KPCs of the 6.0 Å SOAP representation and the decomposed environment volumes and energies in the 10,000-structure sample. The relative variance in the KPCs at each of the first 50 components is also plotted. The correlation coefficients and relative variance of the first three components are highlighted with open symbols, as these are the KPCs that correlate most strongly with the decomposed environment energies and volumes. (b)–(g) Kernel density estimation of all environments in KPC–property space with a linear least squares fit to the data to show the correlations in more detail (red line). The value of the Pearson correlation coefficient is given above each contour set.

defined by the first three KPCs of the dataset. Each point (environment) is colored according to its energy contribution (low = white, medium = blue/purple, high = orange) and is sized according to the volume contributions to its corresponding zeolite framework. Several environments are highlighted to provide examples of the “building blocks” present in this new cloud atlas. In particular, we show in Fig. 6 the lowest-, median-, and highest-energy SOAP environments, as well as the lowest-, median-, and highest-volume SOAP environments. The majority of environment volumes in Fig. 6 range from 30–80 Å³/Si, corresponding to a range of framework densities of 12.5–33.3 Si atoms per 1000 Å³. (Actual zeolites exhibit framework densities in the range of 12.5–21.0 Si atoms per 1000 Å³ [2], although the Deem database extends to larger framework densities than those of actual zeolites.) Most environment energies fall within the range

of –20 kJ/mol Si to +30 kJ/mol Si. We note that the notions of environment energy and volume relate herein to the machine learning process and not necessarily to physico-chemical notions of energy or volume. Nevertheless, as we will show later, this data-driven decomposition ends up being consistent with physical considerations, as was already observed in similar ML models of materials and molecular properties [25, 72]. In each highlighted environment, the atoms included within 6.0 Å SOAP are shown in yellow (Si) and red (O), while the rest of the zeolite framework is depicted as corner-sharing tetrahedra. Contour plots showing the distribution of environments are projected onto the *xy*-, *yz*-, and *xz*-planes in Fig. 6. These contour plots reveal that the statistical distribution of zeolite environments is unimodal and rather broadly peaked, indicating a remarkably uniform distribution of structural motifs. As such, we find no special region in 3D KPC-space that is particularly well stocked with building blocks for making hypothetical zeolites. The broad distribution of environments shown in Fig. 6 suggests that the algorithm for producing these hypothetical zeolite frameworks [3, 4] has left no substantial gap in environment space.

A complementary way to visualize information contained in the Deem dataset is by constructing energy–volume contour plots for both frameworks and their constituent environments; this is shown in the middle panel of Fig. 7, with environment contours in solid black and framework contours in dotted red. Each line in these contour plots indicates a region of constant probability for finding a framework, or an environment, with a certain combination of molar energy and volume. These contour plots show an excellent overlap between framework and environment spaces for energies between –20 kJ/mol Si to +10 kJ/mol Si. Indeed, Fig. 7 shows that most local volumes obtained by machine learning map very well onto the actual space of framework volumes in the Deem dataset. Furthermore, the lower edges of both contour plots (highlighted by the thick red line) follow the energy–density correlation known for actual zeolites [15]. This finding has already been noted by Deem and coworkers [4] for the hypothetical frameworks; we find it remarkable that the SOAP-based environments also exhibit the signature of this energy–volume correlation.

The contour plots in Fig. 7 also show a significant space of environments with local energies greater than 10 kJ/mol Si, i.e., greater than that of any framework in our sample of the Deem database, which was constructed to cover an energy range no more than 30 kJ/mol Si above the molar energy of α -quartz [4]. This extension to higher local energies raises the question of how high-energy local environments manifest in overall framework structures. Such a question on the connection between local environments and overall frameworks can also be raised for high-volume environments, which presumably frame large pores. The top and bottom panels in Fig. 7 answer these questions for two hypothetical frameworks labeled A (top) and B (bottom). Each panel has two

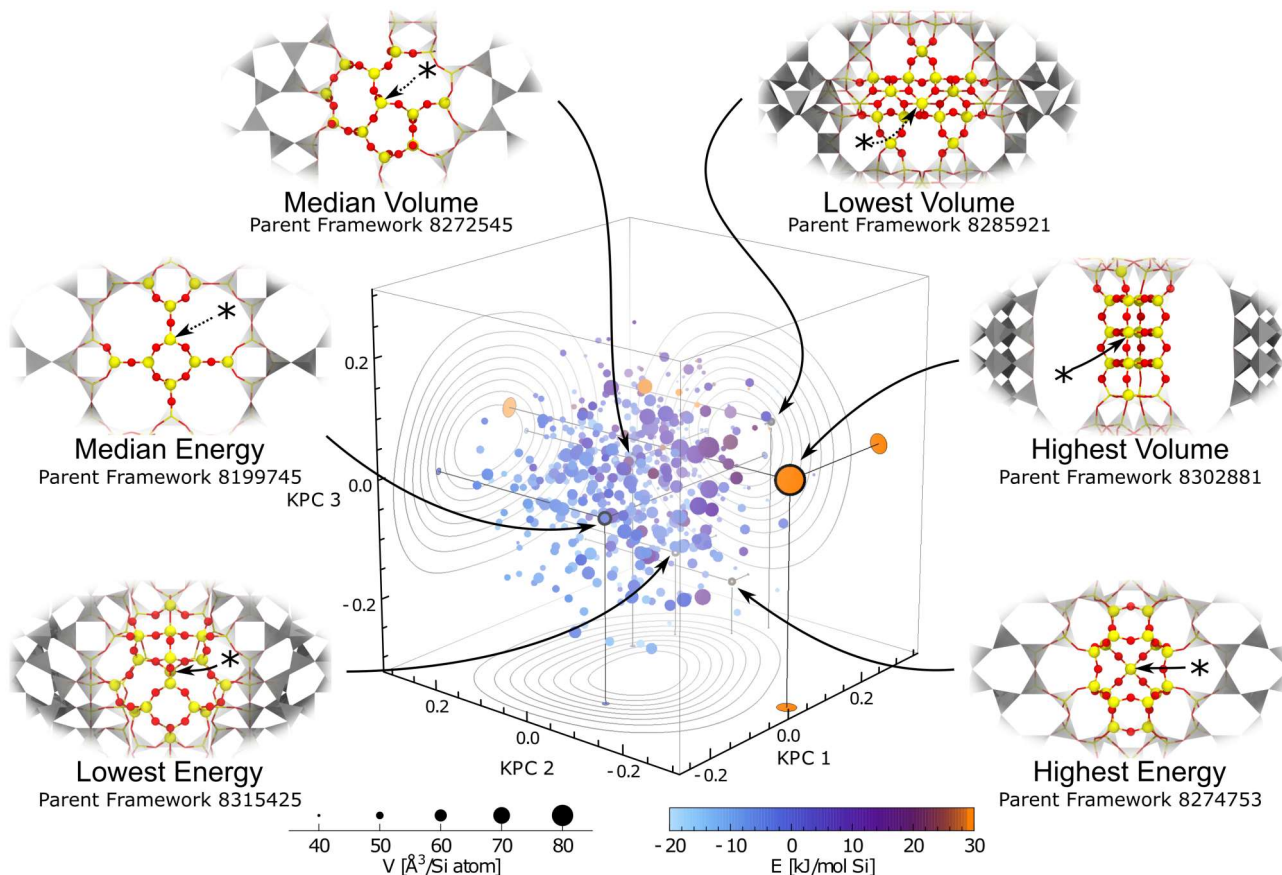


FIG. 6. The new atlas of zeolite building blocks, where every 2,000-th environment of the 10,000-framework subset is plotted as a point in the three-dimensional space formed by the first three kernel principal components of the SOAP construction using a 6.0 Å cutoff. The points are colored and sized according to the energy and volume contribution of the corresponding environment. The environments with the highest and lowest energies are highlighted along with environments contributing energies and volumes close to the median of the dataset. Note that there exist some (extreme) outliers: the highest-energy environment contributes more than 380 kJ/mol Si, and the lowest below -30 kJ/mol Si. The highest-volume environment contributes more than $90 \text{ \AA}^3/\text{Si atom}$, and the lowest less than $30 \text{ \AA}^3/\text{Si atom}$. Energies falling outside the range of the scalebar are assigned to the color at the nearest extreme of the colorscale. Environment centers are indicated by the asterisks and their associated arrows; a dotted arrow signifies that the central atom is hidden behind the foremost atom visible in the atomic snapshot. In each snapshot, the atomic environment is represented as a ball-and-stick model; the surrounding zeolite structure is represented as SiO_2 tetrahedra. Overall, we see a remarkably uniform distribution of environments.

copies of each framework showing local volumes (left) and local energies (right), both color-coded according to low = white, medium = blue, and high = orange. Framework A (top) is shown because it contains the median-energy environment in the Deem dataset, corresponding to the solid blue dot in the contour plot, while its parent framework (A) corresponds to the open blue dot. Framework B (bottom) was chosen because it contains the highest-volume environment in the Deem dataset, corresponding to the solid green dot in the contour plot, while its parent framework (B) corresponds to the open green dot.

These energy-volume color-coded images provide a qualitatively new way to visualize zeolite frameworks, yielding several interesting insights. First, by compar-

ing the color distributions in frameworks A and B, it is readily apparent that A (a medium-pore framework) is much more homogeneous in both local energy and volume than is B (an ultra-large-pore framework). Second, we see that the indicated environment in B (centered on the orange-colored Si atom) does indeed frame the ultra-large pore. Surprisingly, though, the other Si atoms that frame this same ultra-large pore exhibit lower local volumes, presumably because these other Si atoms anchor environments closer to denser regions of the framework and its pore walls. Third, a visual comparison of the volume (left) and energy (right) representations of framework B shows clear correlations between local volumes and energies, reminiscent of the correlation for actual ze-

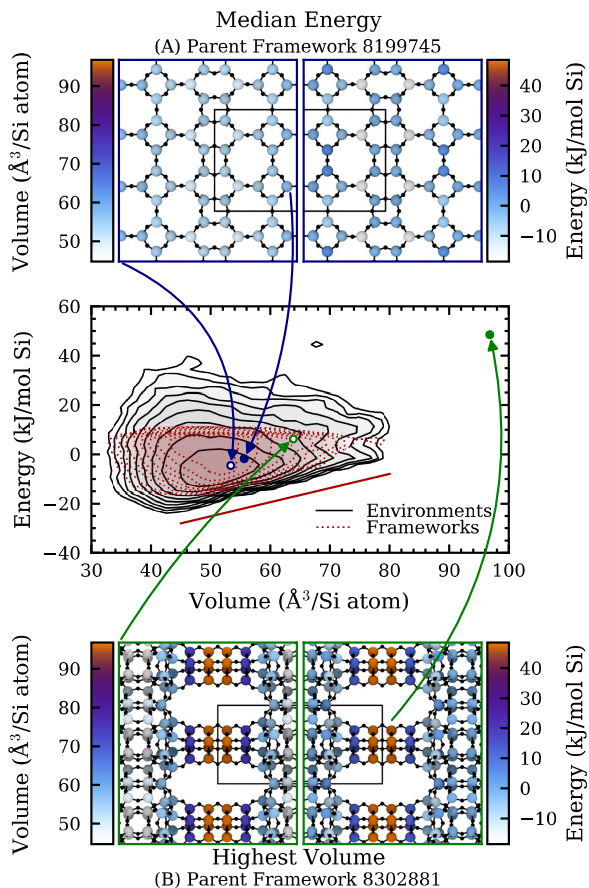


FIG. 7. Middle panel: contour plots showing probabilities of finding frameworks (red dotted lines) and environments (solid black lines) with certain values of molar energy and volume in the Deem dataset, using the SOAP 6.0 \AA representation on the 10,000-structure sample for the environment analysis. Comparing framework and environment contours shows similar structure below 10 kJ/mol Si, with the same energy-volume correlation highlighted by a thick red line. Top and bottom panels: hypothetical frameworks labeled A and B, with median-energy (in framework A) and highest-volume (in framework B) environments, corresponding to solid blue and green dots, respectively, in the contour plot. Parent frameworks A and B are labeled by open blue and green dots, respectively, in the contour plot. Left and right panels show local volume and energy distributions, respectively, for each framework. Visual inspection shows that framework B exhibits much more heterogeneity in both local energy and volume, with strong local energy-volume correlations.

olites reported by Henson and coworkers [15]. Such a correlation is harder to discern for framework A because of its homogeneous character. Overall, these color-coded images, especially those for framework B, show that local energies and volumes vary smoothly across a given framework, possibly allowing the visualization of framework strain and other materials and chemical properties, to be studied in a forthcoming publication.

V. CONCLUSIONS

We have presented an analysis of the structural motifs found within a database of hypothetical zeolites, the Deem SLC-PCOD database [4]. We have investigated whether the structural diversity found in the Deem database can be well-represented by classical descriptors such as Si-Si distances, Si-O-Si angles, and ring sizes, or whether a more general-purpose geometric representation of atomic structure, herein afforded by the SOAP method, is required. We assessed the quality of each descriptor by determining its ability to inform the machine-learning of molar energy and volume for each hypothetical framework in the dataset. We have found that a SOAP representation with a cutoff-length of 6 \AA , which captures local environments beyond near-neighbor tetrahedra, best describes the structural diversity in the Deem database as measured by the accurate predictions of energy and volume afforded by this approach. The performance gain of SOAP can be traced mainly to its description of all inter-atomic correlations within an atom-centered environment. The SOAP approach was found to maintain its lead even when reducing the dimensionality of SOAP feature vectors, through kernel principal component analysis, to be comparable with those of Si-Si distances, Si-O-Si angles, and ring sizes. We thus find that SOAP provides a powerful tool for the statistical analysis of local zeolite environments, i.e., the building blocks of zeolites.

SOAP-KPCA also showed that the first three kernel principal components capture the main variability in the data set, allowing a 3D point cloud visualization of local environments in the Deem database. We refer to this visualization as a “cloud atlas” of zeolite building blocks. This cloud atlas was found to show good correlations with the contribution of a given motif to the density and stability of its parent framework. Overall, the cloud atlas of the Deem database shows a uniform distribution of structural motifs, suggesting that the algorithm for producing the Deem database covered structure space in a thorough and egalitarian fashion. Local volume and energy maps constructed from the SOAP/machine-learning analyses provide new ways of visualizing zeolites that reveal structural heterogeneity, smooth variations of local volumes and energies across a given framework, and correlations between volume and energy of local environments within a given framework.

The present work raises many questions in both data science and zeolite science. From a data science perspective, we wonder whether implementing SOAP with multiple length scales [73] or with radial scaling [74] could optimize the packaging of zeolite structural information, perhaps at the expense of introducing further optimization parameters and obscuring some of the insights discovered above. From a zeolite science perspective, several questions arise, including whether the framework density/stability maps introduced above can predict important zeolite properties like acid-site strength across a

given framework. It is also interesting to consider applying the SOAP/machine-learning methods to actual zeolites in the IZA database to determine similarities and differences between actual- and hypothetical-zeolite data sets. Finally, it is interesting to consider applying the machinery of SOAP to simulations of zeolite formation [75, 76] to determine the local environments that eventually lead to zeolite crystallization. These topics will be addressed in forthcoming contributions.

ACKNOWLEDGMENTS

BAH and MC were supported by the European Research Council under the European Union’s Horizon 2020 research and innovation programme (Grant Agreement No. 677013-HBMAP). RS acknowledges support by the SCCER Efficiency of Industrial Processes. G.P. is thankful to Regione Autonoma della Sardegna for the Ph.D. fellowship provided under the POR-F.S.E. 2014-2020 program. All the Authors would like to acknowledge The National Centre of Competence in Research (NCCR) “Materials’ Revolution: Computational Design and Discovery of Novel Materials (MARVEL)” of the Swiss National Science Foundation (SNSF), which supported a visit of SMA to EPFL, during which this collaboration was initiated.

-
- [1] S. M. Auerbach, K. A. Carrado, and P. K. Dutta, *Handbook of Zeolite Science and Technology* (CRC Press, 2003).
- [2] C. Baerlocher, L. McCusker, and D. Olson, *Atlas of Zeolite Framework Types* (Elsevier, 2007).
- [3] M. W. Deem, R. Pophale, P. A. Cheeseman, and D. J. Earl, *Journal of Physical Chemistry C* **113**, 21353 (2009).
- [4] R. Pophale, P. A. Cheeseman, and M. W. Deem, *Physical Chemistry Chemical Physics* **13**, 12407 (2011), the SLC-PCOD database can be accessed at: http://www.hypotheticalzeolites.net/DATABASE/DEEM/DEEM_PCOD/index.php.
- [5] M. Treacy, K. Randall, S. Rao, J. Perry, and D. Chadi, *Zeit. Krist.* **212**, 768 (1997).
- [6] M. Treacy, I. Rivin, E. Balkovsky, K. Randall, and M. Foster, *Micropor. Mesopor. Mater.* **74**, 121 (2004).
- [7] Y. Li, J. H. Yu, and R. R. Xu, *Angewandte Chemie International Ed.* **52**, 1673 (2013).
- [8] L. Lin, A. Berger, R. Martin, J. Kim, J. Swisher, K. Jarimala, C. Rycroft, A. Bhowan, M. Deem, M. Haranczyk, and B. Smit, *Nat. Mater.* **11**, 633 (2012).
- [9] A. Lupulescu and J. Rimer, *Science* **344**, 729 (2014).
- [10] M. Kumar, M. Choudhary, and J. Rimer, *Nat. Comm* **9**, 2129 (2018).
- [11] X. Zhu, M. Goesten, A. Koekkoek, B. Mezari, N. Kosinov, G. Filonenko, H. Friedrich, R. Rohling, B. Szyja, J. Gascon, F. Kapteijn, and E. Hensen, *Chem. Sci.* **7**, 65066513 (2016).
- [12] V. A. Blatov, G. D. Ilyushin, and D. M. Proserpio, *Chemistry of Materials* **25**, 412 (2013).
- [13] A. P. Bartók, R. Kondor, and G. Csányi, *Phys. Rev. B* **87**, 184115 (2013).
- [14] S. De, A. A. P. Bartók, G. Csányi, and M. Ceriotti, *Phys. Chem. Chem. Phys.* **18**, 13754 (2016).
- [15] N. Henson, A. Cheetham, and J. Gale, *Chem. Mater.* **6**, 1647 (1994).
- [16] I. Petrovic, A. Navrotsky, M. Davis, and S. Zones, *Chem. Mater.* **5**, 1805 (1993).
- [17] *Database of Zeolite Structures – Explanatory Notes – Composite Building Units* (International Zeolite Association, 2018).
- [18] M. J. Sanders, M. Leslie, and C. R. A. Catlow, *Journal of the Chemical Society, Chemical Communications* **0**, 1271 (1984).
- [19] K.-P. Schröder, J. Sauer, M. Leslie, C. R. A. Catlow, and J. M. Thomas, *Chemical Physics Letters* **188**, 320 (1992).
- [20] J. D. Gale and A. L. Rohl, *Molecular Simulation* **29**, 291 (2003).
- [21] S. Zones, *Microporous and Mesoporous Materials* **144**, 1 (2011).
- [22] D. J. Earl and M. W. Deem, *Industrial and Engineering Chemistry Research* **45**, 5449 (2006).
- [23] R. A. Curtis and M. W. Deem, *Journal of Physical Chemistry B* **107**, 8612 (2003).
- [24] M. Ceriotti, *J. Chem. Phys.* **150**, 150901 (2019).
- [25] A. P. Bartók, S. De, C. Poelking, N. Bernstein, J. R. Kermode, G. Csányi, and M. Ceriotti, *Science Advances* **3**, e1701816 (2017).
- [26] L. Ward, R. Liu, A. Krishna, V. I. Hegde, A. Agrawal, A. Choudhary, and C. Wolverton, *Physical Review B* **96**, 024104 (2017).
- [27] F. Faber, A. Lindmaa, O. A. von Lilienfeld, and R. Armiento, *International Journal of Quantum Chemistry* **115**, 1094 (2015).
- [28] F. A. Faber, A. Lindmaa, O. A. von Lilienfeld, and R. Armiento, *Physical Review Letters* **117**, 135502 (2016).
- [29] F. A. Faber, L. Hutchison, B. Huang, J. Gilmer, S. S. Schoenholz, G. E. Dahl, O. Vinyals, S. Kearnes, P. F. Riley, and O. A. von Lilienfeld, *Journal of Chemical Theory and Computation* **13**, 5255 (2017).
- [30] K. Hansen, G. Montavon, F. Biegler, S. Fazli, M. Rupp, M. Scheffler, O. A. von Lilienfeld, A. Tkatchenko, and K.-R. Müller, *Journal of Chemical Theory and Computation* **9**, 3404 (2013).
- [31] K. Hansen, F. Biegler, R. Ramakrishnan, W. Pronobis, O. A. von Lilienfeld, K.-R. Müller, and A. Tkatchenko, *The Journal of Physical Chemistry Letters* **6**, 2326 (2015).
- [32] K. T. Schütt, H. Glawe, F. Brockherde, A. Sanna, K. R. Müller, and E. K. U. Gross, *Physical Review B* **89**, 205118 (2014).
- [33] M. Rupp, A. Tkatchenko, K.-R. Müller, and O. A. von

- Lilienfeld, *Physical Review Letters* **108**, 058301 (2012).
- [34] O. A. von Lilienfeld, R. Ramakrishnan, M. Rupp, and A. Knoll, *International Journal of Quantum Chemistry* **115**, 1084 (2015).
- [35] R. Ramprasad, R. Batra, G. Pilania, A. Mannodi-Kanakithodi, and C. Kim, *npj Computational Materials* **3**, 54 (2017).
- [36] O. Isayev, C. Oses, C. Toher, E. Gossett, S. Curtarolo, and A. Tropsha, *Nature Communications* **8**, 15679 (2017).
- [37] M. Gastegger, L. Schwiedrzik, M. Bittermann, F. Berzsenyi, and P. Marquetand, *The Journal of Chemical Physics* **148**, 241709 (2018).
- [38] J. Behler, *The Journal of Chemical Physics* **134**, 074106 (2011).
- [39] A. P. Bartók, M. C. Payne, R. Kondor, and G. Csányi, *Phys. Rev. Lett.* **104**, 136403 (2010).
- [40] A. Glielmo, C. Zeni, and A. De Vita, *Phys. Rev. B* **97**, 184307 (2018).
- [41] M. J. Willatt, F. Musil, and M. Ceriotti, *J. Chem. Phys.* **150**, 154110 (2019).
- [42] R. Drautz, *Phys. Rev. B* **99**, 014104 (2019).
- [43] S. Yang, M. Lach-hab, I. I. Vaisman, and E. Blaisten-Barojas, *The Journal of Physical Chemistry C* **113**, 21721 (2009).
- [44] J. D. Evans and F.-X. Coudert, *Chemistry of Materials* **29**, 7833 (2017).
- [45] A. P. Bartók and G. Csányi, *International Journal of Quantum Chemistry* **115**, 1051 (2015).
- [46] A. P. B. Wojciech J. Szlachta and G. Csányi, *Physical Review B* **90**, 104108 (2014).
- [47] A. P. Bartók, J. Kermode, N. Bernstein, and G. Csányi, *Physical Review X* **8**, 041048 (2018).
- [48] A. P. Bartók, M. C. Payne, R. Kondor, and G. Csányi, *Physical Review Letters* **104**, 136403 (2010).
- [49] J.-B. Maillet, C. Denoual, and G. Csányi, *AIP Conference Proceedings* **1979**, 050011 (2018).
- [50] V. L. Deringer and G. Csányi, *Physical Review B* **95**, 094203 (2017).
- [51] D. Dragoni, T. D. Daff, G. Csányi, and N. Marzari, *Physical Review Materials* **2**, 013808 (2018).
- [52] F. Musil, S. De, J. Yang, J. E. Campbell, G. M. Day, and M. Ceriotti, *Chemical Science* **9**, 1289 (2018).
- [53] B. A. Helfrecht, P. Gasparotto, F. Giberti, and M. Ceriotti, *Frontiers in Molecular Biosciences* **6**, 1 (2019).
- [54] J. V. Smith, *Mineralogical Society of America* **1**, 281 (1963).
- [55] S. L. Roux and P. Jund, *Computational Materials Science* **49**, 70 (2010).
- [56] A. M. Stoneham and J. H. Harding, *Annual Review of Physical Chemistry* **37**, 53 (1986).
- [57] H. Balamane, T. Halicioglu, and W. A. Tiller, *Physical Review B* **46**, 2250 (1992).
- [58] S. V. King, *Nature* **213**, 1112 (1967).
- [59] L. Guttman, *Journal of Non-Crystalline Solids* **116**, 145 (1990).
- [60] C. S. Marians and L. W. Hobbs, *Journal of Non-Crystalline Solids* **124**, 242 (1990).
- [61] D. S. Franzblau, *Physical Review B* **44**, 4925 (1991).
- [62] “R.I.N.G.S. code,” <http://rings-code.sourceforge.net/>.
- [63] B. Huang and O. A. von Lilienfeld, *The Journal of Chemical Physics* **145**, 161102 (2016).
- [64] A. J. Smola and B. Schölkopf, in *Proceedings of the Seventeenth International Conference on Machine Learning*, ICML '00 (Morgan Kaufmann Publishers Inc., San Francisco, CA, USA, 2000) pp. 911–918.
- [65] C. E. Rasmussen (MIT Press, 2006).
- [66] M. Ceriotti, M. J. Willatt, and G. Csányi, in *Handbook of Materials Modeling*, edited by W. Andreoni and S. Yip (Springer, Cham, 2018).
- [67] Y. Eldar, M. Lindenbaum, M. Porat, and Y. Y. Zeevi, *IEEE transactions on image processing : a publication of the IEEE Signal Processing Society* **6**, 1305 (1997).
- [68] M. Ceriotti, G. A. Tribello, and M. Parrinello, *J. Chem. Theory Comput.* **9**, 1521 (2013).
- [69] G. Imbalzano, A. Anelli, D. Giofré, S. Klees, J. Behler, and M. Ceriotti, *Journal of Chemical Physics* **148**, 241730 (2018).
- [70] For some of the ring descriptors, the dataset contains fewer than 2,000 unique feature vectors, in which case only the unique representations were considered as representative environments.
- [71] B. Schölkopf, A. J. Smola, and K.-R. Müller, *Neural Computation* **10**, 1299 (1998).
- [72] D. M. Wilkins, A. Grisafi, Y. Yang, K. U. Lao, R. A. DiStasio, and M. Ceriotti, *Proc. Natl. Acad. Sci. USA* **116**, 3401 (2019).
- [73] A. A. P. Bartók, S. De, C. Poelking, N. Bernstein, J. R. J. Kermode, G. Csányi, and M. Ceriotti, *Science Advances* **3**, e1701816 (2017).
- [74] M. J. Willatt, F. Musil, and M. Ceriotti, *Physical Chemistry Chemical Physics* **20**, 29661 (2018).
- [75] S. M. Auerbach, W. Fan, and P. A. Monson, *International Reviews in Physical Chemistry* **34**, 35 (2015).
- [76] C. Bores, S. M. Auerbach, and P. A. Monson, *The Journal of Physical Chemistry Letters* **9**, 3703 (2018).

Methods

I. SOAP PARAMETERS

Both the 3.5 Å and 6.0 Å SOAP descriptors employed 12 radial basis functions, a spherical harmonics band limit of 9, a cutoff transition width of 0.3, and an atomic Gaussian width of 0.3 (see the Quippy library reference for further information regarding the SOAP parameters, available at <https://libatoms.github.io/QUIP/descriptors.html>).

II. COMPUTATIONAL TOOLS

The data analysis and visualization for this study was performed in Python [1, 2] with the aid of the SciPy [3], NumPy [4, 5], scikit-learn [6], Atomic Simulation Environment (ASE) [7], quippy [8], and Matplotlib [9] packages in addition to Wolfram Mathematica 11.1 [10]. Atomic structure snapshots were created with Ovito [11] and Visual Molecular Dynamics (VMD) [12] with the Tachyon [13] rendering utility.

III. SPARSE KERNEL METHODS

The sparse kernel methods that we use (KRR and KPCA) both approximate the full $N \times N$ kernel matrix \mathbf{K}_{NN} with a reduced matrix \mathbf{K}_{NM} that includes only a subset of the columns of \mathbf{K}_{NN} in addition to \mathbf{K}_{MM} , the kernel matrix between only the M “representative” feature vectors.

A. Sparse Kernel Principal Component Analysis

One can approximate the full kernel \mathbf{K}_{NN} using the eigendecomposition of its low rank approximation $\tilde{\mathbf{K}}_{NN}$ [14] [15],

$$\mathbf{K}_{NN} \approx \tilde{\mathbf{K}}_{NN} = \tilde{\mathbf{U}}_{NN} \tilde{\mathbf{\Lambda}}_{NN} \tilde{\mathbf{U}}_{NN}^T \quad (1)$$

Under the Nyström approximation [14, 16], one can rewrite $\tilde{\mathbf{K}}_{NN}$ as [14],

$$\tilde{\mathbf{K}}_{NN} = \mathbf{K}_{NM} \mathbf{K}_{MM}^{-1} \mathbf{K}_{MN}. \quad (2)$$

Substituting \mathbf{K}_{MM} for its eigendecomposition [17], one obtains,

$$\tilde{\mathbf{K}}_{NN} = \mathbf{K}_{NM} \mathbf{U}_{MM} \mathbf{\Lambda}_{MM}^{-1} \mathbf{U}_{MM}^T \mathbf{K}_{MN}, \quad (3)$$

which can be rewritten as

$$\tilde{\mathbf{K}}_{NN} = \mathbf{G}_{NM} \mathbf{G}_{MN}, \quad (4)$$

where $\mathbf{G}_{NM} = \mathbf{K}_{NM} \mathbf{U}_{MM} \mathbf{\Lambda}_{MM}^{-1/2} \mathbf{U}_{MM}^T$. Centering \mathbf{G}_{NM} relative to its column means yields an approximation to the centered kernel matrix via Eqn. 4. Therefore, in the following, \mathbf{G}_{NM} and $\tilde{\mathbf{K}}_{NN}$ will be assumed to be centered. The eigendecomposition of $\mathbf{G}_{MN} \mathbf{G}_{NM} = \mathbf{V}_{MM} \mathbf{W}_{MM} \mathbf{V}_{MM}^T$ can then be used to approximate the M largest eigenvectors of $\tilde{\mathbf{K}}_{NN}$, i.e. [18, 19],

$$\tilde{\mathbf{U}}_{NM} \approx \mathbf{G}_{NM} \mathbf{V}_{MM} \mathbf{W}_{MM}^{-1}. \quad (5)$$

Therefore, the KPCA projections can be computed as [20]

$$\tilde{\mathbf{K}}_{NN} \tilde{\mathbf{U}}_{NM} = \tilde{\mathbf{K}}_{NN} \mathbf{G}_{NM} \mathbf{V}_{MM} \mathbf{W}_{MM}^{-1} \quad (6)$$

$$= \mathbf{G}_{NM} \mathbf{V}_{MM} \quad (7)$$

$$= \tilde{\mathbf{U}}_{NM} \mathbf{W}_{MM} \quad (8)$$

B. Sparse Kernel Ridge Regression

A sparse version of KRR can also be constructed from a subselected kernel matrix \mathbf{K}_{NM} . In this case, the solution to the KRR problem can be found by minimizing [21]

$$\frac{1}{2M} \|\mathbf{y}_N - \mathbf{K}_{NM} \mathbf{w}_M\|^2 + \lambda_1 \mathbf{w}_M^T \mathbf{K}_{MM} \mathbf{w}_M, \quad (9)$$

with respect to the weights \mathbf{w}_M , for which the solution is [21]

$$\mathbf{w}_M = (M\lambda_1 \mathbf{K}_{MM} + \mathbf{K}_{MN} \mathbf{K}_{NM})^{-1} \mathbf{K}_{MN} \mathbf{y}_N, \quad (10)$$

where \mathbf{y}_N is a vector of the target properties. In our case, a structural kernel \mathbf{K}_{LM} plays the role of the subselected kernel matrix, where each row \mathbf{k}_M is the summation over a structure S of the kernel between the environments comprising the structure and the M representative

environments across the whole dataset [22],

$$\mathbf{k}_M = \sum_{\mathcal{A} \in S} \mathbf{k}_M^{\mathcal{A}}, \quad (11)$$

where $\mathcal{A} \in S$ denotes the atomic environments that belong to structure S . Furthermore, we apply a scaling $\delta = M \times \text{Var}(\mathbf{y}_L) / \text{Tr}(\mathbf{K}_{MM})$ to the target property vector \mathbf{y}_L and to each kernel matrix, and we include an additional ‘‘jitter’’ parameter $\lambda_2 \mathbf{I}_{MM}$ so that the solution to our sparse KRR model is [22],

$$\mathbf{w}_M = (\lambda_1^2 \delta \mathbf{K}_{MM} + \lambda_2 \mathbf{I}_{MM} + \delta^2 \mathbf{K}_{ML} \mathbf{K}_{LM})^{-1} \delta^2 \mathbf{K}_{ML} \mathbf{y}_L, \quad (12)$$

where λ_1 and λ_2 are regularization parameters, and \mathbf{y}_L is a vector containing the known structural properties for the L structures. The parameter λ_1 was optimized along with the kernel width (in the case of Gaussian kernels) via five-fold cross validation; λ_2 was set to $10^{-16} \times \sigma$, where σ is the largest eigenvalue of $\lambda_1^2 \delta \mathbf{K}_{MM} + \delta^2 \mathbf{K}_{ML} \mathbf{K}_{LM}$.

The optimal weights \mathbf{w}_M can also be used to ‘‘decompose’’ the known structural property values \mathbf{y}_L into contributions from the N individual atomic environments $\hat{\mathbf{y}}_N$ across the whole dataset [22],

$$\hat{\mathbf{y}}_N = \mathbf{K}_{NM} \mathbf{w}_M. \quad (13)$$

IV. RESULTS FOR THE 1,000-STRUCTURE SUBSET

Fig. S1 shows the learning curves for the classical and SOAP descriptors for the 1,000-structure subset, analogous to Fig. 3 in the main text. Fig. S2 gives a comparison of the SOAP-based and classical descriptors at comparable dimensionality for predicting the unit cell volume per atom and the energy per mol Si for the 1,000-structure subset. Fig. S2 is analogous to Fig. 4 in the main text. The learning curves of the 1,000-structure subset are similar to the results of the 10,000-structure subset for the first 1,000 training points.




FIG. S1. Learning curves of the classical and SOAP descriptors for predictions of (a) volume per Si atom and (b) energy per mol Si for the 1,000-structure subset.




FIG. S2. Learning curves of classical and SOAP-KPCA descriptors with similar dimensionality for (a) predictions of volume per Si atom and (b) energy per mol Si for the 1,000-structure subset.

V. SOAP KPCA

Figs. S3, S4, S5, and S6 show the learning curves for the prediction of zeolite volume per Si atom and energy per mol Si using the SOAP-KPCA representation with different numbers of principal components. As the number of principal components composing the representation is increased, the prediction becomes more accurate.

In the case of predicting the volume per atom, a SOAP-KPCA representation including 50 principal components performs similarly to a representation using all 500 of the FPS SOAP components (marked in the graph as “Original”). In the case of predicting the energy per mol Si, upwards of 100 principal components are required to match the prediction accuracy of the representation containing all 500 FPS SOAP vector elements.

The convergence of the prediction accuracy to that of the full FPS SOAP vector as more information (more principal components) are included into the KPCA representation also serves as a validation of the method: the KPCA-based representation can emulate the diversity of the SOAP vector and thus the local chemical environment.

A comparison can also be made between the prediction accuracy of the SOAP-KPCA representation and that of a classical descriptor with comparable information content (dimensionality). In this paradigm, a four-component SOAP-KPCA with 3.5 Å cutoff would contain roughly the same amount of information as the Si–O distance and Si–O–Si angle descriptors; a 10-component SOAP-KPCA with 6.0 Å cutoff would contain approximately the same amount of information as the ring-based descriptor, as ring sizes in our dataset range from 3–12.

When comparing the different representations in this manner, one finds that the performance of the distance- and angle-based descriptors is comparable to, or slightly worse than, the performance of a SOAP-KPCA descriptor including less than five principal components. The same is true for the ring-based descriptor in predicting the energy, but even a single-component 6.0 Å SOAP-KPCA descriptor outperforms the ring descriptor in predicting the volume per Si atom.




FIG. S3. Learning curves for the KPCA-SOAP descriptor for cutoff radii of 3.5 Å and 6.0 Å including different numbers of principal components for the 10,000-structure sample. Increasing the amount of information embedded into the descriptor (increasing the number of principal components) results in a better prediction of the average volume per Si atom




FIG. S4. Learning curves for the KPCA-SOAP descriptor for cutoff radii of 3.5 Å and 6.0 Å including different numbers of principal components for the 10,000-structure sample. Increasing the amount of information embedded into the descriptor (increasing the number of principal components) results in a better prediction of the average energy per mol Si

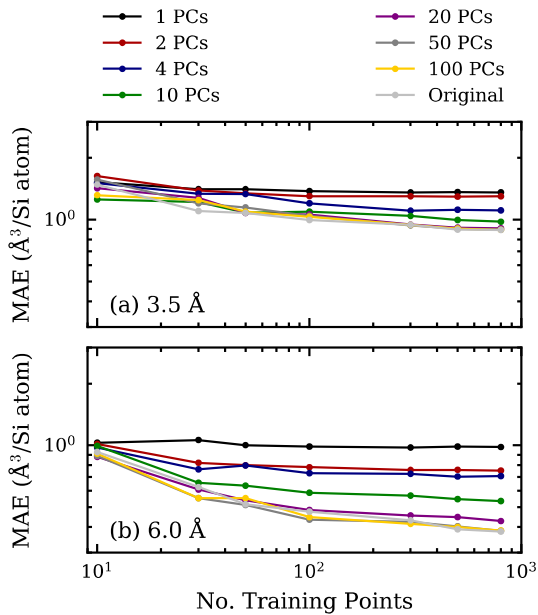


FIG. S5. Learning curves for the KPCA-SOAP descriptor for cutoff radii of 3.5 Å and 6.0 Å including different numbers of principal components for the 1,000-structure sample. Increasing the amount of information embedded into the descriptor (increasing the number of principal components) results in a better prediction of the average volume per Si atom




FIG. S6. Learning curves for the KPCA-SOAP descriptor for cutoff radii of 3.5 Å and 6.0 Å including different numbers of principal components for the 1,000-structure sample. Increasing the amount of information embedded into the descriptor (increasing the number of principal components) results in a better prediction of the average energy per mol Si

VI. RING-BASED DESCRIPTORS

In terms of predicting the molar volume and energy, both King’s definition and the shortest path definition perform very similarly. Figs. S7 and S8 show the learning curves for the ring descriptors built on the 1,000- and 10,000-structure samples. Two variations of the rings descriptor are examined. The “Distribution” (“Dist.”) variant is the descriptor described in the main text: the s -th element of the feature vector for a given Si-centered environment is how many rings of size s that include the central Si. The “Binary” (“Bin.”) variant is a boolean version of the “Distribution” representation: the s -th element of the feature vector is 1 if the central Si participates in at least one ring of size s and is 0 otherwise. As noted in the main article, the FPS of the ring descriptor often results in fewer than 2,000 unique environments. In these cases, only the unique feature vectors serve as representatives in the learning models. Consequently, the models based on the “binary” variants of the King and shortest path ring descriptors use 109 and 53 representatives (for the 1,000-structure sample) and 239 and 94 representatives (for the 10,000-structure sample), respectively; the model for the “distribution” variant of the shortest path ring counts in the 1,000-structure sample uses 763 representatives.

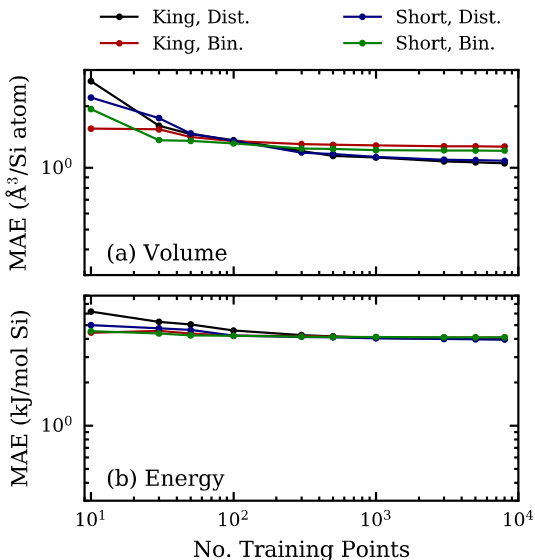


FIG. S7. Learning curves for the ring-based descriptors from the sample of 10,000 structures for predicting the volume per Si atom and the energy per mol Si.




FIG. S8. Learning curves for the ring-based descriptors for the sample of 1,000 structures for predicting the volume per Si atom and the energy per mol Si.

VII. PROPERTY CORRELATIONS

Figs. S9 is a plot of the Pearson correlation coefficients between the first 50 KPCs and the energy and volume as in Fig. 5 in the main text, but using a KPCA of SOAP representation with a 3.5 Å cutoff (instead of 6.0 Å). Similarly, Figs. S10 and S11 show the volume and energy correlations with the 3.5 Å and 6.0 Å SOAP-KPCA representations for the 1,000-structure subset.

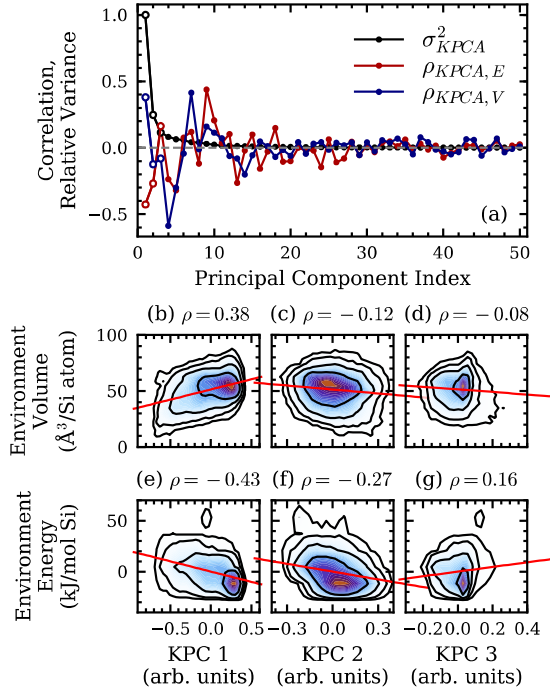


FIG. S9. (a) Pearson correlation coefficients between the first 50 KPCs of the 3.5 Å SOAP representation for the 10,000-structure sample and the decomposed environment volumes and energies. The relative variance in the KPCs at each of the first 50 components is also plotted. The correlation coefficients and relative variance of the first three components are highlighted with open symbols. (b)–(g) Kernel density estimation of all environments in KPC–property space with a least squares fit to the data to show the correlations in more detail. The value of the Pearson correlation coefficient is given above each contour set.

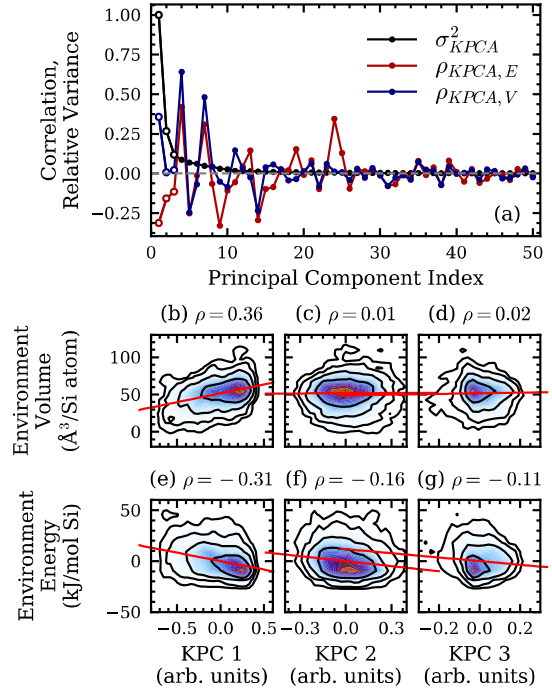


FIG. S10. (a) Pearson correlation coefficients between the first 50 KPCs of the 3.5 Å SOAP representation for the 1,000-structure sample and the decomposed environment volumes and energies. The relative variance in the KPCs at each of the first 50 components is also plotted. The correlation coefficients and relative variance of the first three components are highlighted with open symbols. (b)–(g) Kernel density estimation of all environments in KPC–property space with a least squares fit to the data to show the correlations in more detail. The value of the Pearson correlation coefficient is given above each contour set.

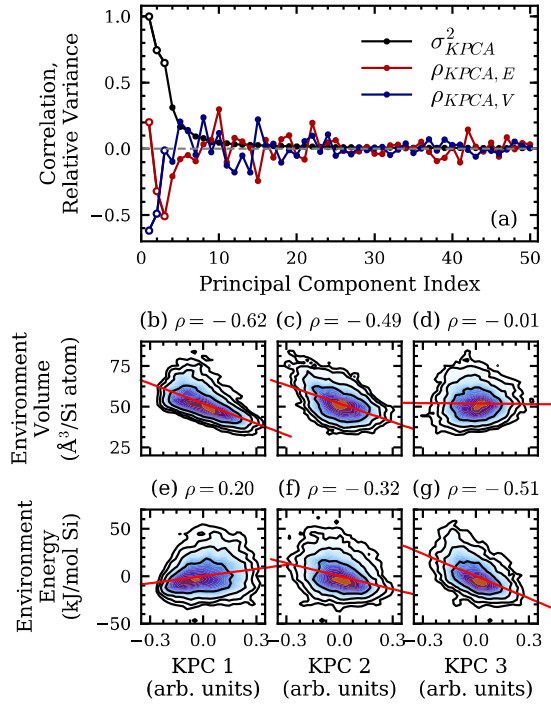


FIG. S11. (a) Pearson correlation coefficients between the first 50 KPCs of the 6.0 \AA SOAP representation for the 1,000-structure sample and the decomposed environment volumes and energies. The relative variance in the KPCs at each of the first 50 components is also plotted. The correlation coefficients and relative variance of the first three components are highlighted with open symbols. (b)–(g) Kernel density estimation of all environments in KPC–property space with a least squares fit to the data to show the correlations in more detail. The value of the Pearson correlation coefficient is given above each contour set.

-
- [1] T. E. Oliphant, *Computing in Science and Engineering* **9**, 10 (2007).
- [2] K. J. Millman and M. Aivazis, *Computing in Science and Engineering* **13**, 9 (2011).
- [3] E. Jones, T. Oliphant, P. Peterson, *et al.*, “SciPy: Open source scientific tools for Python,” (2001–).
- [4] T. E. Oliphant, *A Guide to NumPy* (Trelgol Publishing, USA, 2006).
- [5] S. van der Walt, S. C. Colbert, and G. Varoquaux, *Computing in Science and Engineering* **13**, 22 (2011).
- [6] F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, J. Vanderplas, A. Passos, D. Cournapeau, M. Brucher, M. Perrot, and Édouard Duchesnay, *Journal of Machine Learning Research* **12**, 2825 (2011).
- [7] A. H. Larsen, J. J. Mortensen, J. Blomqvist, I. E. Castelli, R. Christensen, M. Duak, J. Friis, M. N. Groves, B. Hammer, C. Hargus, E. D. Hermes, P. C. Jennings, P. B. Jensen, J. Kermode, J. R. Kitchin, E. L. Kolsbjerg, J. Kubal, K. Kaasbjerg, S. Lysgaard, J. B. Maronsson, T. Maxson, T. Olsen, L. Pastewka, A. Peterson, C. Rostgaard, J. Schitz, O. Schtt, M. Strange, K. S. Thygesen, T. Vegge, L. Vilhelmsen, M. Walter, Z. Zeng, and K. W. Jacobsen, *Journal of Physics: Condensed Matter* **29**, 273002 (2017).
- [8] N. Bernstein, G. Csányi, J. Kermode, A. Bartók-Pártay, L. Bartók-Pártay, F. Bianchini, A. Butenuth, M. Caccin, S. Cereda, A. Comisso, T. Daff, S. T. John, C. Gattinoni, G. Moras, L. Mones, A. Nichol, D. Packwood, L. Pastewka, G. Peralta, I. Solt, O. Strickson, W. Szlachta, C. Varnai, and S. Winfield, “libatoms/quip,” <https://github.com/libAtoms/QUIP>.
- [9] J. D. Hunter, *Computing in Science and Engineering* **9**, 90 (2007).
- [10] W. R. Inc., “Mathematica, Version 11.1,” Champaign, IL, 2017.
- [11] A. Stukowski, *Modelling and Simulation in Materials Science and Engineering* **18**, 015012 (2009).
- [12] W. Humphrey, A. Dalke, and K. Schulten, *Journal of Molecular Graphics* **14**, 33 (1996).
- [13] J. Stone, *An Efficient Library for Parallel Ray Tracing and Animation*, Master’s thesis, Computer Science Department, University of Missouri-Rolla (1998).
- [14] C. K. I. Williams and M. Seeger, in *Advances in Neural Information and Processing Systems 13*, edited by T. K. Leen, T. G. Dietterich, and V. Tresp (MIT Press, 2001) pp. 682–688.