

Visualizing Energy Landscapes through Manifold Learning

Benjamin W. B. Shires^{1,*} and Chris J. Pickard^{1,2,†}

¹*Department of Materials Science and Metallurgy, University of Cambridge,
27 Charles Babbage Road, Cambridge CB3 0FS, United Kingdom*

²*Advanced Institute for Materials Research, Tohoku University 2-1-1 Katahira,
Aoba, Sendai, 980-8577, Japan*

 (Received 5 May 2020; revised 6 August 2021; accepted 27 August 2021; published 5 November 2021)

Energy landscapes provide a conceptual framework for structure prediction, and a detailed understanding of their topological features is necessary to develop efficient methods for their exploration. The ability to visualize these surfaces is essential, but the high dimensionality of the corresponding configuration spaces makes this visualization difficult. Here, we present stochastic hyperspace embedding and projection (SHEAP), a method for energy landscape visualization inspired by state-of-the-art algorithms for dimensionality reduction through manifold learning, such as *t*-SNE and UMAP. The performance of SHEAP is demonstrated through its application to the energy landscapes of Lennard-Jones clusters, solid-state carbon, and the quaternary system C + H + N + O. It produces meaningful and interpretable low-dimensional representations of these landscapes, reproducing well-known topological features such as funnels and providing fresh insight into their layouts. In particular, an intrinsic low dimensionality in the distribution of local minima across configuration space is revealed.

DOI: [10.1103/PhysRevX.11.041026](https://doi.org/10.1103/PhysRevX.11.041026)

Subject Areas: Computational Physics,
Materials Science

I. INTRODUCTION

Potential (and free) energy surfaces (PESs) provide a conceptual and mathematical framework for describing the structures, kinetics, and thermodynamics of systems of interacting atoms [1–4]. The exploration of these surfaces forms the basis for the field of (computational) structure prediction, where the central task is the location of low-lying minima [5]. In order to develop efficient methods for their exploration, the ability to visualize these landscapes is crucial, enhancing our understanding of their complex topologies and providing insight into how we can improve techniques for structure prediction and global optimization. However, the potential energy of a system of atoms varies with the relative positions of all of its nuclei, meaning that for all but the simplest of systems the true surface cannot be depicted in the three spatial dimensions that we can easily visualize.

There are many successful approaches to structure prediction, from minima and basin hopping [6–8], particle swarm [9,10], and evolutionary algorithms [11,12] to

ab initio random structure searching (AIRSS) [13,14]. These methods generate a diverse range of structures for any given system, and AIRSS, in particular, emphasizes a broad sampling of the energy landscape. This approach has been successful in identifying novel and interesting phases of crystalline systems, particularly at high pressures [15–17]. Such a random search for local minima produces meaningful samples of configuration space, relating directly to the underlying distribution of the basins of attraction. However, these samples are generated in huge quantities (around 100 000 for modern application), and the resulting datasets can be hard to fully comprehend, making it difficult to determine how a given search might be improved. A means to extract and depict information contained in these large datasets of local minima, such that insight into the landscape being searched can be obtained, would be valuable.

The desire to visualize high-dimensional objects is not restricted to energy landscapes. A similar challenge is encountered for loss functions defining other optimization tasks, such as statistical regression schemes [18], and machine learning algorithms for classification, clustering, and neural networks [19–22]. For any such cost function surface, it would be revealing if an approximate low-dimensional representation could be produced which preserves as faithfully as possible the key features of the landscape, such as the location of stationary points. For the PESs sampled by methods for structure prediction, the important features include the relative layout of the local

*bs511@cam.ac.uk

†cjp20@cam.ac.uk

Published by the American Physical Society under the terms of the Creative Commons Attribution 4.0 International license. Further distribution of this work must maintain attribution to the author(s) and the published article's title, journal citation, and DOI.

minima, the connectivity of the corresponding basins of attraction, the relative volumes of the basins, and the presence of “superbasins” or funnels [2,3,23–27]. Here, we propose a method for energy landscape visualization which captures these features.

PES visualization can be viewed as a data science task in which we would like to project a high-dimensional source dataset (consisting of vectors describing structures in the configuration space of interest) into one, two, or three Euclidean dimensions. We base our approach on state-of-the-art algorithms for dimensionality reduction for visualization, such as t -distributed stochastic neighbor embedding (t -SNE) [28] and uniform manifold approximation and projection (UMAP) [29,30], drawing inspiration from their increasing application to datasets arising in physics [31–33], materials science [34,35], and cell biology [36–39]. We call this approach stochastic hyperspace embedding and projection (SHEAP).

There are a number of existing tools which allow the depiction of certain features of high-dimensional PESs in lower dimensions, including (metric) disconnectivity graphs [27,40–43] and “sketch maps” [44–46]. These approaches do not focus on faithfully preserving the spatial layout of minima on the true PES but rather on understanding the system dynamics.

Disconnectivity graphs display the connectivity of local minima in terms of minimum energy pathways between them, illustrating which structures are likely to be accessible and interchangeable at a given temperature [27,40,41]. They do not indicate the proximity of minima in the configuration space according to some defined distance metric, which basins are actually adjacent on the PES, or give a clear indication of the relative basin volumes [42].

Metric disconnectivity graphs retain some structural information about the minima by arranging them along one or two metric axes [42,43]. However, since they use one of the two or three plot dimensions to represent the energies of the structures depicted, there is one fewer dimension which can be used to represent the relative positions of the minima on the surface. Furthermore, techniques such as t -SNE and UMAP have been demonstrated to outperform the dimensionality reduction methods considered for metric disconnectivity graphs in the context of visualization [28,29].

Sketch maps follow a similar framework to the state-of-the-art methods that we build on here to produce SHEAP (see Sec. II). The algorithmic choices made within this framework deal with issues specific to the high-dimensional data produced in molecular dynamics trajectories, such as poor sampling at the transition states, and the high-dimensional nature of thermal fluctuations. Sketch maps ensure that data points from the same basin are close and separates clusters of structures from different basins [44,46]. However, it puts less emphasis on accurately reproducing the distances between the structures from different basins. Thus, the approach is less suited to dealing

with other forms of structural data, such as the minima produced from a structure search, for which there are no thermal fluctuations (although a sketch map has been applied to this type of data with success [47]).

In the next section, we introduce dimensionality reduction in the context of visualization, focusing on manifold learning approaches. We then outline our method, SHEAP, designed for PES visualization. We explore the performance of SHEAP by applying it first to model systems of atoms interacting through a Lennard-Jones (LJ) pair potential and then to structures of solid-state carbon, as well as the quaternary system C + H + N + O. We show that SHEAP yields meaningful low-dimensional representations of these landscapes, for example, reproducing well-known topological features such as funnels, using only the structural information contained in minima data obtained from random structure searches. Furthermore, and importantly, we reveal an intrinsic low dimensionality to the distribution of minima across these surfaces.

II. DIMENSIONALITY REDUCTION BY MANIFOLD LEARNING

Dimensionality reduction refers to the mapping of a high-dimensional dataset into a space of lower dimensionality [48]. Here, we refer to a source dataset $X = \{\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_N\} \subset \mathbb{R}^D$ and a low-dimensional representation $Y = \{\mathbf{y}_1, \mathbf{y}_2, \dots, \mathbf{y}_N\} \subset \mathbb{R}^s$, where N is the number of data points and D and s are the dimensionalities of the corresponding spaces. If the dimensionality of the map is low enough ($s \leq 3$), the so-called “map points” can be visualized in a scatter plot. The aim is to preserve latent features of the source data in the reduced representation.

Manifold learning is a nonlinear approach to dimensionality reduction which assumes that the source data lie on, or close to, some low-dimensional manifold embedded within the original high-dimensional space [49,50]. The idea is to produce a low-dimensional representation of the source data by deducing its projected layout across the manifold—this principle is illustrated in Fig. 1. This approach has advantages over linear methods such as principal component analysis (PCA) [51] in that it allows the manifold onto which the data are projected to adopt a folded geometry within the original space, making no assumptions about this geometry. The faithfulness of the projection depends on how close the source data are to actually lying on some embedded manifold of the desired dimension, how good the cover of the source data across this manifold is, and how well the given algorithm is able to capture the key features of the manifold from the spread of the source data.

There are many manifold learning algorithms, generally differing in the geometrical properties of the underlying manifold that they seek to preserve. Some aim to maintain distances or similarities across the entire dataset, e.g., isomap [52] or kernel PCA [53]. Others prioritize the preservation of local proximities, e.g., Sammon mapping [54],




FIG. 1. Data dimensionality reduction by manifold learning. (A) shows a three-dimensional dataset that is close to lying on a two-dimensional surface which is embedded within \mathbb{R}^3 . (B) shows the projection of the source data onto this embedded manifold, which is folded into a Swiss roll in the original 3D space. (C) shows the distribution of the data across the embedded manifold, providing a 2D visual representation of the source data.

locally linear embedding [55], or Laplacian eigenmaps [56]. For source data that are nonlinear, it is usually more important to keep the low-dimensional representations of similar points close together rather than trying to accurately reproduce a given metric at all length scales [28]. The t -SNE [28], LargeVis [57], and UMAP [29] algorithms also follow this principle but are able to patch together the local representations in such a way that global features are captured more faithfully than with other local methods [28,29]. These approaches proceed by constructing a weighted graph describing the structure of the source data and then projecting this description into a space of fewer dimensions through the optimization of a nonconvex cost function. SHEAP is formulated within the same framework, the outline of which is provided next, followed by a discussion of the motivation for the development of SHEAP.

A. Manifold learning through weighted graph matching

A weighted graph is constructed, consisting of similarities p_{ij} between each pair of data points i and j , corresponding to a measure of how close they are in the original space \mathbb{R}^D according to some distance metric $d: X \times X \rightarrow \mathbb{R}$. It is usually beneficial to an algorithm's speed and scalability to compute weights only to the points in the region local to each data point (e.g., the k nearest neighbors), with all other weights being set to zero—as is done in LargeVis [57], UMAP [29], and accelerated (approximate) implementations of t -SNE [58,59]. An initial distribution of N data points is then produced in

a space \mathbb{R}^s of many fewer dimensions (typically 2 or 3). This distribution may be generated at random, possibly according to some bias(es) and/or constraint(s), or through the application of some more primitive algorithm for dimensionality reduction, such as PCA or random projection (see the Appendix C). Another graph of weights q_{ij} is then constructed for the nodes defined by these map points, Y , possibly taking a different functional form from those defined between the high-dimensional points. The low-dimensional representation of the data is then computed by adjusting the layout of the map points such that the faithfulness with which the weights q_{ij} model each p_{ij} is optimized. An optimal layout is achieved by minimizing some cost function $C(Y; X)$ designed to penalize any mismatch. For a nonconvex cost function, local optimization schemes are necessary. Since there is usually some element of randomness in the initialization of the map points and/or the optimization of the cost function, different local minima can be obtained. Hence, these methods are considered stochastic. Note that the resulting projection is not associated with any physically meaningful set of axes (unlike PCA, for example)—it is only the relative positioning and clustering of the mapped data that are significant.

Because there are numerous stages at which there is scope for different possible and justifiable algorithmic choices, many algorithms are possible. These choices include how the weighted graphs describing the connectivity in the high and low dimensions are constructed, what cost function is optimized, how the map points are initialized, and what optimization scheme is used.

B. Motivation for SHEAP

t -SNE has become perhaps the most widely used manifold learning algorithm for visualization. However, in many test cases, the more recent methods LargeVis (2016) and UMAP (2018) appear to offer improvements in terms of the quality of the visualization, run-time speed, and scalability [29,57]. Approximate implementations, such as Barnes-Hut-SNE [58] and fast Fourier transform accelerated interpolation-based t -SNE (Fit-SNE) [59] dramatically accelerate the computation of t -SNE, though these too appear inferior to UMAP for many test datasets [29].

Here, we focus on the relative visualization qualities of these algorithms, leaving speed and scaling to future work. As demonstrated in Ref. [29], UMAP produces a local structure that is comparable to the projections of t -SNE (and LargeVis) but appears able to provide a more faithful global representation. This result is likely because of a better cost function—UMAP's choice provides a more sophisticated treatment of the repulsive contribution to the interaction between map points during optimization by introducing greater dependency on the layout of the source data. This choice is addressed further in Sec. III and discussed in detail in Appendix A.

However, UMAP does not truly optimize the cost function it defines. Rather, the algorithm relies on its specific choice for the optimization procedure, an approximate implementation of stochastic gradient descent (SGD) which uses probabilistic edge sampling and negative sampling [29,57], to deexaggerate the repulsive part of the cost function gradient. This reliance on an approximate optimization scheme is somewhat unsatisfactory and limits the flexibility of the algorithm. For example, UMAP is incompatible with our approach to incorporating basin volumes into the maps—this approach involves representing the structures in the map as hard spheres or circles rather than single points (see Sec. III B) and is straightforward to incorporate within t -SNE. Thus, we develop SHEAP, a hybrid of UMAP and t -SNE which captures some of the advancements made by UMAP, while still producing an algorithm that (i) truly optimizes a well-defined and theoretically justified cost function and (ii) is compatible with the hard-sphere scheme we introduce in Sec. III B.

III. SHEAP

In SHEAP, the weighted graph in the high-dimensional space is constructed according to the same scheme as in t -SNE. Following Ref. [28], the weight p_{ij} between source data points i and j is defined as

$$p_{ij} = \frac{p_{j|i} + p_{i|j}}{2N}, \quad (1)$$

where

$$p_{j|i} = \frac{\exp[-d(\mathbf{x}_i, \mathbf{x}_j)^2/2\sigma_i^2]}{\sum_{k \neq i} \exp[-d(\mathbf{x}_i, \mathbf{x}_k)^2/2\sigma_i^2]}. \quad (2)$$

Here, $d(\mathbf{x}_i, \mathbf{x}_j)$ is some metric for computing the distance between \mathbf{x}_i and \mathbf{x}_j in the high-dimensional space. t -SNE uses the Euclidean distance: $d(\mathbf{x}_i, \mathbf{x}_j) = \|\mathbf{x}_i - \mathbf{x}_j\|_2$. The value of σ_i is chosen such that the so-called perplexity of P_i takes some desired value, where P_i is the conditional probability of all other data points given \mathbf{x}_i . The perplexity of P_i , $\text{Perp}(P_i)$, is computed as

$$\text{Perp}(P_i) = 2^{H(P_i)}, \quad (3)$$

where

$$H(P_i) = -\sum_j p_{j|i} \log_2 p_{j|i}. \quad (4)$$

The perplexity $\text{Perp}(P_i)$ can be interpreted as a smooth measure of the effective number of neighbors of i and is set to the same value for all data points. It can be thought of as defining how far each point “looks out” to describe its local environment. A smaller perplexity favors a

faithful representation of the local connectivity of the data, while a larger value favors the global distribution. A “good” choice for this parameter results from a compromise between these two competing goals. A potential drawback arises in that the perplexity is a global parameter; there may be no single value that is suitable for all source data points and their local regions. It may be the case that a range of perplexities must be explored in order to gain the greatest insight from these projections. Which single perplexity value leads to the most complete picture can then be assessed.

Given a user input perplexity, each σ_i is computed using a straightforward binary search. Because the distributions of any two data points i and j will likely require different σ 's in order to correspond to the same perplexity, the conditional probabilities $p_{j|i}$ and $p_{i|j}$ will, in general, be different. To obtain a symmetric measure for the weight p_{ij} between the pair of data points, the joint probability of $p_{j|i}$ and $p_{i|j}$ is computed, as in Eq. (1).

In the low-dimensional space, the weight between i and j is computed according to a Student t distribution, similarly to in t -SNE but using a different normalization:

$$q_{ij} = \frac{(1 + \|\mathbf{y}_i - \mathbf{y}_j\|_2^2)^{-1}}{N(N-1)}. \quad (5)$$

As outlined in Ref. [28], the use of a heavier-tailed distribution for the low-dimensional weights allows moderate distances in the high-dimensional space to be faithfully represented by much larger distances in the projection, encouraging (larger) gaps to form in the low-dimensional map between the natural clusters present in the source data, alleviating the so-called crowding problem.

The key difference of the SHEAP algorithm from t -SNE is the choice of the cost function:

$$C = \sum_{i,j} p_{ij} \log\left(\frac{p_{ij}}{q_{ij}}\right) + (1 - p_{ij}) \log\left(\frac{1 - p_{ij}}{1 - q_{ij}}\right). \quad (6)$$

The above is the fuzzy set cross-entropy between the two sets of weights, as in UMAP, rather than Kullback-Leibler (KL) divergence [Eq. (A2)], used in t -SNE. Recall that we attribute the improved global structure in UMAP's projections over t -SNE's to this choice of the cost function, due to a better treatment of the repulsive contribution to the interaction between the map points (see Appendix A). In t -SNE, this repulsion arises due to the constraint that all of the weights q_{ij} sum to 1, providing a uniform repulsion between the map points. In UMAP, the repulsive contribution arises directly from the second, additional term in the cost function and is much more dependent on the structure of the source data. This contribution also allows us to discard the computationally

costly normalization appearing in the weights q_{ij} in t -SNE, which must be recomputed at every step during optimization. The constant normalization factor that replaces this term here is chosen such that weights q_{ij} are consistent with the weights p_{ij} in the limit that all data points lie on top of one another in both the source data and the projection, giving $p_{ij} = q_{ij} = [1/N(N-1)]$ for all i and j .

Given the above choices, the cost function gradient in SHEAP is given by

$$\frac{\partial C}{\partial \mathbf{y}_i} = 4N(N-1) \sum_j \left(p_{ij} - q_{ij} \frac{(1-p_{ij})}{(1-q_{ij})} \right) q_{ij} (\mathbf{y}_i - \mathbf{y}_j). \quad (7)$$

Optimization of the low-dimensional layout is achieved using the two-point steepest descent (TPSD) scheme of Barzilai and Borwein [60], and the option of early exaggeration is also implemented, as in t -SNE [28,61]. We can replace the use of early exaggeration with an initial use of SGD, though this alternative scheme does not appear to improve performance.

The layout of map points is initialized with an algorithm for random projection—see Appendix C. Initializing in this way tremendously accelerates the cost function optimization, relative to a random (confined) distribution of points.

Here, we produce low-dimensional depictions of energy landscapes by applying this algorithm to sets of structures in the relevant configuration space. Primarily, we consider datasets containing only local minima, obtained through random structure searching (RSS) [14,15] to promote broad sampling of the space. In Sec. IV, we demonstrate that these data are sufficient for capturing PES features of interest in the context of structure prediction. In Sec. IV C, we also explore the projection of unrelaxed structures obtained through random sampling.

A. Representing structural data

The configuration of a system of atoms can be represented in terms of a structure descriptor—a vector of real-valued functions of the atomic coordinates which describes their relative arrangement [62]. To allow comparisons between structures, a good descriptor is invariant with respect to permutations of equivalent atoms, as well as global translations, rotations, and reflections, and uniquely determines the structure up to these symmetries.

Two descriptors are considered here. The first is a sorted list of the pairwise distances between atoms. It is known that this descriptor is not complete [62]. However, for the simple systems of LJ particles considered (for which the energy depends only on pairwise distances), we find it to capture a sufficient proportion of the structural information and variation.

The second descriptor considered is the smooth overlap of atomic positions (SOAP) [62]. SOAP encodes relative nuclear geometry in the neighborhood of a given atom using an atomic density field, represented by 3D atom-centered Gaussian functions expanded in a basis of spherical harmonics and radial basis functions. A global SOAP descriptor for a given structure can be computed by taking the average SOAP fingerprint between all atomic environments [63]. When multiple species are present, one can either compute a single average incorporating all species or average separately for each species and then concatenate the species-specific averaged vectors [64]—here, we adopt the former approach. Recent work by Cheng *et al.* [64] finds success in pairing SOAP descriptors with kernel PCA to produce low-dimensional maps for a wide variety of structural datasets.

The distance metric d used here is the Euclidean distance, as in t -SNE [28]. This choice implicitly assumes that the source data are locally linear on the underlying manifold [28]. The SHEAP map produced from a given set of structures is expected to depend on both the chosen descriptor and distance metric. Here, we explore the former through a comparison of SHEAP maps produced from sorted lists of distances and from SOAP, for a single LJ system. An investigation into other distance metrics will be the subject of future work.

B. Basin volumes conserved using hard spheres

A shortcoming of using only the spread of local minima to depict a PES is that it gives no sense of how large each of the basins of attraction are. To address this limitation, we extend the methodology outlined in Secs. II A and III by representing each local minimum in the map as a hard n -sphere ($n = s - 1$) rather than a point, with volume (area in 2D) proportional to the (approximate) volume of its basin.

The number of times that structure i is located in a search is labeled c_i . We consider structures to be equivalent if the distance between their descriptor vectors falls below an appropriate similarity threshold. For a suitably large search, these counts should give representative estimates for the relative volumes of the corresponding basins. We represent each structure in the map by a sphere of radius

$$R_i = (c_i)^{1/s} R_0, \quad (8)$$

such that the volume $V_i \propto R_i^s$ is proportional to the count c_i . The quantity R_0 is the minimum sphere radius appearing in the map, taken by any structure that has been found just once—the choice of this parameter is important and is discussed in detail in Appendix (B1).

The hard spheres are incorporated into the framework outlined in Sec. II A through the addition of an extra term to the cost function. This hard-sphere potential is defined as

$$U_{\text{HS}} = \frac{1}{2} \gamma_{\text{HS}} \sum_i \sum_{j \neq i} \alpha_{ij}^2, \quad (9)$$

where γ_{HS} is the hard core strength and

$$\alpha_{ij} = \begin{cases} \frac{R_i + R_j - \|\mathbf{y}_i - \mathbf{y}_j\|_2}{2} & \text{if } \|\mathbf{y}_i - \mathbf{y}_j\|_2 < R_i + R_j, \\ 0 & \text{otherwise.} \end{cases} \quad (10)$$

For a sufficiently large value of γ_{HS} , the repulsion between any overlapping spheres (when present) is the dominant contribution to the gradient, meaning the layout is forced to contain no overlaps.

The inclusion of the hard-sphere potential from the start of the optimization is problematic. The spheres become jammed as they try to move across one another, particularly for densely packed spheres, resulting in convergence being drastically slow or not attainable at all. This problem can be avoided by allowing the spheres to access additional dimensions during the optimization, using the geometry optimization of structures from hyperspace method [65]. Alternatively, an efficient solution is to initially run the optimization without the hard-sphere contribution, switching it on only when the algorithm is close to convergence. With this approach, convergence is more easily attained if the hard spheres are introduced at a fraction of their final value and expanded over the course of a few hundred optimization steps. In this scheme, the hard spheres have minimal impact on the optimization, essentially just distorting the optimal layout once it has already been reached.

Again, we point out that this approach is incompatible with the SGD optimization scheme relied on by the UMAP algorithm—large overlaps between spheres are produced due to the omission of certain contributions to the gradient, resulting in instability in the optimization. A similar problem is encountered with the use of hard spheres at the same time as early exaggeration, though this issue can straightforwardly be circumvented by ensuring that the hard-sphere interaction is not switched on until after the exaggeration period has ended.

C. Energy represented with color

Energy is depicted across our plots through the coloration of the circles or spheres representing each structure. In Sec. IV, we also demonstrate that insight can be gained by illustrating the variation of other quantities across the projected configuration space.

IV. RESULTS

The results of applying SHEAP to various systems are now presented. First, we consider the PESs defined by LJ clusters, with pairwise interactions governed by

$$v_{\text{LJ}}(r) = 4\epsilon_{\text{LJ}} \left(\frac{\sigma_{\text{LJ}}^{12}}{r^{12}} - \frac{\sigma_{\text{LJ}}^6}{r^6} \right), \quad (11)$$

where ϵ_{LJ} is the depth of the potential energy well for a given particle pair and σ_{LJ} is the finite distance at which the interparticle potential is equal to zero. We set $\epsilon_{\text{LJ}} = 1$ and $\sigma_{\text{LJ}} = 2$ throughout. Clusters of 13, 38, and 55 atoms are considered.

For each dataset of LJ clusters, structures are represented as a sorted list of all pairwise distances between atoms, unless explicitly stated otherwise. In Appendix E, we also explore the use of a sorted list of inverse pairwise distances for LJ₁₃.

Following the study of these model systems, datasets of realistic solid-state systems, modeled by first-principles density-functional theory calculations, are addressed.

The examples considered illustrate the applicability of dimensionality reduction by manifold learning to minima data obtained from RSS and demonstrate the insight that can be gained from the resulting maps. A key finding is the intrinsic low dimensionality in the distributions of minima across the corresponding PESs.

A. Computational details

SHEAP is implemented in an OpenMP parallelized Fortran package, available under the GPL2 license [66]. Also implemented in this package are the *t*-SNE algorithm, as well as a simplified version of UMAP—the qualitative performance of SHEAP on standard datasets is compared against these in Appendix D. The maps presented in this work are plotted with Ovito, the open visualization tool [67]. The structural data presented in each map are generated using the AIRSS structure prediction package [68]. SOAP vectors are computed using the implementation in the Automatic Selection And Prediction (ASAP) package [64].

SHEAP is stochastic, so the final map differs for independent projections of the same source dataset, even for identical parameters. However, we find that, for an appropriately chosen parameter set, distinct runs consistently provide qualitatively equivalent results in all cases—for each dataset, we present the map for a single random seed.

B. LJ₁₃

LJ₁₃ has a relatively simple energy landscape. It is described as having a “single funnel” topology, meaning there is a general decrease in the energy of local minima with increasing structural similarity to the global minimum. This ground state structure is a filled 12-vertex icosahedron, and all local minima are within three rearrangements of it [69].

1. Relaxed versus unrelaxed structures

Figure 2 shows a 2D SHEAP map produced from local minima on the LJ₁₃ PES. The search, which comprises of 50 000 samples, uses a confinement sphere of radius $5.0\sigma_{\text{LJ}}$ for the initial random structures, with an imposed minimum

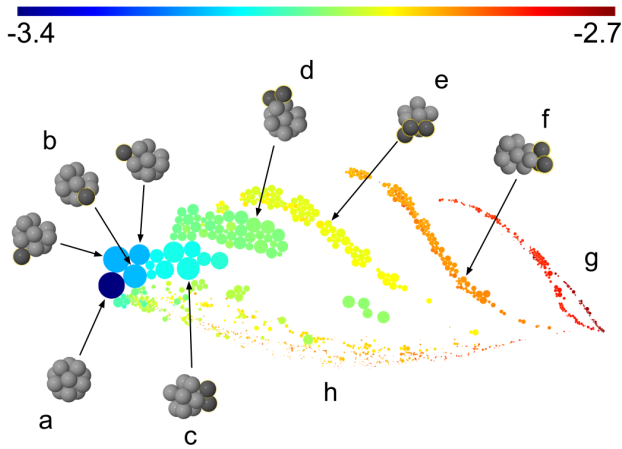


FIG. 2. Labeled 2D SHEAP map produced from a set of 1108 distinct minima on the LJ_{13} PES, located with RSS, described by sorted lists of all pairwise distances. The projection uses a perplexity of 30, with a minimum sphere radius $R_0 = 0.005$. The circles representing each structure are colored according to energy per atom in Lennard-Jones units; see the color bar at the top of the figure. Basin volumes are represented by the area of each circle. (a)–(g) denote clusters of structures of similar energy and configuration; representative structures from each are shown. Structures in (h) are less obviously clustered and tend to be comprised of less-ordered packings of atoms.

separation of $2.0\sigma_{LJ}$. Structures are relaxed to nearby local minima using TPSD [60], resulting in a dataset of 1108 distinct minima.

The funneled structure of the landscape is apparent, with minima energies correlating well with their distance from the global minimum (a) in the map. In addition, SHEAP clusters together structures of similar energies and arrangements. For example, the structures labeled (b) all result from one external atom of the icosahedron being moved from its ground state position to a different surface site. Structures in the cluster labeled (c) result from two external atoms from the icosahedron being moved to different surface sites.

The success observed in the application of SHEAP to the dataset of minima is in stark contrast to the map presented




FIG. 3. 2D SHEAP map produced from 50 000 unrelaxed, randomly sampled structures on the LJ_{13} PES, described by sorted lists of all pairwise distances. The projection uses a perplexity of 30, with a minimum sphere radius $R_0 = 0.01$. The circles representing each structure are colored according to energy per atom in Lennard-Jones units; see the color bar at the top of the figure.

in Fig. 3, which displays the 2D SHEAP projection of the 50 000 unrelaxed structures. Figure 3 contains no obvious clustering. The structures are arranged in a featureless, elongated “blob,” roughly ordered according to energy. Similar results are observed when projecting into 3D.

The disparity between Figs. 2 and 3 is striking. It suggests that the PES itself is truly high dimensional, and one cannot produce a meaningful representation of the landscape from a projection of points drawn randomly from it. In the language of Sec. II, the randomly sampled points do not appear to lie on or close to any low-dimensional manifold. The distribution of minima on this surface, on the other hand, does appear to have some inherent low dimensionality, according to the choice of structure descriptor and distance metric.

2. Projecting relaxed and unrelaxed structures together

The question then is, where do these unrelaxed structures lie relative to the minima? One might intuitively predict that they should fall into the gaps between the minima, but this prediction is not what is observed using SHEAP.

Figure 4 shows a 2D SHEAP map constructed from a dataset comprised of local minima and unrelaxed structures on the LJ_{13} PES. The dataset contains the same relaxed structures that are presented in Fig. 2, as well as a series of unrelaxed random structures generated according to different constraints. These are (i) a specified number of symmetry operations, from 2 up to 5, with a fixed confinement sphere radius of $5.0\sigma_{LJ}$, and (ii) no symmetry constraints, with varying confinement sphere radius, from $3.0\sigma_{LJ}$ up to $5.0\sigma_{LJ}$ in increments of $0.5\sigma_{LJ}$. There are 1000 structures included for each set of constraints. Note that the distribution of minima is distorted by the addition of the unrelaxed structures. We attribute this change to a reduced sensitivity in the relations between the minima, due to having to represent similarities to structures not on the manifold occupied by the minima. A lower perplexity is found to alleviate the distortion somewhat.

The majority of unrelaxed structures lie well away from the minima in a single, elongated cluster. As with Fig. 3, these unrelaxed structures are arranged roughly according to energy. Here, we also see that denser packings (smaller confinement sphere radius) tend to lie toward the low-energy end of the cluster, closest to the minima, while the higher-symmetry packings tend to lie toward the edges. Interestingly, a few of the unrelaxed structures actually cluster with the minima, despite being of significantly higher energies. Each of these structures is found to possess five symmetry operations, and all contain at least 11 atoms that are close to the ideal icosahedral close packing. This result demonstrates one aspect of the AIRSS philosophy—that sensible random structure generation does much of the work in efficiently locating local minima.

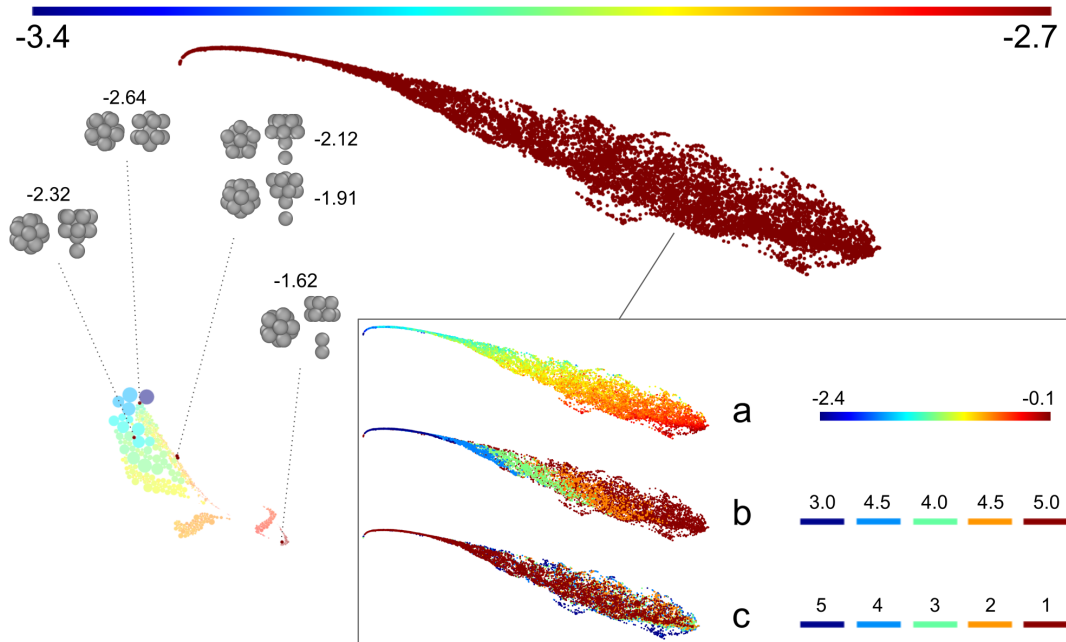


FIG. 4. 2D SHEAP map constructed from a dataset of minima and unrelaxed structures on the LJ_{13} PES, described by sorted lists of all pairwise distances. The projection uses a perplexity of 8. Circles in the main figure are colored according to energy per atom in Lennard-Jones units, given by the color bar at the top of the figure. Minima (all in the bottom-left cluster) are faded, with basin volumes represented by the area of each circle, using a minimum sphere radius $R_0 = 0.005$. Circles representing unrelaxed structures are enlarged to highlight them. The figure inset provides different colorations of the top cluster, which contains only unrelaxed structures. Coloring is according to (a) energy per atom in Lennard-Jones units, (b) the confinement radius in Lennard-Jones units, and (c) the number of imposed symmetry operations—see corresponding color bars to the right of the inset. The few unrelaxed structures that cluster with the minima (all possessing five symmetry operations) are labeled with their energy per atom in Lennard-Jones units, alongside two perpendicular perspectives of their configuration.

3. Projecting a steepest descent pathway

In order to reveal the connection between the prerelaxed structures and the central layout of local minima in Fig. 4, a SHEAP map is produced for a source dataset consisting of the same set of minima, plus selected structures visited along the pathway of steepest descent from an unrelaxed structure generated without symmetry constraint, down to the global minimum. The resulting projection is provided in Fig. 5.

We observe that this pathway traverses the space in the map between the unrelaxed structures and the minima and passes by many local minima on its way to the icosahedral ground state. The presence of the descent pathway leads to slight distortions in the distribution of minima in the map, relative to Fig. 2, likely due to oversampling of the regions on the PES visited during the relaxation.

These SHEAP maps for LJ_{13} provide an interesting perspective on the structure of its energy landscape. They suggest that the surface consists of clusters of minima which all sit very close to one another, with the unrelaxed structures predominantly residing outside the region occupied by these minima. Furthermore, they imply the existence of some inherent low dimensionality in the distribution of these minima across the surface.




FIG. 5. 2D SHEAP map produced from a dataset consisting of local minima on the LJ_{13} PES, plus selected structures visited along the pathway of steepest descent from an unrelaxed structure generated without symmetry constraint, down to the global minimum. Structures are described by sorted lists of all pairwise distances. The projection uses a perplexity of 30. Circles are colored according to energy per atom in Lennard-Jones units, given by the color bar at the top of the figure. For the minima, basin volumes are represented by the area of each circle, with a minimum hard-sphere radius $R_0 = 0.005$. Circles representing structures visited during the relaxation are enlarged to highlight them.

C. LJ_{38}

In previous work, LJ_{38} is described as having an archetypal double funnel energy landscape; this arrangement of its local minima is depicted clearly in disconnectivity graphs, such as in Ref. [70]. The 38-atom system stands out in the context of LJ clusters due to its octahedral global minimum—most other cluster sizes exhibit ground states based on icosahedra [71]. While the octahedron is the lowest-energy arrangement for LJ_{38} , there is an icosahedral metastable configuration that is only very slightly higher in energy. These two minima reside at the bottom of the two separate “funnels” in the landscape, with the funnel

corresponding to the icosahedral structure being much broader than that corresponding to the global minimum, i.e., containing many more minima [72]. For this reason, LJ_{38} is well known to be a relatively difficult test for structure searching algorithms.

1. Two-dimensional map of minima data

Figure 6(A) shows a 2D SHEAP map produced from minima on the LJ_{38} PES. The search consists of 100 800 samples and uses the same parameters and relaxation scheme as for the 13-atom system, resulting in a dataset of 89464 distinct minima.

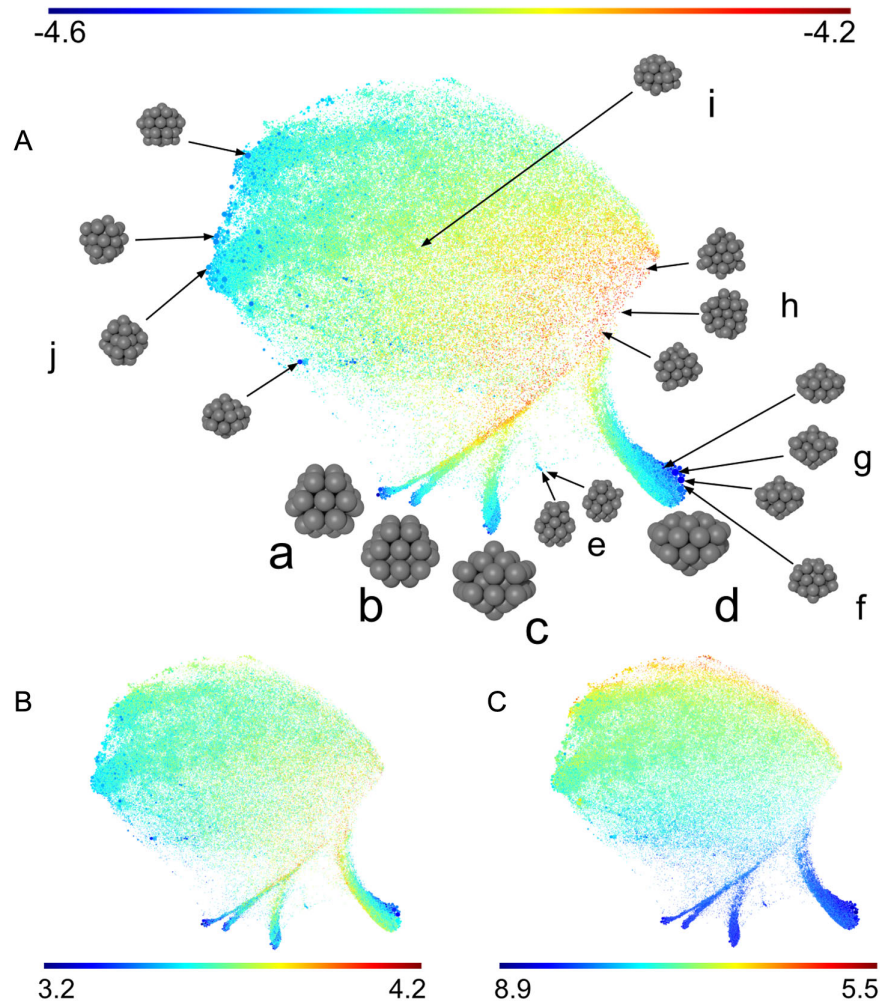


FIG. 6. (A) shows a 2D SHEAP map for 89464 minima on the LJ_{38} PES, generated with RSS, described by sorted lists of all pairwise distances. The projection uses a perplexity of 30, with a minimum sphere radius $R_0 = 0.008$. The circles representing each structure are colored according to energy per atom in Lennard-Jones units; see the color bar at the top of the figure. Basin volumes are represented by the area of each circle. (a)–(d) label proposed funnels in the LJ_{38} PES; structures shown alongside correspond to the lowest-energy minima residing in each. (a) contains the octahedral global minimum; (b) and (c) could be considered subfunnels of this minimum. (d) contains the icosahedral lowest metastable structure. (e) labels a small island of structures occupying the space between (a)–(c) and (d). (f) and (g) label low-energy minima residing in funnel (d) and are closely related to icosahedral structure. (h)–(j) label representative structures from the large, “brainlike” cluster of minima which contains the majority of structures—these are predominantly of low symmetry. (B) and (C) depict the same SHEAP map as in (A) but colored according to different physical parameters. Circles in (B) are colored according to the convex hull volume of each cluster. Circles in (C) are colored according to the number of contacts per atom in each cluster (see the main text for details). Corresponding color bars are provided below these subfigures.

LJ_{38} has many more degrees of freedom than LJ_{13} , leading to a far more complex landscape topology, with many more local minima. This difference is reflected in the searches, with almost 90 times more distinct minima being located for the larger cluster from little over twice as many samples and a much lower disparity between the largest and smallest basin volumes. The map for LJ_{38} exhibits greater connectivity than that for LJ_{13} , with the former showing much less clear-cut clustering of structures and a more continuous variation in the energy. We attribute this difference to the vastly greater number of minima compared to LJ_{13} , with the variation between minima being much more incremental for the larger system. For example, starting with the global minimum and moving one atom from its ground state position to a surface site leads to a relatively much greater deviation from the ground state structure for LJ_{13} than for LJ_{38} .

The standout features in Fig. 6(A) are the large “brainlike” cluster of minima (mostly low-symmetry, amorphous structures) and the strands of minima leading away from this cluster with progressively decreasing energies. With regards to the latter, there is an obvious connection to be drawn to the supposed funnels. We find that the octahedral global minimum and the icosahedral lowest metastable minimum reside at the ends of two of these different strands, labeled (a) and (d), respectively. Preceding them are closely related structures, e.g., those labeled (g), which all result from a single external atom of the icosahedral structure being relocated to a different surface site. Satisfyingly, the existence of these “funnels” arises purely from an analysis of the minima, through the structural information captured in the descriptor vectors. In the construction of the corresponding disconnectivity graph, funneling is inferred through an examination of transition state data and an evaluation of minima connectivity in the context of minimum energy pathways via these transition states. However, note that one can construct approximate disconnectivity graphs based solely on the structural relationships between minima (using fingerprint distances), circumventing the expensive transition state calculations by computing estimates for the energy barriers in the system [73].

The following question arises: What are the remaining branches of structures appearing in the SHEAP map? In Fig. 6, the most apparent of these additional strands of structures are labeled (b) and (c), alongside illustrations of the lowest-energy structure residing in each. We propose these to be previously unreported funnels in the LJ_{38} PES, both of which could potentially be considered subfunnels of a superfunnel also containing (a)—they appear to be part of the same branch, eventually separating off to comprise independent funnels.

2. Other color schemes

To better interpret the layout of structures across the map, two further schemes for coloring the circles are considered, in addition to the energy per atom. These

are the convex hull volume of the structure per atom and the number of contacts in each structure per atom, provided in Figs. 6(B) and 6(C), respectively.

The convex hull volume shows a very similar distribution to the energy, both in the main cluster of structures (albeit with slightly less variation) and in the funnels. We attribute this similarity to the LJ energy being entirely dependent on the pairwise distances between atoms; there are no three-body or higher terms. Hence, the energy is expected to be closely correlated to the packing density within the structures, which, in turn, should correlate well with the map position.

In order to produce Fig. 6(C), a contact is defined as any pairwise separation that is less than $2.3\sigma_{LJ}$, 2.5% above the minimum separation in the LJ pair potential. We observe that the funnels in the map correspond to regions of very high numbers of contacts per atom, and the number of contacts steadily reduces with increasing separation from these funnels, roughly from the bottom to the top of the main cluster.

3. Projecting into the third dimension

So far, SHEAP maps have been presented for only two map dimensions. However, we find that often there is a non-negligible improvement in the quality of visualization by accessing an extra dimension and projecting into 3D. To illustrate this point, in Fig. 7, we present three perpendicular perspectives of a 3D SHEAP map produced




FIG. 7. Three perpendicular perspectives of a 3D SHEAP map produced from the same set of LJ_{38} minima that are presented in Fig. 6. The projection uses a perplexity of 30, with a minimum sphere radius $R_0 = 0.008$. The spheres representing each structure are colored according to energy per atom in Lennard-Jones units; see the color bar at the top of the figure. Basin volumes are represented by the volume of each sphere. Each viewpoint looks down one of the axes displayed at the bottom right of the figure; the axes have no other physical meaning.

for the same dataset of LJ_{38} minima. Note that the axis system is arbitrary.

The 3D map is qualitatively very similar to that in 2D; the 2D map looks almost identical to a projection of the 3D distribution onto the x - y plane. However, one can see from the other two perspectives that there are features protruding out of this plane that are not represented faithfully in 2D. Most notably, the separation between the icosahedral funnel (d) and the three other large (sub) funnels (a)–(c) is more pronounced with the extra dimension. The classification of (a)–(c) as subfunnels of a single, larger “superfunnel” is also more obvious in this projection.

4. Using a different descriptor: SOAP

The distance-based descriptor employed so far is appropriate for LJ systems, since their energy is determined entirely by pairwise separations. However, the picture of an energy landscape obtained with SHEAP should be qualitatively robust to the choice of descriptor, as long as the relevant structural information is captured. To test this hypothesis, we compare the above results for LJ_{38} , which use sorted lists of distances, with the use of SOAP vectors [62].

Unlike the sorted list of distances (for which the only parameter is the cutoff distance), the SOAP vector for a given structure is dependent on many parameter choices:

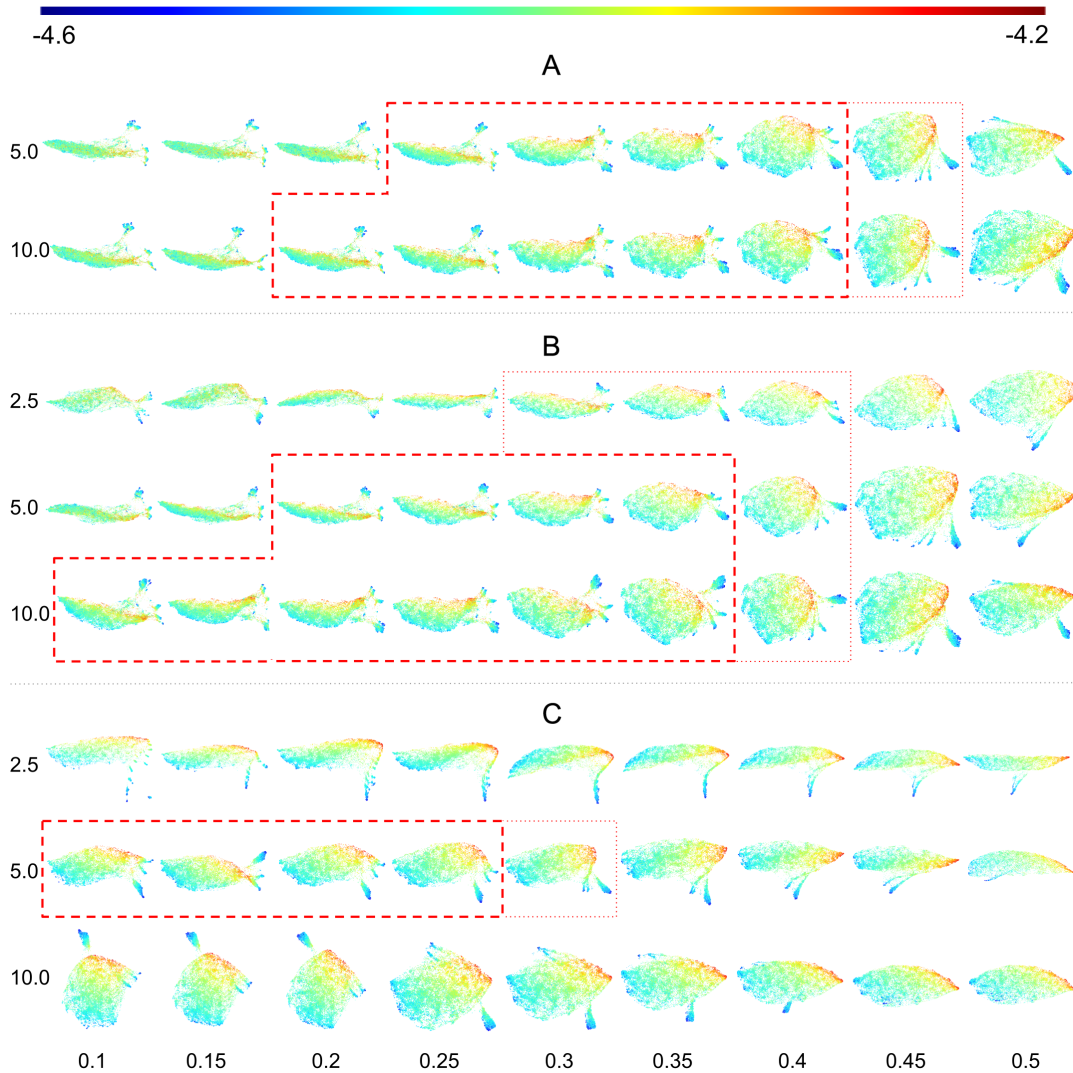


FIG. 8. 2D SHEAP maps for 10 000 minima (prior to similarity check) on the LJ_{38} PES, described using SOAP with various parameters. (A), (B), and (C) show results for $n_{\max}, l_{\max} = 15, 9, 12, 6$, and $8, 4$, respectively. r_{cut} is consistent across each row, values given on the far left in σ_{LJ} . σ is consistent across each column, values given at the bottom in σ_{LJ} . Projections use a perplexity of 15, with R_0 computed by Eq. (B1). The circles representing each structure are colored according to energy per atom in Lennard-Jones units; see the color bar at the top of the figure. Basin volumes are represented by the area of each circle. Maps in the bold, dashed, red perimeters represent “suitable” SOAP parameter choices, producing the same key features as Fig. 6. Maps in the faded, dashed, red perimeters also show these features but less clearly.

r_{cut} specifies the distance beyond which neighbor atoms are not taken into account, n_{max} specifies the number of radial basis functions to be used for each local descriptor, l_{max} specifies the maximum degree of the spherical harmonics, and σ specifies the standard deviation of the Gaussians. This level of freedom can be beneficial, allowing the descriptor to be tailored to better suit a given system, but it demands more input on the part of the user. The choice of radial basis functions is also important but is not addressed in detail here. By default, ASAP uses spherical Gaussian-type orbitals [74] rather than the original polynomial radial basis set [62], allowing much faster analytic computation [75].

Figure 8 shows 2D SHEAP maps for a subset of the LJ₃₈ dataset studied in Sec. IV C, described by SOAP vectors using a range of parameters. A reduced dataset (10 000 structures) is used to lessen the computational cost.

ASAP permits up to 15 radial and nine spherical harmonic SOAP basis functions. Figure 8(A) illustrates the effect of varying σ and r_{cut} for these values of n_{max} and l_{max} . Note that no structure in the dataset contains a separation greater than $10.0 \sigma_{\text{LJ}}$. There are no maps for $r_{\text{cut}} = 2.5 \sigma_{\text{LJ}}$ (as with other n_{max} and l_{max} considered), since this cutoff is too low relative to the number of radial basis functions, preventing normalization of the latter. While the exact structure of the map varies with each set of parameters, there is a broad range (inside the bold, dashed, red perimeter) for which the key features are consistent with the maps produced using sorted lists

of distances. Within this range of parameters, the large cluster of low-symmetry structures that is present in Figs. 6 and 7 is clear, the largest funnel [(d) in Fig. 6] is clearly revealed, and at least two other funnels are also discernible. For σ values too large, the definition of the map diminishes, and the funnels are absorbed into the main cluster. For σ values too small, the smooth variation in energy across the cluster of low-symmetry structures is lost.

Figures 8(B) and 8(C) depict results using fewer basis functions: $n_{\text{max}}, l_{\text{max}} = 12, 6$ and $8, 4$, respectively. Comparing these to Fig. 8(A), we see that the band of suitable σ shifts to lower values with lower n_{max} and l_{max} . We also observe that the scope of acceptable parameters is of similar size for Figs. 8(A) and 8(B) but is significantly restricted in Fig. 8(C), for which only the midrange r_{cut} is able to produce the expected map features for any σ .

In general, the funnels labeled (a) and (b) in Fig. 6 are not as well separated in the maps in Fig. 8. We attribute this difference to the smaller sample size used for these projections—2D and 3D maps produced using SOAP from the full dataset (Fig. 9, for which $r_{\text{cut}} = 10.0 \sigma_{\text{LJ}}$, $n_{\text{max}} = 15$, $l_{\text{max}} = 9$, and $\sigma = 0.3 \sigma_{\text{LJ}}$) provide clear separation of these funnels.

We conclude that, with the right choice of SOAP parameters, the picture of the LJ₃₈ energy landscape obtained through SHEAP is qualitatively equivalent for the two descriptors. Furthermore, we find that the range of suitable SOAP parameters is relatively broad, and easy to




FIG. 9. 2D and 3D SHEAP maps produced from the same set of LJ₃₈ minima that are presented in Fig. 6, described using SOAP with parameters $r_{\text{cut}} = 10.0 \sigma_{\text{LJ}}$, $n_{\text{max}} = 15$, $l_{\text{max}} = 9$, and $\sigma = 0.3 \sigma_{\text{LJ}}$. The projections use a perplexity of 30, with a minimum sphere radius $R_0 = 0.006$ for the 2D map and $R_0 = 0.01$ for 3D. The circles and spheres representing each structure are colored according to energy per atom in Lennard-Jones units; see the color bar at the top of the figure. Basin volumes are represented by the area (volume) of each circle (sphere). Each viewpoint of 3D map looks down one of the axes displayed, as labeled; the axes have no other physical meaning. (a)–(d) in the 2D map label proposed funnels in the LJ₃₈ PES, analogous to those labeled in Fig. 6, with the same low-energy minima residing at the ends.

identify for this system, via a scan of possible values, visualized through SHEAP.

D. LJ₅₅

LJ₅₅ is referred to as a structure seeker, meaning it has a very distinct global minimum that is easily located in a structure search, with no competing morphologies separated by high barriers [76]. Here, we present a representation of this landscape obtained with SHEAP, highlighting its contrast to that of LJ₃₈.

Figure 10 shows 2D and 3D SHEAP maps produced from minima on the LJ₅₅ PES. The search consists of 50 000 samples and uses the same parameters and relaxation scheme as the above LJ systems, giving 49813 distinct minima.

There is clearly a funneled structure to the arrangement of minima, terminating with the icosahedron and closely related configurations, such as those labeled (b), which all have the structure of the I_h ground state but with one external atom moved to a different surface site. However, the maps also suggest that describing the landscape as a single funnel would be an oversimplification. There is a second strand of structures which branches off from the funnel to the global minimum, leading to a relatively low-energy D_{5h} configuration. This structure is related to the I_h

configuration via a twist about the fivefold axis [2]. It stands out in the corresponding disconnectivity graph due to its relatively low energy and high barrier for transition to the global minimum [69]; this picture is complemented by the detail in the structural similarity and subfunnel arrangement provided by SHEAP.

Another band of (sub)funnels is present, well separated from those leading to the I_h and D_{5h} structures in the 3D map [in a similar manner to the separation of (d) from (a)–(c) in Fig. 6]. However, these are all of relatively high energy [unlike (d) in Fig. 6] and so are not competitive with the global minimum.

Despite some similarities, it is apparent from the SHEAP maps why LJ₅₅ is a more straightforward problem than LJ₃₈ in terms of global optimization. With LJ₅₅, the competing funnel(s) or strand(s) are shallow compared to that of the global minimum, and there are no similarly low-energy structures outside of the main funnel—the D_{5h} structure exceeds the global minimum by approximately 0.1ϵ per atom. With LJ₃₈, competing funnels are deep and terminate with structures similar in energy to the global minimum (only approximately 0.01ϵ per atom higher). Furthermore, for LJ₃₈ the most voluminous funnel is not the one containing the global minimum, and the maps also possess a region of low-energy structures within the amorphous bulk of configurations, which is not the case for LJ₅₅.




FIG. 10. 2D and 3D SHEAP maps for 49813 minima on the LJ₅₅ PES, generated with RSS, described by sorted lists of all pairwise distances. Projections use a perplexity of 30, with a minimum sphere radius 0.01. The circles representing each structure are colored according to energy per atom in Lennard-Jones units; see the color bar at the top of the figure. Basin volumes are represented by the area of each circle. (a)–(f) label selected structures in the landscape: (a) I_h , global minimum; (b) C_s , all have one surface atom relocated relative to global minimum; (c) D_{5h} , lowest-energy structure in funnel branching off from that to global minimum; (d) C_1 , one atom relocated relative to D_{5h} ; (e) C_{5v} and (f) C_5 , high-symmetry structures outside the main (sub)funnels to I_h and D_{5h} . Each viewpoint of 3D map looks down one of the axes displayed, as labeled; the axes have no other physical meaning.

E. Carbon

To demonstrate the applicability of SHEAP to realistic periodic systems, we consider a dataset of solid-state carbon structures, taken from the Materials Cloud Archive [77]. The set contains 11370 structures obtained with AIRSS performed at 10 GPa for a range of cell sizes (up to 24 atoms), with 2–4 symmetry operations imposed on the prerelaxed structures. All structures are based on $sp/sp^2/sp^3$ hybridized carbon, but they exhibit a wide range of topologies. The interactions in this system contain significant three-body contributions [78], meaning the sorted list of pairwise distances is not a suitable structure descriptor. Hence, we use SOAP, which is capable of describing these relationships completely [62].

As demonstrated with LJ₃₈, the SOAP descriptor has a number of parameters which effect the appearance of a SHEAP map. With that model system, we are able to obtain suitable values relatively easily via a scan of the parameter space. However, it would, of course, be preferable to have a systematic way to obtain good parameter choices for any given system. For systems of real atoms, we find a satisfying solution is the heuristic approach of Cheng *et al.* [64], providing a set of so-called “universal” parameters. These come in three flavors: “smart,” which computes a suitable set of parameters by considering the typical and minimum bond lengths of each of the species involved, “minimal,” which is similar to smart but uses fewer basis functions, and “long-range,” which again is similar to smart

but aims to capture more of the long-range structure. For the smart and long-range cases, two sets of SOAP parameters are computed, at slightly different length scales. A single structure descriptor is constructed by concatenating the global SOAP vectors for each set. Here, we provide a SHEAP map for the above carbon dataset produced using the smart parameters—see Fig. 11.

Most structures lie within a single, elongated arrangement. Several have distinctly large basins; most of these are low energy and lie toward the edge of the distribution. Many of these correspond to well-known carbon arrangements, including diamond (a) and graphite (g), which reside at opposite ends of the map, as well as lonsdaleite (b) and the chiral framework structure (d) of Ref. [79]. These and other standout clusters are labeled to provide a representative picture of the distribution of structures.

Comparing the SHEAP map for carbon to those for the LJ systems highlights key differences in their energy landscapes. With the LJ systems, neighboring minima tend to be close in energy. With carbon, there is a large degree of “mixing” between high- and low-energy structures; basins of high-energy minima are often surrounded by those of much lower energy, and vice versa. Furthermore, the variation in the energy of minima across the carbon map shows less clear long-range trends than for the LJ clusters. While the LJ landscapes possess clear funnels of minima, such features do not appear to be prevalent in the carbon PES. Through this comparison, it is easy to see why LJ

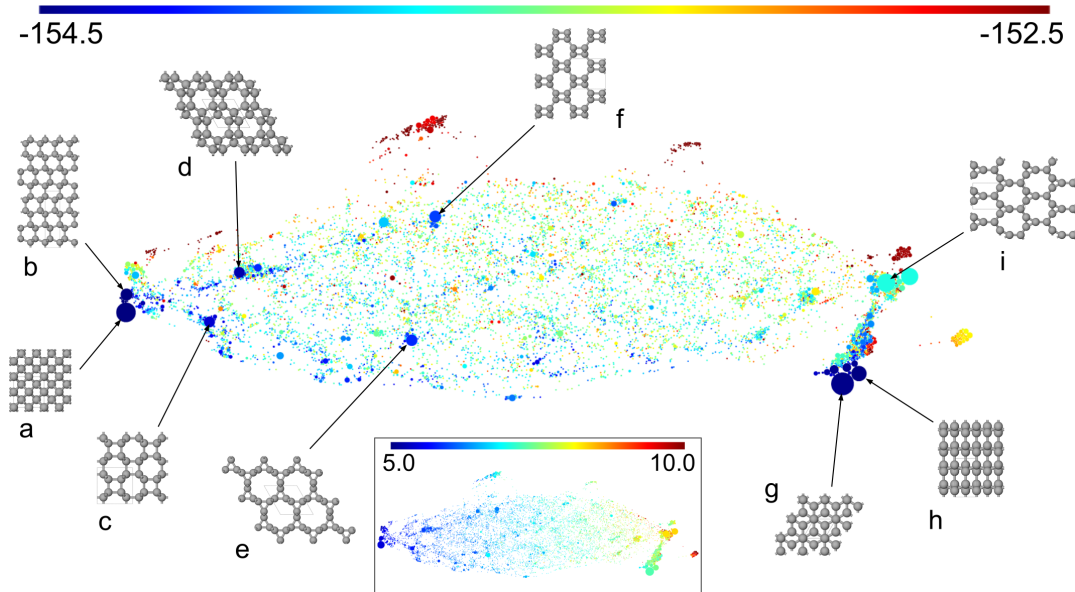


FIG. 11. 2D SHEAP map for 11370 structures of solid-state carbon, obtained with AIRSS at 10 GPa for a range of cell sizes—dataset taken from the Materials Cloud Archive [77]. Structures in the source-data represented as concatenated global SOAP vectors using universal “smart” parameters: $r_{\text{cut}}, n_{\text{max}}, l_{\text{max}}, \sigma = \{2.4 \text{ \AA}, 8, 4, 0.3 \text{ \AA}\}, \{2.0 \text{ \AA}, 8, 4, 0.25 \text{ \AA}\}$. The projection uses a perplexity of 30, with $R_0 = 0.0147$, as computed by Eq. (B1). The circles representing each structure are colored according to their energy per atom in eV; see the color bar at the top of the figure. Basin volumes are represented by the area of each circle. (a)–(n) label selected structures, including (a) diamond, (b) lonsdaleite, (d) the chiral framework structure of Ref. [79], (g) graphite, and (h) graphene layers with nongraphite packing. The inset shows the same map, colored according to volume per atom in \AA^3 ; see the color bar at the top of the inset.

landscapes lend themselves very well to structure prediction algorithms such as simulated annealing [80] and basin hopping [6,7], which utilize the presence of funnels to locate progressively lower-energy minima, while the study of carbon requires more heuristic approaches, such as in Ref. [81].

As illustrated in the inset in Fig. 11, in contrast to the energy variation, the volume per atom of each carbon structure follows a clear global trend across the map. These results demonstrate the additional complexity in the energy landscape resulting from three-body interactions, leading to less correlation between the energy and density than is seen for a simple pairwise interaction.

F. C + H + N + O

As a final demonstration of SHEAP, we present results obtained for a dataset of structures containing multiple species. The set considered consists of 50000 solid-state structures containing carbon, hydrogen, nitrogen, and oxygen (the four most common elements in organic chemistry and biology) at varying stoichiometries, which reduces to 21239 distinct structures following a similarity check. These configurations, taken from a larger set on the Materials Cloud Archive [77], were generated with AIRSS performed at 1 GPa. Structures in the source dataset are represented as concatenated global SOAP vectors, using the smart universal parameters, including the crossover

terms between different species. Presented in Fig. 12 are 2D and 3D SHEAP maps for this system.

The 2D map contains four distinct corners, each corresponding to arrangements of one of the four elements, as labeled. One can easily identify clusters of basins corresponding to low-energy configurations of binary molecules, most notably H_2O , CO_2 , and NH_3 . Intuitively, each of these molecular structures resides closest to the corner corresponding to the element it contains with the highest stoichiometry. The edges of the map represent structures containing predominantly the elements of the adjacent corners, also in agreement with intuition. The same is true for the structures occupying the regions between opposite corners (H, N and C, O). The lowest-energy structures containing each element in equal stoichiometry correspond to polymers of isocyanic acid ($\text{O} = \text{C} = \text{NH}$), which reside close to the center of the map, as labeled. A few other standout clusters of basins are labeled, according to the molecule(s) whose polymorphs are mapped there, to provide a representative picture of the distribution of structures and stoichiometries.

The 3D map also contains a distinct corner for each of the four elements. However, the extra dimension facilitates the emergence of an edge between every pair and provides clearer separation of different stoichiometries, particularly for structures residing toward the center of the 2D map. The result resembles an irregular tetrahedron, with slight

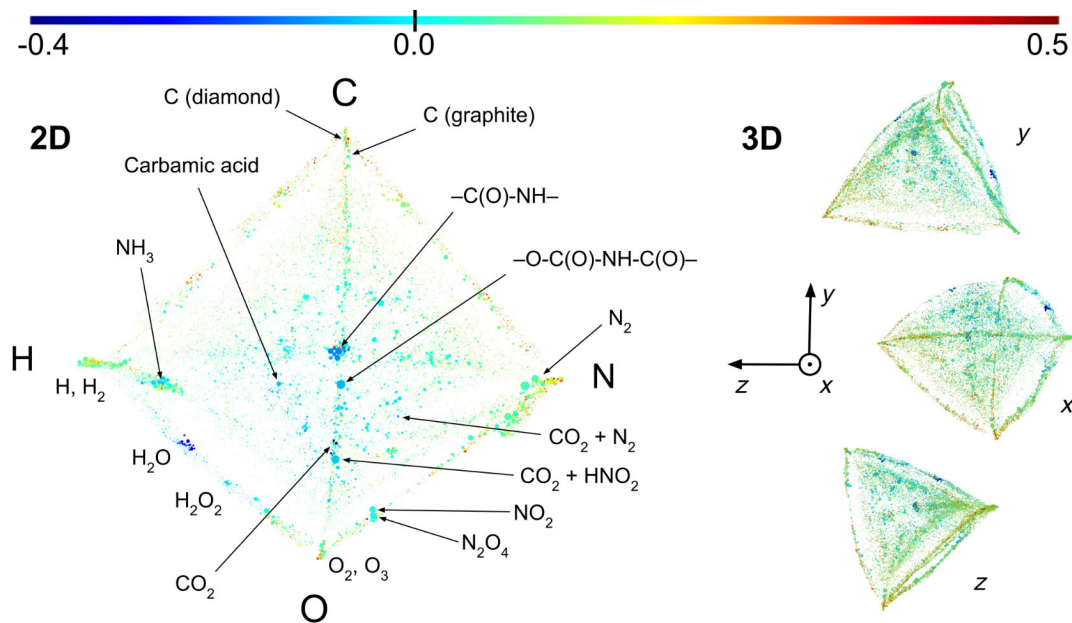


FIG. 12. 2D and 3D SHEAP maps for 21239 structures containing carbon, hydrogen, nitrogen, and oxygen at varying stoichiometries, obtained with AIRSS performed at 1 GPa, taken from the Materials Cloud Archive [77]. Structures in the source data are represented as concatenated global SOAP vectors using smart universal parameters: $r_{\text{cut}}, n_{\text{max}}, l_{\text{max}}, \sigma = \{2.7 \text{ \AA}, 8, 4, 0.33 \text{ \AA}\}, \{2.0 \text{ \AA}, 8, 4, 0.25 \text{ \AA}\}$. Projections use a perplexity of 30, with R_0 computed by Eq. (B1). Circles and spheres representing each structure are colored according to their formation energy per atom in eV; see the color bar at the top of the figure. Basin volumes are represented by the area (volume) of each circle (sphere). Each viewpoint of the 3D map looks down one of the axes displayed, as labeled; the axes have no other physical meaning. Regions containing polymorphs of various molecular species are labeled in the 2D map.

outward curvature of the faces and edges. This layout is reminiscent of a quaternary convex hull plot [82]; it is striking that the similarities to this construction emerge naturally from our projection of the energy landscape.

Both maps depict a higher density of structures at the edges (and between opposite corners in 2D), where the binary (and unary) configurations reside. This result demonstrates that SHEAP is able to highlight nonuniformity or biasing in the sampling of a configuration space.

G. Cost versus dimensionality

In Secs. IV C, IV D, and IV F, we demonstrate that the visualization quality produced by SHEAP can be improved by projecting into three dimensions rather than two. How do we know that there are no important features in a given minima distribution that would be revealed if we were able to view a projection into higher dimensions? We address this question by considering the optimal cost (in the absence of the hard-sphere interaction) achieved by projections into a range of map dimensionalities.

This analysis is conducted for datasets of 10 000 local minima of each of the systems studied above as well as datasets consisting of 10 000 points drawn uniformly at random inside the unit hypercubes of 2–6 dimensions. For the structural data, ten independent SHEAP projections are carried out, without switching on the hard-sphere interaction, for 1–6 map dimensions s . For the points inside hypercubes, SHEAP projections are conducted for up to eight dimensions. The lowest cost obtained at each map dimension, for each dataset, is plotted against s in Fig. 13.

Note that the cost is restricted from approaching zero as $s \rightarrow D$ (and beyond) by the use of different functional forms for the weights in the high and low dimensions; there is an inherent cost associated with the use of different distributions, even when $s \geq D$.

For all datasets of structural minima, we observe a dramatic reduction in the (lowest) relative cost in going from 1D to 2D. We also see a non-negligible reduction from increasing the map dimensionality from two to three (except in the case of carbon), which is most pronounced for C + H + N + O. However, it appears that in each case there is nothing of significance to be gained in going beyond 3D.

We compare to the plots for uniformly drawn random points within hypercubes of varying dimensionality. According to SHEAP, in order to faithfully represent the distribution of points in the hypercube of dimensionality D (which must contain long-range D -dimensional features), at least $s = D$ map dimensions are required. For each cube, using fewer than D dimensions in the map leads to a penalty in the optimal cost that is obtained, which increases with decreasing map dimensionality.

These results suggest that the distribution of minima on the landscapes we consider are inherently low dimensional (according to the given combination of structure descriptor

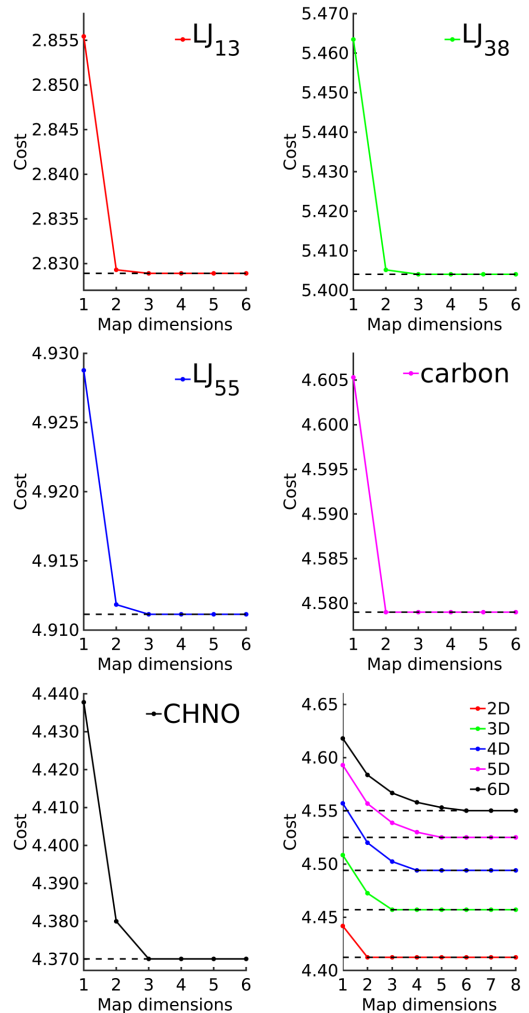


FIG. 13. Lowest cost obtained through SHEAP projections (without hard-sphere interactions) of datasets of 10 000 local minima for five systems, plotted against the number of map dimensions. Also shown are (lowest) cost versus map dimension plots for 10 000 points drawn uniformly at random inside the unit hypercubes of 2–6 dimensions, on a single set of axes. Each SHEAP projection uses a perplexity of 30.

and distance metric), providing validation for our approach to producing visual representations of high-dimensional PESs.

V. DISCUSSION

The striking observation of local minima being approximately distributed across low-dimensional manifolds within configuration space could have profound consequences for structure prediction. The ability to sample directly from this manifold would be a highly efficient means of searching for stable structures. Unfortunately, SHEAP (as with all algorithms in the same family) does not yield a transformation matrix between the original and reduced spaces, and the latter is not defined with respect to a meaningful set of axes. Thus, there is no clear way to generate new structures on the manifold learned by

SHEAP. However, our findings are promising for the use of other machine learning methods such as variational autoencoders [83] and generative adversarial networks [84] for structure prediction (see Ref. [85]) and suggest that learning from previously located minima would be an effective strategy, e.g., by pairing with a non-learning-based approach such as AIRSS. SHEAP could be deployed to illuminate differences in the sampling of a given surface by different approaches or by different parameter choices within a given approach.

Other pioneering work in developing a greater understanding of energy landscapes through visualization is Ref. [3], which also analyzes similarities between local minima to reveal topological features. Oganov and Valle employ a fingerprint structure descriptor based on pairwise separations of atoms [86] and compute the similarities between fingerprints with a cosine distance metric. They use this similarity measure to construct plots of energy versus distance to the global minimum for local minima on a given landscape, as well as mappings of the minima distributions which aim to produce “distances between points in the graph that are maximally close to the distances between the corresponding fingerprints” (see multidimensional scaling [87]). Their analysis reveals funnels in the landscapes of crystalline systems of few atoms per unit cell. While Ref. [3] does not focus on the quality of the mapping between their description of a given PES and the low-dimensional space it is visualized in and deals with only relatively small sample sizes (a few thousand structures), the tools developed there have been demonstrated to facilitate improved structure searching [88]. Dimensionality reduction by SHEAP furthers the potential of such an approach.

In this work, we represent real crystalline systems using SOAP. However, it is important to note that the development of good structure descriptors is an active field, and there are numerous justifiable choices we could have made for the descriptor-metric pairing. Of course, the entire form of a map is dependent on the choice of descriptor and distance metric, and a map’s features cannot be interpreted without knowledge of the descriptor-metric pair used. With SOAP, a global representation of a given structure is constructed by averaging the local fingerprints of each atomic environment. As with any instance of averaging, this step introduces a loss of structural information, which becomes severe for large unit cells. A potential solution is to build on the work of Ref. [89]. Rather than comparing averaged SOAP descriptors via a Euclidean distance metric, one could use the Hungarian algorithm to best match the atomic environments in a pair of structures, summing over the individual distances between matching environments. Any such, more sophisticated approach would need to be assessed with regard to the balance between increased fidelity of the map versus the computational cost of the algorithm.

VI. CONCLUSION

In conclusion, SHEAP enables us to produce meaningful and interpretable low-dimensional representations of a range of PESs, requiring only the structural information contained in minima data obtained from random structure searches as input, providing fresh insight into their topologies. The maps presented contain distributions and clusters of structures that we have been able to rationalize and reveal long-range topological features, including funnels (or the lack thereof).

By evaluating the costs at higher map dimensionalities, we have assessed the validity of SHEAP projections into the low number of dimensions that we can comprehend. In doing so, we have revealed an intrinsic low dimensionality to the distributions of minima, as represented by the descriptors and metrics considered, across both model PESs, and those corresponding to realistic, solid-state systems. We observe that, with all systems studied, no more than three map dimensions are required to represent all of the structural variation captured in our descriptions of the local minima, with two map dimensions also enabling useful projections.

We have demonstrated that SHEAP has the potential to be a useful tool in the field of structure prediction (and beyond), providing a convenient and concise way to visually represent and compare the large quantities of data that are generated in structure searches. The insight gained from these projections could provide inspiration for approaches to improving the sampling of a given landscape for a desired purpose.

ACKNOWLEDGMENTS

We thank Daan Frenkel, Joseph Nelson, and Bonan Zhu for useful discussions. B. W. B. S. acknowledges the EPSRC CDT in Computational Methods for Materials Science for funding under Grant No. EP/L015552/1. C. J. P. has been supported by the Royal Society through a Royal Society Wolfson Research Merit award. This work was performed using resources provided by the Cambridge Service for Data Driven Discovery (CSD3), operated by the University of Cambridge Research Computing Service.

APPENDIX A: UMAP VERSUS t -SNE

For algorithms that follow the framework outlined in Sec. II A, the optimization procedure is usually dictated by a cost-function gradient that can be separated into a sum of attractive and repulsive pairwise interactions between map points:

$$\frac{dC}{dY} = \sum_i \sum_{j>i} F_{\text{att}}(\mathbf{y}_i, \mathbf{y}_j) - F_{\text{rep}}(\mathbf{y}_i, \mathbf{y}_j). \quad (\text{A1})$$

The t -SNE algorithm exploits this separation with its use of early exaggeration to accelerate the optimization [28]. This technique works by the temporary enhancement of the

attractive contribution to the gradient, achieved through multiplying these terms by some positive scalar $\alpha > 1$ for the first few hundred iterations of optimization, promoting the formation of clusters in the projected data.

In t -SNE, the cost function measuring the faithfulness with which the probabilities q_{ij} model their corresponding p_{ij} is the KL divergence of the former distribution from the latter [28]:

$$C_{t\text{-SNE}} = \sum_{i,j} p_{ij} \log\left(\frac{p_{ij}}{q_{ij}}\right). \quad (\text{A2})$$

The attractive contribution to the gradient in t -SNE is a direct result of the form of the cost function and is what drives the algorithm to accurately reproduce the local connectivity of the source data in the projected map. The repulsive contribution is a result of the constraint that the low-dimensional weights q_{ij} sum to 1. This constraint can be incorporated either through the treatment of the optimization with an additional Lagrangian multiplier term or directly in the definition of the weights. Without this constraint, the KL divergence would be minimized by all points in the map collapsing onto one another.

UMAP chooses to minimize the fuzzy set cross-entropy between the two distributions [29]:

$$C_{\text{UMAP}} = \sum_{i,j} p_{ij} \log\left(\frac{p_{ij}}{q_{ij}}\right) + (1 - p_{ij}) \log\left(\frac{1 - p_{ij}}{1 - q_{ij}}\right). \quad (\text{A3})$$

The first term is identical to the KL divergence used by t -SNE and produces the attractive contribution to the gradient. In UMAP, this term acts only between map points that are considered to be nearest neighbors in the source data, meaning $p_{ij} = 0$ for all pairs of points i and j for which neither is considered one of the k th nearest neighbors (k is a user-specified hyperparameter) of the other, according to the defined distance metric. The second, additional term introduces a repulsion between every pair of map points that is greater the further apart the corresponding points are in the source data. This choice of cost function is the key reason that UMAP is able to provide improved performance over t -SNE.

First, in t -SNE, the repulsive term in the gradient appears only because of the constraint that the low-dimensional weights q_{ij} sum to one, which requires the renormalization of all of these weights at each iteration of the optimization procedure. In UMAP, the cost function introduces a repulsive contribution to the gradient without the need for this constraint, meaning the costly normalization step can be avoided.

Second, the repulsive term contribution to the gradient in UMAP is more dependent on the actual data than that in t -SNE. Thus, UMAP tends to do a better job than t -SNE in

producing a faithful representation of the global structure of the source data.

Finally, the removal of the need to normalize the weights q_{ij} means that the cost function is amenable to optimization via SGD, which replaces the true gradient with an estimate that requires only a stochastically chosen subset of the components to be evaluated. UMAP implements an efficient approximate SGD algorithm which makes use of probabilistic edge sampling and negative sampling [29]. The improvements in speed and scaling achieved by UMAP over t -SNE are largely a result of the use of SGD rather than some other full gradient descent method.

However, a fairly subtle caveat arises with UMAP's implementation of SGD. The repulsive part of the interaction between map points is treated with negative sampling [90]. While this approach vastly reduces the number of repulsive interactions that need to be evaluated, it also has the effect of deexaggerating the repulsion between the map points, relative to their interaction as described by the true cost function. We find that this effect is crucial to UMAP producing a meaningful projection. However, this effect also means that UMAP's optimization scheme does not actually minimise the defined cost function, and the algorithm "converges" only because the step size is forced to go to zero. Thus, while UMAP is very efficient and produces low-dimensional projections that appear to improve on those by the t -SNE algorithm, it is rather unsatisfactory in that the cost function defined within the algorithm is not truly minimized in the construction of the final map.

APPENDIX B: CHOOSING R_0

The choice of the minimum circle or sphere radius R_0 is critical. If, in the absence of the hard spheres, the separation between a given pair of map points is smaller than the sum of their hard-sphere radii, then the introduction of the hard-sphere interaction forces these structures apart. Thus, if R_0 is large enough for this separation to be the case for a single pair of structures, the introduction of hard spheres results in a distortion of the layout of the map. On the other hand, if R_0 is too small relative to the spread of data in the map, some (or all) of the spheres may be barely visible. Hence, a good choice for R_0 compromises between ensuring that all of the structures in the map are clearly visible, while minimizing the distortion of the layout of structures.

By default, $R_0 = 0.01$. Also implemented is the option to compute a suitable R_0 in the same iteration at which the hard-sphere interaction is switched on, as a fraction of the average separation of the map points from their center of mass:

$$R_0 = \frac{2^s}{10} \left(\frac{N}{\sum_{i=1}^N c_i} \right)^{1/s} \left\langle \|\mathbf{y}_i - \mathbf{y}_{\text{c.m.}}\|_2 \right\rangle_{\text{HS}}, \quad (\text{B1})$$

where

$$\left\langle \|y_i - y_{c.m.}\|_2 \right\rangle_{HS} = \frac{1}{N} \sum_{i=1}^N \|y_i - y_{c.m.}\|_2, \quad (\text{B2})$$

evaluated at the onset of the hard-sphere interaction. The premultiplying factor is selected by empirical investigation and may be altered in the future. Alternatively, the value of R_0 can be input by the user, allowing interactive feedback.

APPENDIX C: RANDOM PROJECTION

Random projection relies on the Johnson-Lindenstrauss lemma [91], which states that any set of N points in a high-dimensional Euclidean space can be embedded into a $k \geq \mathcal{O} \log(N/\epsilon^2)$ -dimensional Euclidean subspace without distorting the distances between any pair of points by more than a factor of $1 \pm \epsilon$, for any $0 < \epsilon < 1$. It can be proven that, if the k -dimensional subspace is constructed by the generation of k random but mutually orthogonal and normalized vectors in the high-dimensional space, then this is achieved with high probability. Furthermore, it is found that, in practice, even if embedding into a low-dimensional space of dimension $s \ll \mathcal{O} \log(N/\epsilon^2)$, one can often still obtain something useful from a random projection. In our implementation of random projection, a source data point \mathbf{x} existing in the original D -dimensional space is projected onto an s -dimensional subspace by transformation via a $s \times D$ matrix R whose s rows are orthonormal vectors in the original space, to produce the projected data point:

$$\mathbf{y}^{RP} = R\mathbf{x}. \quad (\text{C1})$$

In addition, the descriptor vector of sorted pairwise distances contains only positive entries, meaning all of the data points exist in the positive orthant of the D -dimensional hyperspace—a $1/2^d$ fraction of the total space. Hence, for this descriptor, we choose the row vectors comprising R to contain only positive entries. The same is true for the sorted list of inverse distances descriptor—see Appendix E.

APPENDIX D: QUALITATIVE COMPARISON OF SHEAP TO t -SNE AND UMAP

The visualization quality of SHEAP is compared against our implementations of t -SNE and UMAP on two standard image datasets: the Columbia University Image Library (COIL-20) [92] and the Modified National Institute of Standards and Technology database (MNIST) [93]. In general, the improvements offered by UMAP over t -SNE are maintained with SHEAP.

The COIL-20 [92] dataset contains 1440 images, produced by individually placing 20 objects on a motorized turntable (against a black background) and photographing them at 5° intervals, up to a complete rotation. Each image is grayscale, containing 128×128 pixels, and is treated as a 16384-dimensional vector. This dataset is expected to contain well-defined clusters, corresponding to the sets of images of each object. Furthermore, the submanifold on which each cluster lies should have a relatively intuitive form—objects containing no symmetry about the axis of rotation are expected correspond to a closed loop of points, with consecutive images of a given object neighboring one another. In Fig. 14, both UMAP and SHEAP do a better

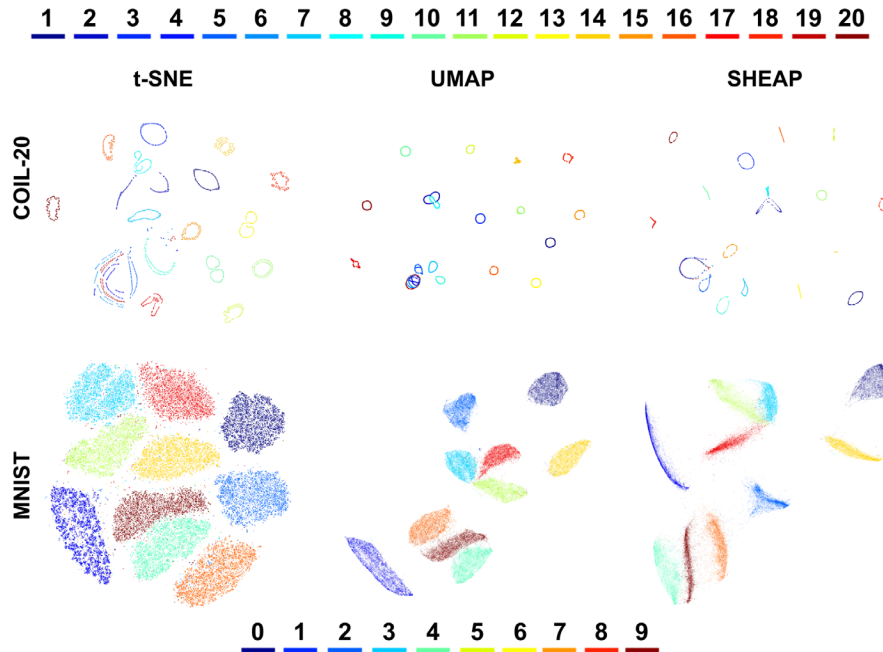


FIG. 14. Projections of the COIL-20 [92] and MNIST [93] datasets made by our implementations of the t -SNE, UMAP, and SHEAP algorithms. The legend at the top of the figure labels distinct objects in the COIL-20 dataset; the legend at the bottom labels distinct digits in the MNIST dataset.

job than t -SNE at forming the clusters corresponding to each object and keep more of the loops intact. The biggest difference between the UMAP and SHEAP projections is that SHEAP produces larger and more frequent distortions of these closed loops—we argue that these distortions make intuitive sense, occurring when there exist symmetries in the objects being imaged.

The MNIST [93] (training) dataset contains 60 000 grayscale 28×28 images of handwritten digits (0–9), each treated as a 784-dimensional vector. A given data point is expected to be neighbored by other writings of the same digit, and the dataset should contain well-defined clusters corresponding to the distinct digits. These are expected to be arranged such that similar digits reside close together and dissimilar ones are well separated (a handwritten 3 is often hard to distinguish from an 8 but never mistaken for a 4). In Fig. 14, all three algorithms do a good job of clustering the images corresponding to each digit. However, UMAP and SHEAP capture more of the global relationships between the clusters, providing much clearer separations between them and grouping together clusters of similar digits (3,5,8 and 4,7,9) in similar layouts. A noticeable difference between UMAP and SHEAP's projections is that the latter seems to contain thinner, denser bands of points in clusters that are particularly elongated. Nonetheless, the two plots are qualitatively very similar.

APPENDIX E: LJ_{13} DESCRIBED WITH INVERSE DISTANCES

Figure 15 shows a 2D SHEAP map for the same set of LJ_{13} minima as considered in Sec. IV B. The difference here is that, rather than describing structures with sorted lists of all distances between pairs of atoms, we use sorted lists of all inverse distances. We observe similar clustering of structures as in Fig. 2; groups of structures in Fig. 15 are labeled according to the corresponding cluster in Fig. 2. However, the map obtained from inverse distances does not result in a consistent global layout of the minima, failing to reproduce the expected single funnel distribution across all minima. We speculatively attribute this difference to the emphasis the descriptor places on the shortest distances in the structures, resulting in a poorer representation of global information.

Also in Fig. 15 is a plot of the lowest cost (in the absence of hard-sphere interactions) achieved with projections into a range of map dimensionalities, similar to the plots in Fig. 13. Again, this analysis is conducted for a reduced dataset of 10 000 minima, with ten independent SHEAP projections carried out for each number of map dimensions. As with the other structural datasets considered, no more than three map dimensions are required for convergence. However, here, three map dimensions do not produce a lower cost than two, whereas a slight reduction is observed with the sorted list of distances descriptor for the same

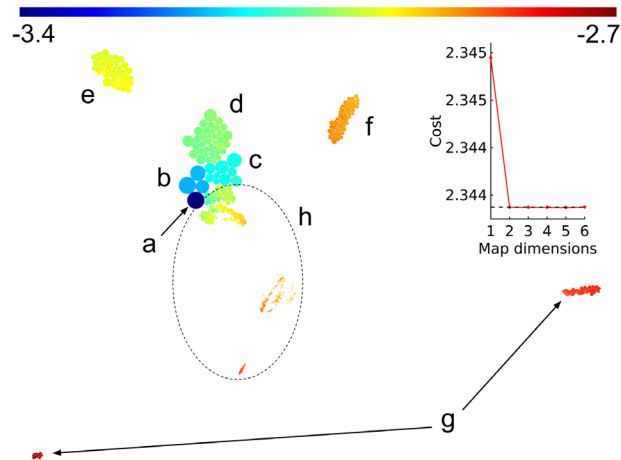


FIG. 15. Labeled 2D SHEAP map produced from local minima on the LJ_{13} PES (the same set as depicted in Fig. 2), described by sorted lists of inverse pairwise distances. The projection uses a perplexity of 30, with a minimum sphere radius $R_0 = 0.01$. The circles representing each structure are colored according to energy per atom in Lennard-Jones units; see the color bar at the top of the figure. Basin volumes are represented by the area of each circle. Groups of structures are labeled according to their clustering in Fig. 2; the dashed line encloses structures collectively labeled (h). Also shown is a plot of the lowest cost (without hard-sphere interactions) obtained by projections of a subset of this dataset into a range of map dimensionalities.

system. This result is likely due to the less well represented global features of the landscape.

- [1] M. Born and R. Oppenheimer, *Zur Quantentheorie der Molekeln*, *Ann. Phys. (Berlin)* **389**, 457 (1927).
- [2] D. J. Wales, *Energy Landscapes: Applications to Clusters, Biomolecules and Glasses* (Cambridge University Press, Cambridge, England, 2003).
- [3] A. R. Oganov and M. Valle, *How to Quantify Energy Landscapes of Solids*, *J. Chem. Phys.* **130**, 104504 (2009).
- [4] D. J. Wales, *Decoding the Energy Landscape: Extracting Structure, Dynamics and Thermodynamics*, *Phil. Trans. R. Soc. A* **370**, 2877 (2012).
- [5] A. R. Oganov, C. J. Pickard, Q. Zhu, and R. J. Needs, *Structure Prediction Drives Materials Discovery*, *Nat. Rev. Mater.* **4**, 331 (2019).
- [6] Z. Li and H. A. Scheraga, *Monte Carlo-Minimization Approach to the Multiple-Minima Problem in Protein Folding*, *Proc. Natl. Acad. Sci. U.S.A.* **84**, 6611 (1987).
- [7] D. J. Wales and J. P. Doye, *Global Optimization by Basin-Hopping and the Lowest Energy Structures of Lennard-Jones Clusters Containing up to 110 Atoms*, *J. Phys. Chem. A* **101**, 5111 (1997).
- [8] M. Amsler and S. Goedecker, *Crystal Structure Prediction Using the Minima Hopping Method*, *J. Chem. Phys.* **133**, 224104 (2010).
- [9] Y. Wang, J. Lv, L. Zhu, and Y. Ma, *Crystal Structure Prediction via Particle-Swarm Optimization*, *Phys. Rev. B* **82**, 094116 (2010).

- [10] Y. Wang, J. Lv, L. Zhu, and Y. Ma, *Calypso: A Method for Crystal Structure Prediction*, *Comput. Phys. Commun.* **183**, 2063 (2012).
- [11] A. R. Oganov and C. W. Glass, *Crystal Structure Prediction Using Ab Initio Evolutionary Techniques: Principles and Applications*, *J. Chem. Phys.* **124**, 244704 (2006).
- [12] C. W. Glass, A. R. Oganov, and N. Hansen, *Uspexevolutionary Crystal Structure Prediction*, *Comput. Phys. Commun.* **175**, 713 (2006).
- [13] C. J. Pickard and R. J. Needs, *High-Pressure Phases of Silane*, *Phys. Rev. Lett.* **97**, 045504 (2006).
- [14] C. J. Pickard and R. J. Needs, *Ab Initio Random Structure Searching*, *J. Phys. Condens. Matter* **23**, 053201 (2011).
- [15] C. J. Pickard and R. J. Needs, *Structure of Phase III of Solid Hydrogen*, *Nat. Phys.* **3**, 473 (2007).
- [16] C. J. Pickard and R. J. Needs, *Highly Compressed Ammonia Forms an Ionic Crystal*, *Nat. Mater.* **7**, 775 (2008).
- [17] C. J. Pickard and R. J. Needs, *Aluminium at Terapascal Pressures*, *Nat. Mater.* **9**, 624 (2010).
- [18] D. G. Kleinbaum, L. L. Kupper, K. E. Muller, and A. Nizam, *Applied Regression Analysis and Other Multivariable Methods* (Duxbury Press, Belmont, CA, 1988), Vol. 601.
- [19] A. K. Jain, M. N. Murty, and P. J. Flynn, *Data Clustering: A Review*, *ACM Comput. Surv. (CSUR)* **31**, 264 (1999).
- [20] C. C. Aggarwal, *Data Classification: Algorithms and Applications* (CRC Press, Boca Raton, FL, 2014).
- [21] M. A. Nielsen, *Neural Networks and Deep Learning* (Determination Press, San Francisco, 2015), Vol. 25.
- [22] S. R. Chitturi, P. C. Verpoort, A. A. Lee, and D. J. Wales, *Perspective: New Insights from Loss Function Landscapes of Neural Networks*, *Mach. Learn. Sci. Technol.* **1**, 023002 (2020).
- [23] F. H. Stillinger and T. A. Weber, *Packing Structures and Transitions in Liquids and Solids*, *Science* **225**, 983 (1984).
- [24] J. P. K. Doye and C. P. Massen, *Characterizing the Network Topology of the Energy Landscapes of Atomic Clusters*, *J. Chem. Phys.* **122**, 084105 (2005).
- [25] C. P. Massen and J. P. K. Doye, *Power-Law Distributions for the Areas of the Basins of Attraction on a Potential Energy Landscape*, *Phys. Rev. E* **75**, 037101 (2007).
- [26] J. D. Bryngelson, J. N. Onuchic, N. D. Socci, and P. G. Wolynes, *Funnels, Pathways, and the Energy Landscape of Protein Folding: A Synthesis*, *Proteins: Struct., Funct., Bioinf.* **21**, 167 (1995).
- [27] O. M. Becker and M. Karplus, *The Topology of Multidimensional Potential Energy Surfaces: Theory and Application to Peptide Structure and Kinetics*, *J. Chem. Phys.* **106**, 1495 (1997).
- [28] L. van der Maaten and G. Hinton, *Visualizing Data Using t-SNE*, *J. Mach. Learn. Res.* **9**, 2579 (2008), <http://www.jmlr.org/papers/v9/vandermaaten08a.html>.
- [29] L. McInnes, J. Healy, and J. Melville, *Umap: Uniform Manifold Approximation and Projection for Dimension Reduction*, *arXiv:1802.03426*.
- [30] L. McInnes, J. Healy, N. Saul, and Großberger, *Umap: Uniform Manifold Approximation and Projection*, *J. Open Source Softw.* **3**, 861 (2018).
- [31] T. Fösel, P. Tighineanu, T. Weiss, and F. Marquardt, *Reinforcement Learning with Neural Networks for Quantum Feedback*, *Phys. Rev. X* **8**, 031084 (2018).
- [32] X. Li, O. E. Dyck, M. P. Oxley, A. R. Lupini, L. McInnes, J. Healy, S. Jesse, and S. V. Kalinin, *Manifold Learning of Four-Dimensional Scanning Transmission Electron Microscopy*, *npj Comput. Mater.* **5**, 5 (2019).
- [33] V. Spiwok and P. Kříž, *Time-Lagged t-Distributed Stochastic Neighbor Embedding (t-SNE) of Molecular Simulation Trajectories*, *Front. Mol. Biosci.* **7**, 132 (2020).
- [34] B. Olsthoorn, R. M. Geilhufe, S. S. Borysov, and A. V. Balatsky, *Band Gap Prediction for Large Organic Crystal Structures with Machine Learning*, *Adv. Quantum Technol.* **2**, 1900023 (2019).
- [35] C. Royse, S. Wolter, and J. A. Greenberg, *Emergence and Distinction of Classes in XRD Data via Machine Learning, in Anomaly Detection and Imaging with X-Rays (ADIX) IV* (International Society for Optics and Photonics, Bellingham, WA, 2019), Vol. 10999, p. 109990D.
- [36] M. Ghandi, F. W. Huang, J. Jané-Valbuena, G. V. Kryukov, C. C. Lo, E. R. McDonald, J. Barretina, E. T. Gelfand, C. M. Bielski, H. Li *et al.*, *Next-Generation Characterization of the Cancer Cell Line Encyclopedia*, *Nature (London)* **569**, 503 (2019).
- [37] C.-A. Dutertre, E. Becht, S. E. Irac, A. Khalilnezhad, V. Narang, S. Khalilnezhad, P. Y. Ng, L. L. van den Hoogen, J. Y. Leong, B. Lee *et al.*, *Single-Cell Analysis of Human Mononuclear Phagocytes Reveals Subset-Defining Markers and Identifies Circulating Inflammatory Dendritic Cells*, *Immunity* **51**, 573 (2019).
- [38] S. Hemmers, M. Schizas, E. Azizi, S. Dikiy, Y. Zhong, Y. Feng, G. Altan-Bonnet, and A. Y. Rudensky, *Il-2 Production by Self-Reactive cd4 Thymocytes Scales Regulatory t Cell Generation in the Thymus*, *J. Exp. Med.* **216**, 2466 (2019).
- [39] E. Becht, L. McInnes, J. Healy, C.-A. Dutertre, I. W. Kwok, L. G. Ng, F. Ginhoux, and E. W. Newell, *Dimensionality Reduction for Visualizing Single-Cell Data Using Umap*, *Nat. Biotechnol.* **37**, 38 (2019).
- [40] R. Czerminski and R. Elber, *Reaction Path Study of Conformational Transitions in Flexible Systems: Applications to Peptides*, *J. Chem. Phys.* **92**, 5580 (1990).
- [41] D. J. Wales, *The Energy Landscape as a Unifying Theme in Molecular Science*, *Phil. Trans. R. Soc. A* **363**, 357 (2005).
- [42] G. J. Rylance, R. L. Johnston, Y. Matsunaga, C.-B. Li, A. Baba, and T. Komatsuzaki, *Topographical Complexity of Multidimensional Energy Landscapes*, *Proc. Natl. Acad. Sci. U.S.A.* **103**, 18551 (2006).
- [43] L. C. Smeeton, M. T. Oakley, and R. L. Johnston, *Visualizing Energy Landscapes with Metric Disconnectivity Graphs*, *J. Comput. Chem.* **35**, 1481 (2014).
- [44] M. Ceriotti, G. A. Tribello, and M. Parrinello, *Simplifying the Representation of Complex Free-Energy Landscapes Using Sketch-Map*, *Proc. Natl. Acad. Sci. U.S.A.* **108**, 13023 (2011).
- [45] G. A. Tribello, M. Ceriotti, and M. Parrinello, *Using Sketch-Map Coordinates to Analyze and Bias Molecular Dynamics Simulations*, *Proc. Natl. Acad. Sci. U.S.A.* **109**, 5196 (2012).
- [46] M. Ceriotti, G. A. Tribello, and M. Parrinello, *Demonstrating the Transferability and the Descriptive Power of Sketch-Map*, *J. Chem. Theory Comput.* **9**, 1521 (2013).

- [47] E. A. Engel, A. Anelli, M. Ceriotti, C. J. Pickard, and R. J. Needs, *Mapping Uncharted Territory in Ice from Zeolite Networks to Ice Structures*, *Nat. Commun.* **9**, 2173 (2018).
- [48] L. van der Maaten, E. Postma, and J. van den Herik, *Dimensionality Reduction: A Comparative Review*, *J. Mach. Learn. Res.* **10**, 66 (2009).
- [49] J. A. Lee and M. Verleysen, *Nonlinear Dimensionality Reduction* (Springer, New York, 2007).
- [50] T. Lin and H. Zha, *Riemannian Manifold Learning*, *IEEE Trans. Pattern Anal. Mach. Intell.* **30**, 796 (2008).
- [51] H. Hotelling, *Analysis of a Complex of Statistical Variables into Principal Components*, *J. Educat. Psychol.* **24**, 417 (1933).
- [52] J. B. Tenenbaum, V. De Silva, and J. C. Langford, *A Global Geometric Framework for Nonlinear Dimensionality Reduction*, *Science* **290**, 2319 (2000).
- [53] B. Schölkopf, A. Smola, and K.-R. Müller, *Nonlinear Component Analysis as a Kernel Eigenvalue Problem*, *Neural Comput.* **10**, 1299 (1998).
- [54] J. W. Sammon, *A Nonlinear Mapping for Data Structure Analysis*, *IEEE Trans. Comput.* **C-18**, 401 (1969).
- [55] S. T. Roweis and L. K. Saul, *Nonlinear Dimensionality Reduction by Locally Linear Embedding*, *Science* **290**, 2323 (2000).
- [56] M. Belkin and P. Niyogi, *Laplacian Eigenmaps and Spectral Techniques for Embedding and Clustering*, in *Proceedings of the 14th International Conference on Neural Information Processing Systems: Natural and Synthetic, Vancouver, BC, Canada, 2001* (MIT Press, Cambridge, MA, USA, 2002), pp. 585–591.
- [57] J. Tang, J. Liu, M. Zhang, and Q. Mei, *Visualizing Large-Scale and High-Dimensional Data*, in *Proceedings of the 25th International Conference on World Wide Web* (International World Wide Web Conferences Steering Committee, Geneva, 2016), pp. 287–297.
- [58] L. van der Maaten, *Barnes-hut-SNE*, [arXiv:1301.3342](https://arxiv.org/abs/1301.3342).
- [59] G. C. Linderman, M. Rachh, J. G. Hoskins, S. Steinerberger, and Y. Kluger, *Efficient Algorithms for t -Distributed Stochastic Neighborhood Embedding*, [arXiv:1712.09005](https://arxiv.org/abs/1712.09005).
- [60] J. Barzilay and J. M. Borwein, *Two-Point Step Size Gradient Methods*, *IMA J. Numer. Anal.* **8**, 141 (1988).
- [61] G. C. Linderman and S. Steinerberger, *Clustering with t -SNE, Provably*, [arXiv:1706.02582](https://arxiv.org/abs/1706.02582).
- [62] A. P. Bartók, R. Kondor, and G. Csányi, *On Representing Chemical Environments*, *Phys. Rev. B* **87**, 184115 (2013).
- [63] S. De, A. P. Bartók, G. Csányi, and M. Ceriotti, *Comparing Molecules and Solids across Structural and Alchemical Space*, *Phys. Chem. Chem. Phys.* **18**, 13754 (2016).
- [64] B. Cheng, R.-R. Griffiths, S. Wengert, C. Kunkel, T. Stenczel, B. Zhu, V. L. Deringer, N. Bernstein, J. T. Margraf, K. Reuter *et al.*, *Mapping Materials and Molecules*, *Acc. Chem. Res.* **53**, 12697 (2020).
- [65] C. J. Pickard, *Hyperspatial Optimization of Structures*, *Phys. Rev. B* **99**, 054102 (2019).
- [66] <https://bitbucket.org/bshires/sheap/src/master/>.
- [67] A. Stukowski, *Visualization and Analysis of Atomistic Simulation Data with Ovito—The Open Visualization Tool*, *Model. Simul. Mater. Sci. Eng.* **18**, 015012 (2009).
- [68] <https://www.mtg.msm.cam.ac.uk/Codes/AIRSS>
- [69] J. P. K. Doye, M. A. Miller, and D. J. Wales, *Evolution of the Potential Energy Surface with Size for Lennard-Jones Clusters*, *J. Chem. Phys.* **111**, 8417 (1999).
- [70] D. J. Wales, *Perspective: Insight into Reaction Coordinates and Dynamics from the Potential Energy Landscape*, *J. Chem. Phys.* **142**, 130901 (2015).
- [71] R. H. Leary, *Global Optima of Lennard-Jones Clusters*, *J. Global Optim.* **11**, 35 (1997).
- [72] J. P. Neirotti, F. Calvo, D. L. Freeman, and J. D. Doll, *Phase Changes in 38-Atom Lennard-Jones Clusters. I. A Parallel Tempering Study in the Canonical Ensemble*, *J. Chem. Phys.* **112**, 10340 (2000).
- [73] B. Schaefer and S. Goedecker, *Computationally Efficient Characterization of Potential Energy Surfaces Based on Fingerprint Distances*, *J. Chem. Phys.* **145**, 034101 (2016).
- [74] M. O. Jäger, E. V. Morooka, F. F. Canova, L. Himanen, and A. S. Foster, *Machine Learning Hydrogen Adsorption on Nanoclusters through Structural Descriptors*, *npj Comput. Mater.* **4**, 37 (2018).
- [75] L. Himanen, M. O. J. Jäger, E. V. Morooka, F. F. Canova, Y. S. Ranawat, D. Z. Gao, P. Rinke, and A. S. Foster, *Describe: Library of Descriptors for Machine Learning in Materials Science*, *Comput. Phys. Commun.* **247**, 106949 (2020).
- [76] M. Sicher, S. Mohr, and S. Goedecker, *Efficient Moves for Global Geometry Optimization Methods and Their Application to Binary Systems*, *J. Chem. Phys.* **134**, 044106 (2011).
- [77] C. J. Pickard, *AIRSS Data for Carbon at 10 GPa and the C + N + H + O System at 1 GPa*, *Materials Cloud Archive*, <https://doi.org/10.24435/materialscloud:2020.0026/v1>.
- [78] V. L. Deringer and G. Csányi, *Machine Learning Based Interatomic Potential for Amorphous Carbon*, *Phys. Rev. B* **95**, 094203 (2017).
- [79] C. J. Pickard and R. J. Needs, *Hypothetical Low-Energy Chiral Framework Structure of Group 14 Elements*, *Phys. Rev. B* **81**, 014106 (2010).
- [80] S. Kirkpatrick, C. D. Gelatt, and M. P. Vecchi, *Optimization by Simulated Annealing*, *Science* **220**, 671 (1983).
- [81] X. Shi, C. He, C. J. Pickard, C. Tang, and J. Zhong, *Stochastic Generation of Complex Crystal Structures Combining Group and Graph Theory with Application to Carbon*, *Phys. Rev. B* **97**, 014104 (2018).
- [82] J. A. Flores-Livas, L. Boeri, A. Sanna, G. Profeta, R. Arita, and M. Eremets, *A Perspective on Conventional High-Temperature Superconductors at High Pressure: Methods and Materials*, *Phys. Rep.* **856**, 1 (2020).
- [83] D. P. Kingma and M. Welling, *Auto-Encoding Variational Bayes*, [arXiv:1312.6114](https://arxiv.org/abs/1312.6114).
- [84] I. J. Goodfellow, J. Pouget-Abadie, M. Mirza, B. Xu, D. Warde-Farley, S. Ozair, A. Courville, and Y. Bengio, *Generative Adversarial Nets*, in *Proceedings of the 27th International Conference on Neural Information Processing Systems—Volume 2, Montreal, Canada, 2014* (MIT Press, Cambridge, MA, USA, 2014), pp. 2672–2680.

- [85] S. Kim, J. Noh, G. H. Gu, A. Aspuru-Guzik, and Y. Jung, *Generative Adversarial Networks for Crystal Structure Prediction*, [arXiv:2004.01396](#).
- [86] M. Valle and A. R. Oganov, *Crystal Structures Classifier for an Evolutionary Algorithm Structure Predictor*, in *Proceedings of the 2008 IEEE Symposium on Visual Analytics Science and Technology* (IEEE, New York, 2008), pp. 11–18.
- [87] W. S. Torgerson, *Multidimensional Scaling: I. Theory and Method*, *Psychometrika* **17**, 401 (1952).
- [88] A. O. Lyakhov, A. R. Oganov, and M. Valle, *How to Predict Very Large and Complex Crystal Structures*, *Comput. Phys. Commun.* **181**, 1623 (2010).
- [89] L. Zhu, M. Amsler, T. Fuhrer, B. Schaefer, S. Faraji, S. Rostami, S. A. Ghasemi, A. Sadeghi, M. Grauzinyte, C. Wolverton, and S. Goedecker, *A Fingerprint Based Metric for Measuring Similarities of Crystalline Structures*, *J. Chem. Phys.* **144**, 034203 (2016).
- [90] T. Mikolov, I. Sutskever, K. Chen, G. S. Corrado, and J. Dean, *Distributed Representations of Words and Phrases and Their Compositionality*, in *Proceedings of the 26th International Conference on Neural Information Processing Systems—Volume 2, Lake Tahoe, Nevada, 2013* (Curran Associates Inc., Red Hook, NY, USA, 2013), pp. 3111–3119.
- [91] W. B. Johnson and J. Lindenstrauss, *Extensions of Lipschitz Mappings into a Hilbert Space*, *Contemp. Math.* **26**, 189 (1984).
- [92] S. A. Nene, S. K. Nayar, and H. Murase, *Columbia Object Image Library (COIL-20)*, Technical Report No. CUCS-005-96, Columbia University, 1996.
- [93] Y. LeCun, L. Bottou, Y. Bengio, and P. Haffner, *Gradient-Based Learning Applied to Document Recognition*, *Proc. IEEE* **86**, 2278 (1998).