



PCCP

Unsupervised learning of representative local atomic arrangements in molecular dynamics data

Journal:	<i>Physical Chemistry Chemical Physics</i>
Manuscript ID	CP-ART-02-2023-000525.R1
Article Type:	Paper
Date Submitted by the Author:	23-Mar-2023
Complete List of Authors:	Roncoroni, Fabrice; Lawrence Berkeley Laboratory, Molecular Foundry; ETH Zurich, Sanz-Matias, Ana; Lawrence Berkeley National Laboratory, Molecular Foundry Sundararaman, Siddharth; Lawrence Berkeley National Laboratory, Molecular Foundry Prendergast, David; Lawrence Berkeley National Laboratory, Molecular Foundry

SCHOLARONE™
Manuscripts

Cite this: DOI: 00.0000/xxxxxxxxxx

Unsupervised learning of representative local atomic arrangements in molecular dynamics data

Fabrice Roncoroni,^{*a} Ana Sanz-Matias,^a Siddharth Sundararaman,^a and David Prendergast^a

Received Date

Accepted Date

DOI: 00.0000/xxxxxxxxxx

Abstract

Molecular dynamics (MD) simulations present a data-mining challenge, given that they can generate a considerable amount of data but often rely on limited or biased human interpretation to examine their information content. By not asking the right questions of MD data we may miss critical information hidden within it. We combine dimensionality reduction (UMAP) and unsupervised hierarchical clustering (HDBSCAN) to quantitatively characterize prevalent coordination environments of chemical species within MD data. By focusing on local coordination, we significantly reduce the amount of data to be analyzed by extracting all distinct molecular formulas within a given coordination sphere. We then efficiently combine UMAP and HDBSCAN with alignment or shape-matching algorithms to partition these formulas into structural isomer families indicating their relative populations. The method was employed to reveal details of cation coordination in electrolytes based on molecular liquids.

Introduction

Molecular dynamics (MD) is a vital tool in gaining molecular-scale insight on the properties and functional behavior of complex systems and interfaces. There is immense value and inspiration in providing the community with visualizations of molecular configurations and their dynamics, particularly with respect to the identification of molecular-scale bottlenecks in functional processes and the rational design of new chemistries to avoid them^{1–6}.

Analysis of MD data sets is frequently driven by simplistic metrics (e.g., density profiles, pair distribution functions) and human intuition on what questions to ask of the data. Insight can be limited significantly by human imagination or past experience when it comes to designing bespoke analysis of such data sets. The underlying assumption, which we presume here also, is that the relevant information content in large data sets has a much smaller dimension than the entire data set, or that there is a relatively simple, low-dimensional underlying probability density that can describe behavior in the system. In an ideal scenario, automated data analysis should provide us with an unbiased path to dimensionality reduction with full disclosure of embedded details of the distribution/density in terms of distinct classification or grouping of data, while the task of labeling these groups (via interpreta-

tion or contextualization) is more suited to human intuition and experience.

To this end, we have developed an unsupervised data-mining approach to extract details of distinct motifs of local coordination from MD simulations. We divide the problem of assessing prevalent local coordination environments into three main steps. First, we identify the species of interest and the expected size of their local coordination environment. Then, we extract local atomic clusters from the MD data set and group them by chemical formula or composition. Finally for each chemical formula, we perform dimensionality reduction, hierarchical clustering and alignment procedures to determine the structural distribution of coordination isomers. We focus on defining a computationally efficient approach, which can quickly analyze large MD data sets to provide details of conformations within minutes or less to facilitate on-the-fly analysis. To this end, the combination of efficient dimensionality reduction (UMAP⁷) with hierarchical clustering (HDBSCAN⁸) is aided significantly by our use of fast alignment algorithms (FASTOVERLAP⁹) that focus on relative alignment with respect to automatically identified exemplars of isolated data clusters. Furthermore, we demonstrate that this approach can be used to "learn" how to classify previously unseen data and perform "on-the-fly" unsupervised analysis of MD data as it is generated. The outcome of this approach is a detailed molecular-scale understanding of the local coordination environment of particular species, with a significant reduction in data

^a Joint Center for Energy Storage Research, the Molecular Foundry, Lawrence Berkeley National Laboratory, Berkeley, California 94720, United States

* E-mail: roncoroni@lbl.gov

dimensionality, that is informed more by the data itself than by human expectation (i.e., minimizing external bias).

In what follows, we provide a case study for mining the coordination environments of ions in liquid electrolytes as provided by extensive molecular dynamics sampling using empirical force fields, computed using LAMMPS¹⁰. The sampling may result from simple trajectories for well-defined thermodynamic ensembles, or it may be biased with respect to certain collective variables, as in umbrella sampling¹¹ or metadynamics^{12–14}, to explore the free energy landscape within a defined low-dimensional collective variable space.

Our chosen example focuses on multivalent ions in nonaqueous electrolytes as a particularly challenging case. In our experience, the distribution of coordination environments about multivalent ions is definitively multimodal – one should not speak of *the* solvation environment, but rather admit that there may be multiple solvation environments¹⁵. This presents several challenges to theoretical modeling, especially when particular ion coordinations may persist in deep local minima of free energy, hindering a full sampling of the available configuration space. Although such challenges have been addressed using collective variables that span coordination number,^{16,17} we may lack validation that the sampled behavior is not limited by the choice of collective variable. In addition, experimental characterization, such as spectroscopy, requires molecular models for interpretation, especially if multiple isomers may contribute to the measured spectra^{3,18–23}. For future work on this complex problem, we claim that unsupervised classification can ultimately deconvolute the distribution of coordination environments into well-separated unimodal, normal distributions, which can be provided as motifs for defining optimal collective variables to sample the free energy landscape effectively and to make unambiguous interpretations of experimental measurements.

Methodology

The behavior of electrolyte, in bulk or in the presence of an interface, can be studied with MD simulations^{24–26}. Because of the high amount of data generated even by a simple MD trajectory, the task to analyze and identify relevant aspects of the calculation can become cumbersome. In the field of electrolyte science, often we are interested in the unique chemical environments that are present around specific solvated ions. In particular, the first solvation shell is critical in determining important properties, such as solubility of the salt, transport properties, stability and reactivity with the electrodes^{27–34}. In the following section we present a workflow we developed to assist with the analysis of large MD datasets. Our goal is to understand which solvation environments are sampled by the MD trajectory. Ideally, we want to minimally rely on arbitrary inputs provided by a user and instead make use of a robust, unbiased procedure that can leverage as much data as is available to us.

Local structure sampling

Since we are interested in the chemical environment around solvated ions in electrolytes, the first step is to define a spherical

region around the ions of interest (e.g. the Ca²⁺ cation) with an element-specific cut off radius and remove the atoms outside it. In this way, we can effectively isolate what we refer in this paper as the *local atomic arrangement* around the ion of interest. An effective choice to define the cut off is to use element-specific radial distribution functions and estimate the extent of the first solvation shell. Extracting the local atomic arrangement out of an entire MD frame dramatically reduces the size of the trajectory to be analyzed. Instead of frames of thousands of atoms we are left with a small cluster of atoms around the species of interest. If there is more than one ion of interest in the simulation box, multiple local atomic arrangements can exist in each frame. In cases where their solvation shells may overlap, we can further decide to merge them into dimers, trimers, etc. and consider them as a single unit.

Alignment and Permutation

Structure alignment and permutation is a central component of the procedure we have assembled below. To determine if two complex atomic arrangements sampled from an MD trajectory are similar is easily complicated by the strong likelihood that their centers of mass are displaced or translated from one another and that they may have quite different orientations with respect to a given set of Cartesian axes. In structural biology, where large, complex macromolecules may have quite low symmetry, determining an optimal translation and rotation to best overlap two structures can be accomplished by minimizing the root mean square deviation (RMSD) of the overall atomic structure - a Euclidean metric in the multi-atomic three-dimensional coordinates³⁵:

$$\text{RMSD}(\mathbf{X}, \mathbf{Y}) = \sqrt{\frac{1}{N} \sum_{i=1}^N \sum_{\alpha=x,y,z} |X_{\alpha}^{(i)} - Y_{\alpha}^{(i)}|^2},$$

where \mathbf{X} and \mathbf{Y} are N -dimensional sets of atomic position vectors, assumed to have the same index ordering with respect to i . Optimal alignment is achieved by finding the parameters that define a translation vector and rotation matrix that minimize the RMSD.

Surprisingly, aligning much smaller clusters can be more challenging due to the strong possibility that they exhibit somewhat random ordering of the atomic coordinates that define the cluster due to their assembly being driven by non-bonding interactions – Coulombic attraction and/or van der Waals forces, most likely. This possible reordering of the atomic species can discontinuously alter the RMSD and prevent gradient-based minimization algorithms from optimally aligning pairs of structures. Multiple approaches have been developed in the literature to overcome this issue. Reducing atomic structural information into orientation-independent features (bond lengths, angles, dihedrals, Coulomb matrices³⁶, bag of bonds³⁷, etc.) can help, but still suffers from the possible random ordering of individual elements³⁸. Binning or histogramming the atomic data as a density field³⁹ or within some other reduced-dimensional feature vector⁴⁰ can remove

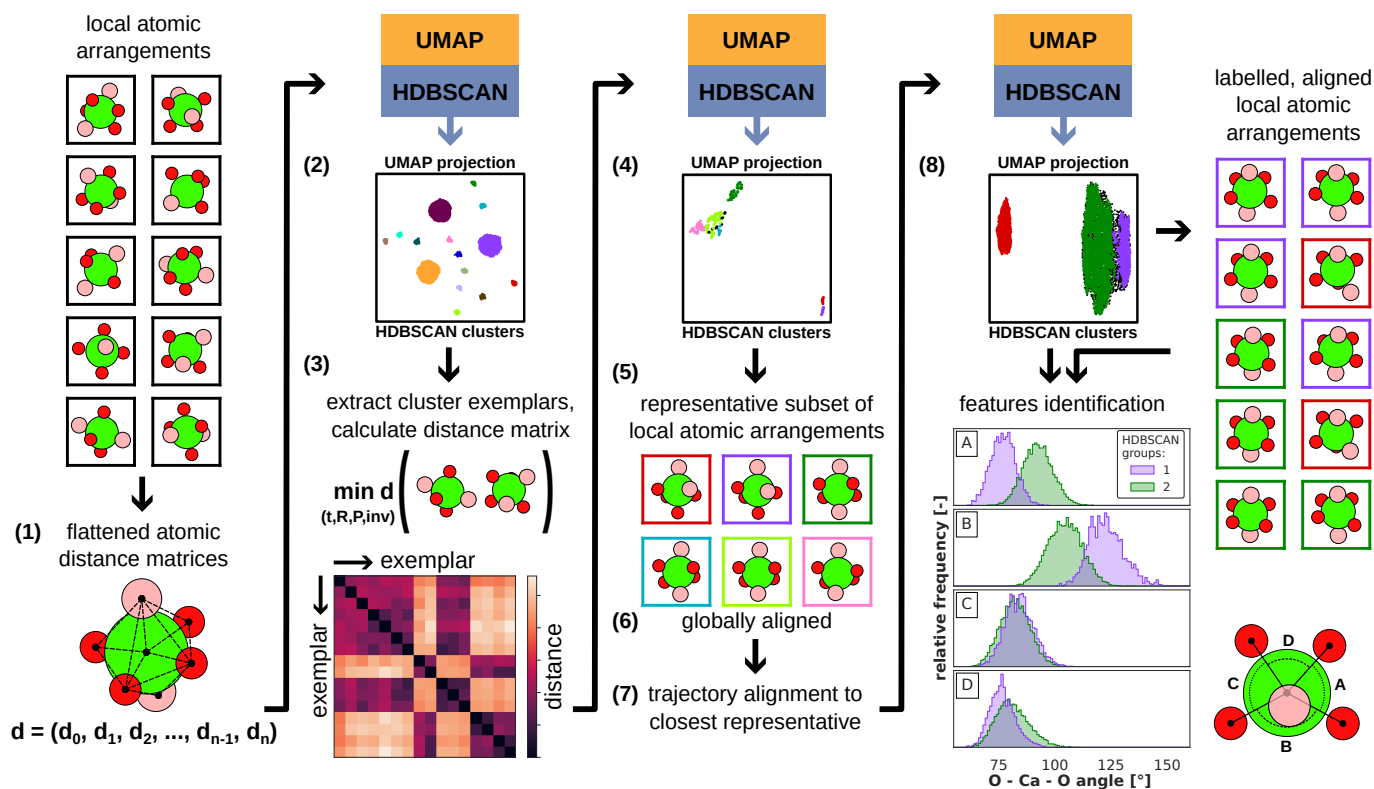


Fig. 1 Schematic described in the text highlighting the multi-step process by which local atomic arrangements for a given chemical formula are ultimately grouped into clusters of aligned atomic arrangements through successive application of dimensionality reduction (UMAP), hierarchical density-based clustering (HDBSCAN) and alignment via minimization of the root mean square deviation d while including translation (t), rotation (R), permutation (P), and inversion (inv) of the atomic coordinates (FASTOVERLAP), followed by specific distinguishing feature identification that helps to define interpretable labeling of the final clusters.

this issue but may require large fields to store the information or may be dependent on the grid-spacing of the density/histogram representation.

Due to the factorial scaling of permutations with respect to the number of identical atoms in the local atomic arrangements, a brute force approach to enumerating and comparing all possibilities may become inefficient or at worst intractable. The so-called Hungarian method⁴¹, and approaches derived from it such as the shortest augmenting path^{42,43} provide a solution to this combinatorial optimization problem in polynomial time. We make use of FASTOVERLAP's branch and bound alignment algorithm⁹ to align pairs of atomic arrangements by minimizing their RMSD with respect to permutation (P) and inversion (inv), in addition to translation (t) and rotation (R) of the 3D structures:

$$d_{\min}(\mathbf{X}, \mathbf{Y}) = \min_{t, R, P, \text{inv}} \text{RMSD}(\mathbf{X}, \mathbf{Y}'),$$

where \mathbf{Y}' is the translated, rotated, permuted, inverted version of \mathbf{Y} .

The final output of optimally aligned structures greatly facilitates a fast appreciation of the salient features of a given cluster of atomic arrangements, so that they might be tagged with physically interpretable labels.

Computing d_{\min} between every pair of atomic arrangements to

define a huge distance matrix (scaling as the square of the number of arrangements) would definitely reveal details of distinct groups within the data set, since this metric defines which arrangements are "close" to one another. This distance matrix could be considered as a representation of the underlying distribution or density of the data set. However, this metric scales with the square of the number of local atomic arrangements. Which is still an unwieldy object to compute and to analyze efficiently and automatically for a large dataset.

Clustering algorithm

In order to classify the local atomic environments into distinct groups, we decide to rely on a procedure that leverages a combination of: the dimensionality reduction technique UMAP⁷, to provide a low dimensional projection of large data sets while maintaining their topological structure; and the hierarchical clustering algorithm HDBSCAN⁸, to organize the low dimensional data into similar groups. Pre-processing the data with UMAP greatly facilitates the clustering procedure, since HDBSCAN understandably struggles to effectively identify dense regions in spaces of higher dimensionality⁸. Our approach is designed to significantly reduce the computational overhead for comparisons via structural alignment and permutation.

A scheme for the analysis procedure is shown in Figure 1. After

extracting the local atomic arrangements from the trajectory, we separate them into groups, each with a distinct chemical formula (e.g. B_2CaO_4), which we analyze separately.

(1) The first step is to calculate the flattened upper triangular atomic distance matrix for every local atomic arrangement in the trajectory. This results in a row vector containing the pairwise distance between every pair of atoms in the local atomic arrangement. Relying on the flattened distance matrix instead of the raw atomic positions is advantageous, since it is invariant to arbitrary translations or rotations of the local atomic arrangements.

(2) Using UMAP, we generate a low-dimensional projection of the input data – a reduction from $n = N(N - 1)/2$ to 2, where N is the total number of atoms in the local arrangement. Points that are close to each other in the UMAP space have a similar flattened upper triangular distance matrix. This densification in the reduced dimensions of the UMAP space facilitates clustering by HDBSCAN into labelled groups. Arrangements that have been assigned a different label due to significant differences in their distance matrix vectors could actually comprise similar structures that differ only because of an arbitrary permutation of their atomic indexes.

(3) We extract from each HDBSCAN cluster some exemplar local atomic arrangements, viz., particular arrangements that come from high density regions in each cluster. Then, using FASTOVERLAP, we obtain the minimal distance d_{\min} between each pair of exemplars (including permutations) and build a corresponding exemplar distance matrix.

(4) We provide UMAP with this exemplar distance matrix as a feature vector to generate another low-dimensional projection, which we then organize into clusters with HDBSCAN. Since d_{\min} obtained with FASTOVERLAP now accounts for permutation of the atomic indexes, as well as inversion, rotation and translation of the 3D structure, we would expect the total number of groups to reduce significantly at this step, as symmetry equivalent exemplars should merge.

(5) We extract new exemplars from this second UMAP projection space to select an even smaller subset that will constitute an alignment basis set. With the assumption that the sampled trajectory is complete – i.e., it contains all the expected local atomic arrangements – the alignment basis set should also be complete.

(6) We self-consistently align the elements of this basis to one another to more easily compare them upon visualization. Obtaining an optimal global alignment between multiple structures can be challenging. Normally, pairwise alignment algorithms can efficiently align structures two-by-two by minimizing their RMSD. However, minimizing the total RMSD of multiple structures at the same time adds an additional level of complexity. It is important to ensure that the main features of the coordination environment are properly aligned in a similar way. As an example, if we were unfortunate enough to try to align all structures to a single reference that coincidentally has a well-defined symmetry axis, it is possible that structures with equivalent structural features end up aligned differently. The goal of the global optimal alignment is to reduce a cost function consisting of the average RMSD between the exemplar arrangements using our chosen structural alignment procedure. This process is iterative, creating a biased random

linkage between the exemplars and sequentially looping through the linkage to align each structure pairwise. At each step, the new alignment is accepted only if the new average RMSD between all exemplars is lower than the previous one.

(7) We now use the globally aligned exemplar atomic arrangements as reference points: we loop through the entire trajectory and try to align each local atomic arrangement to the exemplars, only keeping that alignment with the smallest d_{\min} . This step ensures that all arrangements have a consistent atomic indexing with respect to each other. Only the alignment with the smallest d_{\min} is kept for each local atomic arrangement, but we subsequently compute its RMSD (i.e., without alignment or permutation) with respect to all elements of the alignment basis to define a new feature vector.

(8) At this point, to perform a final clustering iteration, we are left with three possibilities. Since now the atomic indexing is consistent, we can calculate the flattened upper triangular atomic distance matrix once more and provide it to UMAP. Alternatively, since the vector of d_{\min} between each local atomic arrangement and the alignment basis is unique, it can be used as a direct input for dimensionality reduction. Finally, a combination where we provide both the flattened upper triangular atomic distance matrix and the vector of d_{\min} to the alignment basis set can be used. In all three cases, the low dimensional projection obtained with UMAP is then given to the HDBSCAN algorithm. The final product is distinct clusters of structures that have similar local atomic arrangements and that have been optimally aligned with each other.

Computational Details

The model system of this study comprises one Ca^{2+} cation and two BH_4^- anions dissolved in THF next to a graphite interface. THF is an organic solvent, a molecular liquid with the formula C_4H_8O , with the heavy atoms arranged in a pentagonal ring and two hydrogen atoms bound to each carbon. There are 336 THF molecules and two layers of 416 carbon atoms defining the graphite surface. The initial configuration was generated by building a simulation box of size $a = 34.08 \text{ \AA}$, $b = 31.97 \text{ \AA}$, $c = 50.00 \text{ \AA}$, where c is perpendicular to the plane of the graphite surface, with the help of the software PACKMOL⁴⁴.

All classical molecular dynamics simulations were ran with a timestep of 1 fs and the OPLSAA force field⁴⁵. The choice of BH_4^- charges was based on comparisons with Density Functional Theory interaction energies between a Ca^{2+} cation and a BH_4^- anion at face, vertex and edge sites. NBO charges provided the best agreement. Additionally, charge scaling by a factor of 0.7 was applied in order to correctly reproduce the ab-initio free-energy profile of Ca^{2+} in THF as a function of coordination^{46,47} using the methods described below. The particle-mesh Ewald (PME) method with a 1.0 nm cut-off distance and 10^{-5} grid spacing in k -space were used to treat long-range electrostatic interactions. The cut-off for the Lennard-Jones interactions was 1.0 nm and the spline ranged from 0.9 nm to 1.0 nm.

To equilibrate the initial simulation box, we performed multiple steps. First, the structure was relaxed using steepest descent followed by conjugate gradient minimization with a tolerance of 10^{-4} kcal mol $^{-1}$ for the energies and 10^{-4} kcal mol $^{-1}$ Å $^{-1}$ for the forces. Then the system was heated to 298 K during 10,000 steps (10 ps) of NVT simulation. NPT dynamics with an isotropic pressure of 1 atm were then run for 1 ns using a Nose/Hoover thermostat⁴⁸ and a Nose/Hoover barostat⁴⁹. The temperature coupling constant was 0.1 ps and the pressure piston constant was 2.0 ps. The equilibrated lattice parameters were obtained by averaging the box size over the last 0.5 ns of NPT trajectory. Finally, we performed an additional thermalization at 298 K with the newly calculated lattice parameters under NVT dynamics for 5.6 ns.

The evaluation of the free energy profile as a function of the two collective variables (CV) – the coordination number (CN) between calcium and the oxygen of the THF molecules and the distance from the outermost graphite layer (dZ) – was performed using the COLVARS module¹³ as implemented in LAMMPS¹⁰. The initial configuration consisted of equilibrated simulation cells of solvent molecules with ions as described above, and the same MD parameters were used as in the final equilibration step. The cut-off radius for the coordination number was 3.6 Å. To prevent the calcium ion from moving too far away from the graphite, a harmonic potential wall was set at $dZ = 24$ Å. The free energy surface is stored on a discrete grid with a spacing of 0.05 (unitless and Å for the two CVs, respectively). Biasing was applied every 200 time steps by adding Gaussian hills of weight 0.02 kcal/mol and width twice the grid spacing for the two CVs. To accelerate the sampling we used multiple (14) walker metadynamics⁵⁰ communicating every 25,000 timesteps. The trajectory of each replica comprises 300 ns, totalling approximately 4.2 μ s.


Umbrella sampling was performed using the same MD parameters as during the equilibration and metadynamics run by adding a single harmonic umbrella potential at a well defined point on the CV surface. At each point, a trajectory length of 70 ns was collected. The trajectory was saved every 5 ps for further analysis.

Dimensionality reduction with UMAP⁷ was performed first by removing the mean and scaling to unit variance the data⁵¹ and subsequently embedding the scaled data in a two-dimensional space. The size of the local neighborhood was constrained to 15 neighbors with a minimum distance between points of 0.0 and the distance was computed with the Euclidean metric. Hierarchical clustering with HDBSCAN⁸ was performed on the low-dimensional projection with standard parameters and allowing for the possibility of there being only a single cluster. In some cases, parameters were tweaked iteratively depending on visual feedback to obtain reasonable clustering. Namely, varying the number of neighbours in the UMAP projection between 10 and 30 and increasing the value of the smallest size grouping for the HDBSCAN clusters.


Structural alignment was performed with the branch and bound algorithm implemented in the FASTOVERLAP package⁹.

Our clustering algorithm was implemented in Python, leveraging available open-source data science modules (NumPy⁵², sklearn⁵¹) and integrates with the Atomic Simulation Environment (ASE)⁵³. The plots were generated with the help of Matplotlib⁵⁴ and seaborn⁵⁵.

Results



(a) Free energy surface



(b) RDF plots

Fig. 2 (a) Two-dimensional free energy surface for the solvation of Ca^{2+} in THF at a graphite interface as a function of the collective variables: x) coordination number between the Ca and the O of the THF atoms, y) distance between the Ca and the outermost graphite layer. The numbers 1-5 correspond to the region sampled with umbrella sampling and used in the results section. (b) Normalized O and B RDF plots around Ca for trajectories 1-4 (left) and trajectory 5 (right). The vertical dashed line corresponds to the cutoff radii for the analysis discussed later.

Molecular Dynamics Sampling

We provide here some representative examples of applying the clustering procedure to an MD trajectory. Our chosen application is to a cutting-edge electrolyte: calcium borohydride [$\text{Ca}(\text{BH}_4)_2$] dissolved in tetrahydrofuran (THF). We first run a metadynamics simulation to calculate the two-dimensional free energy landscape as a function of two collective variables: the Ca-O coordination number and the distance of the calcium ion from the graphite interface (see Computational Details). The free energy surface shown in Figure 2a exhibits multiple distinct minima, with the most prominent for Ca-O coordination numbers of approximately 3 and 4 that extend as valleys at higher distances from the graphite. Energy barriers on the order of 5-10 $k_B T$ at 300 K separate the minima in these valleys. The lowest point on the energy surface is point 2. Points 1, 3, 4 and 5

have a free energy of $3.1 k_B T$, $0.7 k_B T$, $0.9 k_B T$ and $6.9 k_B T$ higher respectively. We are interested in identifying similarities and differences in the calcium coordination environment at different locations across the free energy landscape. With our initial metadynamics map of the free energy landscape, we next employ umbrella sampling, a constrained MD protocol biased by an additional harmonic potential, centered at particular coordinates on the space spanned by our collective variables. The harmonic potential ensures a local sampling of the MD trajectory close to the free energy surface and allows us to extract quantitative information regarding the relative local ionic populations (see Computational Details).

We choose five distinct regions where we run umbrella sampling. Locations 1, 2, 3 and 4 are at different distances from the graphite interface, but all in the free energy valley with Ca-O coordination number of approximately 4. Location 5 is at the transition point between two minima with coordination numbers 3 and 4, respectively, that are the closest to the graphite. Depending on the location of the harmonic constraint with respect to these collective variables, distinct chemical arrangements around Ca^{2+} can be observed. We are particularly interested in the first coordination shell around the ion. Being positively charged, Ca ions attract the oxygen atom in THF and the BH_4^- anions. In each case, the local atomic arrangement can be assigned a chemical formula that indicates its composition in terms of the number of each atomic species.

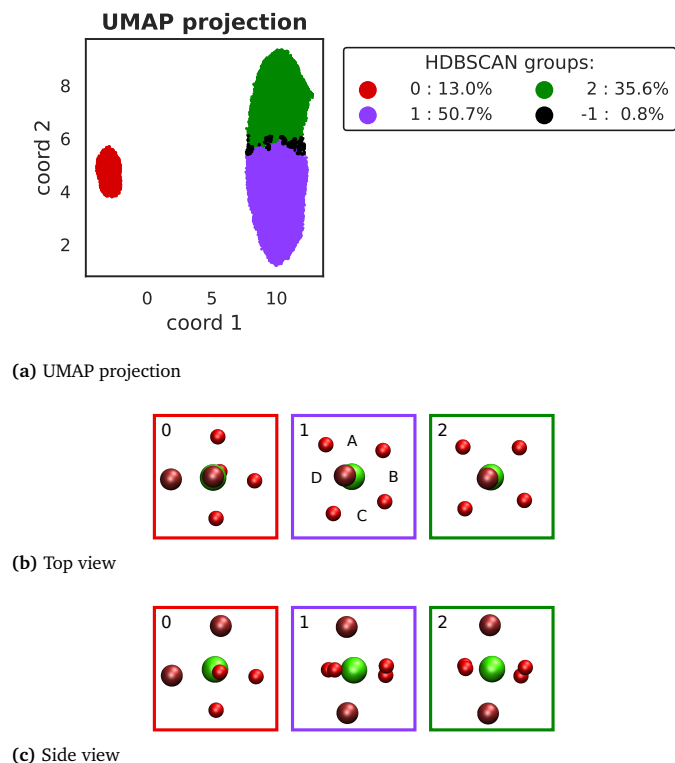


Fig. 3 (a) UMAP projection colored according to HDBSCAN labelling of B_2CaO_4 local atomic arrangements of trajectories 1-4. Points classified as noise are labelled -1 (black). (b) Top and (c) side view of the averaged local atomic arrangements for each HDBSCAN group.

Local Atomic Arrangements

First, we analyze minima 1-4 along the valley with a coordination number of 4. We are interested in knowing the difference in local atomic arrangement as a function of the distance from the graphite interface. Initially, we limit our analysis to only three atomic species (Ca, O, and B), ignoring details of the solvent molecule or anion orientation. To select an appropriate cutoff radius and extract the local atomic arrangements from the trajectories, we use the radial distribution functions (left RDF plot in Figure 2b). The RDFs of O and B with respect to Ca show two narrow peaks at 2.35 \AA and 2.45 \AA , respectively, indicating the extent of the first solvation sphere around the cation. Beyond these first peaks, the RDF is almost flat until approximately 5 \AA , where the oxygen atoms of the second solvation sphere appear. By choosing a cutoff radius of 4 \AA for both B and O, we can effectively isolate the local atomic environment around Ca^{2+} and ensure we only include the first solvation shell. The predominant total coordination number (including boron and oxygen) for calcium in this environment is 6, but other coordination numbers could occur. For example, at the neighboring minima with a coordination number of 3 the maximum total coordination number allowed is 5 (since there are only two available borohydride ions in the entire simulation box). In the umbrella sampling, we constrain 4 oxygen atoms to be coordinated to the calcium and since the two available borohydride anions are strongly attracted by the divalent cation, we observe only one chemical formula at the four sampled points, namely B_2CaO_4 . After extracting the local atomic arrangements, each sampling point provides 12,000 molecular snapshots, totalling 48,000 atomic arrangements. We then proceed to employ the clustering algorithm to obtain the main coordination environments and compare their occurrence as a function of their distance from the graphite interface without previous assumptions on their arrangements.

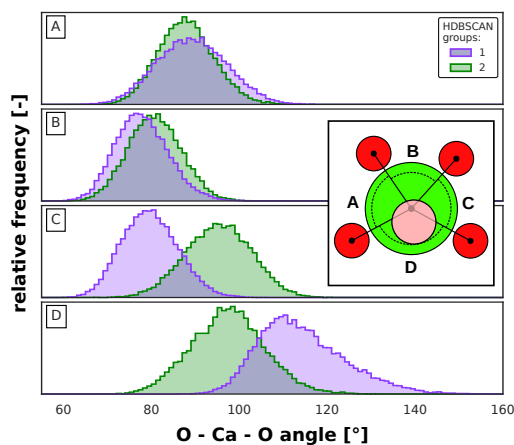


Fig. 4 Distribution of angles between neighbouring O atoms with center on the Ca for the groups 1 and 2. The detail within the figure shows the ordering of the angles A, B, C and D.




Fig. 5 Distribution of species of B₂CaO₄ as a function of distance from the interface (z).

Clustering Analysis

The clustering algorithm identifies 3 main local atomic arrangements across the four distances sampled. The final projection of the UMAP space on 2 dimensions with the relative clusters labelled with HDBSCAN is shown in Figure 3a.

The projected coordinates are mainly divided into two regions in the UMAP space. On the right, a big cluster of points is divided into two groups (1 and 2). On the left, there is the isolated group 0. The black points labeled -1 which HDBSCAN did not classify as belonging to any group are considered noise, although they are most likely intermediate arrangements between groups 1 and 2. The most occurring atomic arrangement, group 1, accounts for roughly 50.7% of the 48,000 arrangements analyzed. The second most occurring, group 2, accounts for 35.6% of the arrangements. Group 0 consists of 13.0% of the arrangements. Only 0.8% of the data was labelled as noise. The reference atomic structures of each labelled group are shown in Figures 3b and 3c.

To try to understand why the local atomic arrangements were divided into different groups, we look at the angles between the oxygen atoms and the boron around the central calcium atom. The biggest difference comes from the bent or linear arrangement of the two B atoms. In group 0, the B atoms are adjacent, as proposed in²¹, while in groups 1 and 2 they are located axially on opposite sides of the Ca ion. Further differences of configuration between groups 1 and 2 arise from differences in the O-Ca-O angles.

Figure 4 shows the four O-Ca-O angles of groups 1 and 2. It becomes evident that group 1 consists of local atomic arrangements where one of the O-Ca-O angles is much larger than the average. Specifically, the mean angle D is $114.3^\circ \pm 9.8^\circ$ and $97.5^\circ \pm 8.3^\circ$, in groups 1 and 2, respectively. The fact that there seem to be normal distribution of angles with distinct means for the different groups indicates that those features play a role in defining the final UMAP projection of the local atomic arrangements. Specifically, since we only provide UMAP with the distances between atoms in a local atomic arrangement, the O-O distances responsible for the O-Ca-O angle are distinct between different groups.


We can see from Figure 5 that depending on the distance of the calcium ion from the graphite, the relative population of the three coordination environments changes. Closest to the interface, only 4.9% of the coordination structures belong

to group 0 (bent configuration). This number increases to 16.9% in the second layer, reaches the highest value of 18.1% at 15 Å from the graphite and finally decreases to 11.9% in the bulk. Interestingly, the ratio between local arrangements 1 and 2 is highest at the interface, where their prevalence is almost equal. Most likely, local atomic arrangement 2 is favoured by the presence of the graphite. At all distances sampled, approximately 50% of the structures belongs to group 1.


The clustering procedure involving UMAP and HDBSCAN was able to separate not only atomic arrangements that are clearly different (adjacent vs opposite boron arrangement), but also atomic arrangements with finer differences of angles between the coordinating atoms (O-Ca-O angle). However, we noticed by repeating the analysis that it is not always possible to distinguish the atomic arrangements with small O-Ca-O angle differences (groups 1 and 2). The main reason behind this is the randomness embedded in the dimensionality reduction and clustering pipeline. UMAP is a stochastic algorithm and therefore some variance in the results is to be expected, although a consistent random seed can be set to guarantee reproducibility for a given data set. There is also some stochasticity in our iterative global alignment procedure, based on its choice of random linkage. In general we observed the low dimensionality projection to be quite stable. However, mostly when working with smaller feature vectors (considering less atoms) and fine structural differences in atomic arrangement, a different final clustering could be obtained. Despite this, the differentiation of the main feature – the B-Ca-B angle – is always retrieved, which is a good indicator that the clustering algorithm functions when the local atomic arrangements are clearly distinct.

Increasing Complexity


In order to obtain a more detailed picture of the chemistry around calcium in the MD trajectory, we decided to increase the size of the local atomic arrangements by including the carbon atoms of the coordinated THF (while still ignoring the hydrogen atoms of the borohydride anion and THF). Here we do not explicitly use a cutoff radius for the carbon atoms, instead we use the molecular bonding information provided to LAMMPS to reconstruct the THF molecules with an oxygen atom within the cut-off radius of 4 Å from the specified calcium ion. After re-running the clustering procedure, we can observe even richer information on the local coordination environment (Figure 6a).



(a) UMAP projection



(b) Top view



(c) Side view

Fig. 6 (a) UMAP projection colored according to HDBSCAN labelling of B_2CaO_4 local atomic arrangements that include THF backbones of trajectories 1-4. Points classified as noise are labelled -1 (black). (b) Top and (c) side view of the averaged local atomic arrangements for each HDBSCAN group. The four THF molecules in the configurations with the boron axially arranged (groups 1-7) are B: back, F: front, L: left, R: right.

This time we identify 8 main groups using the clustering procedure. One group (group 0) is distant from the other ones on the UMAP projection space. The local atomic arrangement of group 0 exhibits two adjacent B atoms, while groups 1-7 all have the B atoms axially arranged on opposite sides of the Ca cation. The total occurrence of group 0 is 12.9%, while in the previous analysis, where we still were excluding the carbon atoms, it was 13.0%. This indicates that both with and without the THF carbon atoms the UMAP projection effectively isolates this structural characteristic – a bent B-Ca-B configuration. The remaining 87.1% of the atomic arrangements belong to cases where the boron atoms are on opposite sides of the calcium ion. In this configuration, the THF molecules are equatorially arranged on a plane around the cation. These local atomic arrangements are divided by our al-

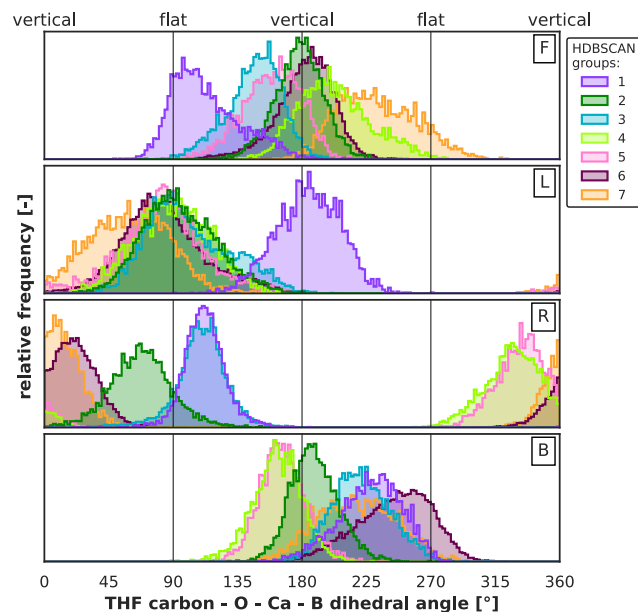


Fig. 7 Distribution of the dihedrals between the C atoms of the THF molecules and the axis spanned by the two B atoms in the case they are axially arranged (groups 1-7 in Figure 6). F: front, L: left, R: right, B: back THF molecule.

gorithm into 7 separated groups in the UMAP space. The main differences between those groups lie in the orientation of the THF molecules around the calcium.

The mean structure of the local atomic arrangements belonging to each HDBSCAN group is shown in Figures 6b and 6c. Apart from the clear difference between group 0 and the rest, we can see how the clustering procedure succeeds in revealing different arrangements of the THF molecules around the calcium ion. From the visualization of the mean structures, we learn how the COCaB dihedral angle formed by one carbon atom adjacent to an oxygen atom in the same THF molecule and the the Ca-B axis seems to be a relevant metric. The distribution of these COCaB dihedral angles is shown in Figure 7. The presence of approximately normal COCaB dihedral distributions (in that they are unimodal) indicates that indeed there are well defined mean COCaB dihedral angles for each group of local atomic arrangements. Furthermore, this validates the choice of averaging the atomic positions to obtain a mean structure, given that the mean structures appear physically reasonable. This is a strong indication that the clustering process has revealed and separated normally distributed data with distinct means.

Since groups 1-7 all have B atoms arranged axially on opposite sides of the calcium ion, then a COCaB dihedral angle of 0° or 180° corresponds to the THF molecule with its pentagonal plane aligned with this axis, which we label as vertical. A COCaB dihedral angle of 90° or 270° we will label as a flat alignment of the THF molecular plane. Angles in between these limiting cases can be referred to as "tilted". Finally, we notice that in most cases the THF molecules are arranged radially with respect to the Ca ion at the center of the atomic arrangement, with the THF dipole




Fig. 8 (a) UMAP projection of group 1 from Figure 6. The labels 0 and 1 correspond to the peaks in density of the heatmap in (b). (c) Contribution of the two regions to the first dihedral angle of group 1 (F) in Figure 7. (d) Dihedral of the F THF molecule with contribution from the two high density regions.

vector antiparallel to a radial vector. In those rarer cases where this is not so (dihedral R in 1 and 3), we can label the alignment as non-radial.

The first thing to be noticed is that no dihedral distribution is similar for the four THF molecules at the same time, which indicates that each group found by the clustering procedure is unique. All configurations with the axial boron have a mix of flat, tilted and vertical coordinating THF molecules. Since the local atomic arrangements have been aligned with each other, the key differences are easily recognizable. Groups 1 and 3, both show a flat, non-radial right THF molecule and only differ from the orientation of the left THF which is vertical and flat respectively. Groups 4 and 5 differ mainly from the orientation of the front and back THF, which are tilted towards the flat THF on the left in one case and towards the vertical THF on the right otherwise. Group 2 has alternating flat and vertical THF molecules, while groups 6 and 7 have a mixture of tilted and flat THFs. Interestingly, we can note how there is no configuration with all four flat THF molecules (or even two adjacent ones). Indeed, ultimately it is steric hindering that dictates which configurations are possible.

In some cases we notice how some of the COCaB dihedral angle distributions have tails. For example, the dihedral angle of the front THF molecule (F) of group 1 extends from its peak near the flat orientation towards the vertical orientation. To investigate this particular feature, we isolate the local atomic arrangements of this group to perform a more detailed analysis. In Figure 8, we can see indeed how there is a finer substructure within the UMAP projection space that was not previously captured when we ran HDBSCAN on the entire low-dimensional space. Instead, if we run HDBSCAN uniquely on a portion of the UMAP projection – namely only the data belonging to group 1 – we are able to

separate the former group into two subgroups. Indeed, there are two regions of high data density in the projected space that can clearly be seen by eye in the associated heatmap (Figure 8b). By partitioning the original cluster into these two regions, and plotting the F dihedral again, we can see how this time the corresponding distribution is separated into two distinct groups that appear unimodal.

As we have shown for the "coarser" clustering on the B_2CaO_4 , the relative population of the groups along the low energy valley of coordination number 4 varies depending on the distance to the graphite. Figure 9b shows the percentage population of each group at the four umbrella sampled points. Again, we can see how the electrolyte/graphite interface shows characteristic atomic arrangements that are more dominant in its immediate vicinity than in regions only 8 Å away or greater. Specifically, the indicated groups 1 and 3 are almost exclusively found in the 5 Å data. When comparing this finding with the mean structures in Figures 6b and 6c, we notice how these two local atomic arrangements have one THF molecule (R) inclined so that it is no longer radially oriented with respect to the Ca ion. Figure 9a shows representative local atomic arrangements for groups 1 (top) and 3 (bottom) and their relative orientation with respect to the graphite surface. The local atomic arrangement is rotated such that the non-radial THF molecule lies parallel to the surface plane. The presence of the interface indeed presents a hard physical barrier to the local atomic arrangements. Clearly, there is some flexibility of the coordination shell around calcium that allows the THF molecules and the borohydride to adjust their arrangement to better accommodate the interface. However, the distortion of the THF coordination sphere costs energy to the system. These non-radial arrangements may be the origin of the overall increase in free energy of 3.1 kT at point 1 (5 Å from the graphite surface) with respect to the other sampled points (2-4) within the same valley in Figure 2a.

Nonequilibrium Sampling

In a second example, we analyze the composition of the umbrella sampled trajectory 5 from Figure 2a. This particular region in the collective variables space is at a transition point close to the graphite surface between two local free energy minima with Ca-O coordination numbers of 3 and 4, respectively. This region has a free energy only ~ 5 kT higher than the lowest of the two minima (at a coordination number of 3). From the radial distribution function around calcium shown in Figure 2b (right) we see that similarly to the previous case there are two narrow peaks at 2.35 Å and 2.45 Å for oxygen and boron respectively. Additionally, in the Ca-O RDF there is a second peak with lower intensity centered at 3.75 Å. The second peak comes with no surprise: by constraining the coordination number to be approximately 3.5, we effectively apply a bias to the system that forces one of the THF molecules farther away from the calcium. In this case, three THF molecules maintain a distance similar to the configuration with coordination number of 4, while the fourth is pushed away, giving rise to the peak at 3.75 Å. To include this

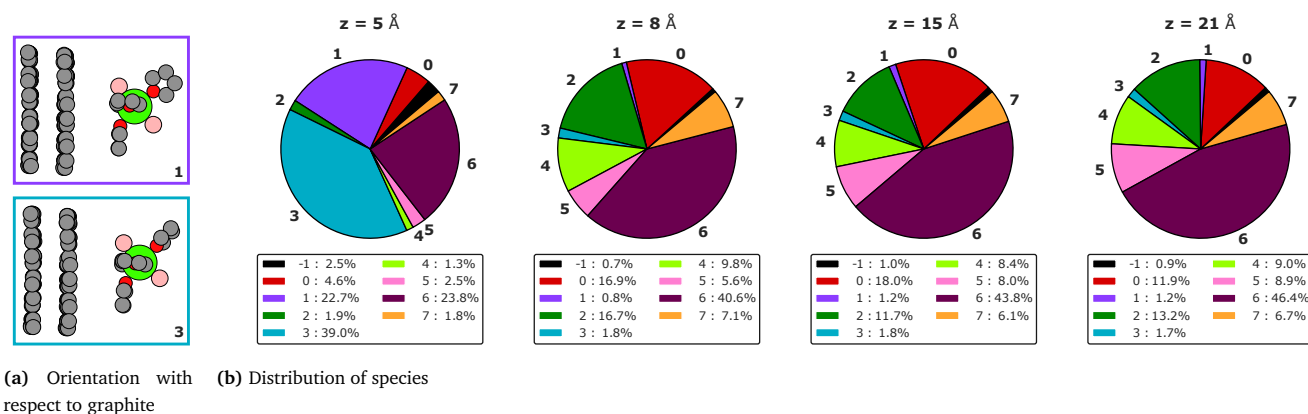


Fig. 9 (a) Orientation of local atomic arrangements of isomers 1 and 3 with respect to the graphite layer. (b) Distribution of species of B_2CaO_4 with THF carbons as a function of the distance to the interface (z). Note that isomers 1 and 3 appear almost uniquely closest to the surface.

second peak in our analysis, we decided to set the cutoff radius of the local atomic arrangement at 4.5 \AA . After isolating the environment around the calcium, we find two distinct chemical formulas: 70.4% of the structures are $C_{16}B_2CaO_4$ and 29.6% are $C_{12}B_2CaO_3$. When multiple chemical compositions are found, a separate clustering procedure is required for each.

Starting from the local atomic arrangements with 3 coordinating THF molecules, we identify 7 different HDBSCAN groups (Figure 10). Only groups 2 and 3, 24.1% of the total population, have a local atomic arrangement with two axially aligned boron atoms, while the majority have a bent boron arrangement. This trend is opposite to what has been observed until now, where the majority of the local atomic arrangements had the two boron atoms on opposite sides of the calcium ion. Again, we see how UMAP and HDBSCAN identify structures with different orientations of the THF molecules.

The structures with 4 coordinating THF molecules were divided into 10 groups and are shown in Figure 11. Also in this case, there is a higher percentage of arrangements with adjacent B atoms with respect to trajectories 1-4, as 40.4% of the local atomic arrangements belong to groups 1, 2, 3, 5 and 6 that show the bent configuration.

Due to the fact that trajectory 5 was sampled with a biasing potential that forces the calcium-oxygen coordination number to be 3.5, the four coordinating THF molecules cannot be at the same distance from the calcium as in trajectories 1-4. Indeed, from the averaged local atomic arrangements we notice how in all cases only one THF molecule is farther away than the others. This is what gives rise to the second peak we observed in the Ca-O RDF plot in Figure 2b. This also is a strong indication that for the system it is energetically favourable to have three THF molecules closer and only one farther away rather than multiple.

In Figure 12a we show the projection of the data from the transition state on the same UMAP embedding used to cluster the data from trajectories 1-4 shown in Figure 6a. This was done by aligning the local atomic arrangements of trajectory 5 to those con-

stituting the alignment basis set of trajectories 1-4 and then using the distance/flattened distance matrix as input to the trained UMAP pipeline. Interestingly, most of the newly projected data lies at the edges of the existing clusters or in between clusters. This indicates that the structures found in trajectory 5 are not a complete departure from points 1-4. On the other hand, the additional data could be used to better understand the link between transition states and stable states found in the MD trajectory. As an example, we focus on the data that has been projected between groups 2 and 3 in Figure 6a. Isolating those data and clustering them with HDBSCAN reveals two separate groups, one closer to group 3 and one closer to group 2. The overlap of the former projections (groups 2 and 3 from trajectory 1-4) with the new data is shown in Figure 12b.

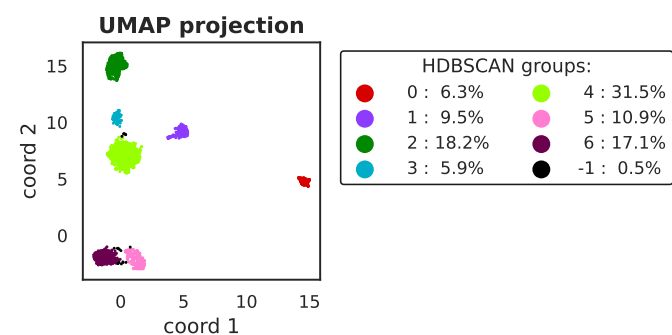
A snapshot of the representative structures of each group is shown at the right of the UMAP plot. We clearly see from the two bottom structures that there is one THF molecule farther away from the calcium, as expected by the umbrella sampling biasing potential in trajectory 5. Moreover, we notice how structures belonging to groups that lie closer on the UMAP space have more structural similarity. Group 1 indeed looks like group 2 with an elongated THF and group 0 is closer to group 3. The RMSD value between the representative structures supports this finding:

Table 1 RMSD values (in \AA) between the four structures in Figure 12b


	group 0	group 1	group 2	group 3
group 0	0.000	0.530	0.450	0.693
group 1	0.530	0.000	0.581	0.445
group 2	0.450	0.581	0.000	0.440
group 3	0.693	0.445	0.440	0.000

Outlook and Discussion


The main power of this unsupervised classification of local atomic arrangements is the great help it provides to the researcher to perform further analysis. Since each local atomic arrangement has been aligned to its representative and the representatives to one another, effective visualization of the atomic structures is easy. Due to the alignment, we can compute the mean structure of



(a) UMAP projection

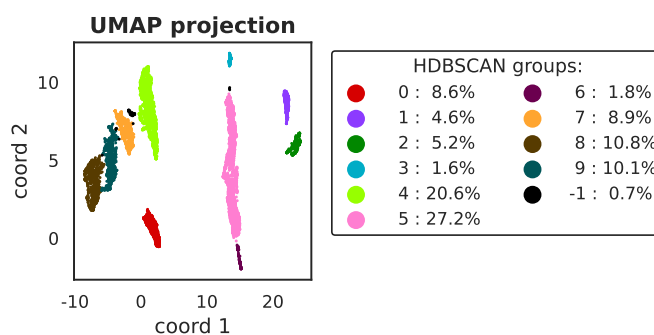


(b) Top view




(c) Side view


Fig. 10 (a) UMAP projection colored according to HDBSCAN labelling of $C_{12}B_2CaO_3$ local atomic arrangements of trajectory 5. Points classified as noise are labelled -1 (black). (b) Top and (c) side view of the averaged local atomic arrangements with chemical formula $C_{12}B_2CaO_3$ for each HDBSCAN group.



(a) UMAP projection



(b) Top view



(c) Side view

Fig. 11 (a) UMAP projection colored according to HDBSCAN labelling of $C_{16}B_2CaO_4$ local atomic arrangements of trajectory 5. Points classified as noise are labelled -1 (black). (b) Top and (c) side view of the averaged local atomic arrangements with chemical formula $C_{16}B_2CaO_4$ for each HDBSCAN group.

each group by averaging the atomic positions within each group. Although the averaged atomic positions may not have physical meaning (the bond lengths may not be maintained, for example), it is helpful for visualization purposes. And, as we have seen, the partition of the sampled data into clusters that exhibit normal distributions for key structural characteristics indicates that the mean structure may indeed be physical.

If the isolated clusters do emerge as normal distributions, then a promising future direction would be to explore the principal components of the variance in these distributions. Methods such as principal component analysis (PCA, for example implemented in the scikit-learn Python package⁵¹) can reveal which vibrational modes are dominant in each group and perhaps serve as input for a reduced dimensional model of the system free energy with temperature dependent vibrational entropic components.

The mean or other representatives (such as exemplars) of each group can be used as inputs for more expensive calculations of ob-

servables that might reference experiment. For example, *ab-initio* electronic structure calculations can provide details relevant to both vibrational, NMR, and X-ray spectroscopy. The additional ease with which we can understand subtle variations in the distribution of representative structures as a function of the local environment (distance from an interface, for example) can also be used in complex interfacial studies to interpret differences in measurements that are specifically related to the interfacial environment – be they structural, chemical, or merely orientational in origin.

The choice of alignment algorithm within our method is not a limitation. For the purposes of this work, we mainly used FASTOVERLAP with its Branch and Bound method for structural alignment including permutations and inversion. However, because of the modular structure of the code, other alignment algorithms could easily be included. The shape matching IRA method was also tested⁵⁶. In our experience, FASTOVERLAP has

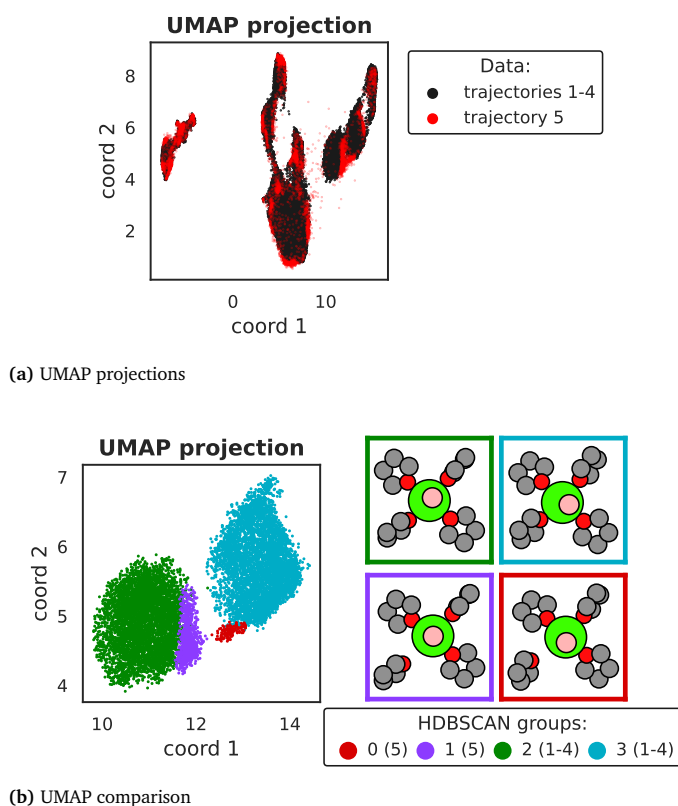


Fig. 12 (a) Projection of the data from trajectory 5 on the UMAP pipeline of trajectories 1-4. Most of the newly projected data lies at the ridges or between previously existing clusters. (b) Overlap of the UMAP projection of groups 2 and 3 from trajectories 1-4 (Figure 6a) and UMAP projection of data from trajectory 5 that is close to it and has been divided by HDBSCAN in two groups (0 and 1).

proven to be more reliable in correctly aligning the structures, with little to no misalignment. With IRA on the other hand we observed occasionally a misalignment to the reference structure. The probability of failure of IRA increases when the structures are not exactly or nearly congruent. Therefore, having a complete sample of reference structures that we try to align to is crucial to ensure that every coordination structure has a near-congruent representative. The best alignment results and a good trade-off in computational performance were obtained with the Branch and Bound alignment algorithm of FASTOVERLAP.

Aligning the structures to a subset of representatives instead of calculating a full distance-distance matrix has the advantage of reducing the size of the matrix that needs to be calculated. Instead of performing a $N \times N$ alignment problem (with N being the total number of structures in the sample), we do a $N \times n$ alignment (with n being the much smaller number of representatives). With a good choice of representatives, this reduced dataset is normally enough to differentiate the main features of the coordination environment. This leads to the fast and reliable post-processing capabilities of our algorithm. On easily available computing resources, for example a regular desktop workstation, analyzing tens of thousands of local atomic arrangements was done within minutes. The rapidity at which it is possible to obtain insights on

large MD datasets further promotes an iterative approach to the analysis, where hidden relevant metrics are learned on the fly.

The advantages of this methodology are not limited merely to unsupervised dimensionality reduction. By providing previously calculated labels, for example, from the clustering of a training set, it is possible to perform supervised dimensionality reduction and metric learning. We can use this approach to train a UMAP model on a smaller or more diverse data set (maybe the meta-dynamics trajectory), and use the model to calculate the reduced coordinates of new structures. Once the learned and supervised metric reduction is done, the new data is given to HDBSCAN and labelled according to the previous labels. Ideally, this can be implemented into a pipeline to obtain on-the-fly classification of the coordinating environment as the simulation goes on. In the worst case, if new data presents examples that cannot be labeled – the entirety of the collective variable space may not have been sampled in the training set – then we are alerted to this deficiency and we can retrain with this new data to augment the number of distinct classes. We demonstrate the efficacy of the learning method on the same trajectory of the example with B_2CaO_4 presented above. Out of the 48,000 structures of the four umbrella sampling trajectory, we randomly took 3,000 as a training set. The clustering procedure recognized the three main coordination groups observed in Figure 3a. By saving a fitted dimensionality reduction pipeline, composed by normal scaling of the input coordinates and reduction of the space by UMAP, we can apply it to the "new" data composed by the 45,000 remaining clusters. The trained model was able to correctly identify 99.5% of the new structures provided. The mislabelling only happened between the two coordination environment where the boron is on opposite sides and only the O-Ca-O angle is different. This example shows the strength of the analysis method, where less than 10% of the data can be used to label the remaining 90%.

Conclusions

In this work, we show how novel dimensionality reduction and hierarchical clustering algorithms can be embedded into a workflow to provide an unbiased description of coordination structure obtained through MD sampling. We isolate, classify and visualize the relevant aspects of the coordination environment of Ca^{2+} ions solvated in an electrolyte comprising BH_4^- anions and THF solvent in the vicinity of a graphite interface. We were able to identify the most prevalent conformational isomers of the first solvation shell of Ca^{2+} and quantitatively estimate their population as a function of the distance from the graphite interface. The method is trivially expandable to larger solvation environments, that can reveal details on electrolyte performance⁵⁷.

The development of new methods to process increasingly complex data sets with as little human intervention as possible is crucial in the rapidly expanding fields of high-throughput computational materials science and chemistry. Dimensionality reduction and hierarchical clustering algorithms proved to be effective tools to facilitate detailed structural analysis and to partition atomic arrangements into distinct structural distributions, which are more amenable to standard statistical analyses and comparison with experiment.

The efficient extraction of representative structural motifs with distinct coordination environments and their associated variance can greatly accelerate the interpretation of various experimental techniques with particular sensitivity to local structure and chemistry. With the approach outlined in this work, measured spectra, and their simulated analogues, may more easily be deconvoluted into contributions from specific coordination environments or regional contributions, especially from the vicinity of active interfaces. Analysis of existing molecular dynamics sampling data for the interpretation of NMR chemical shifts⁵⁸ or X-ray absorption spectroscopy – either near-edge (NEXAFS/XANES)^{59,60} or extended fine structure (EXAFS)^{61,62} – will surely benefit, especially with recent access to operando measurements with nanometer sensitivity to chemistry at the electrode-electrolyte interface^{1,63}. As a step in this direction, we have just applied this unsupervised learning technique to explore the formation of electroactive species in $\text{Ca}(\text{BH}_4)_2|\text{THF}$ near a graphite electrode⁶⁴.

Author Contributions

All authors helped conceptualize the original idea from DP. ASM and FR developed and tested the clustering algorithm. FR ran the MD calculations, analyzed the data and wrote the original manuscript draft with support from ASM and DP. SS contributed to force field development. DP supervised and managed the project.

Conflicts of interest

There are no conflicts to declare.

Acknowledgements

This work was fully supported by the Joint Center for Energy Storage Research (JCESR), an Energy Innovation Hub funded by the U.S. Department of Energy, Office of Science, Basic Energy Sciences. All simulations were performed at the Molecular Foundry, Lawrence Berkeley National Laboratory.

Notes and references

- J.-J. Velasco-Velez, C. H. Wu, T. A. Pascal, L. F. Wan, J. Guo, D. Prendergast and M. Salmeron, *Science*, 2014, **346**, 831–834.
- Y. Sun, T. Yang, H. Ji, J. Zhou, Z. Wang, T. Qian and C. Yan, *Advanced Energy Materials*, 2020, **10**, 2002373.
- F. Yang, Y.-S. Liu, X. Feng, K. Qian, L. C. Kao, Y. Ha, N. T. Hahn, T. J. Seguin, M. Tsige, W. Yang, K. R. Zavadil, K. A. Persson and J. Guo, *RSC Adv.*, 2020, **10**, 27315–27321.
- S. S. Yamijala, H. Kwon, J. Guo and B. M. Wong, *ACS Applied Materials and Interfaces*, 2021, **13**, 13114–13122.
- J. Young and M. Smeu, *Advanced Theory and Simulations*, 2021, **4**, 2100018.
- N. Yao, X. Chen, Z. H. Fu and Q. Zhang, *Chemical Reviews*, 2022, **122**, 10970–11021.
- L. McInnes, J. Healy and J. Melville, *UMAP: Uniform Manifold Approximation and Projection for Dimension Reduction*, 2018, <https://arxiv.org/abs/1802.03426>.
- L. McInnes, J. Healy and S. Astels, *Journal of Open Source Software*, 2017, **2**, 205.
- M. Griffiths, S. P. Niblett and D. J. Wales, *Journal of Chemical Theory and Computation*, 2017, **13**, 4914–4931.
- A. P. Thompson, H. M. Aktulga, R. Berger, D. S. Bolintineanu, W. M. Brown, P. S. Crozier, P. J. in 't Veld, A. Kohlmeyer, S. G. Moore, T. D. Nguyen, R. Shan, M. J. Stevens, J. Tranchida, C. Trott and S. J. Plimpton, *Comp. Phys. Comm.*, 2022, **271**, 108171.
- G. M. Torrie and J. P. Valleau, *Journal of Computational Physics*, 1977, **23**, 187–199.
- A. Laio and M. Parrinello, *Proceedings of the National Academy of Sciences of the United States of America*, 2002, **99**, 12562.
- G. Fiorin, M. L. Klein and J. Hénin, *Molecular Physics*, 2013, **111**, 3345–3362.
- G. Bussi and A. Laio, *Nature Reviews Physics*, 2020, **2**, 200–212.
- A. Baskin and D. Prendergast, *Journal of Physical Chemistry Letters*, 2019, **10**, 4920–4928.
- S. Roy, M. D. Baer, C. J. Mundy and G. K. Schenter, *Journal of Physical Chemistry C*, 2016, **120**, 7597–7605.
- E. H. Byrne, P. Raiteri and J. D. Gale, *Journal of Physical Chemistry C*, 2017, **121**, 25956–25966.
- L. E. Camacho-Forero, T. W. Smith, S. Bertolini and P. B. Balbuena, *Journal of Physical Chemistry C*, 2015, **119**, 26828–26839.
- M. I. Nandasiri, L. E. Camacho-Forero, A. M. Schwarz, V. Shutthanandan, S. Thevuthasan, P. B. Balbuena, K. T. Mueller and V. Murugesan, *Chemistry of Materials*, 2017, **29**, 4728–4737.
- J. Z. Hu, N. N. Rajput, C. Wan, Y. Shao, X. Deng, N. R. Jaegers, M. Hu, Y. Chen, Y. Shin, J. Monk, Z. Chen, Z. Qin, K. T. Mueller, J. Liu and K. A. Persson, *Nano Energy*, 2018, **46**, 436–446.
- N. T. Hahn, J. Self, T. J. Seguin, D. M. Driscoll, M. A. Rodriguez, M. Balasubramanian, K. A. Persson and K. R. Zavadil, *Journal of Materials Chemistry A*, 2020, **8**, 7235–7244.
- L. C. Kao, X. Feng, Y. Ha, F. Yang, Y.-S. Liu, N. T. Hahn, J. MacDougall, W. Chao, W. Yang, K. R. Zavadil and J. Guo, *Surface Science*, 2020, **702**, 121720.
- G. Agarwal, J. D. Howard, V. Prabhakaran, G. E. Johnson, V. Murugesan, K. T. Mueller, L. A. Curtiss and R. S. Assary, *ACS Applied Materials and Interfaces*, 2021, **13**, 38816–38825.
- K. Xu, *Chemical Reviews*, 2004, **104**, 4303–4417.
- S. Kerisit, E. S. Ilton and S. C. Parker, *Journal of Physical Chemistry B*, 2006, **110**, 20491–20501.
- D. Bedrov, J. P. Piquemal, O. Borodin, A. D. MacKerell, B. Roux and C. Schröder, *Chemical Reviews*, 2019, **119**, 7940–7995.
- T. A. Pham, K. E. Kweon, A. Samanta, V. Lordi and J. E. Pask, *Journal of Physical Chemistry C*, 2017, **121**, 21913–21920.
- N. N. Rajput, T. J. Seguin, B. M. Wood, X. Qu and K. A. Persson, *Topics in Current Chemistry*, 2018, **376**, 19.
- S. Han, *Scientific Reports*, 2019, **9**, 5555.
- H. K. Bezabh, M. C. Tsai, T. T. Hagos, T. T. Beyene, G. B. Berhe, T. M. Hagos, L. H. Abrha, S. F. Chiu, W. N. Su and

- B. J. Hwang, *Electrochemistry Communications*, 2020, **113**, 106685.
- 31 T. Hou, K. D. Fong, J. Wang and K. A. Persson, *Chemical Science*, 2021, **12**, 14740–14751.
- 32 Z. Tian, Y. Zou, G. Liu, Y. Wang, J. Yin, J. Ming and H. N. Alshareef, *Advanced Science*, 2022, **9**, 2201207.
- 33 D. Konstantinovskiy, E. A. Perets, T. Santiago, L. Velarde, S. Hammes-Schiffer and E. C. Yan, *ACS Central Science*, 2022, **8**, 1404–1414.
- 34 Z. Yu, T. R. Juran, X. Liu, K. S. Han, H. Wang, K. T. Mueller, L. Ma, K. Xu, T. Li, L. A. Curtiss and L. Cheng, *Energy and Environmental Materials*, 2022, **5**, 295–304.
- 35 I. Kufareva and R. Abagyan, *Methods in Molecular Biology*, 2012, **857**, 231–257.
- 36 M. Rupp, A. Tkatchenko, K.-R. Müller and O. A. von Lilienfeld, *Phys. Rev. Lett.*, 2012, **108**, 058301.
- 37 K. Hansen, F. Biegler, R. Ramakrishnan, W. Pronobis, O. A. V. Lilienfeld, K. R. Müller and A. Tkatchenko, *Journal of Physical Chemistry Letters*, 2015, **6**, 2326–2331.
- 38 D. C. Elton, Z. Boukouvalas, M. S. Butrico, M. D. Fuge and P. W. Chung, *Scientific Reports*, 2018, **8**, 9059.
- 39 M. Ceriotti, M. J. Willatt and G. Csányi, in *Machine Learning of Atomic-Scale Properties Based on Physical Principles*, ed. W. Andreoni and S. Yip, Springer International Publishing, Cham, 2018, pp. 1–27.
- 40 W. F. Reinhart, *Computational Materials Science*, 2021, **196**, 110511.
- 41 H. W. Kuhn, *Naval Research Logistics Quarterly*, 1955, **2**, 83–97.
- 42 R. Jonker and A. Volgenant, *Computing*, 1987, **38**, 325–340.
- 43 G. Carpaneto, S. Martello and P. Toth, *Annals of Operations Research*, 1988, **13**, 191–223.
- 44 L. Martínez, R. Andrade, E. G. Birgin and J. M. Martínez, *J Comput Chem*, 2009, **30**, 2157–2164.
- 45 S. V. Sambasivarao and O. Acevedo, *Journal of Chemical Theory and Computation*, 2009, **5**, 1038–1050.
- 46 A. Baskin and D. Prendergast, *The Journal of Physical Chemistry Letters*, 2020, **11**, 9336–9343.
- 47 D. M. Driscoll, N. K. Dandu, N. T. Hahn, T. J. Seguin, K. A. Persson, K. R. Zavadil, L. A. Curtiss and M. Balasubramanian, *Journal of The Electrochemical Society*, 2020, **167**, 160512.
- 48 S. Nosé, *The Journal of Chemical Physics*, 1984, **81**, 511–519.
- 49 W. G. Hoover, *Phys. Rev. A*, 1985, **31**, 1695–1697.
- 50 P. Raiteri, A. Laio, F. L. Gervasio, C. Micheletti and M. Parrinello, *Journal of Physical Chemistry B*, 2006, **110**, 3533–3539.
- 51 F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, J. Vanderplas, A. Passos, D. Cournapeau, M. Brucher, M. Perrot and E. Duchesnay, *Journal of Machine Learning Research*, 2011, **12**, 2825–2830.
- 52 C. R. Harris, K. J. Millman, S. J. van der Walt, R. Gommers, P. Virtanen, D. Cournapeau, E. Wieser, J. Taylor, S. Berg, N. J. Smith, R. Kern, M. Picus, S. Hoyer, M. H. van Kerkwijk, M. Brett, A. Haldane, J. F. del Río, M. Wiebe, P. Peterson, P. Gérard-Marchant, K. Sheppard, T. Reddy, W. Weckesser, H. Abbasi, C. Gohlke and T. E. Oliphant, *Nature*, 2020, **585**, 357–362.
- 53 A. H. Larsen, J. J. Mortensen, J. Blomqvist, I. E. Castelli, R. Christensen, M. Dułak, J. Friis, M. N. Groves, B. Hammer, C. Hargus, E. D. Hermes, P. C. Jennings, P. B. Jensen, J. Kermode, J. R. Kitchin, E. L. Kolsbjerg, J. Kubal, K. Kaasbjerg, S. Lysgaard, J. B. Maronsson, T. Maxson, T. Olsen, L. Pastewka, A. Peterson, C. Rostgaard, J. Schiøtz, O. Schütt, M. Strange, K. S. Thygesen, T. Vegge, L. Vilhelmsen, M. Walter, Z. Zeng and K. W. Jacobsen, *Journal of Physics: Condensed Matter*, 2017, **29**, 273002.
- 54 J. D. Hunter, *Computing in Science & Engineering*, 2007, **9**, 90–95.
- 55 M. L. Waskom, *Journal of Open Source Software*, 2021, **6**, 3021.
- 56 M. Gunde, N. Salles, A. Hémerlyck and L. Martin-Samos, *Journal of Chemical Information and Modeling*, 2021, **61**, 5446–5457.
- 57 N. T. Hahn, J. Self, D. M. Driscoll, N. Dandu, K. S. Han, V. Murugesan, K. T. Mueller, L. A. Curtiss, M. Balasubramanian, K. A. Persson and K. R. Zavadil, *Physical Chemistry Chemical Physics*, 2022, **24**, 674–686.
- 58 P. R. Markwick, C. F. Cervantes, B. L. Abel, E. A. Komives, M. Blackledge and J. A. McCammon, *Journal of the American Chemical Society*, 2010, **132**, 1220–1221.
- 59 D. Prendergast and G. Galli, *Physical Review Letters*, 2006, **96**, 215502.
- 60 L. F. Wan and D. Prendergast, *Journal of the American Chemical Society*, 2014, **136**, 14456–14464.
- 61 J. L. Fulton, G. K. Schenter, M. D. Baer, C. J. Mundy, L. X. Dang and M. Balasubramanian, *Journal of Physical Chemistry B*, 2010, **114**, 12926–12937.
- 62 L. X. Dang, G. K. Schenter, V. A. Glezakou and J. L. Fulton, *Journal of Physical Chemistry B*, 2006, **110**, 23644–23654.
- 63 C. H. Wu, T. A. Pascal, A. Baskin, H. Wang, H. T. Fang, Y. S. Liu, Y. H. Lu, J. Guo, D. Prendergast and M. B. Salmeron, *Journal of the American Chemical Society*, 2018, **140**, 16237–16244.
- 64 A. Sanz-Matias, F. Roncoroni, S. Sundararaman and D. Prendergast, *Ca-dimers and solvent layering determine electrochemically active species in Ca(BH₄)₂ in THF*, 2023, arXiv:2303.08261.