

Change Point Detection of Events in Molecular Simulations using **dupin**

Brandon L. Butler^a, Domagoj Fijan^a, and Sharon C. Glotzer^{*a,b}

^aDepartment of Chemical Engineering, University of Michigan, Ann Arbor, MI.

^bBiointerfaces Institute, University of Michigan, Ann Arbor, MI, USA

^{*}Corresponding author, *E-mail address:* sglotzer@umich.edu

August 5, 2024

Abstract

Particle tracking is commonly used to study time-dependent behavior in many different types of physical and chemical systems involving constituents that span many length scales, including atoms, molecules, nanoparticles, granular particles, and even larger objects. Behaviors of interest studied using particle tracking information include disorder-order transitions, thermodynamic phase transitions, structural transitions, protein folding, crystallization, gelation, swarming, avalanches and fracture. A common challenge in studies of these systems involves change detection. Change point detection discerns when a temporal signal undergoes a change in distribution. These changes can be local or global, instantaneous or prolonged, obvious or subtle. Moreover, system-wide changes marking an interesting physical or chemical phenomenon (e.g. crystallization of a liquid) are often preceded by events (e.g. pre-nucleation clusters) that are localized and can occur anywhere at anytime in the system. For these reasons, detecting events in particle trajectories generated by molecular simulation is challenging and typically accomplished via *ad hoc* solutions unique to the behavior and system under study. Consequently, methods for event detection lack generality, and those used in one field are not easily used by scientists in other fields. Here we present a new Python-based tool, **dupin**, that allows for universal event detection from particle trajectory data irrespective of the system details. **dupin** works by creating a signal representing the simulation and partitioning the signal based on events (changes within the trajectory). This approach allows for studies where manual annotating of event boundaries would require a prohibitive amount of time. Furthermore, **dupin** can serve as a tool in automated and reproducible workflows. We demonstrate the application of **dupin** using two examples and discuss its applicability to a wider class of problems.

1 Introduction

Computer simulations of molecular systems from the atomic to colloidal particle scale are a cornerstone of modern materials research. Given ongoing improvements in algorithms and processor speed, newer studies may make routine use of thousands of simulations or more [1, 2]. As a result, much work has been done to automate, streamline or otherwise simplify the management of large-scale simulation studies from software that manages data and workflows [3, 4] to packages used in data pipelines [5, 6, 7, 8, 9, 10].

We continue this trend by developing a software package that detects events (transitions) within point cloud data, the kind of data produced in molecular simulations. Event detection in point cloud data from particle trajectories is difficult because *a priori* information on the nature of the transition, identifying features, or precursors of the transition (if any) is often missing. For studies involving multiple transitions or pathways, as well as large systematic studies, detection is further complicated by the need to automate the task. This paper addresses this problem using approaches collectively known as change point detection (CPD) [11].

The techniques of CPD, which are commonly used in signal processing, have yet to infiltrate materials research [11]. Change points are defined by abrupt changes beyond expected fluctuations in a time-resolved dataset. In this paper, we will use particle trajectory data as our starting dataset. Molecular dynamics simulation as well as experiments using dynamic (e.g. liquid phase) transmission electron microscopy [12, 13] or confocal laser scanning microscopy [14, 15] — in conjunction with particle tracking software — can produce data consisting of time resolved particle positions and orientations.

Many algorithms [11] have been introduced to detect events associated with change points in data for the purpose of monitoring human activity [16], determining useful telemetry in data centers [17] or predicting machine degradation [18]. Change point detection approaches can be categorized as supervised [19, 20, 21] or unsupervised [22, 23, 24, 25, 26] as well as offline or online. Supervised methods require external labelling of events to *train* on, while unsupervised do not. Offline methods require the entire dataset as input. In contrast, online methods analyze the data as a stream while data is generated so that signals are screened in real time. The approach we take casts event detection as an optimization problem where the objective is minimized over change point locations [27].

Current approaches for detection of interesting events in molecular simulations involve the use of system or particle level order parameters. After computing these features, order parameters or environment labels, events are detected using system-specific or problem-specific detection schemes or by visual inspection. Automating the detection of meaningful events would significantly improve current workflows, facilitate “big data” studies in materials research, and provide a consistent approach across studies.

The first step for the approach we take is to generate a set of descriptors (usually order parameters) which capture structural features of the system that might undergo a change during the simulation. Steinhardt order parameters[28] (including its derivatives such as Minkowski structure metrics (MSM) [29]) are an extremely versatile type of order parameters used to describe local structure of matter based on an expansion of bond orientational order in the basis of spherical harmonics. They are used in all fields of molecular simulations, spanning simulations with both atomistic and coarse grained potentials to study liquids,[30] crystals,[28, 31] and other phases. Depending on types of phases studied other order parameters such as nematic order parameter, local density can also be employed. In the atomistic simulations Smooth Overlap of Atomic Position (SOAP) [32] is another commonly used order parameter, including recent additions of time-dependent SOAP variants[33]. These local order parameters can be used to classify particles into environments using machine learning (ML) [34, 35, 36, 37, 38, 39, 40]. Current practices in molecular simulations for event detection involve system or particle level order parameters.

In this paper, we present a new, open-source Python package `dupin` (named after Edgar Allan Poe’s detective C. Auguste Dupin) for generic, autonomous event detection in particle trajectories with local or system-wide transitions. We show how `dupin` can partition a system’s trajectory into regions of transition and stable (or metastable) states through the use of generic order parameters. In the following section, we outline `dupin`’s multi-stage procedure for detecting a set of change points from a system trajectory. We

discuss options available to the user at each stage of the process. We then present two example applications using `dupin` and show its utility in determining the temporal bounds of system-wide structural transitions and particle-level events. We conclude with a general discussion of `dupin` and potential extensions. `dupin` is available on GitHub and distributed through conda-forge and the Python Package Index (PyPI).

2 Results

Primary purpose of `dupin` is detection of events in time correlated point cloud data sets (particle trajectories). We propose and implement an approach in which we translate a set of point cloud data into a set of relevant descriptors representative of each frame of the trajectory which we refer to as the signal. These descriptors should be carefully selected to capture critical aspects of the system’s dynamics, to ensure successful CPD process. Once such a signal is generated, we can apply change point detection (CPD) algorithms to detect events in the trajectory. To do this we have developed a multi-step scheme which prepares the raw point cloud input into a form suitable for event detection. `dupin` introduces novel methodologies for some steps of this workflow while also supporting a broad range of established methods and user-defined functions. This flexibility allows users to tailor the analysis of events to specific requirements. In this section, we will present the components of our multi-step scheme and options available to the user, discussing both the novel and traditional approaches implemented in `dupin`.

2.1 Detection Scheme

`dupin`’s detection scheme is based on CPD [11, 41]. CPD seeks to assign a set of points, K , where a signal, $S : \{s_0, s_1, \dots, s_N\}$, undergoes a change. We denote a single point in a signal (e.g. s_i) as a frame. We define two operators - indexing and slicing - that act on the signal. The indexing operator returns a single value such that $S[i] = s_i$. The slicing operator produces a sub-signal using a semi-closed interval $S[i, j] : \{s_i, s_{i+1}, \dots, s_{j-1}\}$. Using the slicing operation, we can see that the set of points K encompass $|K| + 1$ sub-signals between $[k_i, k_{i+1})$, where $k_0 = 0$ and $k_{|K|}$ is the last frame in the signal. The first and last change points, k_0 and $k_{|K|}$ are trivial, so the number of change points is often written as $|K| - 2$ instead of $|K|$. We adopt this convention from this point on in the paper.

As an example of CPD, imagine a 100-frame trajectory of a protein with two conformers, A and B, and a signal S comprising data representing the conformation of the protein in each frame. If the protein changes conformation at frames 40 and 60, then the signal S can be sliced into three sub-signals (Figure 1): $S[0, 40)$, $S[40, 60)$, and $S[60, 100)$. In this case, there are two change points $K = \{40, 60\}$. This is because the random fluctuations (the fluctuations of a (metastable) equilibrium system around the mean) should be ignored.

We employ a class of CPD algorithms that use the paradigm of optimization and loss/cost functions to select change points from S [41]. A loss or cost function penalizes some notion of error, which is then optimized to minimize the *cost*. We find a set of change points K of size n , such that choice of each single change point $k \in K$ minimizes the cost function C ,

$$K = \arg \min_k \sum_{i=1}^{n+1} C(S[k_{i-1}, k_i)), \quad (1)$$

where $S[k_{i-1}, k_i)$ is the partition of the signal S between potential change points k_{i-1} and k_i . Notice that the

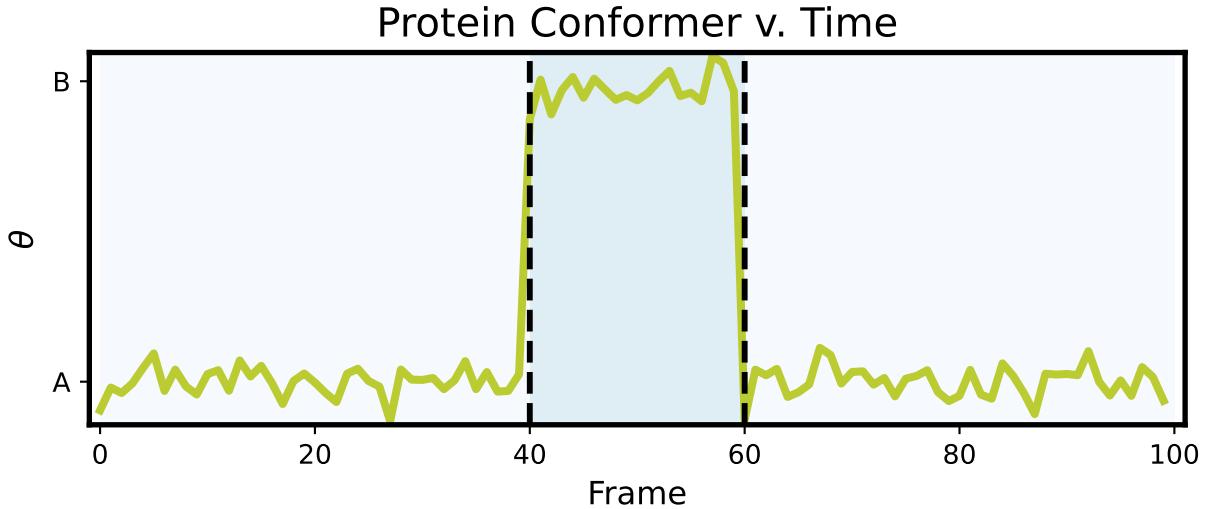


Figure 1: An example detection of a protein transitioning between conformers A and B. At frame 40, the protein goes from conformer A to B and back to A at frame 60. The background colors and dashed black line indicate these two change points, and the corresponding three sub-signals are $S[0, 40]$, $S[40, 60]$, $S[60, 100]$. θ represents a fictitious order parameter along which a structural change happens.

number of change points expected, n , in the signal must be known in advance. A more detailed exploration of this issue is provided in the subsequent discussion. Further elaboration of this class of CPD algorithms is given in Supplementary Material Section 1.

2.1.1 Overview of Method

To develop a generic protocol for event detection in molecular trajectories, we must answer three fundamental questions (i) how to generate a signal from a trajectory, (ii) what cost function(s) best partition the trajectory into sub-signals and (iii) how to determine the correct number of change points in a signal to ensure detection of all events. To address these problems, `dupin` uses a novel approach in which different steps are grouped into stages (see Figure 2). We group the steps into three separate stages: data collection, data augmentation, and detection. The data collection stage includes the *generate*, *map*, *reduce* and *aggregate* steps. The data augmentation stage includes the *transform* step. The detection stage includes the *detect* step. The *generate*, *aggregate* and *detect* steps within the data generation and detection stages are always required. The *reduce* steps can be required or optional depending on the data generated while the *map* and *transform* steps are always optional. The schematic of a typical pipeline for event detection in `dupin` is presented in Figure 2.

2.1.2 Data Collection

We begin by constructing feature vectors for every frame up to a total number of frames N_{frames} . A feature vector is a set of n features used to describe a data point in an n -dimensional space. Order parameters (MSM, local density, etc.) are examples of features. For our case, we refer to features that describe aspects of a given point/time in a trajectory. We combine the feature vectors into a signal S , which is a matrix of size $N_{\text{frames}} \times N_{\text{features}}$. Each matrix element s_{ij} in S contains the value of one of the features in one of the frames of the trajectory. We assume that any change in a molecular system can be adequately described by a

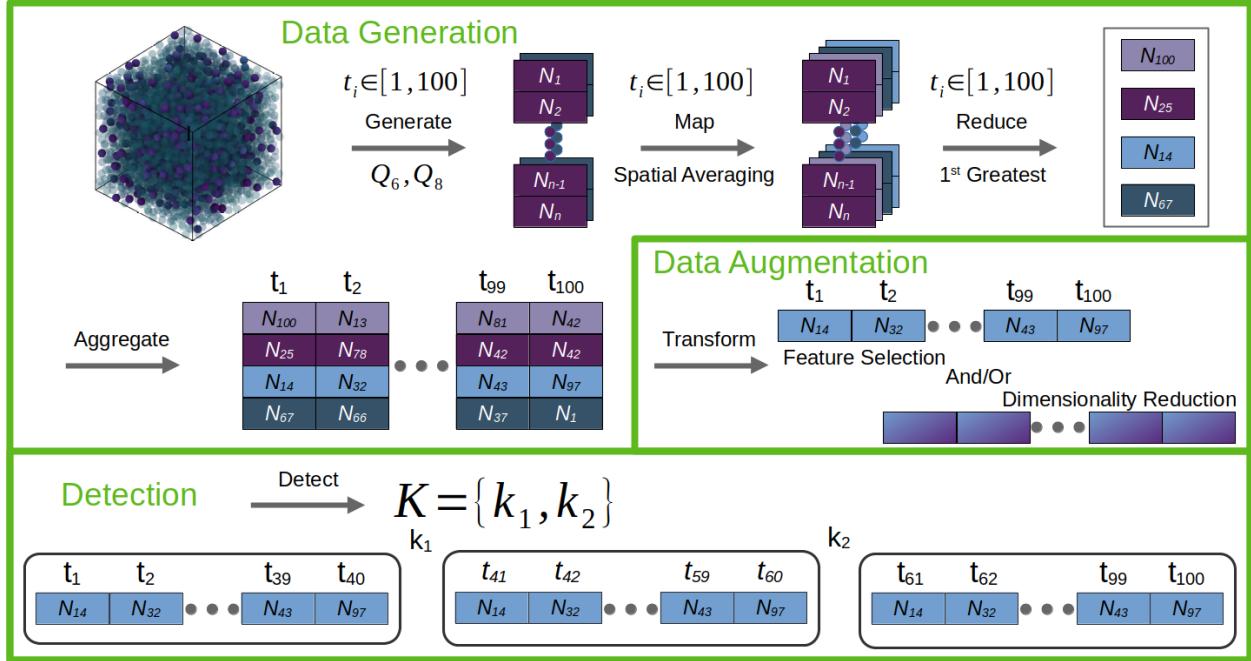


Figure 2: An illustration of the typical pipeline for event detection in a particle trajectory. We use the same numbers as the protein conformer example (change points are 40 and 60 and the total frames are 100). Arrows indicate steps and small colored rectangles represent individual per-particle feature values generated from the trajectory. Green boxes separate steps into three general stages: data collection, data augmentation, and detection. Data is *generated* from the frame and *mapped* to a new distribution and to itself (note the replication of the original features after mapping). The distributions are then *reduced* into one scalar feature each. For the *generate*, *map* and *reduce* steps underneath the arrows, we provide potential functions/applications of the step. The features are then *aggregated* across frames, and *transformed* with either feature selection or dimensionality reduction. Finally, the change points as well as the number of change points are *detected*. The *transform* step is not required, and the *aggregate* step could immediately precede the final step, *detect*. The figure assumes per-particle features only. For global features, the *generate* step immediately precedes the *aggregate* step — they could also optionally be *mapped* first.

such a signal. Changes in the system are thus indicated by changes in the feature vectors over time/frames.

The *generate* step requires us to choose the class of features that will be computed for every frame. Once the feature data is generated, non-scalar or vector features (e.g. per-particle quantities such as Steinhardt order parameters [28]) must first be reduced. Reduction takes a vector feature and converts it through a variety of reducers (functions) to a finite number of scalar features representative of the distribution. Scalar features are never reduced (they are already scalar). The *reduce* step cannot be used on global properties (scalar features) such as system potential energy, but is required for per-particle or high dimensional properties (vector features) like the per-particle potential energy. Examples of scalar features resulting from a *reduce* step include the maximum value, minimum value, mean, mode, median, range, n -th greatest, n -th least, etc. `dupin` offers several novel reducers that are not commonly used in literature, such as the n -th greatest and least reducers. Such reducers aren't useful in settings the descriptors are usually used, such as description of global structure, but are the appropriate choice for detecting transitions in trajectories. This is because initial transition events occur at tails (minima or maxima) of the descriptor distributions.

As an example, consider the utility of the n -th greatest and the n -th least reducers in the context of a study of crystallization. The process of homogeneous nucleation and growth of a crystal from a metastable fluid phase starts when local fluctuations cause a group of neighboring particles to form an ordered cluster. If this cluster reaches or exceeds a critical size, it begins to grow, leading to the production of the crystal phase [42]. This process is known as classical nucleation theory. Ideally, CPD would have access to information on the initial fluctuations conspiring to produce the cluster as well as the cluster's subsequent growth. In the following analysis, we assume (i) we have already constructed or defined features that adequately distinguish the solid and fluid phases, and (ii) the initial fluctuation is small compared to the system size. To detect the nucleation event, we must use reducers closer to the extrema such as the 1st and 10th greatest or least values (e.g. for local densities). Selecting near the extrema is necessary as most particles are fluid when a nucleus forms and selecting reductions that average over the distribution or select near the mean will not register a nucleation event. On the other hand, as the nucleus grows to 10 then 20 then 50 then 100 particles, more particles assume the approximate feature vector of the crystal. Thus, the 100th greatest or least reducers would capture the growth of a cluster up to a size of 100 particles.

While never required, vector features can be mapped to other distributions before aggregating. An example of a useful mapping is spatial averaging of a feature over its neighbors as is sometimes done with Steinhardt order parameters [31].

After a set of scalar features describing a single frame of the trajectory is generated, the process repeats across all trajectory frames and the results are *aggregated* into a single multi-dimensional, time-dependent signal. This signal can be sent directly to the *detect* step or to the *transform* step. The *transformation* step can involve any combination of three tasks: feature selection, dimensionality reduction and signal filtering. Next, we describe each of these data augmentation tasks as they pertain to `dupin`.

2.1.3 Data Augmentation

After features have been reduced, the number of features may easily be in the hundreds. In particle trajectories, each feature incurs noise due to thermal fluctuations — the same fluctuations involved in, e.g., nucleation and growth events in crystallization. As a result, this thermal noise increases the chance of spurious or undetected events for two related reasons. First, one or more features may fluctuate enough to be mistaken as an event. If we assume a baseline chance for such a fluctuation to be 1%, then a signal of 200 features has an 86.6% chance of recording such a spurious event. Second, given n features, if the event

appears in only one feature, then as $n \rightarrow \infty$ the relative reduction in cost to fitting to that one feature decreases drastically. Assuming comparable noise in each feature, then for n features each feature contributes $1/n$ to the cost. For 200 features, the reduction in cost for fitting to a single feature's change point is at most 0.5% percent of the original cost.

To prevent the noise in high feature dimensions from leading to poor event detection, data augmentation through feature selection or dimensionality reduction can improve performance. The amount of improvement depends on the dimension of the feature set, noise level, and other characteristics of the signal. Furthermore, data augmentation drastically improves computational performance of the event detection scheme without diminishing detection performance. The cost of the detection algorithm used in this paper is linear in the dimensionality of the feature set; that is, doubling the number of features reduces the performance by half.

`dupin` currently implements two different approaches to feature selection: (a) “mean-shift” filtering and (b) feature correlation. The main idea of mean-shift filtering is to compare two distributions and determine if they have a significantly different means from each other. For our purposes the distributions considered are the distributions of features in the system at the beginning and the end of the signal. A significant change in the mean of these features suggests that the feature effectively captures an event within the system. To perform this comparison, we analyze the means and standard deviations of both distributions to assess whether the means are equivalent. Based on this comparison and a predefined sensitivity parameter, we then make a *yes/no* decision about whether to include the descriptor for event detection. For more information see Supplementary Material Section 3. Note that the mean-shift filtering is different from a mean-shift cost function detection.

Feature selection via feature correlation is done through spectral clustering [43], where the similarity matrix is computed based on the correlation matrix of the signal. A pre-determined number of features from each cluster is taken based on a provided feature importance score or is randomly selected from a cluster. More information on these methods for feature selection can be found in Supplementary Material Section 3. Other feature selection methods such as forward selection or backward selection seamlessly interoperate with `dupin` and can be used as well. Both methods can be found in scikit-learn [44].

For dimensionality reduction, `dupin` implements a novel machine learning (ML) classifier to reduce the signal to a single dimension based on local signal similarity; this approach is a variation of the approach found in Reference [45]. Signal reduction is accomplished by taking a sliding window across the original signal and using an ensemble (collection) of weak learners to determine the similarity of the window's left and right halves. A weak learner is one with limited ability to discriminate between classes (e.g. a decision stump [46] that classifies based on a single *yes/no* condition). Thus, weak learners cannot distinguish window halves based solely on noise due to their limited discrimination ability. This limitation is a desirable property as we only want the classifier to discriminate on *significant* differences between window halves. Each learner in the ensemble is trained and tested on different data within the window via a stratified shuffle split (from scikit-learn), which reduces the chance that a bad data partition will decrease the test loss. To determine the local signal similarity, we label each frame in the left window half with the class label zero and each frame in the right window half with class label one. The classifiers within the ensemble are then trained on a subset of the data and tested on the remaining data. The testing loss (e.g. the zero-one loss) is used as a metric of dissimilarity. When there is no event, the classifiers should have nothing but random noise to train on, resulting in the classifier being wrong on test data $\sim 50\%$ of the time. However, when an event occurs within the window, a classifier can train on differences between the window halves and accuracy tends towards 100% (i.e. 0.0 loss). For each window position, we take the average loss as our one-dimensional signal (here, the

zero-one loss). Figure 3 shows a graphical depiction of the dimensionality reduction implemented in `dupin`.

The ML classifier has several hyperparameters, including the window size, number of classifiers (decision stumps) and the percentage of the data used for training. The window size must be chosen based on the number of samples (trajectory size). In general, bigger is better unless it leads to only a few windows or if it's much larger relative to the length of the transition. The choice of dumping frequency can also have important implications for the window size. If the dumping frequency of the feature is smaller compared to its relaxation time, the window size should be large enough that it contains several uncorrelated samples. This way we ensure that training is done on at least partially uncorrelated samples (random fluctuations). The number of classifiers determines smoothness of the error. Higher numbers naturally smooth the error across a trajectory and are generally preferable. The highest possible number of classifiers should be half of the non-repeating combinations of training and test set. This limit ensures that for small window sizes enough classifiers are given while being relatively confident that most window data splits will be unique. The test size is the percentage of the data used for testing. Larger test sizes provide better representation of the true representation. A consequence of this is that larger test sizes result in better error estimates and increased smoothness of the error. In general, test size should be smaller than half but not too small, to ensure comparison to correlated samples by chance is avoided.

Other schemes and algorithms for dimensionality reduction such as principal component analysis (PCA), uniform manifold approximation and projection (UMAP) [47], etc. can be used to reduce the number of features into a few information-dense dimensions as well.

Signal filtering is commonly used for smoothing a signal and is similar to the *mapping* step in the data generation stage. Here, the completed signal is transformed along the temporal dimension into a new signal. Signal filtering is commonly used for smoothing a signal. `dupin` has a rolling mean signal filter, which smooths noise by averaging the signal across neighboring frames. Other signal filters are available in packages like SciPy and can be used easily with `dupin`.

2.1.4 Detection

The final stage in `dupin`'s pipeline is event *detection*. To detect local events, two cost functions are available for use in `dupin`. Both are based on piecewise linear fits of time versus signal features, and both have increased cost when this fit has a higher summed p -norm error $|y - f(x)|_p$. In the examples in Section 3, we use the summed and square rooted 2-norm.

The first cost function, C_1 , in `dupin` computes the p -norm loss of the piecewise, least-squared linear fit of each feature as a function of time. We refer to this approach as linear cost function approach. This procedure minimizes the p -norm via a least-squared fit on the slope m and intercept b for the given sub-signal $S[i, j)$. To prevent units or feature magnitudes from contributing to faulty detection, we map all features to the unit square independently so that the range for each feature is $[0, 1]$. After mapping to the unit square, C_1 is given by:

$$C_1(S[i, j)) = \min_{m, b} \sum_{x=i}^j |S[x] - (mx + b)|_p, \quad (2)$$

where $\|\cdot\|_p$ is the p -norm, $S[x]$ is the signal value at frame x , and $mx + b$ is the linear fit of the signal where m is the slope of the linear fit and b is the y-intercept.

The second cost function, C_2 , is similar to the first, but the linear fit is determined simply by drawing a line from the beginning point to the end point of the sub-signal, without any fitting. We refer to this

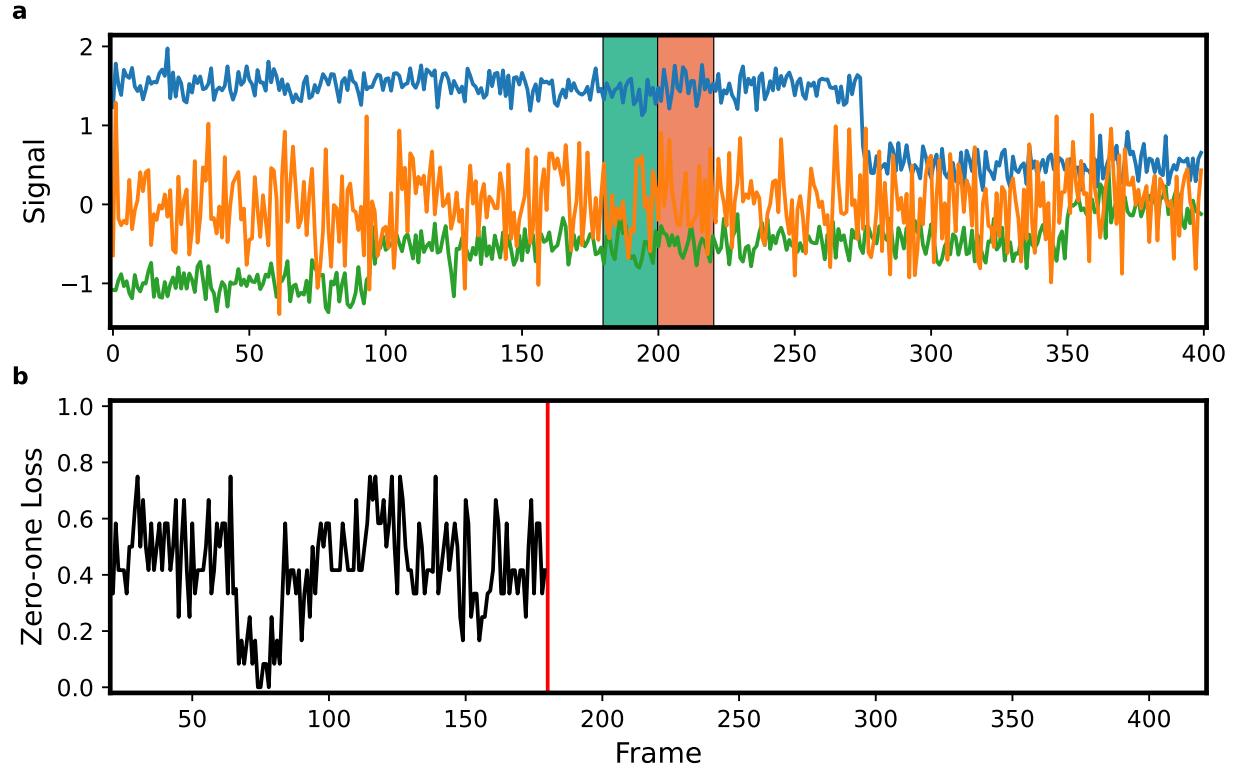


Figure 3: An example highlighting the use of classifiers to reduce a three-dimensional signal to one dimension. The process was stopped with a window center frame of 200 for instructive purposes. (a) A plot of the three-dimensional signal. Change points are located at 75, 275 and 350. The box centered at frame 200 represents the two window halves: $[190, 200)$ and $[200, 210)$. (b) The zero-one loss from the start of the signal to the current window. Notice the drop towards zero loss near frame 75, which corresponds to a mean-shift in (a) (blue line).

approach as biased cost function approach. As a result, this cost function is more sensitive to sudden shifts and changes in slope in the signal compared to C_1 . The C_2 cost function is given by:

$$\begin{aligned} C_2(S[i, j)) &= \sum_{x=i}^j |S[x] - (mx + b)|_p & (3) \\ m &= \frac{S[j] - S[i]}{j - i} \\ b &= S[i] - mi, \end{aligned}$$

The only optimizable parameters in Equation 3 are the locations of the change points. In cases where greater sensitivity is desired, C_2 is a viable alternative.

`dupin` can also detect events by using a simple mean-shift cost function detection approach. In this approach, a mean-shift cost function is computed instead of fitting linear functions to the signal. Such approach is especially effective when the ML classifier is used for the data augmentation step. Note that the mean-shift cost function detection is different from mean-shift filtering. For more details on this approach including the cost function definition see Supplementary Material Section ??.

Any detection algorithm can be used in `dupin`. For this work we use the Python package `ruptures` [41], which implements various algorithms for solving the optimization problem already posed (see Supplementary Material Section 1 for more information). `dupin` provides the implemented cost functions above to the detection class from `ruptures` (we use the dynamic programming solver in this work). Using `ruptures`, we find the optimal change points for a given number of change points.

After detection, we still have one more problem to solve: finding the optimum number of change points. To do this, we find an elbow in the total cost function as a function of the number of change points n . The elbow is defined as the point of maximum curvature, but can be expanded to discrete points. `dupin` can use any elbow detection method that works with discrete points. We choose for this work the kneedle algorithm [48] found in the `kneed` Python package; kneedle behaves well with the test systems studied and provides a hyperparameter allowing for sensitivity control in elbow detection. The sensitivity parameter also means the algorithm can return no elbow, allowing `dupin` to select zero change points as the correct CPD. More details on the kneedle algorithm and its usage in `dupin` can be found in Supplementary Material Section 2.

2.2 Online Detection

Another novelty enabled by `dupin` is the ability to perform online detection of events with approaches outlined above. Online detection is the process of detecting events while the particle trajectories are still being generated. This approach is useful in cases where responses to events are desired or when the generated trajectory is too large to be stored in memory. In online detection, the *detection* step is applied on-the-fly to the set of frames generated up until the moment of detection. This set of frames can extend to the beginning of the trajectory or be a sliding window of the last N generated frames. To use `dupin` online, a predefined set of features is calculated on-the-fly at the desired frequency leading to a continuous application of *generate* to *aggregate* steps. If using a sliding window rather than the whole trajectory, the order parameters are placed into a “first-in first-out” (FIFO) queue, so that the sliding window moves with the time evolution of the data. Sliding window is the preferred approach for online detection, due to improved performance

and easier CPD setup for detection. Using the queue approach allows CPD to be run on only part of the trajectory, which dramatically speeds up detection, making it more viable for online use. Each time a new frame is added to the signal, CPD can be run on the current data in the queue. The algorithm works best when the window size is commensurate with the size of a single event. Either way, `dupin` should be run with the assumption that we are expecting to detect up to one event. However, in practice we run CPD up to a change point set size of approximately four as any elbow detection scheme will need some points beyond the elbow to detect it. To detect multiple events, the queue should be cleared after detecting an event. If this is not done that event will continue to be detected in consecutive runs, leading to undesirable behavior.

3 Example Applications

We now demonstrate `dupin`'s event detection scheme step-by-step using `dupin` for two example systems to highlight its usefulness and versatility. All trajectory data was produced by molecular dynamics (MD) simulations run using HOOMD-blue [49, 50, 51]. Feature vector construction (*generation*) was carried out using `freud` [5, 6] in conjunction with `dupin`. The data was organized and managed using the `signac` framework [3, 52, 53]. System visualization was performed using OVITO [10].

3.1 Binary Lennard-Jones system

Our first example is a binary Lennard-Jones (LJ) system containing particles of type A and type B. This example demonstrates that `dupin` can be used to detect change points even if only a small portion of the system undergoes a transition. The MD simulations were performed in canonical (NVT) ensemble at temperature $T = 1.5$, density $\rho = 0.8$. System contained a total of $N_p = 2916$ particles of which 145 (5% of the system) were particles of type B, and the rest were particles of type A. The simulation starts off with all particles interacting with the same LJ potential ($\epsilon = 1$, $\sigma = 1$ with cutoff $r_{\text{cutoff}} = 2.5$). To prepare the system for production run we first place the particles on the simple cubic lattice, melt it, and equilibrate the system. Figure 4 (a) left shows the randomized and equilibrated system at the beginning of the production run. To generate the trajectory of point cloud data for detection we saved the system snapshot every 1000 time steps. The production part was run in two phases. In the first phase we ran for 1 million time steps. Next, the interaction between particles of type B is changed in such a way that LJ interaction potential parameters are changed to $\epsilon = 2.5$, $\sigma = 0.5$ without changing the cutoff. In second phase we ran for another million time steps. The production run generated $N_f = 2000$ frames in total.

We set up `dupin`'s detection pipeline to compute the Steinhardt order parameter for $l = 6$ (Q_6) with 12 nearest neighbors for each particle in the system as our *generate* step. This step results in $N_p \times N_f$ values of Q_6 in total. We represent this step as a distribution of Q_6 as a function of time in Figure 4 (b) 1. The *map* step is skipped and the *reduce* step is applied to find the 10th, and 50th greatest and least values of the Q_6 distribution. The features produced by the reduce step are kept and aggregated into the final signal which contains $N_f \times 4$ values in total (Figure 4 (b) 2.). The *transform* step was skipped because the dimensionality of the signal is already low. To *detect* change points in the trajectory we used the C_1 linear cost function. Elbow detection was used to compute the optimal number of change points (Figure 4 (b) 3.).

The detection scheme employed resulted in a single change point detected at frame 1030 (Figure 4 (a) middle). The detected change point corresponds to a frame occurring 30000 time steps after the interaction potential between particles of type B was changed. The employed change in the potential between particles of type B caused particles of type B to aggregate into clusters due to much stronger interaction. At frame

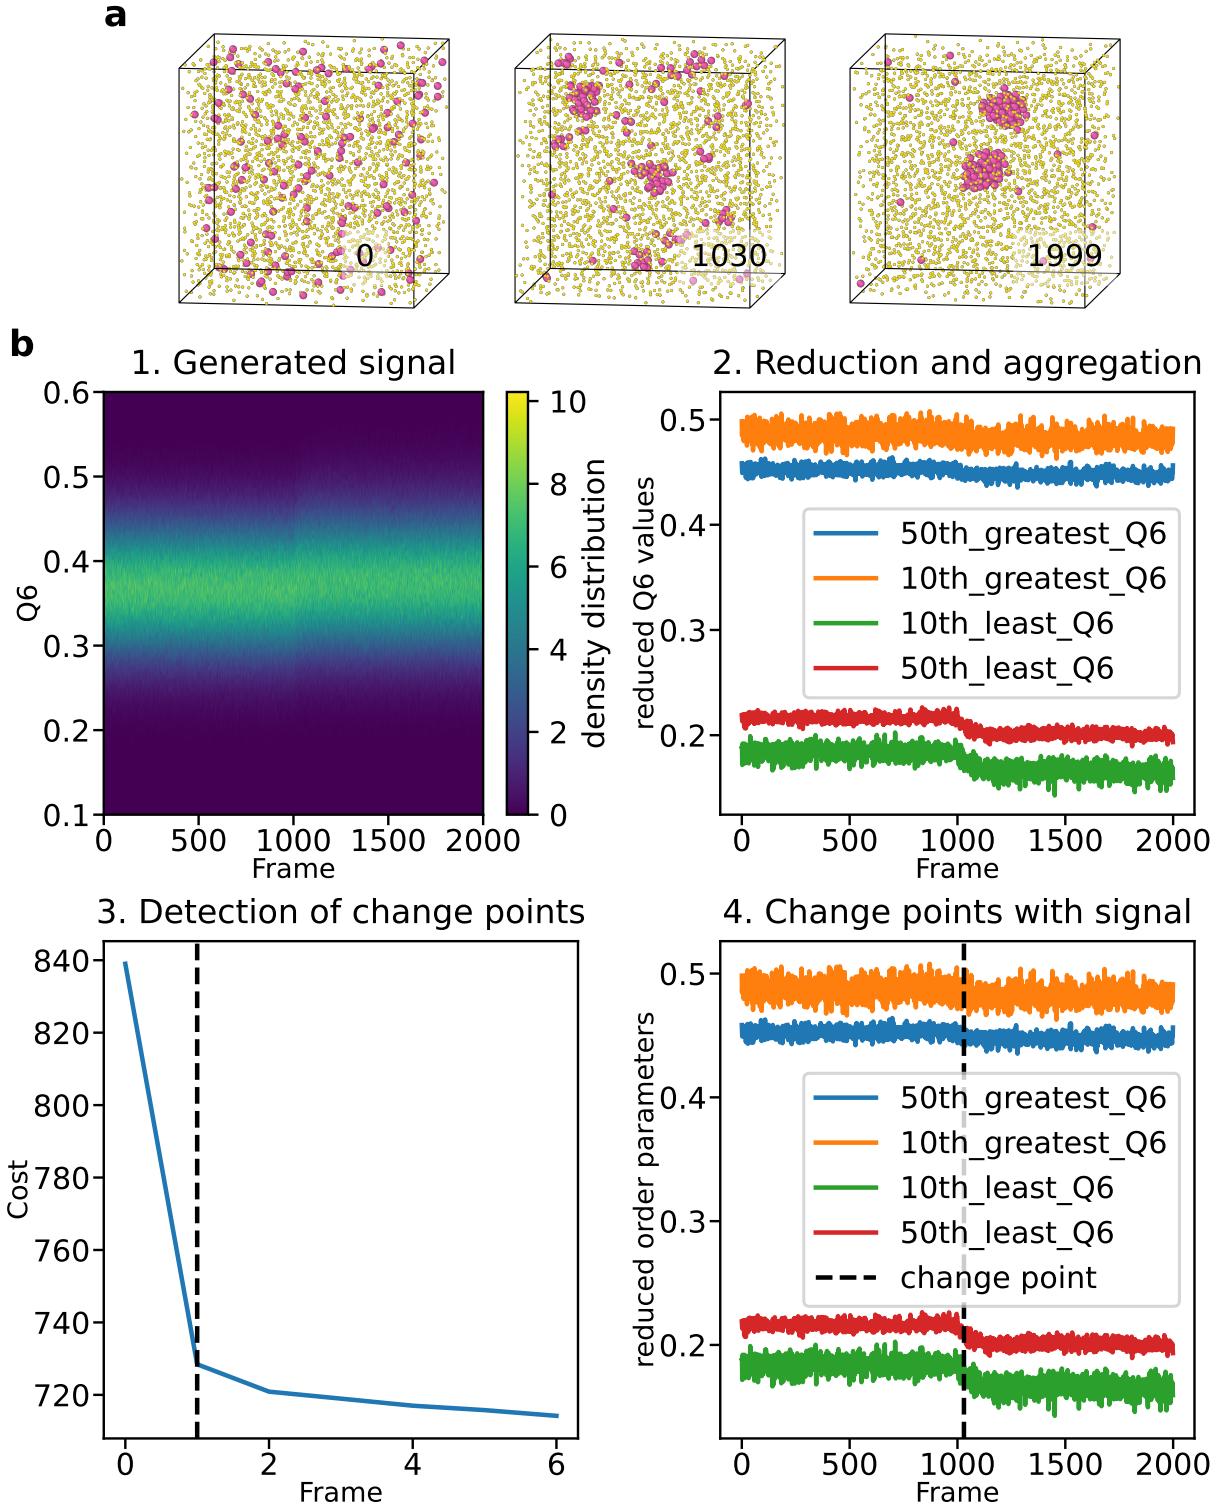


Figure 4: Change point detection applied to system of binary LJ particles. (a) Simulation snapshots with associated frame number (lower right). The middle image shows the system at the detected change point. Particles are colored according to their type (A - yellow; B - pink). (b) 1. The result of the signal generation step is represented as a density distribution of Q_6 for each frame of the trajectory. 2. After applying reduction and aggregation we are left with final signal on which the detection will be performed. 3. The cost plot used to compute the optimal number of change points using elbow detection (black dashed line) for linear change point detection. 4. Resulting change point (black dashed line) plotted on the signal.

1030 (Figure 4 (a) middle) we can clearly see that several smaller clusters of particles of type B have already formed. We can roughly estimate from Figure 4 (b) that the transition started at frame 1000, when the potential was changed and ended at around frame 1100. This places the detected change point squarely in the middle of the transition. The snapshot shown in Figure 4 (a) right, corresponds to the simulated system at the end of the simulation. We can see that that most of particles of type B have aggregated into two larger clusters by the end of the transition. This example showcases how `dupin`'s change point detection capabilities even if events are undergone by a small part of the system.

3.2 Nucleation and growth of a binary crystal of particles

Our second example is a binary system of point particles interacting via the Mie potential [54] ($n = 50$, $m = 25$) with a size ratio of 0.55. The simulation was run using MD for 36.8 million time steps in the isothermal isobaric ensemble (temperature $T = 0.35$, pressure $P = 0.052$, number of particles $N_p = 27,000$, in reduced units) to simulate the solidification of a liquid into a crystal by homogeneous nucleation and growth. The simulation was initialized in a NaCl lattice with random placement of the two species, which quickly dissolves upon thermalizing at slightly higher temperature prior to a quench to the target T .

To *generate* the signal we compute the Voronoi polyhedron volume [55, 56] and MSM [29] for spherical harmonics $l = 2, 4, 6, 8, 10, 12$ for each particle. For each feature, we *map* to two distributions — itself (no transformation) and the Voronoi tessellation neighbor average [31]. We then *reduce* each distribution (raw and averaged) to six features: the 1st, 10th, and 100th greatest and least values. After *reducing*, the MSM for each l produces twelve features, six from the raw distribution and six from spatial averaging. Following *aggregation*, we *transform* the signal via feature selection through a mean-shift filter with sensitivity of $10e-4$. After the mean-shift filter is applied we are left with the signal containing several features which change significantly during the length of the trajectory. Some of the features that pass the mean-shift filter are shown in Figure 5 (b). Next, we showcase two different detection approaches. In the first detection route, we take the filtered signal and *detect* the change points using C_1 (linear cost function) with `rupture`'s dynamic programming algorithm, using `kneed` for elbow detection with a sensitivity of 1, for $|K| \in [1, 10]$. For the second detection route, we take the filtered signal and we first *transform* the signal again by applying the ML classifier dimensionality reduction (window size 80) introduced in Subsection 2.1.3. We used 200 decision stumps (decision trees of depth one) on features selected by the mean-shift filter to classify windows halves. The test set size used was 0.4, which means that 40% of the data in the window was not used for fitting. The zero-one loss is then smoothed over the neighboring three errors on each side with the mean signal filter. The new signal obtained from the ML classifier dimensionality reduction route was then used to *detect* events using a L_1 mean-shift cost function [57].

The first detection route (Figure 5 (d) left) which *detects* using C_1 linear cost function detection results in two change points. The linear cost function detection was dominated by a region of sharp change in several properties ≈ 160 –200 during crystallization. We note that the continuing shift of some Voronoi polyhedra volumes or MSMs at the end of the simulation is roughly linear, meaning C_1 does not penalize grouping them into a single sub-signal.

On the other hand, the second detection route (Figure 5 (c) and (d) right) in which we first add another *transform* step using the ML classifier for dimensionality reduction (Figure 5 (c)) and then *detect* using mean-shift cost function (Figure 5 (d) right) results in a much larger transition event window (i.e. the change points are farther apart). Such a partitioning is expected as the dimensionality reduction scheme picks up on *any* deviation across the window. As a consequence, mean-shift cost function route detected

change points at frames 90 and 210, in contrast to the linear cost function route. We note that the reason for the change point locations can be seen in Figure 5 (c) where the smoothed average zero-one loss is plotted. The system is still undergoing structural changes at frame 210, indicated by the value of average loss which while higher than 0.0 is lower than the expected value outside a transition of 0.5. This behavior is explained by the observed trends in some Voronoi polyhedra volumes and MSMs, which are still not in equilibrium (flat) after frame 210. The final slice of the trajectory (frame 210 till the end) is thus associated with a new phase of the ongoing transition, which is not finished by the end of the trajectory as indicated by the average loss value.

These same schemes can be followed to apply **dupin** to simulations of molecules, nanoparticles, colloids, polymers, or other "particle"-based systems that generate particle positions (and possibly orientations) as a function of time, making **dupin** highly extensible and generalizable.

3.3 Collapse of polymer in poor solvent

Our third example involves detecting the collapse of an isolated polymer chain in an implicit solvent following a change from good to poor solvent using **dupin**. We simulated a polymer comprised of 5000 connected polymer beads interacting via a Lennard-Jones pairwise interaction, with bonds between beads modeled by a simple harmonic potential $U = -0.5k(r - r_0)$. The MD simulation was performed in the canonical ensemble (number of particles $N = 5,000$, system volume $V = 10^6$, and temperature $T = 1$). Following equilibration (1.2 million steps) in good solvent, the simulation was run for 2.5 million steps, where the ε of the Lennard-Jones potential was increased to mimic a change in the Flory χ parameter to a poor solvent. Over the 2.5 million steps, ε was increased from 0.1 to 1.25 over 50 intervals of 30,000 steps, running at the final ε for 1 million steps. Increasing the ε parameter causes the polymer beads to start aggregating. Several stages of aggregation are observed in which the number of lobes decreases in stages until final metastable configuration of two lobes is obtained.

We *generate* the signal using the polymer end-to-end distance, the local density defined as Voronoi polytope volume, and the number of clusters of neighboring beads using **freud**'s clustering algorithm ($r = 1.2$). Because Voronoi polytope volume is the only non-scalar feature, we *reduce* only it. We *reduce* the volumes to six features: the 10st, 100th, and 1,000th greatest and least values. We do not *transform* the data due to low dimensionality. Finally, for offline detection, we *detect* the change points using C_1 with **rupture**'s dynamic programming algorithm, using **kneed** for elbow detection with a sensitivity of 1, for $|K| \in [1, 10]$. We also perform online *detection* using the sliding window approach with a window size of 50. All steps up to *detection* are identical though the signal is fed to the detector as a stream rather than simultaneously. The detector used for the online *detection* uses the same cost function, algorithm, and elbow detection, but only goes to $|K| = 6$. The cost function gets computed on-the-fly for each frame in the window, taking into account previous frames in the window. We clear the window whenever an event is detected.

Figure 6 shows the CPD analysis for this system. The collapse of the polymer due to poor solvent can be seen in Figure 6 (a). The collapse *looks* visually continuous, although the images show a collapse mediated by multiple intermediate, discrete steps, and the shape of the order parameter with time is sigmoidal. Figure 6 (c - left) shows the cost associated with the choice of optimum n (number of change points) for offline detection. The offline detection scheme only detects the transition to the final snapshot of the simulation because fitting to the individual sections of the collapse does not sufficiently decrease the cost function.

Figure 6 (c - right) shows **dupin**'s ability to detect the collapse and its component parts in the case of

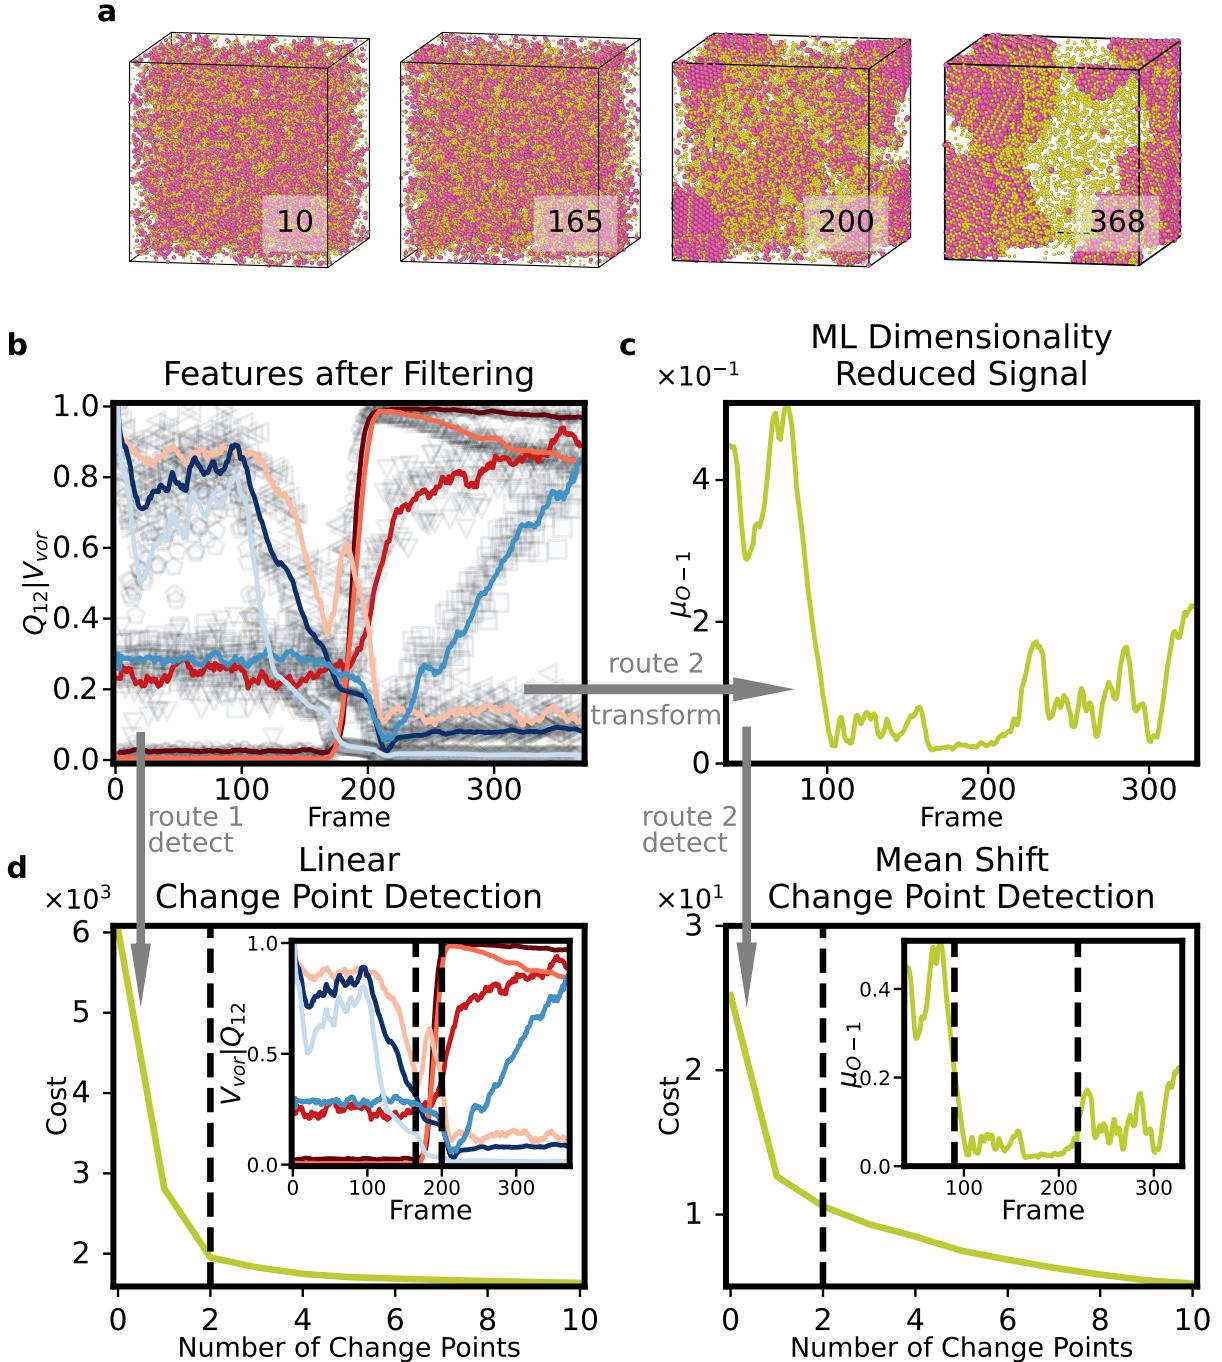


Figure 5: Change point detection applied to system of binary Mie particles. (a) Simulation snapshots with associated frame number (lower right). The two middle images show the system at the two change points detected using linear change point detection. Particles are colored according to their type. (b) Plot of some features obtained after the transform step (100th greatest and least MSM $l = 12$ (left) and Voronoi polyhedra volumes (right)). Solid lines are the (rolling mean) smoothed features and the translucent points are the actual data points used for calculations. Gray arrows and text indicate two different change point detection routes applied. (c) The zero-one loss function computed using the ML classifier for dimensionality reduction applied to features from transform step (b). (d) The two main figures show the cost plots used to compute the optimal number of change points with elbow detection (black dashed line) for linear change point detection (left) and mean-shift change point detection (right). The insets show the change point frames (black dashed line) on the signals used to compute the change points.

online detection. For online detection the elbow cost plot would need to be shown for every frame. Instead, we show the relative improvement of cost inside the current window from adding a second change point in online detection. The formula used is $\xi = (c_1 - c_2)/c_0$ where c_i is the cost for selecting i change points. This approach yields independent cost curves for each window (detected event). When this proxy is high, it indicates that the current window is likely to contain a transition. The “ramp up”, “growth” and “slow down” behavior of the sigmoid leads to three events detected with online detection at frames 65, 95, and 140. This granularity results from the max signal length of 50, which increases the relative reduction in cost for fitting to the three sections of the polymer collapse. Processing the entire trajectory for online detection took 820 milliseconds ($\pm 3.12\mu s$) on a single core of a 3.0 GHz Intel Xeon Gold 6154. This speed is sufficiently fast to use in real time applications such as autonomously triggering simulation protocols during a simulation.

4 Discussion and Conclusions

We’ve demonstrated how the procedure described in Section 2.1 allows for detection of transition points within a simulation trajectory with a high degree of accuracy. The obvious benefit from this approach is the automation of structural transition detection within a study. In studies with hundreds and thousands of simulations the dividends of this approach increase exponentially. We have also demonstrated that **dupin** is successful at detecting events underpinned by a small number of particles in the system (< 5%). Our method does, however, require informative descriptors for the transition. This requirement can be met by selecting a wide variety of descriptors and applying feature selection tools to the signal afterwards.

From the examples presented, we conclude that the ML dimensionality reduction classifier produces a signal which results in a larger partitioning of the detected events compared to results produced using the linear cost function approach. However, the ML dimensionality reduction classifier approach falls short in cases where multiple events are overlapping or in cases where there are abrupt changes at the end of the signal. For instance, if we were to apply the ML dimensionality reduction classifier approach to a system with only one frame into a transition, no event would be detected. The linear approach is better suited for detection in such scenarios and in scenarios where multiple events are expected as showcased in the online polymer example.

dupin opens the doors for new ML applications to phenomena such as crystallization pathways, defect formation and active matter, all of which involve structural transitions. By curating the transition data for ML applications, large scale studies that would have been prohibitively costly in terms of human hours are now accessible. Furthermore, leveraging **dupin** for online event detection holds promise to lessen data storage and processing demands and provide a powerful avenue for real-time control over simulations and experiments.

The source code can be found at GitHub at <https://github.com/glotzerlab/dupin>. The documentation hosted by Read the Docs, <https://dupin.readthedocs.io>, also contains three additional examples of event detection: one for a simple example of WCA spheres forming FCC, one of hard truncated tetrahedra forming cF432, and one of ionic beryllium and chlorine forming a nematic then crystalline phase. The analysis code and data will be available on Deep Blue Documents following publication.

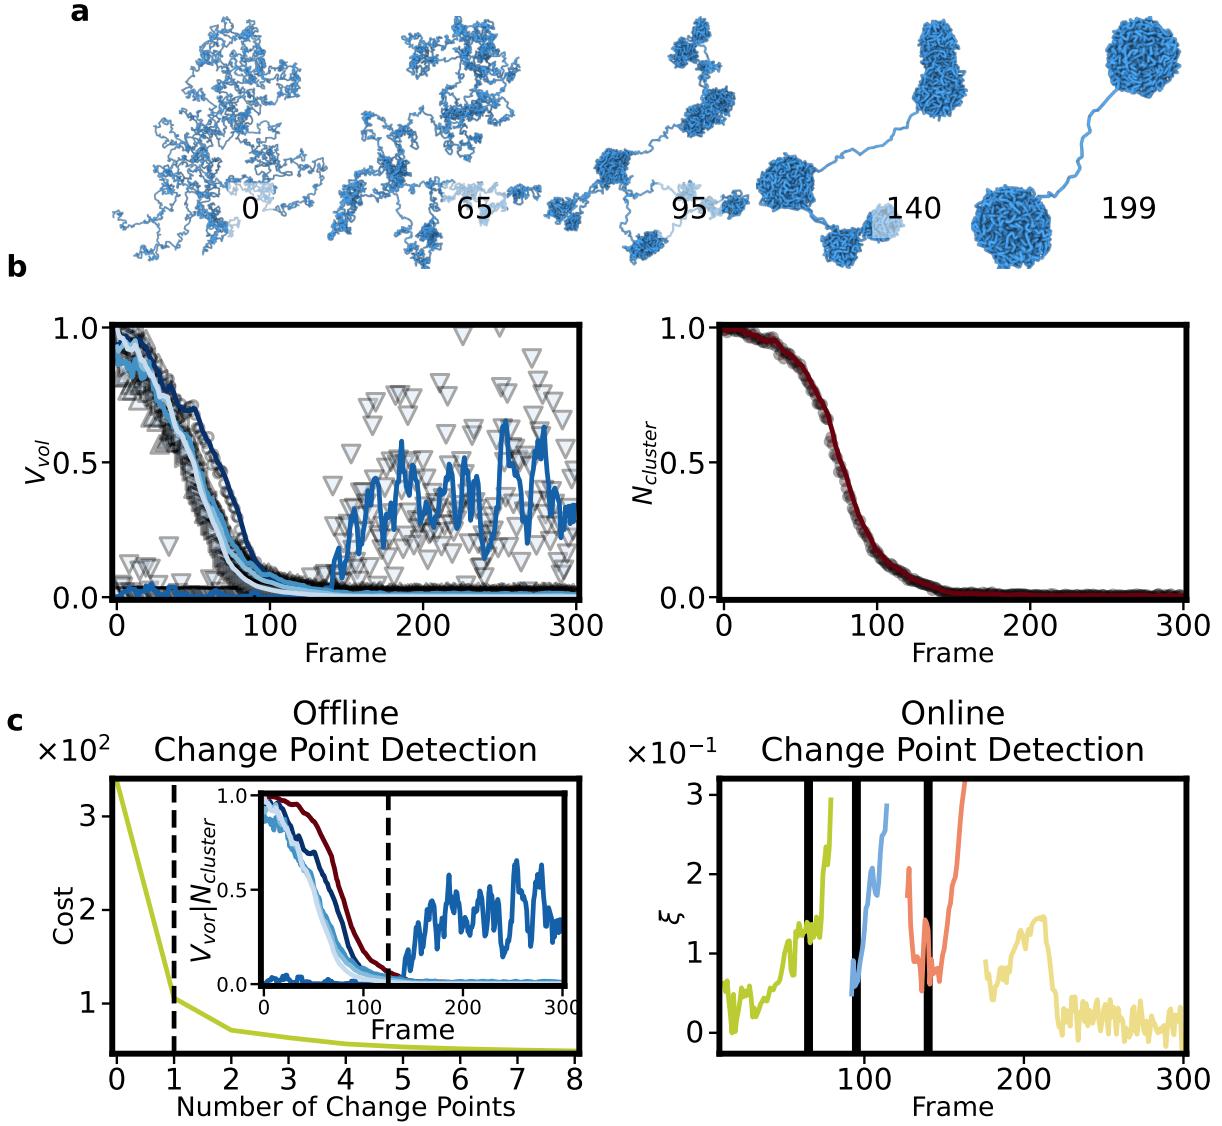


Figure 6: Change point detection applied to polymer trajectory. (a) Images corresponding to the beginning of the simulation, online change points, and end of the simulation. (b) Plots of features for detection: Voronoi polytope volumes (left) and the number of clusters (right). Solid lines are the (rolling mean) smoothed features and the translucent points are the data points on which all calculations were done. (c) Detection of change points using the offline (left) and online approach (right). Offline plot (left) of the cost associated with the optimum n change points computed on all features. The inset for the offline plot (left) shows the features from (b) with the computed change points (black dashed line). Online detection (right): Window cost proxy function (see main text) for online detection and associated change points (black lines).

5 Acknowledgements

This research was supported by the National Science Foundation, Division of Materials Research under a Computational and Data-Enabled Science and Engineering (CDS&E) Award # DMR 2302470 (2023 - 2027) and # DMR 1808342 (2019-2023). This work used the Extreme Science and Engineering Discovery Environment (XSEDE) [58], which is supported by National Science Foundation grant number ACI-1548562; XSEDE award DMR 140129. This research was supported in part through computational resources and services provided by Advanced Research Computing, a division of Information and Technology Services at the University of Michigan, Ann Arbor.

References

- [1] Stephen Thomas et al. “Routine Million-Particle Simulations of Epoxy Curing with Dissipative Particle Dynamics”. In: *J. Theor. Comput. Chem.* 17.03 (May 1, 2018), p. 1840005. ISSN: 0219-6336. DOI: 10.1142/S0219633618400059. URL: <https://www.worldscientific.com/doi/abs/10.1142/S0219633618400059> (visited on 01/31/2022).
- [2] Matthew W. Thompson et al. “Scalable Screening of Soft Matter: A Case Study of Mixtures of Ionic Liquids and Organic Solvents”. In: *J. Phys. Chem. B* 123.6 (Feb. 14, 2019), pp. 1340–1347. ISSN: 1520-6106. DOI: 10.1021/acs.jpcb.8b11527. URL: <https://doi.org/10.1021/acs.jpcb.8b11527> (visited on 11/03/2020).
- [3] Carl S. Adorf et al. “Simple Data and Workflow Management with the Signac Framework”. In: *Computational Materials Science* 146 (Apr. 15, 2018), pp. 220–229. ISSN: 0927-0256. DOI: 10.1016/j.commatsci.2018.01.035. URL: <https://www.sciencedirect.com/science/article/pii/S0927025618300429> (visited on 01/28/2022).
- [4] Sebastiaan P. Huber et al. “AiiDA 1.0, a Scalable Computational Infrastructure for Automated Reproducible Workflows and Data Provenance”. In: *Sci Data* 7.1 (1 Sept. 8, 2020), p. 300. ISSN: 2052-4463. DOI: 10.1038/s41597-020-00638-4. URL: <http://www.nature.com/articles/s41597-020-00638-4> (visited on 04/14/2022).
- [5] Vyas Ramasubramani et al. “Freud: A Software Suite for High Throughput Analysis of Particle Simulation Data”. In: *Computer Physics Communications* (Mar. 17, 2020), p. 107275. ISSN: 0010-4655. DOI: 10.1016/j.cpc.2020.107275. URL: <http://www.sciencedirect.com/science/article/pii/S0010465520300916> (visited on 05/04/2020).
- [6] Bradley D. Dice et al. “Analyzing Particle Systems for Machine Learning and Data Visualization with Freud”. In: *Proceedings of the 18th Python in Science Conference* (2019), pp. 27–33. DOI: 10.25080/Majora-7ddc1dd1-004. URL: http://conference.scipy.org/proceedings/scipy2019/bradley_dice.html (visited on 01/28/2022).
- [7] Robert T. McGibbon et al. “MDTraj: A Modern Open Library for the Analysis of Molecular Dynamics Trajectories”. In: *Biophysical Journal* 109.8 (2015), pp. 1528–1532. DOI: 10.1016/j.bpj.2015.08.015.
- [8] Naveen Michaud-Agrawal et al. “MDAnalysis: A toolkit for the analysis of molecular dynamics simulations”. In: *Journal of Computational Chemistry* 32.10 (Apr. 2011), pp. 2319–2327. DOI: 10.1002/jcc.21787. URL: <https://doi.org/10.1002%2Fjcc.21787>.

- [9] Richard Gowers et al. “MDAnalysis: A Python Package for the Rapid Analysis of Molecular Dynamics Simulations”. In: *Proceedings of the Python in Science Conference*. SciPy, 2016. DOI: 10.25080/majora-629e541a-00e. URL: <https://doi.org/10.25080%2Fmajora-629e541a-00e>.
- [10] Alexander Stukowski. “Visualization and Analysis of Atomistic Simulation Data with OVITO—the Open Visualization Tool”. In: *Modelling Simul. Mater. Sci. Eng.* 18.1 (Dec. 2009), p. 015012. ISSN: 0965-0393. DOI: 10.1088/0965-0393/18/1/015012. URL: <https://doi.org/10.1088%2F0965-0393%2F18%2F1%2F015012> (visited on 12/18/2019).
- [11] Samaneh Aminikhanghahi and Diane J. Cook. “A Survey of Methods for Time Series Change Point Detection”. In: *Knowl Inf Syst* 51.2 (May 1, 2017), pp. 339–367. ISSN: 0219-3116. DOI: 10.1007/s10115-016-0987-z. URL: <https://doi.org/10.1007/s10115-016-0987-z> (visited on 01/28/2022).
- [12] Binbin Luo et al. “Quantifying the Self-Assembly Behavior of Anisotropic Nanoparticles Using Liquid-Phase Transmission Electron Microscopy”. In: *Acc. Chem. Res.* 50.5 (May 16, 2017), pp. 1125–1133. ISSN: 0001-4842. DOI: 10.1021/acs.accounts.7b00048. URL: <https://doi.org/10.1021/acs.accounts.7b00048> (visited on 01/09/2023).
- [13] Zihao Ou et al. “Kinetic Pathways of Crystallization at the Nanoscale”. In: *Nat. Mater.* 19.4 (4 Apr. 2020), pp. 450–455. ISSN: 1476-4660. DOI: 10.1038/s41563-019-0514-1. URL: <http://www.nature.com/articles/s41563-019-0514-1> (visited on 01/09/2023).
- [14] C. Patrick Royall, Ard A. Louis, and Hajime Tanaka. “Measuring Colloidal Interactions with Confocal Microscopy”. In: *J. Chem. Phys.* 127.4 (July 28, 2007), p. 044507. ISSN: 0021-9606. DOI: 10.1063/1.2755962. URL: <http://aip.scitation.org/doi/full/10.1063/1.2755962> (visited on 01/09/2023).
- [15] Salma M. Abdel-Hafez, Rania M. Hathout, and Omaima A. Sammour. “Tracking the Transdermal Penetration Pathways of Optimized Curcumin-Loaded Chitosan Nanoparticles via Confocal Laser Scanning Microscopy”. In: *International Journal of Biological Macromolecules* 108 (Mar. 1, 2018), pp. 753–764. ISSN: 0141-8130. DOI: 10.1016/j.ijbiomac.2017.10.170. URL: <https://www.sciencedirect.com/science/article/pii/S0141813017329100> (visited on 01/09/2023).
- [16] Samaneh Aminikhanghahi and Diane J. Cook. “Using change point detection to automate daily activity segmentation”. In: *2017 IEEE International Conference on Pervasive Computing and Communications Workshops (PerCom Workshops)*. 2017, pp. 262–267. DOI: 10.1109/PERCOMW.2017.7917569.
- [17] Daniel Alves, Katia Obraczka, and Rick Lindberg. “Identifying Relevant Data Center Telemetry Using Change Point Detection”. In: *2020 IEEE 9th International Conference on Cloud Networking (Cloud-Net)*. 2020, pp. 1–4. DOI: 10.1109/CloudNet51028.2020.9335800.
- [18] Zunya Shi and Abdallah Chehade. “A dual-LSTM framework combining change point detection and remaining useful life prediction”. In: *Reliability Engineering & System Safety* 205 (2021), p. 107257. ISSN: 0951-8320. DOI: <https://doi.org/10.1016/j.ress.2020.107257>. URL: <https://www.sciencedirect.com/science/article/pii/S0951832020307572>.
- [19] Kyle D. Feuz et al. “Automated Detection of Activity Transitions for Prompting”. In: *IEEE Transactions on Human-Machine Systems* 45.5 (2015), pp. 575–585. DOI: 10.1109/THMS.2014.2362529.
- [20] Manhyung Han et al. “Comprehensive Context Recognizer Based on Multimodal Sensors in a Smartphone”. In: *Sensors* 12.9 (Sept. 2012), pp. 12588–12605. ISSN: 1424-8220. DOI: 10.3390/s120912588. URL: <http://dx.doi.org/10.3390/s120912588>.

- [21] Sasank Reddy et al. “Using Mobile Phones to Determine Transportation Modes”. In: *ACM Trans. Sen. Netw.* 6.2 (Mar. 2010). ISSN: 1550-4859. DOI: 10.1145/1689239.1689243. URL: <https://doi.org/10.1145/1689239.1689243>.
- [22] Alexander Aue et al. “Break detection in the covariance structure of multivariate time series models”. In: *The Annals of Statistics* 37.6B (2009), pp. 4046–4087. DOI: 10.1214/09-AOS707. URL: <https://doi.org/10.1214/09-AOS707>.
- [23] Daniel Barry and J. A. Hartigan. “A Bayesian Analysis for Change Point Problems”. In: *Journal of the American Statistical Association* 88.421 (1993), pp. 309–319. ISSN: 01621459.
- [24] Yoshinobu Kawahara and Masashi Sugiyama. “Sequential change-point detection based on direct density-ratio estimation”. In: *Statistical Analysis and Data Mining: The ASA Data Science Journal* 5.2 (2012), pp. 114–127.
- [25] Yoshinobu Kawahara, Takehisa Yairi, and Kazuo Machida. “Change-Point Detection in Time-Series Data Based on Subspace Identification”. In: *Seventh IEEE International Conference on Data Mining (ICDM 2007)*. 2007, pp. 559–564. DOI: 10.1109/ICDM.2007.78.
- [26] E. Keogh et al. “An online algorithm for segmenting time series”. In: *Proceedings 2001 IEEE International Conference on Data Mining*. 2001, pp. 289–296. DOI: 10.1109/ICDM.2001.989531.
- [27] Marcel Bosc et al. “Automatic change detection in multimodal serial MRI: application to multiple sclerosis lesion evolution”. In: *NeuroImage* 20.2 (2003), pp. 643–656. ISSN: 1053-8119. DOI: [https://doi.org/10.1016/S1053-8119\(03\)00406-3](https://doi.org/10.1016/S1053-8119(03)00406-3).
- [28] Paul J. Steinhardt, David R. Nelson, and Marco Ronchetti. “Bond-Orientational Order in Liquids and Glasses”. In: *Phys. Rev. B* 28.2 (July 15, 1983), pp. 784–805. DOI: 10.1103/PhysRevB.28.784. URL: <https://link.aps.org/doi/10.1103/PhysRevB.28.784> (visited on 05/11/2020).
- [29] Walter Mickel et al. “Shortcomings of the Bond Orientational Order Parameters for the Analysis of Disordered Particulate Matter”. In: *J. Chem. Phys.* 138.4 (Jan. 22, 2013), p. 044501. ISSN: 0021-9606. DOI: 10.1063/1.4774084. URL: <http://aip.scitation.org/doi/full/10.1063/1.4774084> (visited on 11/12/2020).
- [30] Jeremy C. Palmer et al. “Metastable liquid-liquid transition in a molecular model of water”. In: *Nature* 510 (2014), pp. 385–388. DOI: 10.1038/nature13405.
- [31] Wolfgang Lechner and Christoph Dellago. “Accurate Determination of Crystal Structures Based on Averaged Local Bond Order Parameters”. In: *J. Chem. Phys.* 129.11 (Sept. 21, 2008), p. 114707. ISSN: 0021-9606. DOI: 10.1063/1.2977970. URL: <https://aip.scitation.org/doi/full/10.1063/1.2977970> (visited on 05/05/2022).
- [32] Albert P. Bartók, Risi Kondor, and Gábor Csányi. “On representing chemical environments”. In: *Phys. Rev. B* 87 (18 May 2013), p. 184115. DOI: 10.1103/PhysRevB.87.184115. URL: <https://link.aps.org/doi/10.1103/PhysRevB.87.184115>.
- [33] Cristina Caruso et al. “TimeSOAP: Tracking high-dimensional fluctuations in complex molecular systems via time variations of SOAP spectra”. In: *The Journal of Chemical Physics* 158.21 (June 2023), p. 214302. ISSN: 0021-9606. DOI: 10.1063/5.0147025. URL: <https://doi.org/10.1063/5.0147025>.
- [34] Bradley Dice. “Complex Crystallization Pathways Analyzed in a Continuous Feature Space”. Thesis. 2021. DOI: 10.7302/3103. URL: <http://deepblue.lib.umich.edu/handle/2027.42/170058> (visited on 04/06/2022).

- [35] Emanuele Boattini, Marjolein Dijkstra, and Laura Filion. “Unsupervised Learning for Local Structure Detection in Colloidal Systems”. In: *The Journal of Chemical Physics* 151.15 (Oct. 15, 2019), p. 154901. ISSN: 0021-9606. DOI: 10.1063/1.5118867. URL: <http://aip.scitation.org/doi/abs/10.1063/1.5118867> (visited on 11/23/2020).
- [36] S. S. Schoenholz et al. “A Structural Approach to Relaxation in Glassy Liquids”. In: *Nature Phys* 12.5 (May 2016), pp. 469–471. ISSN: 1745-2481. DOI: 10.1038/nphys3644. URL: <https://www.nature.com/articles/nphys3644> (visited on 12/18/2019).
- [37] Carl S. Adorf et al. “Analysis of Self-Assembly Pathways with Unsupervised Machine Learning Algorithms”. In: *J. Phys. Chem. B* 124.1 (Jan. 9, 2020), pp. 69–78. ISSN: 1520-6106. DOI: 10.1021/acs.jpcb.9b09621. URL: <https://doi.org/10.1021/acs.jpcb.9b09621> (visited on 10/22/2020).
- [38] C. Dietz, T. Kretz, and M. H. Thoma. “Machine-Learning Approach for Local Classification of Crystalline Structures in Multiphase Systems”. In: *Phys. Rev. E* 96.1 (July 19, 2017), p. 011301. DOI: 10.1103/PhysRevE.96.011301. URL: <https://link.aps.org/doi/10.1103/PhysRevE.96.011301> (visited on 10/07/2020).
- [39] Matthew Spellings and Sharon C. Glotzer. “Machine Learning for Crystal Identification and Discovery”. In: *AIChE Journal* 64.6 (2018), pp. 2198–2206. ISSN: 1547-5905. DOI: 10.1002/aic.16157. URL: <https://aiche.onlinelibrary.wiley.com/doi/abs/10.1002/aic.16157> (visited on 11/25/2019).
- [40] Trent Barnard and Gabriele C. Sosso. “Combining machine learning and molecular simulations to predict the stability of amorphous drugs”. In: *The Journal of Chemical Physics* 159.1 (July 2023), p. 014503. ISSN: 0021-9606. DOI: 10.1063/5.0156222. URL: <https://doi.org/10.1063/5.0156222>.
- [41] Charles Truong, Laurent Oudre, and Nicolas Vayatis. “A Review of Change Point Detection Methods.” In: (Jan. 1, 2018). URL: https://openreview.net/forum?id=4_eVexSJ8B (visited on 01/28/2022).
- [42] S. Karthika, T. K. Radhakrishnan, and P. Kalaichelvi. “A Review of Classical and Nonclassical Nucleation Theories”. In: *Crystal Growth & Design* 16.11 (Nov. 2, 2016), pp. 6663–6681. ISSN: 1528-7483. DOI: 10.1021/acs.cgd.6b00794. URL: <https://doi.org/10.1021/acs.cgd.6b00794> (visited on 10/22/2020).
- [43] Andrew Y. Ng, Micheal I. Jorden, and Yair Weiss. “On Spectral Clustering: Analysis and an Algorithm”. In: *Advances in Neural Information Processing Systems*. Neural Information Processing Systems. 2002, pp. 849–856.
- [44] Fabian Pedregosa et al. “Scikit-Learn: Machine Learning in Python”. In: *Journal of Machine Learning Research* 12.85 (2011), pp. 2825–2830. ISSN: 1533-7928. URL: <http://jmlr.org/papers/v12/pedregosa11a.html> (visited on 05/09/2020).
- [45] Shohei Hido et al. “Unsupervised Change Analysis Using Supervised Learning”. In: *Advances in Knowledge Discovery and Data Mining*. Ed. by Takashi Washio et al. Lecture Notes in Computer Science. Berlin, Heidelberg: Springer, 2008, pp. 148–159. ISBN: 978-3-540-68125-0. DOI: 10.1007/978-3-540-68125-0_15.
- [46] Wayne Iba and Pat Langley. “Induction of One-Level Decision Trees”. In: *Machine Learning Proceedings 1992*. Ed. by Derek Sleeman and Peter Edwards. San Francisco (CA): Morgan Kaufmann, 1992, pp. 233–240. ISBN: 978-1-55860-247-2. DOI: <https://doi.org/10.1016/B978-1-55860-247-2.50035-8>. URL: <https://www.sciencedirect.com/science/article/pii/B9781558602472500358>.

- [47] Leland McInnes, John Healy, and James Melville. "UMAP: Uniform Manifold Approximation and Projection for Dimension Reduction". Sept. 17, 2020. arXiv: 1802.03426 [cs, stat]. URL: <http://arxiv.org/abs/1802.03426>.
- [48] Ville Satopaa et al. "Finding a "Kneedle" in a Haystack: Detecting Knee Points in System Behavior". In: *2011 31st International Conference on Distributed Computing Systems Workshops*. 2011 31st International Conference on Distributed Computing Systems Workshops. June 2011, pp. 166–171. DOI: 10.1109/ICDCSW.2011.20.
- [49] Joshua A. Anderson, Jens Glaser, and Sharon C. Glotzer. "HOOMD-blue: A Python Package for High-Performance Molecular Dynamics and Hard Particle Monte Carlo Simulations". In: *Computational Materials Science* 173 (Feb. 15, 2020), p. 109363. ISSN: 0927-0256. DOI: 10.1016/j.commatsci.2019.109363. URL: <http://www.sciencedirect.com/science/article/pii/S0927025619306627> (visited on 05/09/2020).
- [50] Joshua A. Anderson, M. Eric Irrgang, and Sharon C. Glotzer. "Scalable Metropolis Monte Carlo for Simulation of Hard Shapes". In: *Computer Physics Communications* 204 (July 1, 2016), pp. 21–30. ISSN: 0010-4655. DOI: 10.1016/j.cpc.2016.02.024. URL: <https://www.sciencedirect.com/science/article/pii/S001046551630039X> (visited on 01/31/2022).
- [51] Brandon L. Butler et al. "HOOMD-blue version 3.0 A Modern, Extensible, Flexible, Object-Oriented API for Molecular Simulations". In: *Proceedings of the 19th Python in Science Conference*. Ed. by Meghann Agarwal et al. 2020, pp. 24–31. DOI: 10.25080/Majora-342d178e-004.
- [52] Vyas Ramasubramani et al. "Signac: A Python Framework for Data and Workflow Management". In: *Proceedings of the 17th Python in Science Conference* (2018), pp. 152–159. DOI: 10.25080/Majora-4af1f417-016. URL: https://conference.scipy.org/proceedings/scipy2018/vyas_ramasubramani.html (visited on 01/28/2022).
- [53] Bradley D. Dice et al. "Signac: Data Management and Workflows for Computational Researchers". In: *Proceedings of the 20th Python in Science Conference* (2021), pp. 23–32. DOI: 10.25080/majora-1b6fd038-003. URL: https://conference.scipy.org/proceedings/scipy2021/bradley_dice.html (visited on 01/28/2022).
- [54] Gustav Mie. "Zur Kinetischen Theorie Der Einatomigen Körper". In: *Annalen der Physik* 316.8 (1903), pp. 657–697. ISSN: 1521-3889. DOI: 10.1002/andp.19033160802. URL: <http://onlinelibrary.wiley.com/doi/10.1002/andp.19033160802> (visited on 01/09/2023).
- [55] Georges Voronoi. "Nouvelles applications des paramètres continus à la théorie des formes quadratiques. Premier mémoire. Sur quelques propriétés des formes quadratiques positives parfaites." In: *Journal für die reine und angewandte Mathematik (Crelles Journal)* 1908.133 (Jan. 1, 1908), pp. 97–102. ISSN: 1435-5345. DOI: 10.1515/crll.1908.133.97. URL: <https://www.degruyter.com/document/doi/10.1515/crll.1908.133.97/html> (visited on 01/18/2023).
- [56] Georges Voronoi. "Nouvelles applications des paramètres continus à la théorie des formes quadratiques. Deuxième mémoire. Recherches sur les paralléloèdres primitifs." In: *Journal für die reine und angewandte Mathematik (Crelles Journal)* 1908.134 (July 1, 1908), pp. 198–287. ISSN: 1435-5345. DOI: 10.1515/crll.1908.134.198. URL: <https://www.degruyter.com/document/doi/10.1515/crll.1908.134.198/html> (visited on 01/18/2023).

- [57] Venkata Jandhyala et al. “Inference for Single and Multiple Change-Points in Time Series”. In: *Journal of Time Series Analysis* 34.4 (2013), pp. 423–446. ISSN: 1467-9892. DOI: 10.1111/jtsa.12035. URL: <http://onlinelibrary.wiley.com/doi/abs/10.1111/jtsa.12035> (visited on 01/09/2023).
- [58] John Towns et al. “XSEDE: Accelerating Scientific Discovery”. In: *Computing in Science & Engineering* 16.5 (2014), pp. 62–74. DOI: 10.1109/MCSE.2014.80.