RESEARCH ARTICLE | NOVEMBER 20 2018

Unsupervised machine learning for detection of phase transitions in off-lattice systems. I. Foundations

R. B. Jadrich; B. A. Lindquist; T. M. Truskett

J. Chem. Phys. 149, 194109 (2018)
https://doi.org/10.1063/1.5049849
CHORUS

THE JOURNAL OF CHEMICAL PHYSICS **149**, 194109 (2018)

Unsupervised machine learning for detection of phase transitions in off lattice systems. I. Foundations

R. B. Jadrich,[1] B. A. Lindquist,[1] and T. M. Truskett[1,2,a)]
[1]*McKetta Department of Chemical Engineering University of Texas at Austin Austin Texas 78712 USA*
[2]*Department of Physics University of Texas at Austin Austin Texas 78712 USA*

(Received 26 July 2018; accepted 23 October 2018; published online 20 November 2018)

We demonstrate the utility of an unsupervised machine learning tool for the detection of phase transitions in off-lattice systems. We focus on the application of principal component analysis (PCA) to detect the freezing transitions of two-dimensional hard-disk and three-dimensional hard-sphere systems as well as liquid-gas phase separation in a patchy colloid model. As we demonstrate, PCA autonomously discovers order-parameter-like quantities that report on phase transitions, mitigating the need for *a priori* construction or identification of a suitable order parameter—thus streamlining the routine analysis of phase behavior. In a companion paper, we further develop the method established here to explore the detection of phase transitions in various model systems controlled by compositional demixing, liquid crystalline ordering, and non-equilibrium active forces. *Published by AIP Publishing.* https://doi.org/10.1063/1.5049849

I. INTRODUCTION

Phase transitions, diverse in character and ubiquitous in physical and biological systems, result from the correlated response of a near-infinite number of interacting microscopic degrees of freedom to a change in one or more macroscopic variables.[1–6] Here, we focus on off-lattice models of particle-based systems, i.e., atomic/molecular materials and their colloidal analogs, which are known to exhibit a rich variety of equilibrium phase transitions including condensation, freezing, and compositional demixing, as well as formation of microphases including cluster fluids.[7] Phase transitions that occur due to nonequilibrium driving forces (e.g., oscillatory shear,[8–12] time dependent magnetic/electric fields,[13,14] or particle self-propulsion[15–17]) are also possible in these materials. Given the diversity of microscopic degrees of freedom and macroscopic outcomes exhibited by such systems, new types of phase transitions and corresponding states of matter can be challenging to detect or predict from first principles.

Typically, phase transitions are described via an abrupt change in a coarse-grained representation of the system—i.e., an order parameter (OP)—that is sensitive to the specific character of the transformation and can provide insights into its origin.[1–6] Given the sheer number and variety of known phase transitions in such systems, it is perhaps unsurprising that there is no universal choice for the OP. Even for equilibrium transitions with a well-developed statistical mechanical foundation and established OPs, detection and systematic characterization of phase transitions can be computationally inconvenient, requiring simulations in specific ensembles.[18–21] As a result, simple and general strategies for identifying known phase transitions and discovering new ones would be welcome.

Toward this end, one promising approach is the use of dimensionality reduction techniques from machine learning to extract OP-like descriptors for phase transitions from configurational data.[22–25] For example, principal component analysis (PCA) has recently been used to detect phase transitions in the two-dimensional ferromagnetic Ising, XY, and other related spin models.[26–31] Beyond PCA, research into nonlinear machine learning strategies has led to the development of (1) a so-called confusion-based scheme which successfully detected transitions in the Kitaev chain and Ising models, as well as in a disordered quantum spin chain,[32] and (2) neural networks that identified emergence of subtle many-body localized phases as well as those arising in the square-ice model and in Ising lattice gauge theory.[33]

Here, we build upon this prior work and explore the extension of PCA to detect and characterize phase transitions for off-lattice models of particle-based materials. Specifically, we perform PCA on features derived from the particle coordinates, a strategy that is straightforward to implement in simulations but could also be applied in experiments that track particles (e.g., confocal microscopy of colloidal dispersions). In Paper I of a two-part series, we establish general guidelines and practices for using PCA to detect phase transitions in such systems. In Paper II,[54] we test the generality of this strategy for a variety of particle-based systems that exhibit equilibrium and non-equilibrium phase transitions.

The balance of this manuscript is organized as follows. In Sec. II, we describe the simulation protocols utilized in this work, provide a brief review of PCA, and identify and resolve complications that arise from a naïve construction of features from particle coordinate data. In Sec. III, we present the results of PCA of the features for some canonical phase transitions: the freezing transition in monodisperse hard disks and hard spheres and spinodal decomposition in a three-dimensional binary mixture with complementary attractions. We present conclusions for this study in Sec. IV.

a)Electronic mail: truskett@che.utexas.edu

II. COMPUTATIONAL METHODS

A. Simulation and analysis

Hard-disk and hard-sphere simulations were performed using the hard-particle Monte Carlo integrator in the HOOMD-blue simulation suite,[34–36] where $N = 4096$ particles were simulated in periodically replicated square and cubic simulation boxes of side-length L, respectively. These equilibrium hard-particle model systems are athermal and as such are uniquely described by an area or volume fraction (η); $\eta \equiv (N/L^2)\pi d^2/4$ for disks and $\eta \equiv (N/L^3)\pi d^3/6$ for spheres. For the PCA, we generated equilibrium samples over a range of η spanning below and above the freezing transition via a successive compression, equilibration, and sampling scheme. For hard disks, we incremented η from 0.55 to 0.80 in increments of 0.005, performing 1×10^7 HOOMD "time steps" (where a time step is approximately equal to four sweeps of MC moves over all particles) at each state point until $\eta = 0.695$ and 5×10^7 time steps per state point thereafter. Particle translations were restricted to 10 -30 of the particle diameter (d) along any Cartesian direction. Between 10^3 and 10^4 evenly spaced configuration snapshots were stored at each state point to provide sufficient data for the PCA. For hard spheres, samples were generated for packing fractions between $\eta = 0.42$–0.66 in increments of 0.005 with approximately a 23 acceptance rate. 2×10^6 time steps were used for equilibration and production with 40 evenly spaced snapshots sampled from the latter. After compression, the hard sphere configurations were re-expanded (via the same protocol) to capture hysteresis effects arising from the appreciable metastability possessed by the fluid phase beyond freezing.

For the case of spinodal decomposition, we revisit simulation data acquired in Ref. 37, where we studied binary mixtures of spherical particles decorated with six patches and two patches, respectively. While the spherical cores exhibited excluded volume interactions, patches on unlike particle types were mutually attractive. More details on the model and simulation protocol for the patchy particle mixture are described in Ref. 37. In this work, we re-analyze the simulation data for the ratio of the number of two-patch particles to six-patch particles, , equal to 1.3. We use PCA to detect liquid-gas phase separation and compare the results to the (η, T) phase diagram reported in Ref. 37 on the basis of $k = 0$ structure factor calculations.

In all cases, PCA was carried out using the Incremental PCA (IPCA) module in Scikit-Learn.[38] Importantly, IPCA allows for training using only small portions of the full data set at once; this is useful given the large size of the simulation trajectories aggregated across state points.

All visualizations were rendered using the Open Visualization Tool (or OVITO).[39] Hard sphere renderings utilized the built-in Polyhedral Template Matching (or PTM) capabilities of OVITO for crystal structure identification.[40]

B. Principal component analysis: Review

Principal component analysis (PCA) is a widely used unsupervised machine learning method that systematically discovers a lower dimensional representation of high dimensional data.[22,23,41] Here, we briefly review the general concepts behind PCA. Interested readers can find more detailed descriptions in the literature.

PCA, like all machine learning models, requires a numerical description for each of the M observations that describe the system. A given observation is represented by a so-called feature vector,

$$\boldsymbol{f} \equiv \left[f_1, \ldots, f_m\right]^T, \tag{1}$$

where each of the features (f_i) is a scalar measurable, T denotes a transpose, and m is the dimensionality of the feature representation. For context, observations in our work are derived from snapshots of particle configurations. The feature vector could describe the entire configuration or a subset of the particles. Without loss of generality, we assume that the elements of the feature vectors are centered (i.e., their mean value is zero). In most applications, $m \gg 1$. To conveniently denote a data set of observations (implicitly of size M), we adopt the notation .

If multiple features are highly correlated across the M observations, then the dataset necessarily contains redundant information. Dimensionality reduction strategies provide a mechanism for "concentrating" this redundant information into fewer dimensions. Toward this end, PCA determines a linear transformation that yields a decorrelated feature representation of the data. The transformation provides a new set of m directional unit vectors $\boldsymbol{w}_i$ (of dimension m) that define a new coordinate system and a corresponding $m \times m$ transformation matrix,

$$\boldsymbol{W} \equiv \left[\boldsymbol{w}_1, \ldots, \boldsymbol{w}_m\right]^T. \tag{2}$$

The projection of $\boldsymbol{f}$ along any of these new directions is simply, p_i $\boldsymbol{w}_i^T \boldsymbol{f}$. This is more succinctly stated in terms of the transformation matrix and the vector of m projections, $\boldsymbol{p} \equiv \left[p_1, p_2, \ldots, p_m\right]^T$, via

$$\boldsymbol{p} \quad \boldsymbol{W}\boldsymbol{f}. \tag{3}$$

For linear decorrelation, $\boldsymbol{W}$ must be of a form that diagonalizes the covariance matrix of $\boldsymbol{p}$,

$$\begin{aligned} \langle \boldsymbol{p}\boldsymbol{p}^T \rangle \quad & \boldsymbol{W} \langle \boldsymbol{f}\boldsymbol{f}^T \rangle \ \boldsymbol{W}^T, \\ \langle p_i p_j \rangle \quad & 0, \ i \quad j, \end{aligned} \tag{4}$$

where $\langle \cdots \rangle_D$ denotes an average over the data.

Given the above prescription, $\boldsymbol{W}$ is still underdetermined. Uniqueness is realized within PCA by requiring that $\boldsymbol{W}$ be an orthogonal matrix ($\boldsymbol{W}^T = \boldsymbol{W}^{\ 1}$, where the right-hand side notation indicates an inverse matrix). Under this constraint, the orthogonal set of unit vectors contained in $\boldsymbol{W}$ is equivalent to the sequential construction of $\boldsymbol{w}_i$ such that the variance projected along them (v_{p_i}) is maximized while maintaining orthogonality with all previous vectors $\boldsymbol{w}_i \cdot \boldsymbol{w}_j = 0$ for all $j < i$. The v_{p_i} are equal to the eigenvalues of the diagonalized covariance matrix $\langle \boldsymbol{p}\boldsymbol{p}^T \rangle$.

The $\boldsymbol{w}_i$ constructed in the aforementioned manner are called Principal Components (PCs), and the corresponding p_i are called PC scores. Each PC contains more information than the next, as quantified by the variance (i.e., $v_{p_1} \geq v_{p_2} \geq \cdots \geq v_{p_m}$). The efficiency with which information is sequestered toward the earlier PCs can be quantified by computing the