

Relevant, Hidden, and Frustrated Information in High-Dimensional Analyses of Complex Dynamical Systems with Internal Noise

Original

Relevant, Hidden, and Frustrated Information in High-Dimensional Analyses of Complex Dynamical Systems with Internal Noise / Lionello, Chiara; Becchi, Matteo; Martino, Simone; Pavan, Giovanni M.. - In: JOURNAL OF CHEMICAL THEORY AND COMPUTATION. - ISSN 1549-9618. - 21:14(2025), pp. 6683-6697. [[10.1021/acs.jctc.5c00374](https://doi.org/10.1021/acs.jctc.5c00374)]

Availability:

This version is available at: 11583/3001515 since: 2025-07-03T09:15:53Z

Publisher:

American Chemical Society

Published

DOI:[10.1021/acs.jctc.5c00374](https://doi.org/10.1021/acs.jctc.5c00374)

Terms of use:

This article is made available under terms and conditions as specified in the corresponding bibliographic description in the repository

Publisher copyright

(Article begins on next page)

Relevant, Hidden, and Frustrated Information in High-Dimensional Analyses of Complex Dynamical Systems with Internal Noise

Chiara Lionello, Matteo Becchi, Simone Martino, and Giovanni M. Pavan*

Cite This: *J. Chem. Theory Comput.* 2025, 21, 6683–6697

Read Online

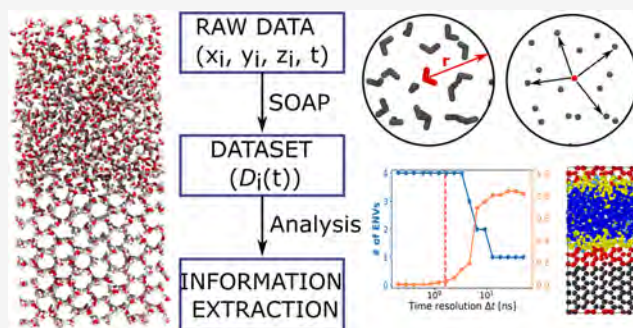
ACCESS |

Metrics & More

Article Recommendations

Supporting Information

ABSTRACT: Extracting from trajectory data meaningful information to understand complex molecular systems might be nontrivial. High-dimensional analyses are typically assumed to be desirable, if not required, to prevent losing important information. But to what extent such high-dimensionality is really needed/beneficial often remains unclear. Here we challenge such a fundamental general problem. As a representative case of a system with internal dynamical complexity, we study atomistic molecular dynamics trajectories of liquid water and ice coexisting in dynamical equilibrium at the solid/liquid transition temperature. To attain an intrinsically high-dimensional analysis, we use as an example an abstract high-dimensional descriptor of local molecular environments (e.g., Smooth Overlap of Atomic Positions, SOAP), obtaining a large dataset containing 2.56×10^6 576-dimensional SOAP spectra that we analyze in various ways. Our results demonstrate how the time-series data contained in one single SOAP dimension accounting only $<0.001\%$ of the total dataset's variance (neglected and discarded in typical variance-based dimensionality reduction approaches) allows resolving a remarkable amount of information, classifying/discriminating the bulk of water and ice phases, as well as two solid-interface and liquid-interface layers as four statistically distinct dynamical molecular environments. Adding more dimensions to this one is found not only ineffective but even detrimental to the analysis due to recurrent negligible-information/non-negligible-noise additions and “frustrated information” phenomena leading to information loss. Such effects are proven general and are observed also in completely different systems and descriptors' combinations. This shows how high-dimensional analyses are not necessarily better than low-dimensional ones to elucidate the internal complexity of physical/chemical systems, especially when these are characterized by non-negligible internal noise.



INTRODUCTION

Elucidating the physics of complex dynamical systems is a challenging task and remains one of the most thought-provoking topics in both physics and chemical physics. These systems feature multiple levels of local and collective dynamical events that coexist and interconnect within complex dynamical networks, making their understanding and classification particularly difficult. Phenomena ranging from microscopic events—such as phase transitions, nucleation processes,^{1,2} or the intricate internal dynamics of molecular assemblies^{3,4}—to macroscopic events like the rapid turns and convolutions of large crowds of individuals, bird flocks,^{5,6} and fish schools,⁷ highlight the need for advanced analytical methods to detect and track the diverse dynamical environments that emerge within these systems.

A common assumption when studying complex systems (of any scale), about which little is known *a priori*, is that high-dimensional analyses are desirable, and even necessary, to avoid losing information due to prior assumptions or incomplete characterization. However, selecting a set of physically relevant descriptors capable of providing a complete characterization of a specific system requires a deep knowledge

of the system itself, which is not always available. Moreover, managing high-dimensional analyses and extracting meaningful physical insights from them is often challenging. Fundamental questions arise: What is the intrinsic dimension of the dataset?^{8,9} Which components are the most relevant?

At the molecular scale, high-dimensional analyses can be approached in various ways. One common method is to use descriptors that transform, for example, the simulation trajectories of the system's units into analyzable data formats, such as time-series. Multiple descriptors can be employed under the assumption that each captures different (orthogonal) information. However, the degree to which the information provided by these descriptors is complementary or redundant is not known *a priori*. To ensure information-rich analyses,

Received: March 5, 2025

Revised: June 20, 2025

Accepted: June 25, 2025

Published: July 2, 2025



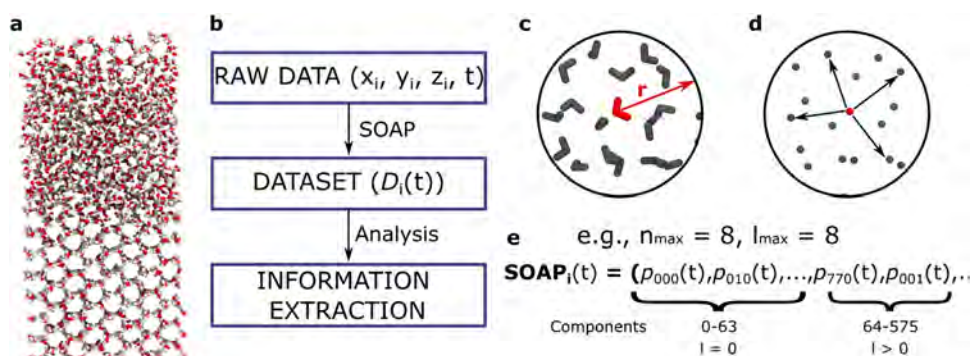


Figure 1. Ice/water coexistence MD simulation. (a) Screenshot of the system studied. (b) Schematic representation of the steps necessary to extract information from simulated systems. (c) Zoom into water molecules, centered in a molecule (red) within a cutoff r . (d) Zoom of the water molecules, considering exclusively oxygen atoms, which are used to calculate the SOAP descriptor. (e) Definition of the Smooth Overlap of Atomic Position (SOAP) descriptor with distinction of components with $l = 0$ and $l > 0$.

descriptors that are intrinsically high-dimensional by their mathematical definition are particularly valuable. A notable example is the Smooth Overlap of Atomic Position (SOAP) descriptor,¹⁰ which, inspired by wave functions and orbital definitions in quantum mechanics, effectively identifies local structural environments in complex molecular systems.

Similarly to other descriptors, such as atomic cluster expansions (ACE)¹¹ or N-body iterative contraction of equivariants (NICE),¹² SOAP provides high-dimensional mathematical representations of local density and order/disorder. Recently, SOAP has been successfully applied to investigate complex molecular systems, including liquid–solid and liquid–liquid coexistence in aqueous systems,^{13–16} and nontrivial dynamics and local disorder events in metal surfaces and nanoparticles,^{17–22} and ionic behaviors in dense environments.²³

In principle, the high-dimensional description provided by SOAP analyses can be extremely useful, especially for studying systems about which little is known *a priori*.^{10,24} However, while high-dimensional descriptions are often assumed to be desirable for such systems, our recent work has shown that in many real cases—including aqueous, solid, and metallic systems of various kinds¹⁷—a monodimensional version of SOAP, called *TimeSOAP*,¹⁷ can extract more meaningful features than the standard high-dimensional SOAP analysis by tracking the variation of the SOAP spectrum of each molecule over time.

This finding raises several important questions: To what extent is high-dimensionality truly necessary to capture the physics of these systems? How can a high-dimensional analysis be less informative than its dimensionally-reduced counterparts? What type of information remains hidden, where is it located, and how can it be extracted? Is the information contained in the orthogonal dimensions of a high-dimensional feature space always additive, does adding a second dimension always contribute new information? These questions are particularly central when using intrinsically high-dimensional descriptors, e.g., SOAP or ACE,¹¹ but they also apply to any combination of descriptors, whether physics-inspired, abstract, or general. At their core, these questions challenge a fundamental issue in physics: How to effectively extract meaningful information from high-dimensional analyses, a problem that is central in many fields.

Most often, high-dimensional dataset extracted from, e.g., molecular dynamics (MD) trajectories (but this is true for any

type of trajectory, including experimentally resolved ones), are studied under an ergodic assumption as a collection of temporally independent data, and clustering methods are used to detect high-density peaks identifying dominant/recurrent SOAP domains.^{25–27} However, we have recently shown that significant information may be hidden in the temporal correlations of the data—such as local, sparse but dominant fluctuations, important local transitions, etc.—that are typically overlooked in pattern recognition approaches which neglect time correlations. We have also developed an efficient method capable of maximizing the extraction and classification of information in time-series data, systematically distinguishing statistically relevant fluctuations from noise as a function of the time resolution used in the analysis.²⁸

In this work, we leverage such approaches to tackle fundamental questions relevant in high-dimensional analyses of molecular (and nonmolecular) systems, as To what extent a high-dimensional analysis approach is really beneficial compared to lower-dimensional ones? How to compare information quantity *vs* information quality (relevance)? As a first prototypical example of an information-rich dataset, we analyze SOAP data extracted from atomistic MD simulation of a periodic box where liquid water and ice coexist in dynamical equilibrium at the melting temperature.^{16–19,28} This provides us with a prototypical system in equilibrium that respects a detailed balance and, at the same time, exhibits nontrivial internal complexity. This is explored here by computing 576-dimensional SOAP spectra for each of the 2048 water molecules in the system every $\delta t = 40$ ps (sampling interval) over 50 ns of MD trajectory. This results in a large high-dimensional dataset containing 2.56×10^6 SOAP power spectra (vectors), each composed of 576 components that we analyze both by ignoring or considering the time correlations between the data (i.e., pattern recognition *vs* time-series analyses).

Our results demonstrate how dimensionality reduction approaches, such as Principal Component Analysis (PCA), fall short when dealing with noisy data. Similar approaches allow discriminating and classifying only two solid ice and liquid water environments: i.e., the two most-populated dominant phases in the system.¹⁸ On the other hand, we show how time-series data contained in a single SOAP component, which accounts remarkably only $<0.001\%$ of the total variance of the SOAP dataset (and that is thus typically discarded in PCA), can provide more information than that

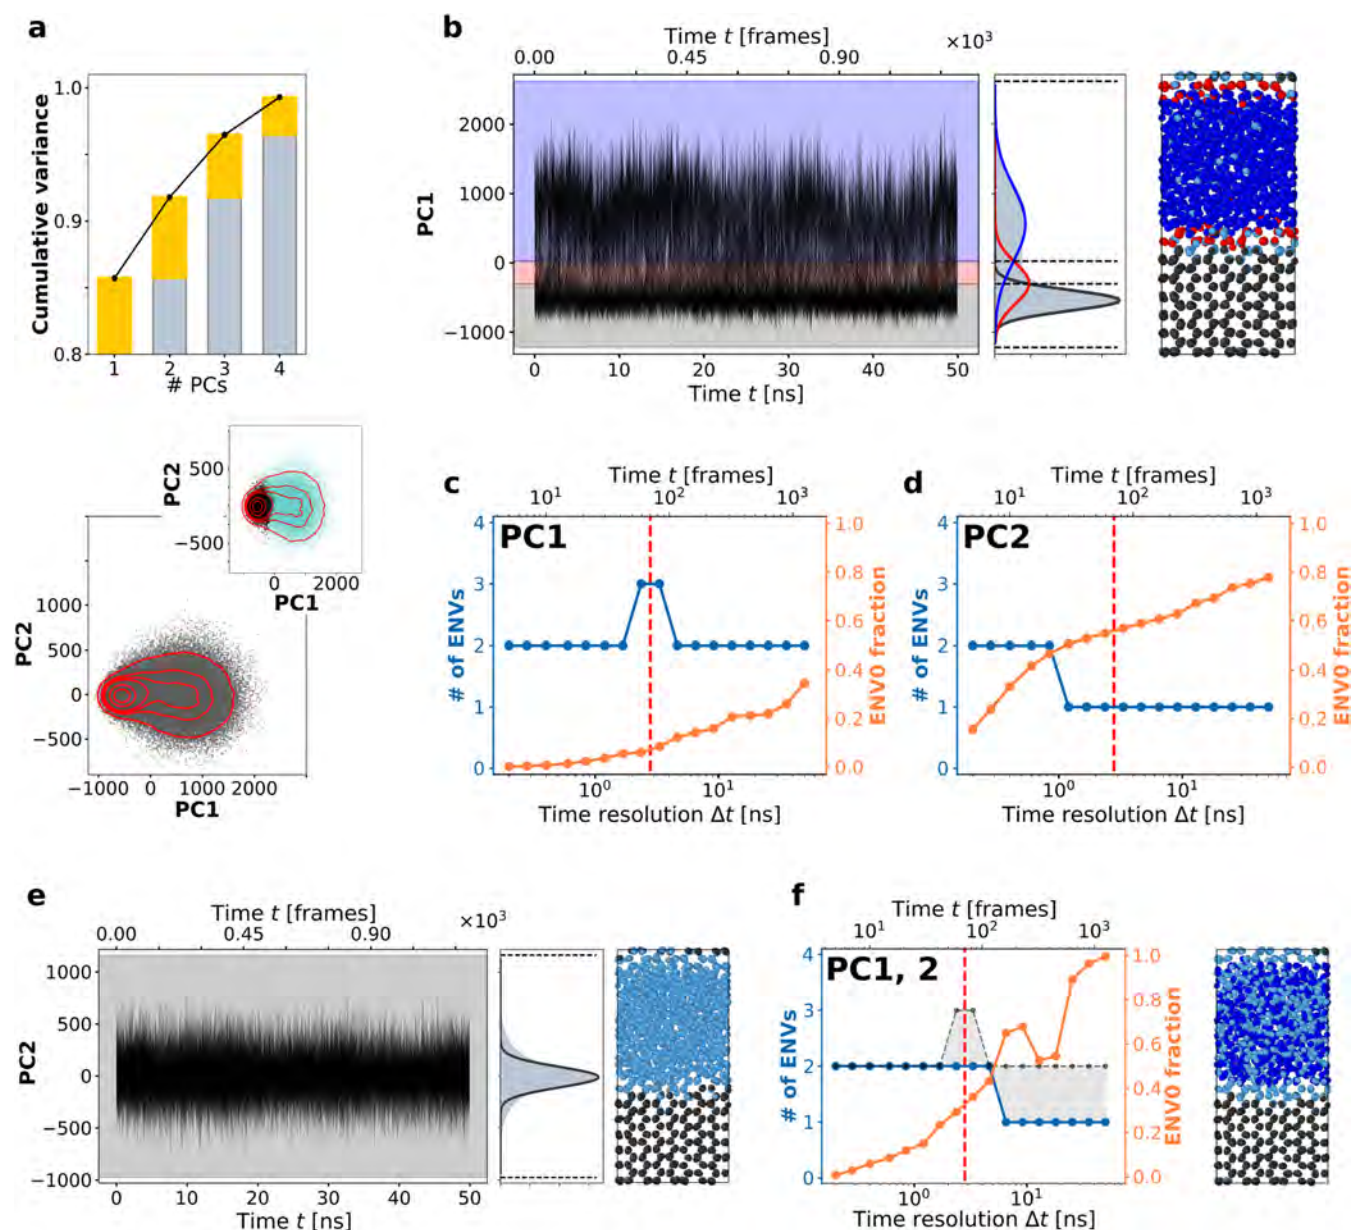


Figure 2. The information nested along the time dimension: Onion Clustering on PC1 and PC2 time-series data. (a) Top: Cumulative data variance explained by the first four PCs. Bottom: SOAP dataset projection onto the first 2 PCs. Red contour isolines help to visualize the data density. Inset: common clustering approaches (hierarchical clustering) distinguish two main clusters (solid ice and liquid water) in the SOAP dataset. (b) Left: PC1 time-series data of all the molecules. Center: kernel density estimate (KDE) of the PC1 time-series data, with the Gaussian environments (solid curves) identified by Onion Clustering.²⁸ The dark gray environment corresponds to solid ice, the red one to the ice/water interface, and the blue one to liquid water. Right: snapshot of the MD simulation, where the water molecules are colored according to the Onion microclusters. (c) Onion output plot, indicating the number of clusters (in blue) identified by Onion Clustering, and the fraction of lost information (in orange: unclassified information (ENV0)—fast SOAP changes—due to insufficient resolution) as a function of the time resolution (Δt) used in the analysis of the PC1 time-series. (d) Same as panel (c), but for the PC2 time-series. (e) Same as panel (b), but for the PC2 time-series. (f) Left: same results as panel (c), obtained by Onion Clustering analyzing a bivariate (bidimensional) PC1, PC2 time-series (classified clusters in blue, unclassifiable information in orange). The gray area shows the information lost by combining both dimensions (gray = $\max(\#_{\text{clust}}(\text{PC1}), \#_{\text{clust}}(\text{PC2})) - \#_{\text{clust}}(\text{PC1}, \text{PC2})$). Right: snapshot of the MD simulation, where the water molecules are colored according to the Onion microclusters detected at the resolution of $\Delta t = 3$ ns (red dotted vertical line) from the bidimensional (PC1, PC2) time-series.

attainable by analyzing all dimensions in the SOAP dataset altogether (consistent results can be obtained with other types of dimensionality reduction techniques, e.g., Time-lagged Independent Component Analysis (tICA),^{29,30} or VAMP-nets,³¹ as reported in Figures S7–S8). Our results reveal that adding more (noisy) components to this single one can be not only nonbeneficial but also detrimental to the analysis, a

phenomenon we describe as “noise-frustrated information”. Finally, we analyze a completely different system—whose trajectories are obtained experimentally and are analyzed by combining physics-based descriptors—showing that such effects are not exclusive of SOAP data, but can arise in any noisy multidimensional analysis. The findings presented here challenge classical paradigms in data analysis, highlighting the

critical importance of focusing on information quality (relevance) rather than quantity to minimize information loss and misinterpretation.

RESULTS

In this work, we use as a case study a 50 ns-long MD trajectory of an atomic system where liquid water and ice coexist in dynamic equilibrium (see Figure 1a). Briefly, the periodic simulation box contains 2048 TIP4P/ICE water molecules,³² initially arranged with 50% in the liquid state and 50% in a hexagonal ice configuration. The molecules are simulated in periodic boundary conditions at the melting temperature (see Methods section for details). This system has been recently proven to exhibit significant internal complexity.^{16–19}

The extraction of information from the trajectories of such a simulated system involves three main steps: (i) acquisition of raw data, which consist in the sequence of the particles' coordinates acquired every 40 ps, over $\tau = 50$ ns of MD simulation; (ii) selection of an appropriate descriptor D to convert the trajectory of each molecule i —given by the time-series data $x_i(t)$, $y_i(t)$, $z_i(t)$ —into a dataset $\{D_i(t)\}$; (iii) analysis of the dataset to extract relevant information on the system's physics (see Figure 1b). The information that can be effectively collected from the system depends on the choice of the descriptor in step (ii). The correct choice of descriptor is crucial for two main reasons. First, different descriptors capture distinct aspects of the system's behavior: for example, a structural descriptor would be useless to study purely dynamical features.¹⁹ Second, each descriptor has an intrinsic signal-to-noise ratio, which influences the clarity and reliability of the extracted information.¹⁶ When the physics of a system is unknown, selecting a general, agnostic—possibly intrinsically high-dimensional—descriptor would be desirable in order to minimize the risk of introducing assumptions that could bias the analysis or result in incomplete information.

Building the High-Dimensional (SOAP) Dataset.

Among the various descriptors used to study the coexistence of ice and liquid water,^{13,16,17,33–35} here we employ SOAP¹⁰ as a representative example of an abstract, general descriptor with an intrinsically high-dimensional nature. Recent studies have demonstrated SOAP's effectiveness in capturing complex phenomena, such as liquid–liquid and liquid–solid coexistence, early phase transition nucleation, and other nontrivial processes.^{13,17,25} SOAP is a representation of the local particle density around each particle in the system through a power spectrum \mathbf{p} , whose components $p_{mn'l}$ are indexed by radial (n and n') and angular (l) indexes. These components encode information about the local density and order, capturing features such as symmetry and the relative spatial arrangements of neighbors around each particle. This allows for information-rich analyses and enables the discrimination of different local environments within the system, evaluation of system homogeneity, and more.

On a practical level, the results of SOAP analyses depend on parameters such as the maximum values n_{\max} and l_{\max} . Selecting appropriate values for these parameters is crucial to balance the retention of information relevant to understand the system's physics while avoiding an unnecessarily large number of components that could make the analysis computationally prohibitive. In the next section, we present results obtained using $n_{\max} = l_{\max} = 8$, the default value in the SOAP computation package *Dscribe*.^{36,37} These values have been extensively used in prior works.^{17–19,25,28} However, consistent

results are also obtained using $n_{\max} = l_{\max} = 4$ (see results in the Figures S1–S2). Using $n_{\max} = l_{\max} = 8$, we computed one 576-components SOAP power spectrum for each of the 2048 molecules in the system every $\Delta t = 40$ ps over the course of $\tau = 50$ ns of MD simulation (1250 frames in total). This resulted in a large high-dimensional dataset containing 2.56×10^6 SOAP spectra, each with 576 components.

The Information Nested in the Time Dimension.

Common approaches to analyze high-dimensional datasets require first to reduce their dimensionality. One of the most widely used techniques is Principal Component Analysis (PCA),^{38–40} based on the concept of variance. PCA reduces a high-dimensional dataset into a space of orthogonal features, ordering these components proportionally to the variance they explain. While PCA, like other dimensionality reduction methods, has intrinsic limitations (it works well only if the intrinsic feature space is linear), we use it here as an initial demonstrative test case to introduce fundamental concepts, which will be then explored using other approaches in subsequent sections. For the SOAP dataset derived from the water/ice trajectories analyzed in this study, the first four PCs (PC1–PC4) account for more than 99.3% of the total variance (see Figure 2a top). The projection of the full SOAP dataset onto the first two principal components, PC1 and PC2, is shown in Figure 2a (bottom). In this projection, the areas of highest density (identified by the red isolines) correspond to recurrent and most-populated patterns—or molecular motifs—characterizing the system. The red contours reveal a density maximum at PC1 ~ -500 and PC2 ~ 0 , with a plateau extending to the right side of the plot. Using an unsupervised clustering method—in this case hierarchical clustering⁴¹—we identified two domains: one encompassing the molecules in the density maximum and another including those in the plateau region (Figure 2c, inset). Assigning the water molecules to these two clusters, based on this analysis, allows distinguishing the molecule in the solid (black) and in the liquid phase (cyan).

It is worth noting that such pattern recognition approaches are based on an ergodic assumption, where temporal correlations between the data are ignored, and all these spectra are treated as part of a single static dataset. However, it is established that relevant information is contained in the time correlation of data. This has been demonstrated by various approaches, such as Markov state models and time-series analyses.^{42,43} Recently, we demonstrated that tracking the temporal sequence of individual signals from each molecule in a system can yield invaluable microscopic-level insights, which are essential for reconstructing both the microscopic and global behavior of the entire system.^{17–19,21,22,28,34,44} In particular, we developed a fully unsupervised and essentially parameter-free clustering method, known as Onion Clustering,²⁸ which can systematically detect statistically relevant fluctuations in noisy time-series data of any type. Briefly, the Onion Clustering identifies and classifies all dynamical environments within time-series data, proceeding with a hierarchical approach. Starting from the most evident features, the method iteratively detects, classifies, and archives them, continuing until no further statistically robust dynamical clusters can be identified.²⁸ For a detailed description of the Onion Clustering method, we refer the interested readers to ref 28.

In this work, we apply Onion Clustering to extract and classify the information contained in the SOAP dataset, which

is studied as an ensemble of time-series rather than of uncorrelated SOAP spectra. We began by analyzing the time-series of the SOAP PC1 for all the 2048 molecules (see Figure 2b (left)). Notably, PC1 alone accounts for $\sim 86\%$ of the cumulative variance of the entire SOAP dataset, as shown in Figure 2a. Onion Clustering successfully identified three statistically distinct clusters within the PC1 time-series, represented by the three Gaussian distributions shown in Figure 2b (center). These clusters correspond to three distinct environments, each characterized by a different average value of the PC1 SOAP spectra and a unique internal variance. These findings highlight that, although focusing solely on PC1 captures only a portion of the information contained in the full dataset, incorporating temporal correlations rather than treating the data as spatially and temporally decorrelated can yield more insightful results. Specifically, while the two main clusters—solid ice and liquid water—are detectable using standard pattern recognition methods applied to the full dataset (Figure 2a, bottom), Onion Clustering additionally resolves a third cluster: the ice–water interface. This observation underscores an essential concept: the quantity of information in a dataset (its completeness) may be less critical than the quality (or relevance) of the information extracted. Here, the temporal correlations within the PC1 time-series contain more pertinent information than the entire SOAP dataset treated as static. This suggests that in certain cases, particularly for noisy datasets, “less may be more”—a principle that will be further demonstrated in the following sections.

Information Loss in Time-Series Analyses: Less May Be More in Time... Another example of the “less may be more” principle becomes evident when considering how changes in time resolution affect clustering results in time-series analyses. By design, Onion Clustering analyzes time-series data across all possible time-resolutions—from the highest resolution (in this case, $\Delta t = 2 \times 40$ ps = 80 ps, corresponding to two trajectory frames), to the lowest resolution (here $\Delta t = 50$ ns, encompassing the entire simulation). The results for PC1 and PC2 are shown in Figure 2c,d, while the results for PC3 and PC4 are reported in Figure S3. In Figure 2c, the blue line shows the number of clusters detected as a function of time resolution, while the orange line depicts the fraction of data points that cannot be classified into a stable environment at the chosen resolution (ENV0), corresponding to molecules undergoing faster transitions.

A key observation is that the third cluster, corresponding to the ice–water interface domain, can only be detected as a distinct environment within a specific time-resolution range $2 \leq \Delta t \leq 4$ ns. The exact Δt window where this third environment can be resolved depends on both the system and the descriptor used.¹⁶ However, in general, as time resolution decreases ($\Delta t > 4$ ns), the number of clusters reduces from three to two. This reduction occurs because the characteristic residence time of a molecule at the ice–water interface is at most ~ 4 ns. When the analysis resolution falls below this threshold, the interface becomes indistinguishable from the other two environments, leading to the reassignment of interfacial molecules to the bulk ice or liquid environment. Conversely, using an excessively high time resolution ($\Delta t < 2$ ns) leads to “oversampling” the time-series, saturating the dataset with information related to molecular vibrations. This oversampling reduces the statistical relevance of sparsely observed solid-to-liquid (and back) transitions. In this case, while subnanosecond molecular vibrations effectively distin-

guish solid ice from liquid water, they fail to resolve molecules at the ice/liquid interface. These interface molecules exhibit longer-scale dynamics and undergo sharp transitions on a nanosecond time scale, which are overlooked when the analysis focuses on finer temporal details.

As expected, these findings demonstrate that being at a time resolution higher than the time scales of the observed dynamics does not yield greater insights or higher-quality information. Instead, particularly when dominated by noise or irrelevant details, it can obscure important transitions, potentially leading to information loss. The results in Figure 2a–c highlight the importance of the temporal component in uncovering relevant information. For this reason, we focus the study on the results obtained by Onion Clustering, which emerges as an optimal method for revealing the hidden dynamics within time-series datasets.

...and in Space: Information Loss in High-Dimensional Analyses. The same approach, used to analyze the PC2 time-series data, leads to different results than those from PC1 (Figure 2e). Unlike PC1, the PC2 time-series data exhibits higher noise levels and present a single density peak. Onion Clustering identifies only two environments in the PC2 time-series, corresponding to ice and liquid water. However, as the time resolution of the analysis decreases, even at relatively high resolutions ($\Delta t > 0.7$ ns), liquid water is not no longer detected as a statistically significant cluster and instead appears as a cluster of unclassified data (Figure 2e, right). Notably, while PC2 accounts for a significant portion of the cumulative variance ($\sim 6\%$) of the dataset, it does not provide information about the ice-liquid water interface environment. Instead, it only retains information about the ice and liquid water environments, which were already prominent in PC1. This raises a fundamental question about the actual benefit of incorporating PC2 into the analysis when its contribution appears redundant compared to the insights already obtained from PC1.

High-dimensional analyses are often thought as beneficial—or even necessary—for avoiding information loss. To explore this, we performed a bidimensional Onion Clustering analysis on a bivariate PC1–PC2 time-series, combining information from both components. Theoretically, the combination of the PC1 and PC2 components should enhance the analysis by increasing the cumulative variance captured, from $\sim 86\%$ to $\sim 92\%$. However, as shown in Figure 2f, the results reveal the opposite effect. The comparison between the blue line, which identifies at most two clusters at all resolutions, and the corresponding data for PC1 in Figure 2c is particularly revealing. The black dashed line in Figure 2f represent the minimum expected advantage when adding PC2 information to PC1. Conceptually, if adding dimensions were truly beneficial, the high-dimensional analysis (PC1, PC2) should never yield less information than the maximum achievable by using either dimension (PC1 or PC2) alone. In this case, since PC1 resolves more clusters than PC2 at every Δt , the two-dimensional analysis should, at minimum, match the clustering results of PC1 alone (represented by the black dashed line in Figure 2f). However, the results of Figure 2f show that the bidimensional PC1–PC2 analysis fails to detect the interface environment altogether, with the number of clusters falling below the black line for all time-resolutions $\Delta t > 1$ ns. This finding provides a concrete example of the “curse of dimensionality”,^{45,46} where the introduction of additional spatial dimensions complicates the resolution of relevant

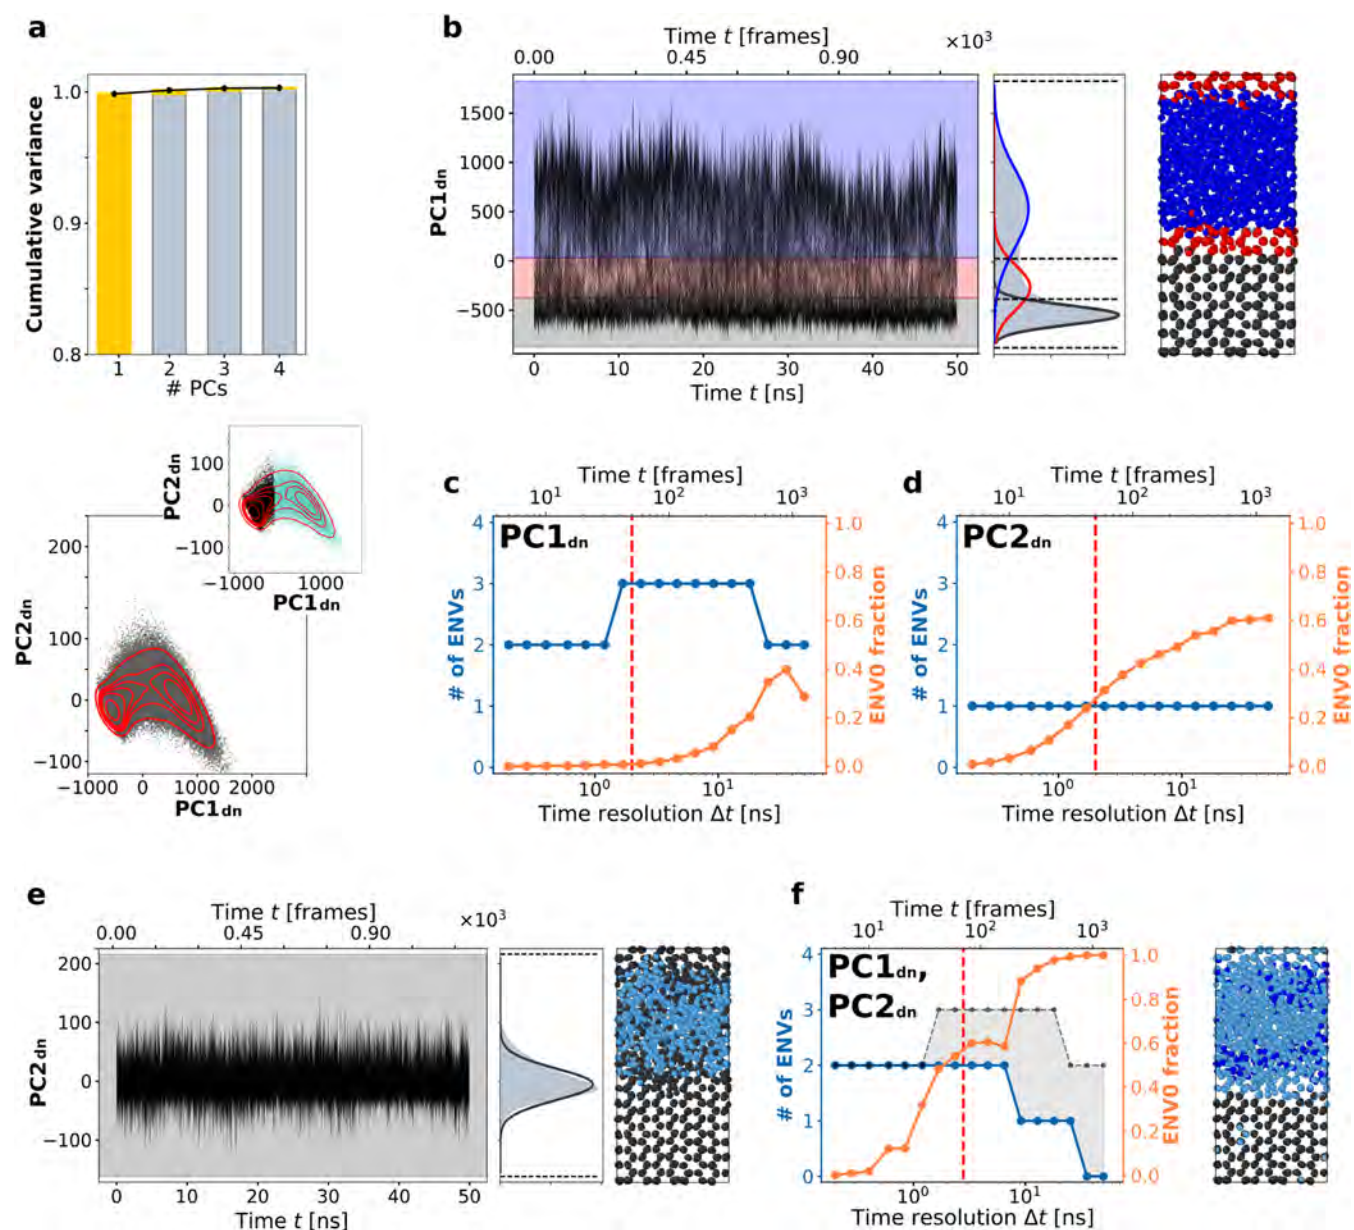


Figure 3. Onion Clustering on denoised PC1 and PC2. The figure follows the same structure of Figure 2, but all the analysis are performed on the smoothed time-series PC1_{dn} and PC2_{dn}. (a) Top: Cumulative data variance explained by the first four PCs. Bottom: SOAP dataset projection onto the first 2 PCs. Inset: hierarchical clustering distinguishes two main clusters (solid ice and liquid water) in the SOAP_{dn} dataset. (b, c) Onion Clustering results for PC1_{dn} and (d, e) Onion Clustering results for PC2_{dn}. (f) Onion Clustering of the bivariate PC1_{dn}, PC2_{dn} time-series, The gray dotted line shows the best result obtained by either of the two single components: the gray area shows the information lost by combining both dimensions. Right: snapshot of the MD simulation, colored according to the microclusters detected by Onion Clustering. The results of PC3_{dn} and PC4_{dn} are reported in Figure S4.

information and ultimately leads to information loss. In the next section, we discuss the physical basis for such effects.

Relevant Information vs Noise. Descriptors like SOAP, rich in “local information”, can be significantly affected by local noise. Recently, Donkor et al. introduced an effective approach to reduce the local noise via spatial average, thus by averaging the SOAP spectrum of each particle with those of its neighbors within the SOAP cutoff.¹³ The result is a denoised SOAP descriptor, with a decreased heterogeneity in SOAP spectra, enhancing the sensitivity of the subsequent analysis.^{13,16} Here, we illustrate the impact of this local denoising on our case study. In our analyses, local-noise reduction in the principal components time-series produces notable effects. As shown in

Figure 3a (top), the cumulative variance of PC1 increases to ~99.5%, compared to the original ~86% observed in Figure 2a. This means that the other PCs contribute only marginal additional information (e.g., PC2 adds only ~0.3%). The projection of PC1 and PC2 in Figure 3a (bottom) further demonstrate the effect of denoising (raw results are reported in Figure 2a, bottom): two distinct density peaks, corresponding to the solid and liquid domains, are clearly visible. These clusters can be easily identified by any clustering method, as shown by the gray and light blue regions in the inset.

It is particularly noteworthy that, in Figure 2a, four PCs were required to reach ~99.5% of the cumulative variance, implying that the dataset was approximately four-dimensional. However,

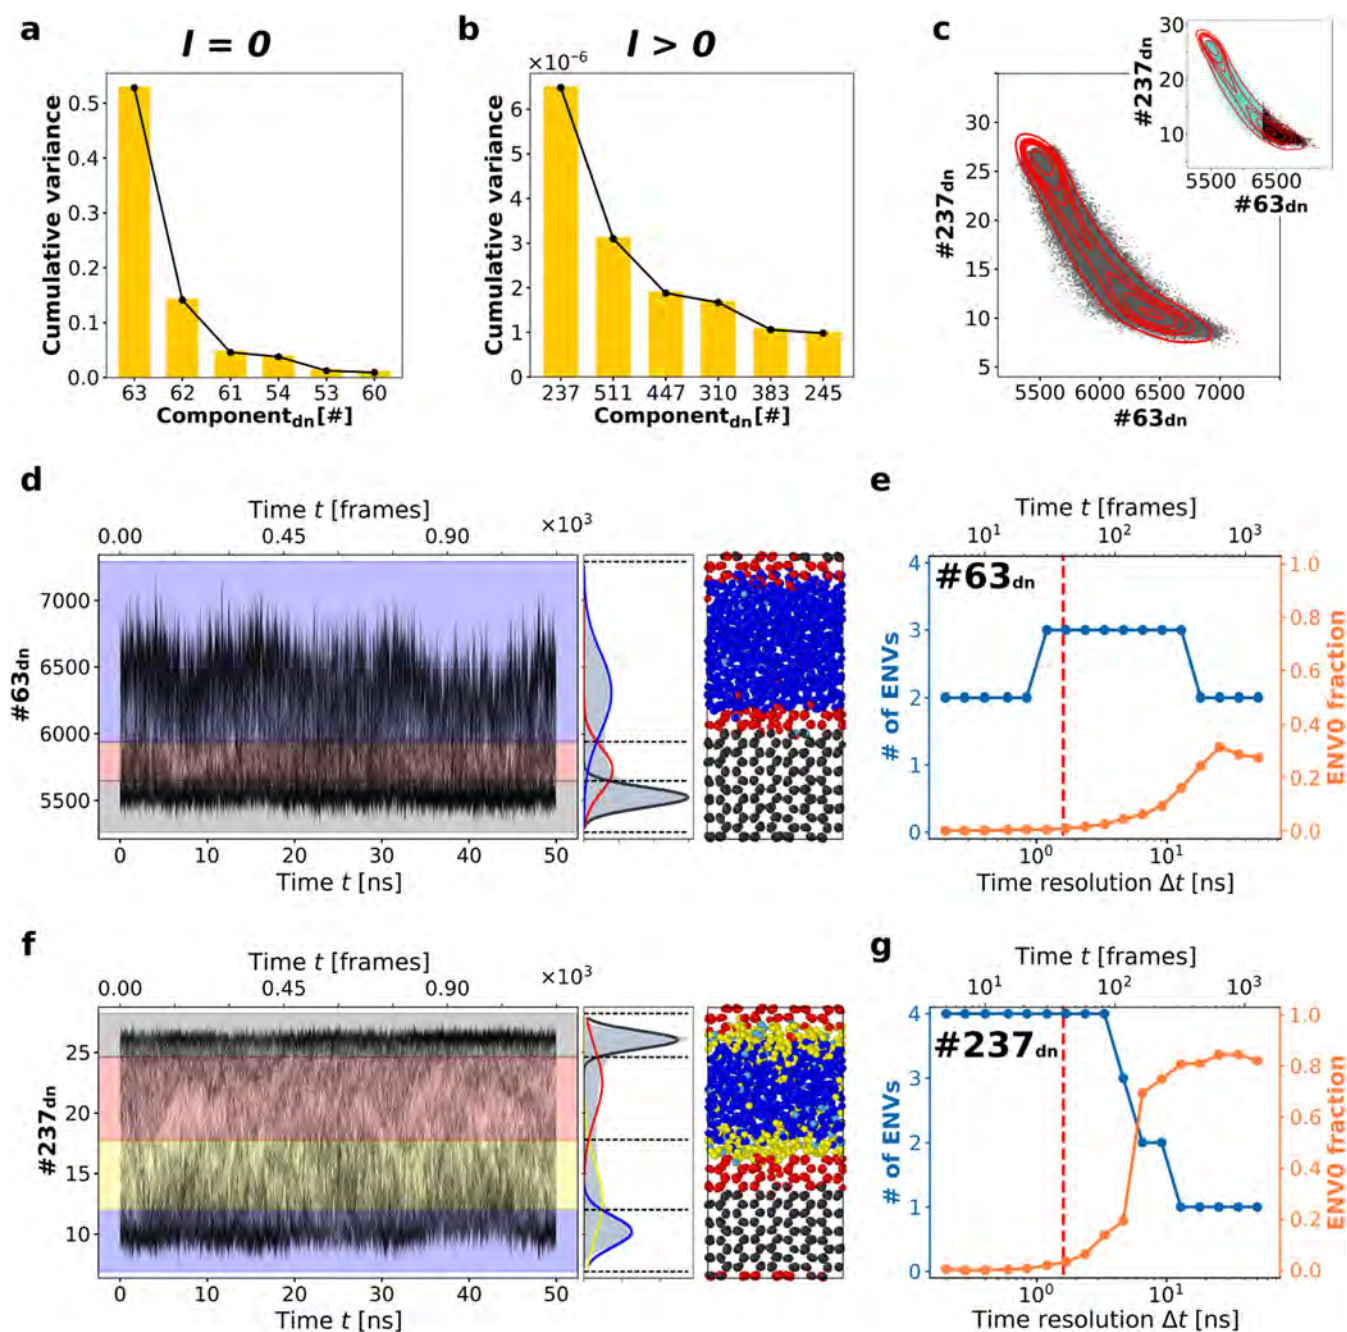


Figure 4. Onion Clustering on denoised components. (a) Variance of the six most significant spherical ($l = 0$) SOAP components. (b) Variance of the six most significant nonspherical ($l > 0$) SOAP components. (c) Denoised dataset projection onto components #63 and #237; red contour lines help visualize the data density; in the inset, static clustering distinguishes 2 environments. (d, e) Onion Clustering results of component #63_{dn}. (f, g) Onion Clustering results of component #237_{dn}.

the local noise reduction allows PC1 alone to account for nearly the entire variance of the dataset, indicating that the “relevant” information, without noise, is essentially mono-dimensional. This suggests that the high-dimensionality observed in the original dataset is primarily a reflection of noise, not of the inherent complexity of the system.

We applied Onion Clustering to analyze the denoised PC1 (PC1_{dn}), which successfully identifies three distinct environments: ice, liquid water, and the liquid/solid interface (Figure 3b). Compared to the results from Figure 2b, the clusters in this case exhibit better separation, with the interface being more clearly detected and appearing thicker. Figure 3c shows

the number of clusters detected as a function of time resolution. Notably, the interface environment becomes detectable at $\Delta t = 1\text{--}2$ ns, as observed in Figure 2, but the reduction of local noise extends the interface’s detectability up to $\Delta t \sim 20$ ns, approximately ten times more detectable in PC1_{dn} than in PC1. This demonstrates the crucial role of local-noise reduction in time-series data, enabling a more stable and robust detection of physically relevant environments and events. In contrast, the analysis of denoised PC2 (PC2_{dn}) identifies only a single environment, corresponding to ice, while all the other molecules and data points are not classifiable.

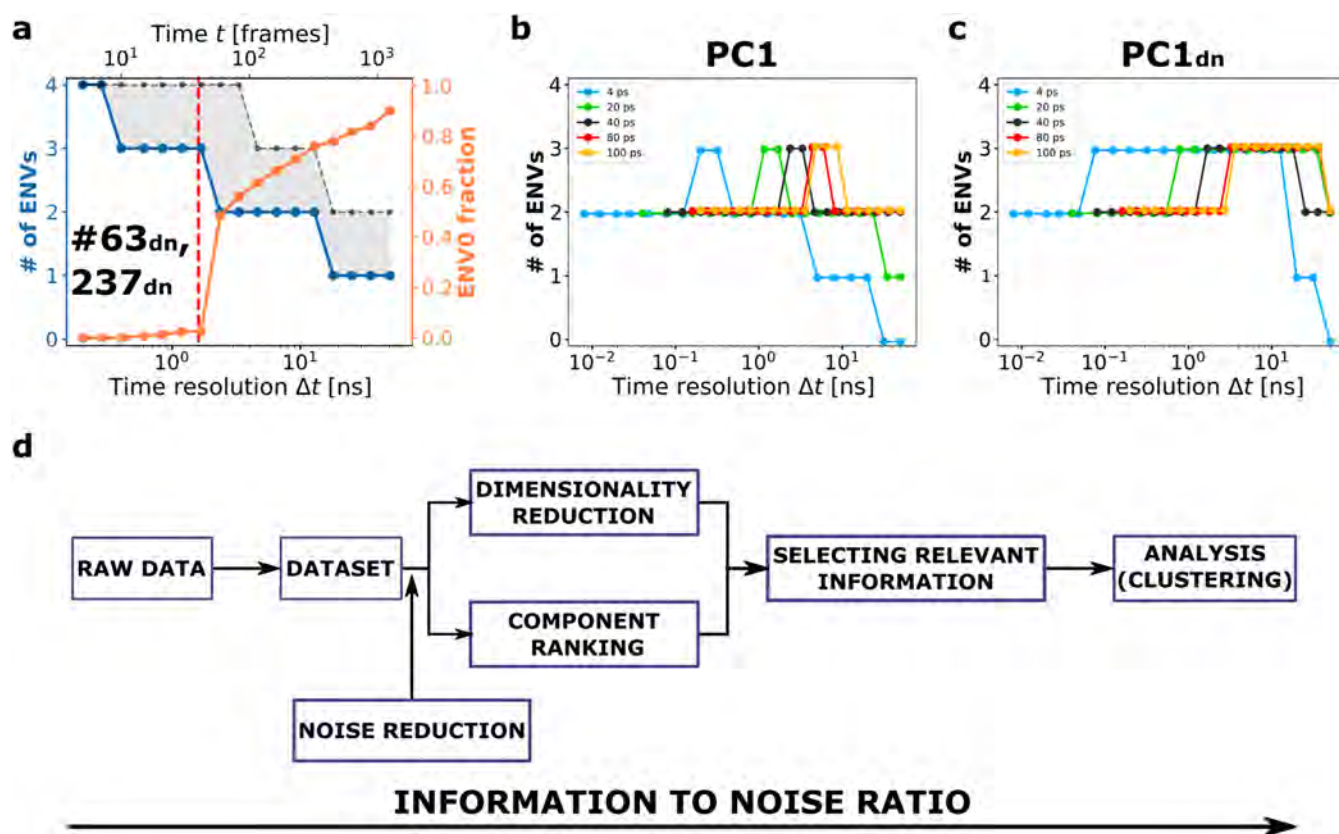


Figure 5. Information frustration: (a) Onion Clustering on bivariate #63_{dn} and #237_{dn} time-series (see Figure 2f for the legend). In gray, the combination of two components leads to considerable information loss, even when they are denoised. (b, c) Dependence of the detection of the ice/water interface on the sampling Δt in the raw MD trajectories onto which the PC1 is calculated in the native/noisy (b) or denoised (c) SOAP dataset. (d) Schematic representation of the analysis process to be followed then analyzing time-series data. Note that, after the noise reduction, the dataset obtained is completely different than the one with noise (now it has more information), thus the dataset should be considered as a new dataset.

Combining PC1_{dn} and PC2_{dn} into a bidimensional time-series introduces similar issues as for the PC1, PC2 case. Also in this case, conducting a high-dimensional analysis not only fails to provide any advantage, but also results in information loss (note the extended gray area in Figure 3f). This suggests that such high-dimensionality-induced information loss is not merely due to the noise of the individual time-series: in fact, the gray area of Figure 3f (indicating information loss upon adding PC2_{dn} to PC1_{dn}) is even larger than that in Figure 2f. Instead, this highlights a critical issue intrinsic to high-dimensional analyses. Adding dimensions does not necessarily introduce relevant information; more often, it guarantees adding noise. In other words, the information in one dimension (PC2_{dn}) manifest as noise relative to PC1_{dn}, degrading the quality of the bidimensional analysis compared to analyzing the PC1 time-series alone.

This evidence rises further question about the effectiveness of PCA as a dimensionality reduction technique. While PCA is widely used, it is known that its linearity of PCA may have limitations in handling high-dimensional datasets. Alternative dimensionality reduction techniques could be considered, each offering distinct advantages and drawbacks.^{13,47} However, in this work we are not interested in the specific results of the SOAP analysis or in how to optimize it, but rather on specific concepts that typically emerge in high-dimensional analyses. For this reason, instead of testing different automatic dimensionality reduction approaches, we investigate the

quantity, quality, and additivity of the information across individual SOAP components.

Information Quantity vs Information Quality in Noisy Datasets. The intrinsic high-dimensionality of the SOAP spectra allow us to directly examine the information contained in its orthogonal n components, without relying on any dimensionality reduction (thus avoiding the related issues). This is the primary reason why we used SOAP as a first example for this general study. PCA operated by selecting the PCs that maximize the explained variance, meaning, the amount (quantity) of captured information. In our analysis, the SOAP spectrum comprises 576 components, with approximately 53% of the variance of the entire SOAP dataset captured by component #63_{dn} ($n = n' = 8, l = 0$), as shown in Figure 4b. The next most significant components, in terms of variance, #62_{dn} and #61_{dn}, together with #63_{dn}, account for ~86% of the total variance. It is already notable that for a system with such complex internal physics (typically assumed to require high-dimensional analysis^{13,17,25,48}) ~86% of the cumulative variance is captured by just three SOAP components. Therefore, these three components are essentially contained within PC1_{dn} (see Figure S5). This is, to some extent, expected as components #63_{dn}, #62_{dn}, and #61_{dn} are in fact spherical components ($l = 0$), meaning that they contain information on the relative distances of the neighbors around every molecule: even though these do not contain any information on their orientations, this is enough to capture

the difference (in density) between the water and ice environments.

The SOAP power spectrum becomes particularly interesting and enriched with relevant information for components with $l > 0$. Unlike components with $l = 0$, which lack of information about the relative orientation of neighboring molecules and encode only their distances, components with $l > 0$ capture details about local symmetries and orientation. Figure 4c highlights several $l > 0$ components, ordered by their explained variance, starting from the highest. Notably, among them, component #237_{dn} is the one with the highest relevance: 6×10^{-6} . In any variance-based dimensionality reduction approach, such as PCA, all $l > 0$ components would likely be discarded due to their negligible statistical weight. However, while it might seem that these components contribute little to the understanding of the system's physics, this assumption is not necessarily accurate.

We investigated the effects of combining the two components with the highest variance from the $l = 0$ and $l > 0$ families, specifically #63_{dn} and #237_{dn}. By definition, these components are orthogonal in the SOAP space. In Figure 4a, these two components are shown. Notably, an Onion Clustering analysis reveals that #63_{dn} alone (Figure 4d) detects the interface more effectively than the PC1_{dn} in previous cases (Figure 3b). The increased separation between the peaks of the red, blue, and black Gaussian distributions demonstrates that the three different environments can be more easily discriminated in this case. As shown in Figure 4e, #63_{dn} achieves comparable stability and robustness in detecting the third cluster to that of PC1_{dn} in Figure 3c, despite capturing only $\sim 53\%$ of the variance compared to PC1_{dn}'s $\sim 98\%$. This underscores the critical insight that variance, typically interpreted as the amount of information, does not necessarily correlate with the relevance or quality of that information. Furthermore, the fact that a single denoised component (#63_{dn}) provides clearer and more discriminative information than PC1_{dn} suggests that the combination of #63_{dn}, #62_{dn}, #61_{dn}, and of the other components contributing to PC1_{dn} increases the noise, thereby diminishing the clarity of the analysis.

It is particularly fascinating to observe the results of Onion Clustering applied to the time-series of component #237_{dn}, a variable with negligible variance (6×10^{-6}) that would typically be overshadowed by the noise of the “heavier” components. In this case, up to four distinct physically relevant environments are robustly detected at resolutions up to $\Delta t \sim 3$ ns. Specifically, the solid/liquid interface splits into two distinct layers: one representing liquid water in contact and exchange with ice, and the other representing ice in contact and exchange with liquid water, alongside the bulk of ice and liquid water phases. This result highlights how a single component (#237_{dn}), despite it seemingly insignificant variance, can provide the most significant insights, underscoring the critical distinction between information quantity (variance) and information quality (relevance).

The Concept of Frustrated Information. In principle, combining two orthogonal components, such as #63_{dn} and #237_{dn}, should not result in signal degradation. At worst, their combination should preserve the maximum information obtainable from the individual components—specifically, detecting four clusters for $\Delta t < 3$ ns (guaranteed by #237_{dn}), followed by a reduction to three and then two clusters at coarser resolutions. However, the bidimensional analysis still

leads to significant information loss (Figure 5a, gray area). Interestingly, this deterioration is not primarily due to superposition of noise from the two components (as shown in Figure S6). In fact, the information loss observed when combining the original, noisy, #63 and #237, components is smaller than that observed in Figure 5a, which results from combining their denoised counterparts. Instead, this loss arises from an “information frustration” phenomenon, where the relevant information captured in one dimension (#63_{dn}) appears as additional noise in the other (#237_{dn}). One explanation of these effects is that, even if the raw components are mathematically orthogonal with each other by definition, this does not mean that the information that these contain refer to physically different events. Therefore, in cases such as this one, where the system is clearly monodimensional (Figure 4a), the information contained in all the individual (SOAP, PC, etc.) components is effectively redundant and correlated, so that they do not add anything to the most informative component (#237_{dn}). *Vice versa*, a typical problem is the opposite one: if a few components out of many are really informative, their combination in a high-dimensional analysis makes them statistically negligible, leading to information loss.

Notably, all our evidence demonstrated that this phenomenon is general and can occur independently on the number of dimensions, their individual noise levels (native *vs* denoised), or the dimensionality reduction approach employed in the analysis. To further support this concept, Figure S9 shows the manifold representations of the SOAP and SOAP_{dn} datasets (computed using UMAP^{13,49}), and the Onion Clustering results on UMAP1 and UMAP2 denoised (Figure S10). This approach operates directly on the full high-dimensional dataset—meaning combining all component information “correctly”—and captures comparable physically relevant information to the one detected by the single component #237_{dn} (Figure 4f,g). This confirms that discriminating the components by their variance is not necessarily the best approach, and better metrics should be used to extract the relevant information.⁵⁰ Differently from PCA, Time-lagged Independent Component Analysis (tICA)^{29,30} performed on the same system allows us to robustly identify the ice–water interface. Onion Clustering analysis of the tIC1 time-series data performs very well in characterizing the system, providing comparable results to, e.g., the Onion analysis of SOAP component #237_{dn} while adding also tIC2 to the analysis, at the same time, still leads to considerable frustration information effects (Figure S7). As one additional test, we have also tested VAMPnets³¹ as another (evolute) method to reduce the dimensionality of the 576-dimensional SOAP dataset. We analyzed the first two more informative components in different ways. As shown in Figure S8, a pattern recognition approach (neglecting the information contained in the time-correlations between the data) allows detecting three separated environments in the data, well visible on the v_1, v_2 plane. However, again, Onion clustering analyses of the time-series data contained along the two v_1 and v_2 components (taken separately) provide more information than a pattern recognition approach on the entire dataset on the v_1, v_2 space, allowing to resolve as many as 4–5 distinct environments. This demonstrates again how there can be more relevant information over time rather than on multiple spatial dimensions combined together in a “static” dataset. Nonetheless, again, Onion clustering analysis of the bivariate v_1, v_2 time-series leads to considerable information frustration

compared to analyzing v_1 or v_2 alone. Also in this case, thus, “less is more” when dealing with time-series analysis (Figure S8d). Together with all other tests conducted in this work, this underlines how, in order to maximize information extraction, one first key point is understanding how much information is nested along the time dimension before choosing the best method to analyze the data.

Information frustration effects are also observed when additional data are added in time (by increasing the sampling frequencies) rather than in space (by adding SOAP components). To investigate the effects of different sampling frequencies, we compared the results obtained analyzing the MD trajectories (raw data) sampled every $\Delta t = 40$ ps (Figures 2, 3, 4) with those obtained by increasing or decreasing the sampling frequency. We therefore resampled the 50 ns MD trajectory at intervals of 4, 20, 80, and 100 ps. The expectation is that Onion Clustering, which extracts relevant information about physical events in the system (for instance, the detectability of the ice/water interface environment up to a certain Δt), should maintain or enhance its ability to resolve this environment when using finer sampling ($\Delta t < 40$ ps). Conversely, coarser sampling ($\Delta t > 40$ ps) might compromise the resolution, potentially preventing proper detection of the interface. However, the results of these analyses reveal issues analogous to those encountered when adding more spatial information. As shown in Figure 5b, the black line represents the reference case ($\Delta t < 40$ ps): the third environment—the ice/water interface—is detectable within the resolution range $2 \text{ ns} < \Delta t < 4 \text{ ns}$ in the PC1 time-series, as seen in Figure 2c. Increasing the sampling time to 4 or 20 ps produces striking effects, as the temporal window within which the interface is detected shifts backward in the time-series. For a 20 ps sampling interval (which doubles the number of frames compared to the reference), the detectable temporal window for the interface shifts to $1 \text{ ns} < \Delta t < 2 \text{ ns}$ (green data). With a 4 ps sampling interval (a 10-fold increase in the number of analyzed frames), the interface is detectable only between 200 and 300 ps, a much shorter period compared to the nanosecond-scale range observed in other cases.

Remarkably, increasing the sampling frequency not only shifts the time window for detecting the same environment backward, but also makes its detection progressively less robust. For instance, doubling the analyzed data ($\Delta t = 20$ ps) reduces the detection robustness by 50% (from 2 ns with sampling $\Delta t = 40$ ps to 1 ns with $\Delta t = 20$ ps), and 20 times less robust when analyzing 10-fold data (100 ps with $\Delta t = 4$ ps). This is a nontrivial outcome. Physically, the environment characterized by molecular exchange between ice and liquid water remains the same in all cases. However, the addition of temporal data highlights a significant dependency on the chosen sampling interval. One might intuitively expect that increasing the sampling frequency would enhance resolution and simply allow earlier detection of the interface. Instead, oversampling leads to unexpected consequences: not only does it diminish the robustness of detecting the event, but it also introduces “data-drive hallucination”, where the same physical event appears to exhibit different dynamics depending on the time resolution. This phenomenon may suggest that the analysis itself is altering the reconstructed physics of the system. The addition of more data can generate an artifact, which appears to be an effect of oversampling and of the system’s local noise.

We repeated the same analysis on the denoised PC 1s (Figure 5c). The reduction of local noise supports such hypothesis: as the sampling interval increases, the detection of the interface shifts earlier, but the final time resolution at which detection concludes remains consistent. In this case, the detection of the interface is more robust across all sampling frequencies, with an average detection range of $0.8 \text{ ns} \leq \Delta t \leq 40 \text{ ns}$. This suggests that increasing the sampling frequency improves the ability to detect this dynamic environment sooner, while the lowest resolution capable of its detection remains consistent across different samplings. This behavior is governed by the inherent physics of the system, particularly the longest residence time of the molecules at the ice/water interface. Notably, the unique offset occurs in the case of a 4 ps sampling case (cyan). Here, even in the absence of noise, the interface is detectable only within a time window of ~ 10 ns, compared to the broader ~ 40 ns observed in the other cases. This provides a clear example of how oversampling, even with denoised data, can lead to information loss.

The data in Figure 5b,c, together with those in Figure 5a and of Figures 2, 3, and 4 highlight a general take-home message: adding more data increases the total amount of information, which can be understood as the sum of relevant information and noise. However, while adding more data does not guarantee an increase in relevant information, it invariably introduces additional noise, potentially leading to oversampling and to information loss. In principle, increasing the amount of information is beneficial only if it is proven *a priori* that the added data contain relevant information—a determination that requires preliminary analysis. It is important to recognize that these issues are particularly inherent in typical pattern recognition approaches applied to “static datasets”, where the temporal information is lost. Furthermore, these phenomena are not exclusive to specific descriptors or to the SOAP dataset. On the opposite, as demonstrated in the next section, the effect of information frustration and information loss can arise (to a smaller or larger extent) whenever two or more features are combined in an attempt to enrich the analysis.

Generality of the Concept: Quincke Rollers as a Second Different Case Study. In the tests of Figures 1–5 we used SOAP as a representative example of a high-dimensional descriptor, providing high-dimensional datasets. However, the same phenomena can occur independently of the specific system or the set of descriptors/variables selected. To illustrate the generality of these concepts, we analyzed a completely different system, characterized by unique complex internal dynamics. In this case, we examined time-series obtained from experimentally acquired trajectories of a system involving the two-dimensional collective motions of colloidal polymeric microparticles, known as Quincke rollers. Readers seeking more details on this specific system are referred to refs 18,28,51. Briefly, Quincke rollers are dielectric polystyrene particles ($\sim 9.9 \mu\text{m}$ in diameter) suspended in a weakly conductive fluid and confined within a two-dimensional cell. Under the influence of a weak vertical DC electric field, these particles can exhibit complex collective motions, such as vortexes and waves. While this is a higher-scale system, completely different from the ice–water molecular-scale one studied in the previous sections, it is useful to show how also in this case we encounter similar “information frustration” effects/issues, which are thus not limited to certain scales,

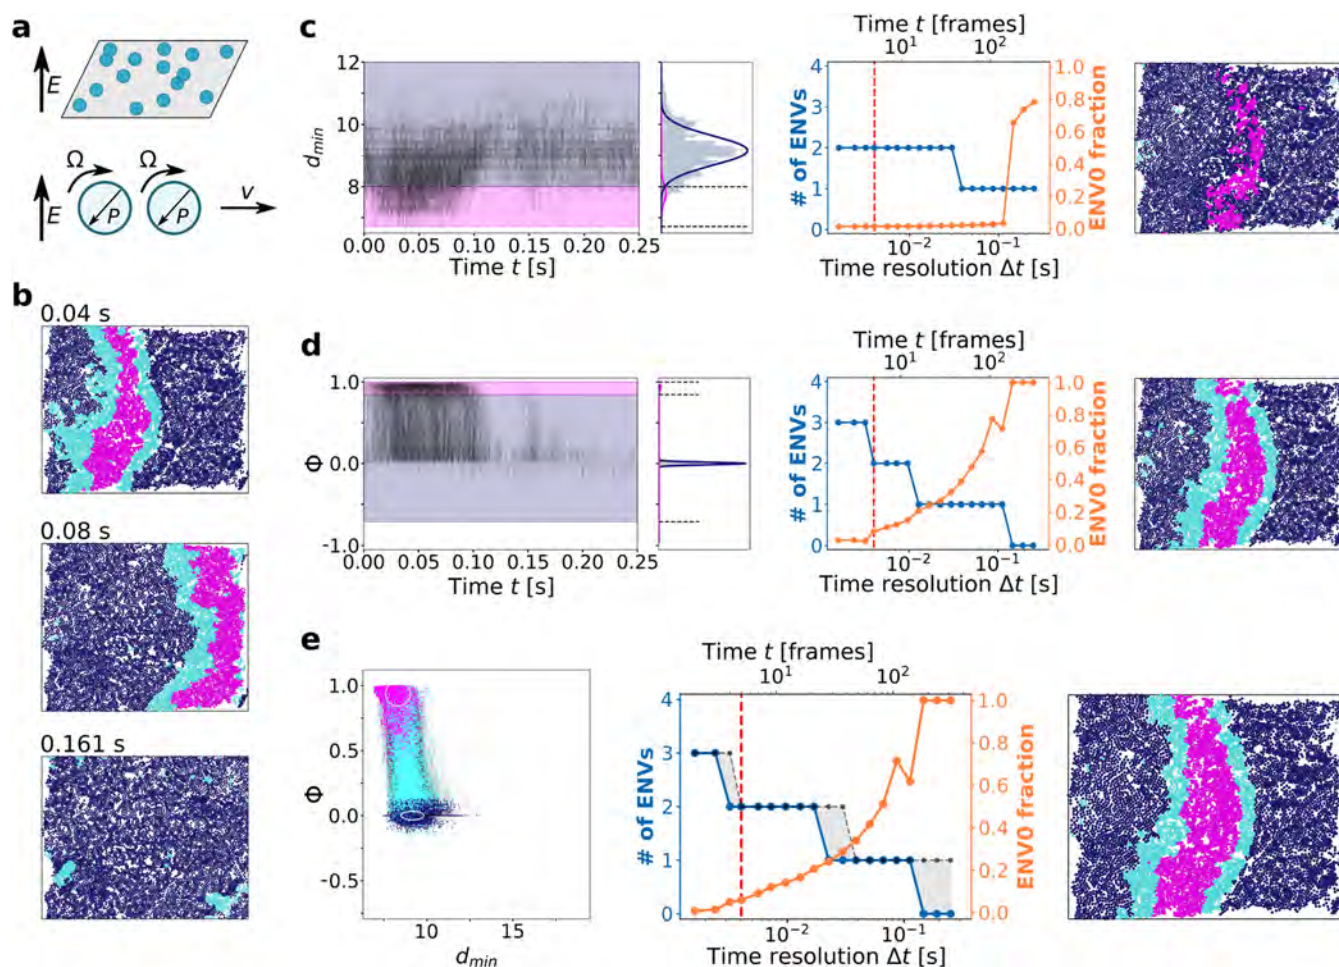


Figure 6. Onion Clustering on experimental trajectories of dynamical Quincke rollers. (a) Schematic representation of Quincke rollers. When suspended in conductive fluid and exposed to a vertical electric field, the particles exhibit collective motions.⁵¹ (b) Three snapshots from the video showing the wave moving from left to right, colored according to clusters detected by Onion Clustering on Φ . (c) Onion Clustering output for spatial averaged d_{\min} . (d) Same as panel (c), but for denoised Φ . (e) Same as in panel (c), but for the clustering on both spatial averaged d_{\min} and Φ .

systems, or descriptors, but are general in high-dimensional analyses of noisy systems.

Here, we analyze the trajectory we resolved for one example Quincke rollers system,²⁸ consisting of 6921 particles moving within a $700 \times 700 \mu\text{m}$ cell over a real-time duration of 0.25 s (310 frames in total). During this interval, a collective density wave is seen to propagate from left to right in the microscopy field (see Figure 6a,b).

For the sake of generality, in this case our analysis uses two human-based descriptors. Recently, we have demonstrated that integrating static descriptors, such as SOAP, with dynamic ones, such as LENS, can enrich high-dimensional analyses.³⁴ In the same spirit—and reflecting common practice in studying collective dynamical events in active matter⁵²—we selected two descriptors: (i) the minimum neighbor distance d_{\min} , as a proxy for the local particle density, and (ii) the local alignment of the particle velocity ϕ . This choice results in a bidimensional analysis (details are provided in the Methods section). We analyzed the denoised mono- and bidimensional time-series of the 6921 particles. Using Onion Clustering, d_{\min} was identified as a less informative descriptor (Figure 6c). From its time-series, two environments are detected up to $\Delta t \leq 30$ ms: one representing the wave's core (having higher density) and the other identifying static particles (with lower density). In

contrast, the particles' local velocity alignment ϕ (Figure 6d) is found to be more informative, capturing the system's collective motion and identifying three distinct environments for resolutions higher than $\Delta t \leq 3$ ms, and two environments for $\Delta t \leq 10$ ms. When three clusters are detected, Onion Clustering distinguishes the wave's core (Figure 6d in fuchsia) and two regions (cyan) where particles transition between static and dynamic states around the wave's core. These regions are shown in the snapshot in Figure 6d, alongside with the static particles' environment (dark violet). For $\Delta t \geq 140$ ms, Onion Clustering fails to resolve any statistically distinct dynamical cluster, indicating that, in the system, all particles change environment at list once during the trajectory. The residence time of particles in any cluster is therefore less than 140 ms.

While the variable ϕ is found more descriptive than d_{\min} for this system, it is interesting to examine the results of combining the two variables in a two-dimensional (ϕ , d_{\min}) time-series analysis. The outcomes are shown in Figure 6e. Comparing the number of clusters identified in the bidimensional analysis to those from the individual monodimensional analyses confirms that ϕ captures all relevant information, while d_{\min} contributes minimally. From this perspective, one might expect the two-dimensional analysis to at least match the

results obtained from the single-variable analysis of ϕ . However, the number of clusters detected in the two-dimensional analysis is diminished compared to the one-dimensional clustering of ϕ , as highlighted by the black line and gray area in Figure 6e, center. Thus, the addition of a second variable (d_{\min}) does not enrich the analysis, but instead leads to a loss of information. While the severity of the issue is somehow mitigated in this case—likely due to the combination of position-dependent (d_{\min}) and velocity-dependent (ϕ) information¹⁹—this demonstrates that combining multiple descriptors can lead to nontrivial phenomena of information frustration and loss. This can occur even in relatively simple cases where the two variables are markedly different. Notably, this problem is not exclusive to specific systems or combinations of descriptors. Rather, it is a general issue in (high-dimensional) analysis, stemming from the fact that information overloading—even when the individual pieces of information are intrinsically relevant—can hinder the ability to discern meaningful details within a dataset.

CONCLUSIONS

Understanding the physics of complex systems is inherently challenging. It is often assumed that high-dimensional analyses are not preferable, but necessary to comprehensively capture the underlining physics of such systems. However, the actual necessity of high-dimensionality, as well as the methods to extract meaningful physical information from such data, often remain unclear. In this work, we address this fundamental issue by systematically exploring the use of high-dimensional abstract descriptors and combinations of multiple physics-inspired ones. Our findings reveal nontrivial and counter-intuitive aspects of high-dimensional analyses, demonstrating that high-dimensionality is not always essential and, in some cases, can lead to frustrated information and even information loss.

As a prototypical case study, we analyzed Smooth Overlap of Atomic Position (SOAP) spectra¹⁰ extracted from MD simulations of coexisting ice and liquid water at dynamic equilibrium at the solid/liquid transition temperature. Typically, SOAP analyses of such systems involve dimensionality reduction methods, such as Principal Component Analysis (PCA) followed by clustering based on a combination of principal components (PCs) that account for high cumulative variance. However, approaches like this ignore the time-correlation in the data; this can result in significant information loss and limit the ability to capture key dynamic events.²⁸ To maximize the information extraction, we opted to analyze the SOAP data in time-series using Onion Clustering,²⁸ which allows resolving the information within individual SOAP PC1 and PC2 time-series, as well as their combination in bivariate time-series. Our results demonstrate that two-dimensional analysis systematically lead to information loss, independently of whether the time-series data are noisy or denoised.¹³ This phenomenon persists even when analyzing two of the original SOAP components, such as #63 and #237 (see Figures 4–5), and thus regardless of the chosen dimensionality reduction approach. We call this effect “frustrated information”, wherein relevant information captured in one dimension manifest as noise in another one. A schematic representation of the process that we followed in this work is shown in Figure 5d.

To ascertain what it is typically meant with “relevant information”, we systematically analyzed the individual SOAP components, separating them into radial ($l = 0$) and angular (l

< 0) contributions. Remarkably, our findings reveal that a specific angular component—#237—which accounts for mere $10^{-3}\%$ in the total variance (and is typically overlooked in variance-based dimensionality reduction approaches), contains more relevant and discernible information than all other components, or even the entire dataset. This shifts the focus from traditional “dimensionality reduction” to a more nuanced “data mining” perspective, highlighting that identifying the component or dimension containing the most relevant information is far more critical than simply prioritizing those with higher variance or combining multiple dimensions.

Distinguishing between the information contained in temporal *vs* spatial dimensions reveals that adding more information than necessary in space (high-dimensionality) can lead to “frustrated information”. Similarly, excessive information in time (oversampling) can result in “information hallucinations”, where the same physical event appears to have different features depending on the analysis (Figure 5). The varying time scales over which the ice/water interface—an environment characterized by water molecules undergoing solid-to-liquid and reverse phase transition—is detected in Figure 5b,c, clearly illustrate the challenges of analyzing noisy data. These findings underscore how the observed physics of a phenomenon can appear to shift dramatically based on the chosen sampling or measurement method.

While we used SOAP as a prototypical descriptor enabling high-dimensional analysis, the results shown in Figure 6, obtained for a system with entirely different scale, dynamics, and descriptors, demonstrate that such issues of nested, hidden, and frustrated information can emerge in virtually any type of multidimensional analysis. These findings provide clear examples of the difference existing between information quantity and quality (relevance), showing that relying solely on variance-based approaches (which prioritize quantity) can be very risky—especially in noisy, high-dimensional dataset, where the noise from one dimension may outweigh the relevant information in another one. At the same time, this work offers useful insights for navigating such complexities and extracting meaningful information from noisy dataset more effectively.

In conclusion, this manuscript aims to demonstrate that high-dimensional analyses do not necessarily provide more information than monodimensional ones. *Vice versa*, they may suffer of noise-addition and information-frustration effects that may hinder the extraction of information from data, especially when dealing with noisy datasets. In such cases, all the results discussed herein clearly demonstrate how it may be much more informative to (i) identify the most informative component/parameter among many and (ii) study it over time as a time-series (thus capturing also the information contained in the time dimension). This may guarantee to extract more information than by, e.g., analyzing all dimensions together in a high-dimensional analysis, but neglecting the information contained in the time-correlations between the data. As a consequence, one single feature can effectively describe a complex system better than a combination of many features, since “information frustration” effects can make them effectively “disturbing” each other. These are fundamental aspects that should be taken into consideration to avoid and minimize information loss when analyzing high-dimensional data in the study of any type of system. Finally, while proving the absolute character of the phenomena discussed here would require testing all possible analysis approaches for all possible

dataset types, it is worth noting that all comparisons and tests reported and discussed herein support the conclusion that this is not exclusive of specific techniques or dataset. The fact that different dimensionality reduction approaches (i.e., components-selection, PCA, tICA, VAMPnets) provide similar evidence indicates that a time-series analysis of one single dimension may provide more information than a pattern recognition approach considering all spatial dimensions together (but neglecting the information contained in the time-correlations of the data). At the same time, such precious information contained in the data—time-correlations may get lost when combining time-series in multivariate analyses. All this suggests important aspects to consider when choosing how to analyze data. A first important step is, for example, to test how much information is nested along the time dimension: this would provide precious indication for choosing between, e.g., a pattern recognition or a time-series analysis when trying to maximize information extraction from the data.

METHODS

Simulations and Data Analysis. Water/Ice MD Simulation. The atomistic system studied herein consists of 2048 TIP4P/ice molecules³² initially set 50% in *Ih* ice and 50% in liquid water phase, that we used to simulate the coexistence of hexagonal ice *Ih* and liquid water.⁵³ To maintain phase coexistence, MD simulations are performed at the melting temperature of 268 K in the TIP4P/ice model.³² The temperature is kept constant using the *v*-rescale thermostat, with a relaxation time of 0.2 ps. The system is equilibrated for 10 ns using the *c*-rescale pressure coupling, with water compressibility $c = 4.5 \times 10^{-5}$ bar, at ambient pressure. The equilibration and production run are carried out under *NPT* and semi-isotropic conditions using the GROMACS software.⁵⁴ The pressure is applied exclusively along the axis perpendicular to the ice/water interface. The production run is 50 ps long, sampled every 1 ps. The trajectory is then resampled based on the chosen sampling intervals (i.e., 4, 20, 40, 80, and 100 ps, see Figure Sb,c). Further details on the MD simulation can also be found in refs^{17,34}

Quincke Rollers Trajectory and Analysis. The trajectory studied in Figure 6 is obtained from microscopy videos from ref 51. The position of each particle at every frame are extracted by applying an image recognition and particle tracking, performed with *Trackpy*.⁵⁵ The raw trajectory consists of 6921 particles in a $700 \times 700 \mu\text{m}$ bidimensional cell over a real-time duration of 0.25 s sampled every $\Delta t = 0.8$ ms, for a total of 310 frames. For each molecule, we calculate the distance from the closest neighbor d_{min} and the local alignment of the velocities ϕ as in eq 1:

$$\phi_i = \frac{1}{n_c^i} \sum_j \frac{\mathbf{v}_i \cdot \mathbf{v}_j}{|\mathbf{v}_i| |\mathbf{v}_j|} \quad (1)$$

where j iterates over the n_c^i particles located within a specific cutoff distance of $r_c = 15$ pixels from particle i .²⁸ \mathbf{v}_i and \mathbf{v}_j are the velocities of particles i and j , respectively. The variable ϕ_i represents the average cosine of the angle between the velocity of particle i and those of its neighboring particles at every instant t . This value ranges from -1 to 1 , reflecting the degree of alignment or orientation similarity between the velocities of the particle and its neighbors. Further details can also be found in ref 28.

Smooth Overlap of Atomic Position (SOAP). The Smooth Overlap of Atomic Positions (SOAP)¹⁰ descriptor provides a high-dimensional description of the molecular arrangement of each atom/molecule by encoding its atomic/molecular environment about a specific central reference entity. Herein, we computed the SOAP descriptor centered on each oxygen atom. At each sampled step, the technique involves the application of a Gaussian smoothed density profile at the selected reference site; the SOAP environment takes into account the contributions of other particles closer than a certain cutoff radius r_{cut} . Then, the SOAP power spectra can be calculated as

$$\mathbf{p}(\mathbf{r})_{ml}^i = \pi \sqrt{\frac{8}{2l+1}} \sum_{m=-l}^l c_{nlm}^i(\mathbf{r}) \times c_{n'lm}^i(\mathbf{r}) \quad (2)$$

where c_{nlm}^i are the expansion coefficients of the particle density surrounding the i -th center.

In this work, we use SOAP¹⁰ as a prototypical example of high-dimensional descriptor. The SOAP power spectrum is calculated for each oxygen atom every 4 ps, using the *Dynsight* software, available from ref 56, considering all other oxygen atoms within a cutoff radius of $r_{\text{cut}} = 10 \text{ \AA}$ (previously demonstrated to be a good cutoff distance for aqueous systems).^{17–19,25,28,34} The SOAP vectors are computed using the *Dscribe* Python package,³⁶ using both $l_{\text{max}} = n_{\text{max}} = 8$ and $l_{\text{max}} = n_{\text{max}} = 4$, for comparison, while the other parameters are kept as default. In the first case, we obtain 576 components for the spectrum of each oxygen atom at each time step, while in the second case, the number of components is reduced to 80, globally obtaining in both cases consistent results (see Figures S1–S2). The PCA to reduce the dimensionality of the dataset (Figures 2, 3) were performed using the *SciPy* Python package.⁵⁷

Time-Series Analysis with Onion Clustering. Onion Clustering is a recently developed algorithm for single-point time-series clustering, which allows to automatically identify stable states in noisy time-series data. The algorithm requires a single parameter, the time resolution, indicated with Δt . This quantity sets the minimum duration of the signal sequences that are classified in a certain state; this ensures the stability of the states. Sequences that cannot be assigned to any state for at least Δt frames are grouped in the “unclassified” cluster (ENV0). Onion Clustering proceeds with an iterative approach. At each iteration, (i) A Gaussian is fitted on the highest peak of the data probability density (ii) All the signal sequences closer than a certain threshold to the mean of the Gaussian are assigned to the state characterized by that Gaussian (iii) These signal sequences are removed from the dataset, changing the data probability density and the next iteration starts. The iterative procedure continues until there are still data points to be classified, and it is possible to assign them to a Gaussian state. In this way, in the end all the signal sequences are either assigned to one of the Gaussian states or left unclassified. In this work, to filter out statistically insignificant states, the ones with population below 1% are removed and their points considered unclassified. Then, to remove the *a priori* (and possibly biased) choice of the time resolution Δt , for every dataset the clustering is performed over a broad range of possible Δt values, ranging from the highest time resolution (Δt double the sampling time) to the lowest possible one (Δt equal to the entire trajectory length). We just underline that performing the analysis at varying time-

resolutions (Δt) helps to automatically identify the optimal Δt that maximizes the subdivision of the time-series (and the system) into distinct statistically relevant microscopic environments and minimizes the fraction of unclassifiable points (ENVO). For further details, we refer the readers to ref 28. for details on the onion method, and to refs 16,18 for other recent applications. The Onion Clustering code is available on the platform Dynsight.⁵⁶

■ ASSOCIATED CONTENT

Data Availability Statement

Complete data and materials pertaining to the trajectories employed, and the analysis conducted herein, are available at: <https://doi.org/10.5281/zenodo.15782997>. Other information needed is available from the corresponding author upon reasonable request.

SI Supporting Information

The Supporting Information is available free of charge at <https://pubs.acs.org/doi/10.1021/acs.jctc.5c00374>.

Data related to Onion clustering results of PC1 and PC2 from SOAP $l_{\max} = n_{\max} = 4$ and on its most relevant components; Onion Clustering on PC3 and PC4 both denoised and raw; Onion Clustering on denoised components #62 and #61; Onion Clustering on components #63 and #237 raw; Onion Clustering on tIC1 and tIC2 raw; Onion Clustering on VAMPnets v1 and v2; hierarchical clustering and Onion Clustering on UMAP1 and UMAP2 (PDF)

■ AUTHOR INFORMATION

Corresponding Author

Giovanni M. Pavan – Department of Applied Science and Technology, Politecnico di Torino, Torino 10129, Italy;
orcid.org/0000-0002-3473-8471;
Email: giovanni.pavan@polito.it

Authors

Chiara Lionello – Department of Applied Science and Technology, Politecnico di Torino, Torino 10129, Italy;
orcid.org/0000-0002-7491-8952

Matteo Becchi – Department of Applied Science and Technology, Politecnico di Torino, Torino 10129, Italy;
orcid.org/0000-0002-6306-5229

Simone Martino – Department of Applied Science and Technology, Politecnico di Torino, Torino 10129, Italy;
orcid.org/0009-0009-7369-3809

Complete contact information is available at: <https://pubs.acs.org/10.1021/acs.jctc.5c00374>

Notes

The authors declare no competing financial interest.

■ ACKNOWLEDGMENTS

G.M.P. acknowledges the support received by the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation program (Grant Agreement no. 818776- DYNAPOL). G.M.P. and C.L. also acknowledge the support received by the ICSC—Centro Nazionale di Ricerca in High Performance Computing, Big Data and Quantum Computing, funded by European Union - NextGenerationEU.

■ REFERENCES

- (1) Wolde, P. R. T.; Frenkel, D. Enhancement of protein crystal nucleation by critical density fluctuations. *Science* **1997**, *277*, 1975–1978.
- (2) Lutsko, J. F. How crystals form: A theory of nucleation pathways. *Sci. Adv.* **2019**, *5*, eaav7399.
- (3) Hagan, M. F.; Grason, G. M. Equilibrium mechanisms of self-limiting assembly. *Rev. Mod. Phys.* **2021**, *93*, 025008.
- (4) Crippa, M.; Perego, C.; de Marco, A. L.; Pavan, G. M. Molecular communications in complex systems of dynamic supramolecular polymers. *Nat. Commun.* **2022**, *13*, 2162.
- (5) Cavagna, A.; Cimarelli, A.; Giardinà, I.; Parisi, G.; Santagati, R.; Stefanini, F.; Viale, M. Scale-free correlations in starling flocks. *Proc. Natl. Acad. Sci. U.S.A.* **2010**, *107*, 11865–11870.
- (6) Nagy, M.; Ákos, Z.; Biro, D.; Vicsek, T. Hierarchical group dynamics in pigeon flocks. *Nature* **2010**, *464*, 890–893.
- (7) Porfiri, M. Inferring causal relationships in zebrafish-robot interactions through transfer entropy: a small lure to catch a big fish. *Anim. Behav. Cognit.* **2018**, *5*, 341–367.
- (8) Facco, E.; d'Errico, M.; Rodriguez, A.; Laio, A. Estimating the intrinsic dimension of datasets by a minimal neighborhood information. *Sci. Rep.* **2017**, *7*, 12140.
- (9) Noia, A. D.; Macocco, I.; Glielmo, A.; Laio, A.; Mira, A. Beyond the noise: intrinsic dimension estimation with optimal neighbourhood identification. *ArXiv* 2024, Accessed: 2024–07–12, DOI: [10.48550/arXiv.2405.15132](https://doi.org/10.48550/arXiv.2405.15132).
- (10) Bartók, A. P.; Kondor, R.; Csányi, G. On representing chemical environments. *Phys. Rev. B* **2013**, *87*, 184115.
- (11) Drautz, R. Atomic cluster expansion for accurate and transferable interatomic potentials. *Phys. Rev. B* **2019**, *99*, 014104.
- (12) Nigam, J.; Pozdnyakov, S.; Ceriotti, M. Recursive evaluation and iterative contraction of N-body equivariant features. *J. Chem. Phys.* **2020**, *153*, 121101.
- (13) Donkor, E. D.; Offei-Danso, A.; Rodriguez, A.; Sciortino, F.; Hassanali, A. Beyond Local Structures in Critical Supercooled Water through Unsupervised Learning. *J. Phys. Chem. Lett.* **2024**, *15*, 3996–4005.
- (14) Offei-Danso, A.; Hassanali, A.; Rodriguez, A. High-Dimensional Fluctuations in Liquid Water: Combining Chemical Intuition with Unsupervised Learning. *J. Chem. Theory Comput.* **2022**, *18*, 3136–3150.
- (15) Ansari, N.; Onat, B.; Sosso, G. C.; Hassanali, A. Insights into the Emerging Networks of Voids in Simulated Supercooled Water. *J. Phys. Chem. B* **2020**, *124*, 2180–2190.
- (16) Martino, S.; Doria, D.; Lionello, C.; Becchi, M.; Pavan, G. M. A data driven approach to classify descriptors based on their efficiency in translating noisy trajectories into physically-relevant information. *ArXiv* 2024, Accessed: 2024–12–15, DOI: [10.48550/arXiv.2411.12570](https://doi.org/10.48550/arXiv.2411.12570).
- (17) Caruso, C.; Cardellini, A.; Crippa, M.; Rapetti, D.; Pavan, G. M. TimeSOAP: Tracking high-dimensional fluctuations in complex molecular systems via time variations of SOAP spectra. *J. Chem. Phys.* **2023**, *158*, 214302.
- (18) Caruso, C.; Crippa, M.; Cardellini, A.; Cioni, M.; Perrone, M.; Piane, M. D.; Pavan, G. M. Classification and spatiotemporal correlation of dominant fluctuations in complex dynamical systems. *PNAS Nexus* **2025**, *4*, pgaf038.
- (19) Crippa, M.; Cardellini, A.; Cioni, M.; Csányi, G.; Pavan, G. M. Machine learning of microscopic structure-dynamics relationships in complex molecular systems. *Mach. Learn.: Sci. Technol.* **2023**, *4*, 045044.
- (20) Rapetti, D.; Delle Piane, M.; Cioni, M.; Polino, D.; Ferrando, R.; Pavan, G. M. Machine learning of atomic dynamics and statistical surface identities in gold nanoparticles. *Commun. Chem.* **2023**, *6*, 143.
- (21) Cioni, M.; Delle Piane, M.; Polino, D.; Rapetti, D.; Crippa, M.; Irmak, E. A.; Van Aert, S.; Bals, S.; Pavan, G. M. Sampling Real-Time Atomic Dynamics in Metal Nanoparticles by Combining Experiments, Simulations, and Machine Learning. *Adv. Sci.* **2024**, *11*, 2307261.

- (22) Perrone, M.; Cioni, M.; Piane, M. D.; Pavan, G. M. Unsupervised Tracking of Local and Collective Defects Dynamics in Metals Under Deformation. *ArXiv* 2024, Accessed: 2024-11-29, DOI: 10.48550/arXiv.2410.20999.
- (23) Lionello, C.; Perego, C.; Gardin, A.; Klajn, R.; Pavan, G. M. Supramolecular Semiconductivity through Emerging Ionic Gates in Ion-Nanoparticle Superlattices. *ACS Nano* 2023, 17, 275–287.
- (24) Lin, A.; Huguenin-Dumittan, K. K.; Cho, Y.-C.; Nigam, J.; Cersonsky, R. K. Expanding density-correlation machine learning representations for anisotropic coarse-grained particles. *J. Chem. Phys.* 2024, 161, 074112.
- (25) Capelli, R.; Muniz-Miranda, F.; Pavan, G. M. Ephemeral ice-like local environments in classical rigid models of liquid water. *J. Chem. Phys.* 2022, 156, 214503.
- (26) Gardin, A.; Perego, C.; Doni, G.; Pavan, G. M. Classifying soft self-assembled materials via unsupervised machine learning of defects. *Commun. Chem.* 2022, 5, 82.
- (27) Gasparotto, P.; Boichichio, D.; Ceriotti, M.; Pavan, G. M. Identifying and Tracking Defects in Dynamic Supramolecular Polymers. *J. Phys. Chem. B* 2020, 124, 589–599.
- (28) Becchi, M.; Fantolino, F.; Pavan, G. M. Layer-by-layer unsupervised clustering of statistically relevant fluctuations in noisy time-series data of complex dynamical systems. *Proc. Natl. Acad. Sci. U.S.A.* 2024, 121, e2403771121.
- (29) Molgedey, L.; Schuster, H. G. Separation of a mixture of independent signals using time delayed correlations. *Phys. Rev. Lett.* 1994, 72, 3634–3637.
- (30) Pérez-Hernández, G.; Paul, F.; Giorgino, T.; De Fabritiis, G.; Noé, F. Identification of slow molecular order parameters for Markov model construction. *J. Chem. Phys.* 2013, 139, 015102.
- (31) Mardt, A.; Pasquali, L.; Wu, H.; Noé, F. VAMPnets for deep learning of molecular kinetics. *Nat. Commun.* 2018, 9, 5.
- (32) Abascal, J. L. F.; Sanz, E.; García Fernández, R.; Vega, C. A potential model for the study of ices and amorphous water: TIP4P/Ice. *J. Chem. Phys.* 2005, 122, 234511.
- (33) Fitzner, M.; Sosso, G. C.; Cox, S. J.; Michaelides, A. Ice is born in low-mobility regions of supercooled liquid water. *Proc. Natl. Acad. Sci. U.S.A.* 2019, 116, 2009–2014.
- (34) Crippa, M.; Cardellini, A.; Caruso, C.; Pavan, G. M. Detecting dynamic domains and local fluctuations in complex molecular systems via timelapse neighbors shuffling. *Proc. Natl. Acad. Sci. U.S.A.* 2023, 120, e2300565120.
- (35) Donkor, E. D.; Laio, A.; Hassanali, A. Do machine-learning atomic descriptors and order parameters tell the same story? The case of liquid water. *J. Chem. Theory Comput.* 2023, 19, 4596–4605.
- (36) Himanen, L.; Jäger, M. O.; Morooka, E. V.; Canova, F. F.; Ranawat, Y. S.; Gao, D. Z.; Rinke, P.; Foster, A. S. DScribe: Library of descriptors for machine learning in materials science. *Comput. Phys. Commun.* 2020, 247, 106949.
- (37) Laakso, J.; Himanen, L.; Himm, H.; Morooka, E. V.; Jäger, M. O.; Todorović, M.; Rinke, P. Updates to the DScribe library: New descriptors and derivatives. *J. Chem. Phys.* 2023, 158, 234802.
- (38) Pearson, K. LIII. On lines and planes of closest fit to systems of points in space. *The London, Edinburgh, and Dublin Philos. Mag. J. Sci.* 1901, 2, 559–572.
- (39) Hotelling, H. Relations Between Two Sets of Variates. *Biometrika* 1936, 28, 321.
- (40) Jolliffe, I. T.; Cadima, J. Principal component analysis: a review and recent developments. *Philos. Trans. R. Soc., A* 2016, 374, 20150202.
- (41) Nielsen, F. *Introduction to HPC with MPI for Data Science*; Springer International Publishing, 2016; pp 195–211.
- (42) Gupta, M.; Gao, J.; Aggarwal, C. C.; Han, J. Outlier detection for temporal data: A survey. *IEEE Trans. Knowl. Data Eng.* 2014, 26, 2250–2267.
- (43) Butler, B. L.; Fijan, D.; Glotzer, S. C. Change point detection of events in molecular simulations using dupin. *Comput. Phys. Commun.* 2024, 304, 109297.
- (44) Empereur-mot, C.; Pedersen, K. B.; Capelli, R.; Crippa, M.; Caruso, C.; Perrone, M.; Souza, P. C. T.; Marrink, S. J.; Pavan, G. M. Automatic Optimization of Lipid Models in the Martini Force Field Using SwarmCG. *J. Chem. Inf. Model.* 2023, 63, 3827–3838.
- (45) Friedman, J. H. On bias, variance, 0/1—loss, and the curse-of-dimensionality. *Data Min. Knowl. Discovery* 1997, 1, 55–77.
- (46) Bishop, C. M.; Nasrabadi, N. M. *Pattern recognition and machine learning*; Springer, 2006; Vol. 4.
- (47) Tribello, G. A.; Gasparotto, P. Using Dimensionality Reduction to Analyze Protein Trajectories. *Front. Mol. Biosci.* 2019, 6, 46.
- (48) Monserrat, B.; Brandenburg, J. G.; Engel, E. A.; Cheng, B. Liquid water contains the building blocks of diverse ice phases. *Nat. Commun.* 2020, 11 (1), 5757.
- (49) McInnes, L.; Healy, J.; Saul, N.; Großberger, L. UMAP: Uniform Manifold Approximation and Projection. *J. Open Source Softw.* 2018, 3, 861.
- (50) Ebrahimi, N.; Maasoumi, E.; Soofi, E. S. Ordering univariate distributions by entropy and variance. *J. Econometrics* 1999, 90, 317–336.
- (51) Liu, Z. T.; Shi, Y.; Zhao, Y.; Chaté, H.; Shi, X.-q.; Zhang, T. H. Activity waves and freestanding vortices in populations of subcritical Quincke rollers. *Proc. Natl. Acad. Sci. U.S.A.* 2021, 118, e2104724118.
- (52) Peruani, F.; Deutsch, A.; Bär, M. A mean-field theory for self-propelled particles interacting by velocity alignment mechanisms. *Eur. Phys. J. Spec. Top.* 2008, 157, 111–122.
- (53) Matsumoto, M.; Yagasaki, T.; Tanaka, H. GenIce: Hydrogen-Disordered Ice Generator. *J. Comput. Chem.* 2018, 39, 61–64.
- (54) Abraham, M. J.; Murtola, T.; Schulz, R.; Páll, S.; Smith, J. C.; Hess, B.; Lindahl, E. GROMACS: High performance molecular simulations through multi-level parallelism from laptops to supercomputers. *SoftwareX* 2015, 1–2, 19–25.
- (55) Allan, D. B.; Caswell, T.; Keim, N. C.; van der Wel, C. M.; Verweij, R. W. *soft-matter/trackpy: v0.6.4*, Accessed: 2023-09-15, 2024.
- (56) Dynsight Github repository. <https://github.com/GMPavanLab/dynsight>, Accessed: 2024-07-15. 2023.
- (57) Virtanen, P.; Gommers, R.; Oliphant, T. E.; Haberland, M.; Reddy, T.; Cournapeau, D.; Burovski, E.; Peterson, P.; Weckesser, W.; Bright, J.; et al. SciPy 1.0: fundamental algorithms for scientific computing in Python. *Nat. Methods* 2020, 17, 261–272.