JCIM JOURNAL OF CHEMICAL INFORMATION AND MODELING

pubs.acs.org/jcim Article

Systematic Study of Machine Learning Classification Algorithms of Zeolitic Imidazolate Framework Polymorphs

Emilio Méndez,[§] Léna Triestram,[§] Dune André, François-Xavier Coudert,* and Rocio Semino*

Cite This: *J. Chem. Inf. Model.* 2026, 66, 10046−10055 Read Online

ACCESS | Metrics & More | Article Recommendations | Supporting Information

ABSTRACT: Zeolitic imidazolate frameworks (ZIFs) are a family of metal−organic frameworks that feature metal centers tetrahedrally linked to imidazole-based ligands and adopt zeolite-like topologies. ZIFs formed by zinc cations and imidazolate linkers exhibit a remarkable degree of polymorphism, which can be modulated by varying synthesis parameters or thermodynamic conditions (i.e., temperature and pressure). Computer simulations provide a unique way of studying these materials and their phase transitions from the microscopic standpoint, revealing their underlying molecular mechanisms. However, studying these mechanisms requires to be able to classify the phase of each molecular entity in an agnostic and automatic fashion, which is particularly challenging when the two phases involved are structurally very similar. In this work, we systematically study neural network classifiers to classify ZIF phases on-the-fly during molecular dynamics simulations. We test a variety of input features, differing both in the dimensionality and nature of the descriptors and in the kind of force field used for building the training/testing database. We reveal that even with low-dimensional descriptors, the classification is highly accurate, while the use of high-dimensional descriptors leads to an even better performance. Training the classifier with configurations coming from different force fields, we can remove force field bias and enhance the classifier performance and general applicability. Finally, we apply our classifiers to reveal mechanistic details of the ZIF-4-cp → ZIF-4-cp-II phase transition.

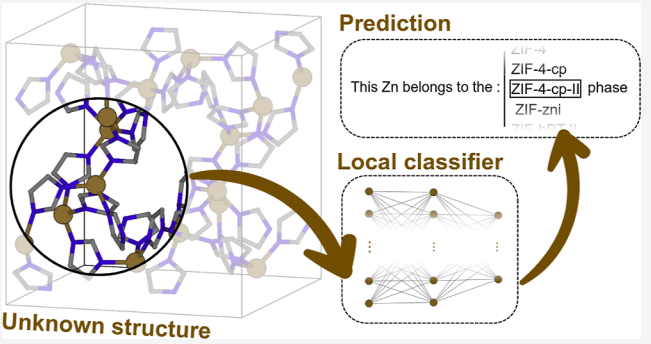

INTRODUCTION

Metal−organic frameworks, or MOFs, are hybrid organic−inorganic supramolecular assemblies based on relatively weak interactions that present a high number of intramolecular degrees of freedom. Due to this combination of factors, they present a higher propensity than classical solids for phenomena such as large-scale flexibility, framework dynamics, stimuli-responsive behavior, the presence of defects and disorder.[1−3] For the same reasons, many MOFs also display polymorphism, i.e., the ability to exhibit two or more crystalline phases that differ in the arrangement of the organic linkers that form them or their conformation within the crystal lattice. Such MOFs can therefore display several phases with the same chemical composition but with differences in their structure, impacting or not the topology of the framework. As is known in the case of molecular crystals, polymorphism has an important impact on both the physical and chemical properties of MOFs[4] and has been widely studied both experimentally and computationally.[5,6]

Zeolitic imidazolate frameworks (ZIFs) are a typical example for the occurrence of polymorphism in MOFs, with dozens of reported Zn(imidazole)$_2$ phases varying in topology as a function of synthesis conditions, thermal and mechanical treatment, and other influences during and after their synthesis.[7,8] This extends even to the existence of disordered phases, including liquids, glasses, gels, and composites thereof.[9−11] However, despite the large amount of experimental evidence gathered on the polymorphism of ZIFs, many important questions remain widely open. The exact pressure−temperature phase diagram of ZIF-4, for example, has not yet been fully solved; although several studies have addressed the question,[8,12] these phase diagrams do not include many of the ZIF polymorphs originally found by Park and co-workers[13] nor are they really representative of thermodynamic stability, since it is well-known that the most stable phases are the dense ones, and porous phases occur due to kinetic considerations. Moreover, the nature and structure of some of the high-pressure phases of ZIF-4, ZIF-8, and ZIF-62 are still not fully solved,[14−16] the differences in the structures of the known phases can be quite subtle and some of the phases have a very narrow thermodynamic stability range.[17]

Received: April 16, 2026
Revised: July 22, 2026
Accepted: July 24, 2026
Published: August 4, 2026

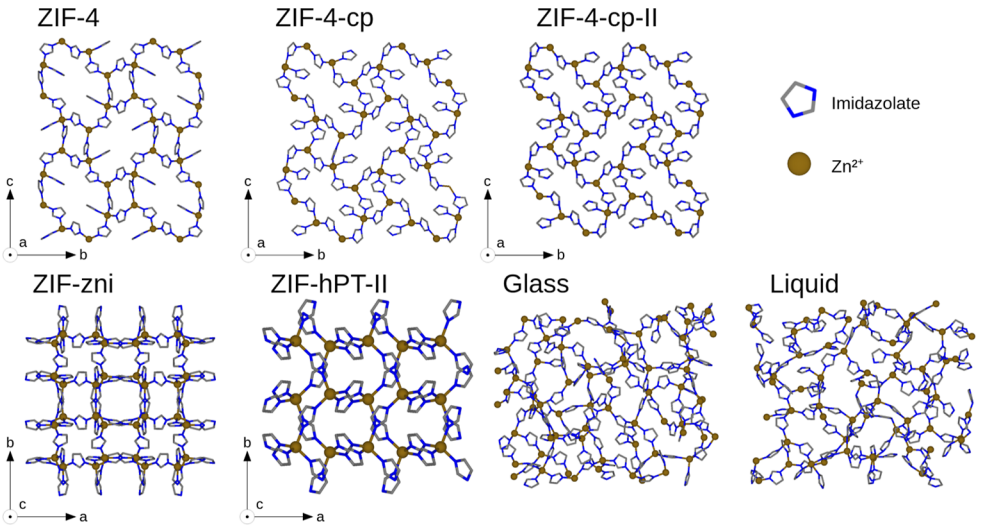

Figure 1. Structures of the polymorphs of ZIF-4 studied. Color code: Zn (ochre), N (blue), C (gray). Hydrogen atoms are not shown for clarity.

Computational techniques give us access to molecular-level resolution, which is key in revealing structural and mechanistic information pertaining to ZIF polymorphs and their phase transformations.[12,18−20] However, simulation trajectories contain large volumes of data which require automation and agnostic data analysis methods to reveal all their hidden potential. For these reasons, there is interest in the development of computational methodologies for the identification of phases in ZIF systems. The challenge is even greater for the assignment of phases in a time-resolved and local manner. Being able to determine, during a molecular dynamics simulation, that a specific atomic environment at a given time is closest to a specific phase, allows for the study of the dynamics of phase transitions and their mechanisms at the atomistic scale. This is made difficult because differences between structures can be subtle, and simple geometrical criteria based on intuition are largely insufficient.[21] Techniques for local classification of phases have been employed previously: for example, geometric methods such as polyhedral template matching[22] have been used to classify amorphous phases of silicon.[23] In a previous study, a database of imidazolate-based ZIF polymorphs was created and used to train a neural-network classifier to identify ZIF polymorphs.[20]

In this work, we extend this previous methodology for local phase identification in ZIF systems in several directions. First, we broaden the classification problem from ambient-pressure phases to a larger family of crystalline and amorphous polymorphs spanning a wide range of thermodynamic conditions, including closely related high-pressure phases that are considerably more difficult to distinguish. Second, we investigate two complementary classes of local descriptors: Behler−Parrinello symmetry functions (BPSF)[24] and Smooth Overlap of Atomic Positions (SOAP),[25] differing both in dimensionality and in the structural information they encode, allowing us to assess the trade-off between computational efficiency and predictive accuracy. Finally, we investigate the effect of combining training data obtained from two different force fields, with the objective of improving the robustness and transferability of the resulting classifiers across interaction potentials. Beyond the development of the classifiers themselves, we demonstrate their usefulness by applying them to the microscopic analysis of the ZIF-4-cp/ZIF-4-cp-II transformation. Unlike the crystal-to-amorphous and melting transitions investigated previously, this transformation occurs between two crystalline polymorphs with very subtle structural differences, providing a significantly more demanding benchmark for local phase identification. We identify a preferential direction in which the phase transition is slower, and correlate it to a change of relative orientation in the rings that are present in both closed-pore structures.

SYSTEMS AND METHODS

ZIF-4 Polymorphs and Their Classification

We aim to develop robust methods of phase classification in ZIFs that are able to identify the phases to which a Zn-centered environment belongs within a series of polymorphs containing both ordered and disordered phases. To do so, we present two phase classifiers based on two general-purpose chemical descriptors.

We focused on seven distinct polymorphs with formula $Zn(Im)_2$, where Im stands for the unsubstituted imidazolate anion: $C_3H_3N_2^-$. These structures include five crystalline phases and two disordered phases which are shown in Figure 1. These polymorphs were previously described in both experimental[8,26] and theoretical[12,18−20] studies and can be obtained by varying the pressure and temperature (see the experimentally measured phase diagram in Figure S14 and in ref 8). Among the crystalline phases, three frameworks share the *cag* topology: ZIF-4 is the open-pore phase, which undergoes a phase transition to the closed-pore phases when pressure is increased (ZIF-4-cp and ZIF-4-cp-II). The two remaining crystal phases ZIF-zni and ZIF-hPT-II have distinct topologies: *zni* and *dia*, respectively. In all of these crystalline phases, the linkers adopt a tetrahedral geometry around the metal ion. We also include two phases that lack long-range order: a low-temperature glass phase and a high-temperature liquid phase.

Simulation Methods

In order to build our classifiers, we first needed to create a database of configurations. To do so, we performed molecular dynamics (MD) simulations of each of the phases studied. To train generalizable classifiers, we produced configurations with two different kinds of force fields: the classical nb-ZIF-FF force field[27] and a machine-learning potential based on the MACE

https://doi.org/10.1021/acs.jcim.6c01175
J. Chem. Inf. Model. 2026, 66, 10046−10055

Table 1. Simulation Conditions of Temperature T, Pressure P, Density ρ, and Number of Zn Atoms N_{Zn} in Each Configuration for the Generation of the Data Set[a]

phase	force field	ensemble	T (K)	P (MPa)	ρ (Zn/nm^3)	N_{Zn}
liquid	nb-ZIF-FF	*NPT*	700	0	4.05	64
liquid	MACE	*NVT*	1500	803	3.65	128
glass	nb-ZIF-FF	*NPT*	300	0	4.75	64
glass	MACE	*NVT*	300	482	3.65	128
ZIF-4	nb-ZIF-FF	*NPT*	300	0	4.03	128
ZIF-4	MACE	*NVT*	300	633	3.65	128
ZIF-4-cp	nb-ZIF-FF	*NPT*	300	20 and 105	4.76 and 4.87	128
ZIF-4-cp	MACE	*NVT*	300	1218 and 1494	4.70 and 4.87	128
ZIF-4-cp-II	nb-ZIF-FF	*NPT*	300	150	5.11	128
ZIF-4-cp-II	MACE	*NVT*	300	1448	4.92	128
ZIF-hPT-II	nb-ZIF-FF	*NPT*	563	810	7.05	64
ZIF-hPT-II	MACE	*NVT*	560	7263	6.54	64
ZIF-zni	nb-ZIF-FF	*NPT*	300	0	4.92	128
ZIF-zni	MACE	*NVT*	300	508	4.43	64

[a]For simulations in the *NPT* ensemble, the density corresponds to the average value along the trajectory.

architecture.[28] On the one hand, the nb-ZIF-FF model is a partially reactive force field in which the Zn-Ligand bonds are modeled via Morse potentials, allowing the possibility of bond breaking and formation along a trajectory. The parameters of this model were optimized to reproduce experimental thermodynamic properties of different ZIF phases as well as *ab initio* Zn−ligand binding energies and it has been successfully applied to modeling phase transitions[12,20] and self-assembly[29,30] of ZIFs. On the other hand, the machine learning potential (MLP) is trained by fitting the energies and forces obtained from electronic structure calculations without imposing any connectivity at all. The potential was based on the MACE code[31,32] (version 0.3.9). It was demonstrated in previous work by Triestram and Coudert[28] that this MLP provides a chemically accurate description of the different phases of ZIF-4 and is transferable across wide temperature and pressure ranges. It has been trained on data sampled from the liquid phase of ZIF-4 through *ab initio* molecular dynamics.[18] Since the two force fields have been constructed by very different methods, we consider that a mixture of configurations obtained by both methods may lead to model-agnostic classifiers.

We aim to classify the phases at the local Zn-centered environment level. In this way, each of the obtained configurations will include a number of data points equal to the number of Zn atoms present in the simulations. Then, from the total set of configurations, we obtained 672000 Zn environments, 96000 for each phase (7 phases are considered) simulated with both force fields. In Table 1, we summarize the thermodynamic conditions at which each simulation was performed. These conditions were selected to lie inside the respective stability regions of each phase.

For the systems run with the MLP, we performed the simulations under the *NVT* ensemble because the force field does not reproduce the correct density in the *NPT* ensemble.[28] The simulation details of each setup are summarized in the Supporting Information. Using both *NVT* and *NPT* for the simulations allows us to obtain more diversity in the training data. Since ZIF-4-cp and ZIF-4-cp-II exhibit very similar local environments, configurations corresponding to two thermodynamic conditions within the stability region of both phases were included in the training data set. This broader sampling increases the structural variability during training and improves the accuracy of the discrimination between these closely related phases. The effect of this data augmentation in the performance of the classifier will be studied in the "Case study: ZIF-4-cp/ZIF-4-cp-II Transitions" subsection of our Results & Discussion.

Classifiers

To generate methods that allow us to analyze structural transformations at a local level, we trained neural network classifiers based on local descriptors. As a result, we obtain as many predictions for a given structure as there are Zn atoms in the simulation box. This allows us to use the classifier in structures of different sizes and analyze processes at the atomic scale. These metal-centered predictions can also be averaged to yield global predictions at the framework structure level.

This approach was implemented in previous work by Méndez and Semino to classify ZIF-4, ZIF-zni, glass ZIF, and liquid ZIF phases and the ZIF-4 ↔ glass phase transition.[20] We now expand the methodology to include the rest of the ZIF phases from Figure 1, thus including structures that are stable at pressures higher than ambient pressure. We also aim to test the use of a more complete set of descriptors of Zn-centered environments and compare this method to a simpler version that is just the extension of that developed in ref 20, which relies on only 12 features per Zn atom.

We thus generated two classifiers that differ in the nature and dimensionality of the descriptors used to encode Zn environments. One of them is based on BPSFs, which are descriptors originally used for the development of machine learning potentials.[24] This method has a low-associated computational cost since only 12 BPSFs were selected to represent Zn-centered environments and only Zn neighboring atoms within the environment were included in the spatial correlations, neglecting any ligand information. As a result, a vector of size 12 that encodes the spatial distribution of Zn atoms in the vicinity of the tagged Zn is obtained for each Zn atom in each configuration. The other classifier is based on the SOAP[25] descriptor. In contrast to the BPSF classifier that only considers the metal ions surrounding the central Zn, the SOAP classifier also considers the atoms that are part of the ligands and computes a 780-size vector, capturing more details of the environment.

https://doi.org/10.1021/acs.jcim.6c01175
J. Chem. Inf. Model. 2026, 66, 10046−10055

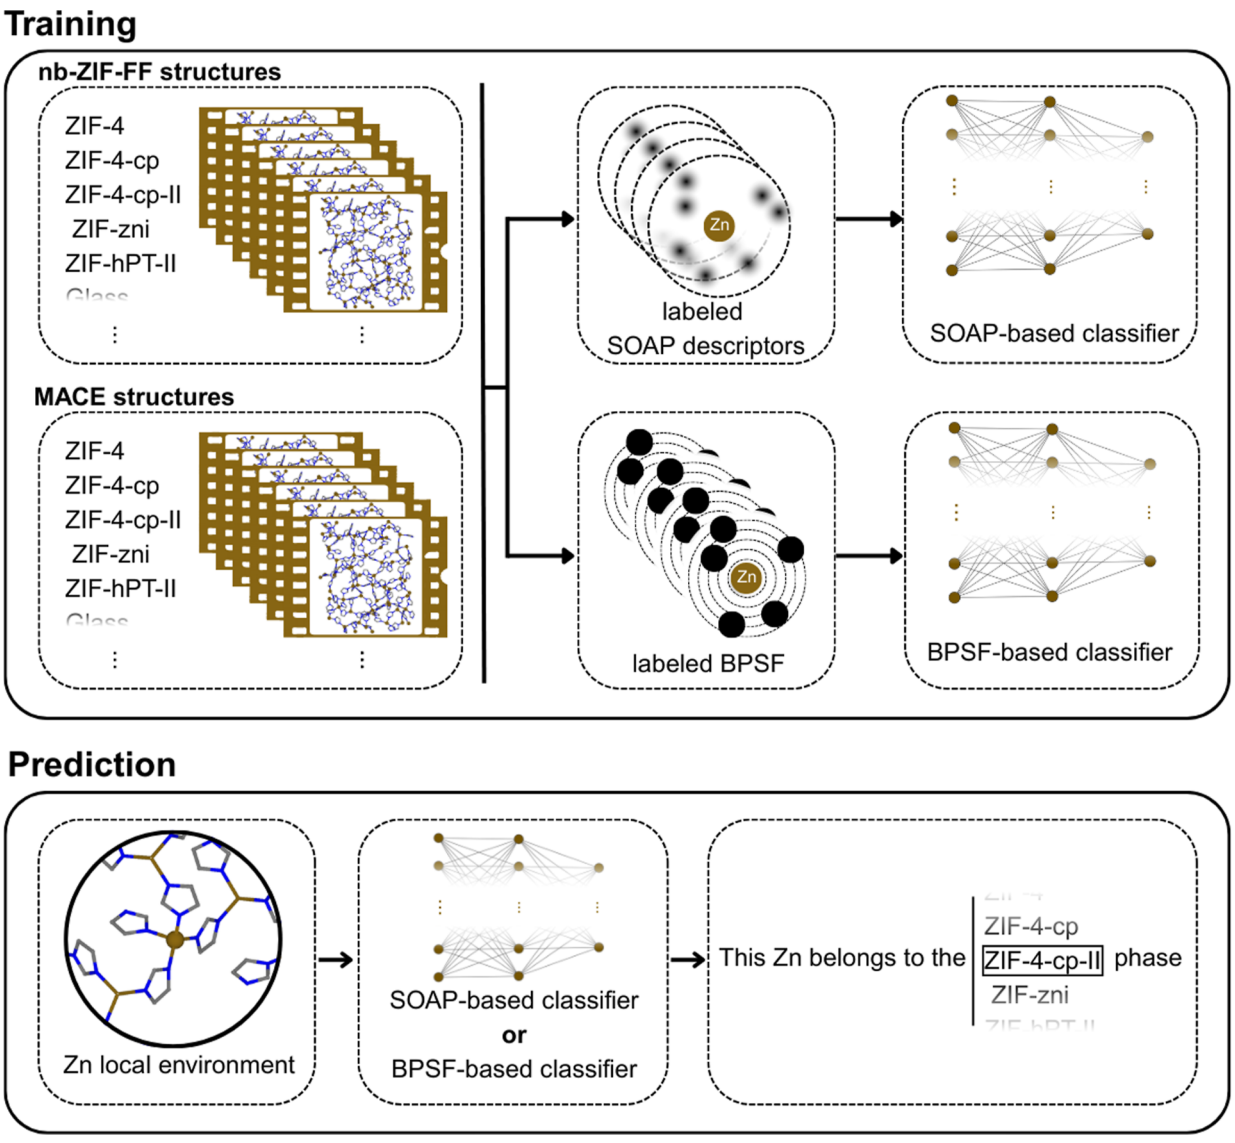

Figure 2. Workflow of the training and classification procedure for both classifiers. Trajectories of ZIF-4 polymorphs were first generated from molecular dynamics simulations with two different force fields. SOAP and BPSF descriptors were then calculated to represent all Zn-centered environments. Subsequently, two neural networks were trained on these two sets of descriptors to generate two separate classifiers, which can then be used independently to classify the phase of a given Zn environment in a structure.

The objective of this study is thus not to perform a direct comparison between SOAP and BPSF descriptors but to present two different strategies for phase classification in ZIFs. Both BPSF and SOAP descriptors are rotationally and translationally invariant, and their resulting vectors have a fixed size that does not depend on the number of neighbors in the local environment.

To train the classifiers, we used the database consisting of MD configurations of the seven phases presented in Table 1 and Figure 1. No distinction was made for the labels of the training data between the two simulation methods. To train the BPSF- and SOAP-based classifiers, we computed the respective descriptors for all the Zn centers in the data set which results in feature vectors labeled according to their phase. For both classifiers, this data was then randomly split into training and test data in an 80−20% manner, and we trained two respective neural networks, which take the features as the input and predict the phase label of the environment. The workflow for the training of the classifiers is represented in Figure 2.

Both classifiers employ fully connected feed-forward neural networks trained as multiclass classifiers. For the BPSF-based classifier, the neural network architecture comprised one input layer of 12 nodes, one for each symmetry function, a single hidden layer of 24 nodes with ReLU activation functions,[33] and 7 output nodes with softmax activation functions. For the SOAP-based classifier, the number of nodes in the first and second layers was set to 780 and 100, respectively. The models were optimized using the Adam optimizer[34] with categorical cross-entropy loss, while the training data were randomly divided into 80% training and 20% test sets. More details of each of the considered descriptors and the neural network training procedures are included in the Supporting Information.

To obtain a dimensionality reduction for the visualization of the SOAP and BPSF descriptors, we use a similar methodology as described in ref 35: we first do a preprocessing of the

https://doi.org/10.1021/acs.jcim.6c01175
J. Chem. Inf. Model. 2026, 66, 10046−10055

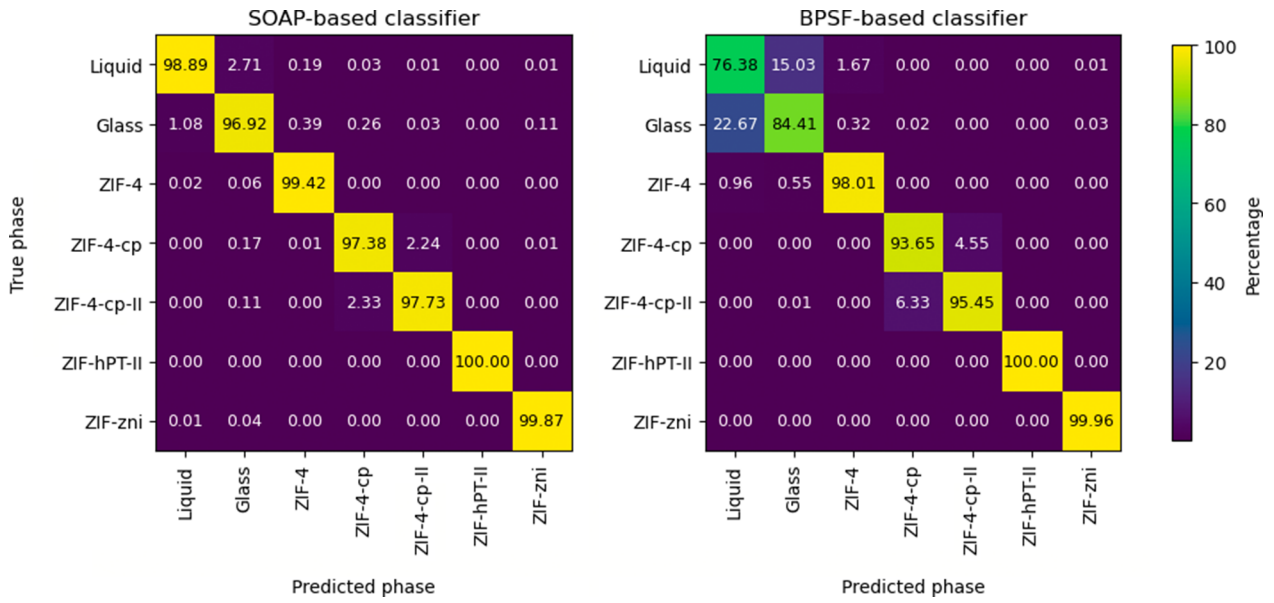

Figure 3. Confusion matrices of the SOAP- (left) and the BPSF-based (right) classifiers on the test set, which includes structures from both MACE and nb-ZIF-FF potentials.

descriptor with PCovC[35] to reduce the array to 5 dimensions, using a mixing parameter α of 0.5 and logistic regression. Subsequently, UMAP[36] is applied to obtain a 2D representation. To make sure that our qualitative analysis remains unchanged with respect to the mixing parameter, we tested the impact of varying this parameter in Figures S1 and S2.

RESULTS AND DISCUSSION

Performance of the Classifiers

Once trained on the local environments of the different phases, the SOAP- and BPSF-based classifiers achieve accuracies on the test set of 98.6% and 92.6%, respectively. The confusion matrices of both classifiers are listed in Figure 3.

In a previous work,[20] a different BPSF-based classifier that was trained on fewer phases (ZIF-4, ZIF-zni, glass, and liquid) was presented, originating from nb-ZIF-FF simulations, which had a classification accuracy of 90.3% over the test set. The accuracy of the previous BPSF-based classifier is slightly lower because most of the confusion arises from the liquid/glass phases. By adding more data of these phases to the training data set, the accuracy increases.

To analyze the confusion matrices and understand why some phases are better predicted than others, it is important to examine the structural differences between phases as well as the representation of the local chemical environments of the metal ions in these phases afforded by the two descriptors employed. The disordered phases (liquid and glass) and the two closed-pore phases (ZIF-4-cp and ZIF-4-cp-II) represent the largest sources of error for both classifiers. This is expected from the structural similarities between these two pairs of phases. Both amorphous phases were already a source of error in the previous BPSF classifier.[20] Higher errors between the ZIF-4-cp and ZIF-4-cp-II phases were also expected: they have the same topology and very close atomic positions, which makes them challenging to differentiate.[12] The ZIF-hPT-II structure, on the other hand, exhibits the best classification performance because of its distinct structure: its density is significantly higher than that of the other phases.

We subsequently performed cross experiments to assess force field dependence. For this, we trained new classifiers with the same methods on one set of structures (nb-ZIF-FF–generated structures) and tested them on a different set of structures (MACE-generated structures), and vice versa. The accuracy of the classifiers trained and tested with different force field databases is detailed in Table 2. The resulting confusion matrices are represented in Figures S4–S7 of the Supporting Information. The objective of these plots is to observe whether the classifiers were able to capture the physical features unique to each phase or if the classification is based on effects arising from the simulation methods, in other words, the objective is to analyze the generalizability of the classifiers.

Table 2. Performance of the Classifiers Trained and Tested with Data Sets Produced by Different Force Fields in %[a]

	test set		
train set	nb-ZIF-FF	MACE	all
nb-ZIF-FF	99.5 (98.4)	76.4 (64.4)	87.9 (81.4)
MACE	84.5 (64.9)	98.6 (90.0)	91.6 (77.4)
all	99.0 (98.0)	98.2 (87.9)	98.6 (92.6)

[a]Results correspond to the SOAP-based classifier followed by the BPSF-based classifier in parentheses.

The combination of both data sets for the training procedure significantly improves the performance of the classifiers. The improved transferability obtained from mixed training data sets is consistent with the fact that each force field samples slightly different configurations of the same underlying structural motifs. Although both force fields reproduce the same crystalline phases, differences in the equilibrium densities, thermal fluctuations, and simulation ensembles lead to variations in the local atomic environments. Training on both data sets therefore exposes the neural networks to a wider range of physically plausible configurations, reducing overfitting to the specific characteristics of a single force field and improving generalization when classifying previously unseen trajectories. For both cross-experiments (trained on MACE,

https://doi.org/10.1021/acs.jcim.6c01175
J. Chem. Inf. Model. 2026, 66, 10046−10055

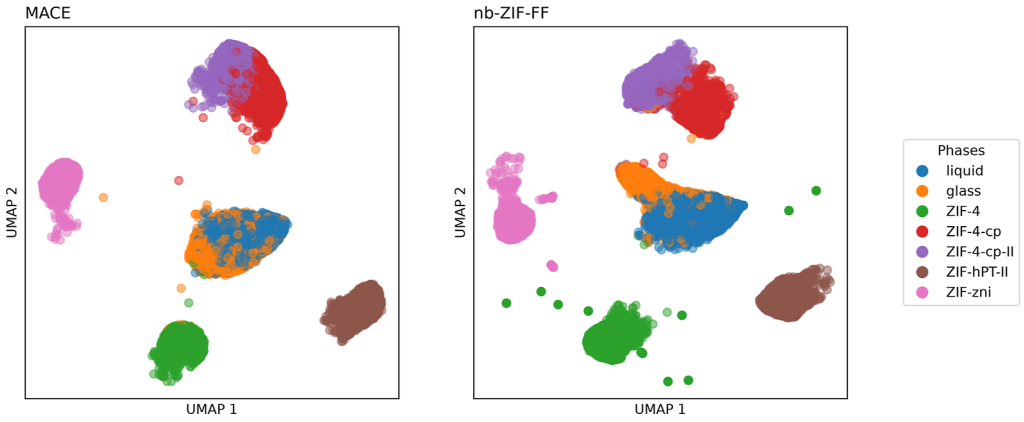

Figure 4. UMAP projection with PCovC preprocessing of the Zn-centered SOAP descriptors for all studied phases, from molecular simulations with the MLP (left) and nb-ZIF-FF (right). The projection space is the same in both plots.

tested on nb-ZIF-FF; trained on nb-ZIF-FF, tested on MACE), the amount of classification errors increases. This seems to be partly caused by the difference between the MACE and nb-ZIF-FF structures and is correlated both to force field differences as well as to the use of different simulation ensembles.

In order to analyze visually the difference in the SOAP vector between phases as well as within the same phase modeled using the two force fields, we plot a 2-dimensional embedding of this high-dimensional vector in Figure 4. The 2-dimensional representation is plotted in the same UMAP space for both force fields (left and right panels). In this plot, the ZIF-4-cp and ZIF-4-cp-II descriptors cluster in the same region, which explains the confusion between these phases in the resulting classifier. This is also the case for the glass and liquid phases. For the SOAP descriptor, the glass phase seems to be the most different across the force fields. We show a similar plot for the BPSF-based classifier in Figure S3, which reveals that this classifier is more sensitive to density: for the MLP, the liquid, glass, and ZIF-4 phases, which share the same density, are grouped in the same cluster. We can also notice larger disparities for the same phases across the force fields. This might indicate that the BPSF-based classifier results in fewer chemically meaningful features than the SOAP-based one, and thus the classification is more density-based. It is important to note that in these figures, we show only one possible 2D representation which captures the main features of these high-dimensional vectors.

Although the BPSF classifier does not contain ligand information and only uses a 12-feature descriptor, it is able to classify phases with high accuracy. The ligand information thus is not strictly necessary for classification. However, using a higher feature descriptor that takes into account the ligands seems to result in a more generalizable classifier and improve the accuracy.

The significant improvement in performance observed for classifiers trained on mixed force-field data sets, compared to those trained on individual force fields, suggests that combining data from multiple force fields is an effective strategy for improving generalization across simulation conditions.

Compared with previous work,[20] the present study systematically evaluates the robustness of local phase classifiers with respect to descriptor choice, simulation methodology, and the complexity of the structural transformation under investigation. In particular, the distinction between the ZIF-4-cp and ZIF-4-cp-II polymorphs represents a considerably more demanding classification task than previously studied crystal−glass or crystal−liquid transformations[20] because both phases share the same topology and differ only through subtle local structural distortions. Successfully resolving this transition demonstrates that the proposed framework is applicable to challenging polymorphic transformations beyond the cases previously reported.

Case Study: ZIF-4-cp/ZIF-4-cp-II Transitions

We probed the capability of our classification algorithms to identify the phase transition between ZIF-4-cp and ZIF-4-cp-II (which we label CP and CPII, respectively). This phase transition was selected for its high reversibility and for the close similarity between the structures of both phases, which makes the classification task more difficult than that in other kinds of transitions. Furthermore, this transition could not yet be observed experimentally, although it was studied in a previous computational work.[12] Nevertheless, the dynamic aspects of the growth of a CPII phase in the bulk of a CP phase and vice versa remain unknown.

For the analysis of the dynamic features of the CP/CPII interconversion, we generated nonequilibrium trajectories in which the system starts at one phase and is simulated under thermodynamic conditions in which the opposite phase is the most stable, so that phase transition occurs spontaneously within a short simulation trajectory. Three independent simulations of CP to CPII transitions were generated, and another three starting from CPII to account for the stochasticity of the interconversion process. Since the transitions involve a change in volume, it was necessary to perform the simulations in the *NPT* ensemble, so the force field used to generate these transitions was nb-ZIF-FF due to its higher stability in NPT simulations and lower computational requirements. The simulated phases comprise $6 \times 6 \times 6$ unit cell systems in order to prevent periodic boundary conditions from affecting the process. The simulation setup used for generating these transitions is detailed in the Supporting Information.

We applied the classification algorithm to each Zn ion in the system over time for each trajectory. In Figure 5, we plot the fraction of Zn atoms classified as CPII as a function of time, using the SOAP-based classifier since it has a higher accuracy in the classification of these two phases. Equivalent plots for the BPSF-based classifier are included in the Supporting

https://doi.org/10.1021/acs.jcim.6c01175
J. Chem. Inf. Model. 2026, 66, 10046−10055

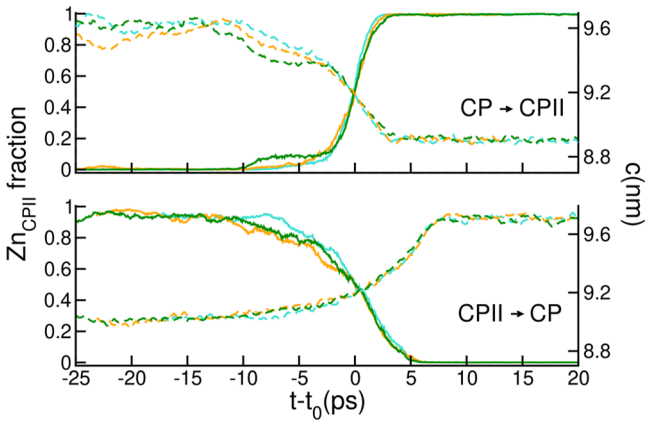

Figure 5. Fraction of Zn ions classified as CPII by the SOAP-based algorithm (left axis, full lines) and simulation box size in the z direction (along the crystallographic c axis) (right axis, dashed lines) vs time for CP to CPII (upper panel) and CPII to CP (lower panel) transitions. All the trajectories are aligned at t_0, which corresponds to the time at which half of the Zn ions are classified as CPII. The different colors correspond to independent transition trajectories.

Information. In all cases, the fraction of Zn ions that were not classified as neither of the two phases was negligible. Since the phase transition is an activated process, the absolute time at which it occurs is different in each trajectory. To align the plots from different simulations, we centered each simulation at time t_0 in which the fraction of Zn atoms belonging to the final phase is 0.5. The plots that correspond to the BPSF-based classifier are centered using the same values of t_0 to allow a clear comparison. We also plot, in dashed lines, the behavior of the c cell parameter, which corresponds to the simulation box size in the z direction. It can be observed that the result of the classification algorithm is well correlated with the change in the c cell parameter, which is characteristic of the CP ↔ CPII phase transition. This demonstrates that the classifiers can detect the phase transition correctly. In some trajectories, frustrated attempts of new phase formation are detected before the full phase transition occurs (see, for example, the behavior of the orange full curve in the upper panel at times between −25 and −20 ps). This kind of events can be associated with the nucleation of a new phase cluster that is not big enough to reach the critical size required to continue growing, and thus it reverts to the original phase.

The classification results present larger fluctuations before the phase transition occurs. This can be explained by the fact that the transition state of the system is metastable before the transition, and the training set of the algorithm was constructed only in thermodynamic conditions of stability. Also it could be attributed to the presence of small, fast-decaying clusters of the stable phase that nucleates in metastable conditions, as explained above. Both classifiers are able to detect the phase transitions properly. The CP to CPII transitions are detected earlier by the BPSF-based classifier while the opposite transitions are detected later. This means that the BPSF-based algorithm has a higher tendency to classify environments as CPII than the SOAP-one does. We also observe that the SOAP-based classifier presents higher fluctuations when applied to the metastable CPII phase, while the BPSF-based algorithm presents more fluctuations when applied to the metastable CP phase. This can be explained by the fact that the BPSF-based algorithm does not contain Zn−ligand spatial correlations, which makes it more sensitive to density fluctuations, which are higher in the CP phase. Again, the BPSF-based classifier overestimates the amount of Zn atoms in the CPII phase, in line with the previous conclusions.

In Figure S13, we present the same results of Figure 5 but for the classifiers trained without the data that comes from structures of CP at $\rho = 4.87$ Zn/nm^3. The comparison between these two plots allows us to probe the influence of the thermodynamic conditions-based data augmentation in the performance of the classifiers. For both classifiers, the transitions happen at slightly different times as a consequence of the higher tendency to classify structures as CPII of the BPSF-based classifier. This is because the pressure conditions used for the generation of CP data are far from the transition pressure (i.e., 20 MPa versus 140 MPa[12]). Interestingly, the amount of fluctuations in the result of the SOAP-based classifier increases before the phase transition events. This corresponds to the situation in which a metastable phase dominates the simulation box. On the contrary, for the BPSF-based classifier, we observe an increase of the fluctuation always in the region where the CP phase dominates. The amount of misclassified structures in this case can go up to 60% in the case of CP to CPII transitions and 30% in the opposite one. This means that the BPSF-based classifier is more sensitive to the lack of CP data at pressures closer to the one at which the transition occurs. This is a consequence of the high sensitivity of this classifier to local density fluctuations as previously discussed, while the SOAP-based classifier, which takes into account ligand orientation information, produces a more robust result.

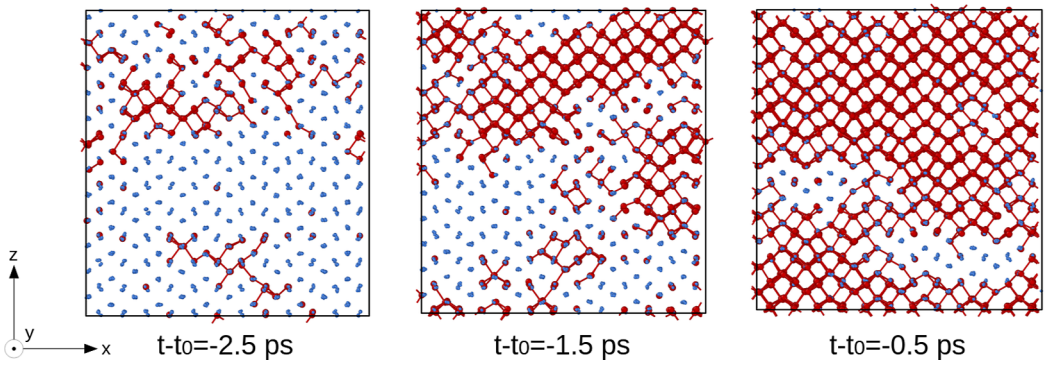

Figure 6. Snapshots of the phase transition from CP to CPII. Only Zn ions are represented and they are classified using the SOAP-based classifier, color-coded depending on the predicted phase, with CP environment in blue and CPII environment in red.

https://doi.org/10.1021/acs.jcim.6c01175
J. Chem. Inf. Model. 2026, 66, 10046−10055

To provide a visual representation of the dynamics of the phase transition, we plot in Figure 6 three snapshots of a trajectory in which a CP to CPII transition takes place. A similar plot for the opposite transition is included in Figure S8. The Zn atoms classified as CP are colored in blue, while the ones classified as CPII are colored in red. In this picture, we can observe how the new phase emerges from the bulk of the previous phase and expands until it occupies the full box. The atoms classified as CPII appear to be clustered and are not randomly distributed. Also, we can observe the tendency of the new phase atoms to form in a layered-like fashion parallel to the (xy) plane. We discuss these aspects more quantitatively below.

The capability of the algorithm to identify phases only using local information opens the possibility of analyzing the geometrical features of the formation of the new phase. We performed a deep-first-search algorithm that identified clusters of Zn ions that belong to the same phase according to the classifier. In the CP to CPII transitions, we analyzed clusters of CPII Zn environments and the opposite in the CPII to CP transitions. In this way, we always considered clusters of the final phase formed. In Figure 7, we plot the average size of the

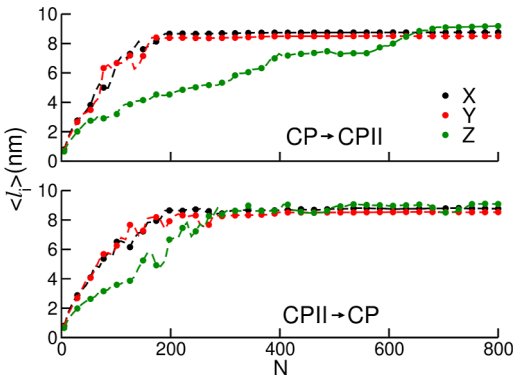

Figure 7. Average cluster size $\langle l_i \rangle$ for directions $i = x$ (black), y (red), and z (green), plotted against the number N of Zn ions comprising the cluster. Upper panel: CP to CPII transitions. Lower panel: CPII to CP transitions. The atoms that form the clusters correspond to the ones classified as belonging to the final phase in each case. Results from three independent simulations were averaged.

clusters composed of neighbor Zn ions that belong to the target phase. The size was measured as the largest distance in the three directions x, y, and z among all pairs of atoms comprising the cluster. We observe that the growth of the new phase is nonisotropic in all the transitions studied. For a given N, the size of the clusters is always larger in the x and y directions than in the z direction. This means that the propagation of the new phase boundary over the simulation box is slower in the z direction than in the other two. This behavior is also identified by the BPSF-based classifier, whose plot is shown in the Supporting Information. Visual inspection of the trajectory depicted in Figure 6 also confirms this trend. The slower growth rate of the final phase in the z direction can be attributed to the fact that the biggest change in cell parameters that occurs during a CP to CPII transition or vice versa is in the value of c, which corresponds to the size of the box in the z direction. This means that during the propagation of a phase boundary in the z direction, a greater change in linear density is produced, in comparison with a growth along x or y. This phenomenon seems to be more pronounced in the CP to CPII transition than in the CPII to CP transition. Nevertheless, the opposite behavior is observed in the results that come from the BPSF-based algorithms. We can attribute this difference to the higher error of the BPSF-based classifier in recognizing CP and CPII phases.

Finally, we sought to interpret the classifier output in physical terms by identifying which features of the local environments distinguish CP from CPII. To do so, we studied the correlation between the classification algorithm results and a local order parameter that varies systematically across the transition. Simple descriptors like average distances or angles between Zn and ligand species do not differ significantly between CP and CPII structures. For that reason we chose the orientation of the four-membered rings of Zn atoms, which had been identified in prior work as a local order parameter able to track the progress of the transition.[12] To measure the orientation, we computed the angle ϕ_a between the vector that is normal to the plane formed by the four-membered ring and the crystallographic vector $\vec{a}$ aligned with the X axis. In Figure 8, we plot the histograms of ϕ_a for each possible amount of Zn

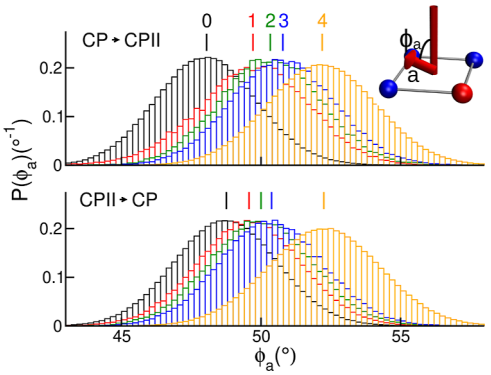

Figure 8. Histograms of the angle ϕ_a between the vector normal to the surface formed by the four-membered rings of Zn atoms and the vector $\vec{a}$ that corresponds to the X axis of the simulation box. A sketch of the angle ϕ_a is depicted in the upper right corner, where each ball corresponds to a Zn atom forming a 4-membered ring, the vertical stick indicates the normal vector to the surface, and the arrowed stick depicts the vector $\vec{a}$. Zn atoms are colored according to the classification result of the SOAP-based algorithm, blue for CP and red for CPII. The histograms are divided by the number of Zn ions in the ring classified as CP. Mean values of each histogram are marked with vertical lines following the same color code. Results obtained from transition trajectories of CP to CPII are plotted in the upper panel and the opposite ones in the lower panel.

ions in the ring that are classified as CPII by the SOAP-based algorithm. A value of zero represents a ring fully classified as CPII, while a value of four represents a fully CP ring. In the upper right corner of the figure, we show a sketch of the definition of ϕ_a, which is explained in the caption. In all cases, results from three independent trajectories were averaged. For computing the normal vector of the plane, we calculated the normal vector to the planes formed by different triplets of atoms in the 4-membered ring and then averaged the result. In most cases, the four atoms almost lie in the same plane, so the values coming from individual triplets present small deviations with respect to the average. All the Zn ions in the crystal structure belong to at least one four-membered ring. In total, there are four different ring orientations, but they all share the same histogram for the angle ϕ_a since they can be interconverted by reflection operations. We can observe in

https://doi.org/10.1021/acs.jcim.6c01175
J. Chem. Inf. Model. 2026, 66, 10046−10055

the plot that there is a correlation between the average value of ϕ_a (marked with vertical lines) and the amount of Zn atoms classified as CPII, which means that the classifier is correctly capturing the change in local geometrical features that characterize the phase transition.[12] The same trend is observed in the plots that correspond to the BPSF-based classifier, which are shown in the Supporting Information. Nevertheless, ϕ_a cannot replace the classifier as order parameter to probe the phase transition since there is a significant overlap in the histograms that correspond to each phase.

The mean values of ϕ_a for each amount of CPII Zn ions are not the same in the case where the transition goes from CP to CPII than in the reverse reaction. This is expected since the pressure values used in both simulations were different, changing the equilibrium values of all the geometrical parameters. Mild differences were detected between the two classification algorithms, being the mean values of ϕ_a closer to each other in the BPSF-based one, especially in the CP to CPII transition. This can be attributed to the fact that this classifier lacks the information about the bonding ligands orientation, which can improve the quality of the classification in cases where the structures of the two phases are similar.

CONCLUSIONS AND PERSPECTIVES

In this work, we developed and evaluated neural network classifiers for the identification of polymorph phases in ZIF-4 systems based on local atomic environments. By combining configurations generated from two distinct force fields and employing both low-dimensional (BPSF) and high-dimensional (SOAP) descriptors, we assessed the accuracy, robustness, and transferability of phase classification methods across a wide range of thermodynamic conditions.

Both classifiers achieve high predictive performance with the SOAP-based approach yielding a higher accuracy. On the other hand, we show that the BPSF approach, even though it contains only metal-centered information and a reduced amount of descriptors, can capture most of the relevant structural features required for phase discrimination. However, the inclusion of ligand information enhances robustness, particularly when comparing phases with similar structures, which are responsible for the greatest source of error.

Training on configurations generated from two different force fields consistently improves predictive robustness relative to models trained on a single data set. These results show that increasing the diversity of the sampled local environments enhances the transferability of the classifiers across the different force fields.

Application of the classifiers to nonequilibrium molecular dynamics trajectories demonstrates their ability to resolve phase transitions in both space and time. The analysis of ZIF-4-cp ↔ ZIF-4-cp-II transitions reveals mechanistic aspects of the process. In particular, we detected an anisotropic growth behavior in the z direction with respect to the x and y axes. We were also able to provide a physical interpretation of the classification result by correlating it with a microscopic order parameter that captures the transition between both phases. The successful application of the algorithms to the challenging ZIF-4-cp ↔ ZIF-4-cp-II transition illustrates that local machine-learning classifiers can resolve subtle crystalline transformations and provide quantitative insight into nucleation and growth processes that are difficult to characterize using conventional structural descriptors alone.

Future work could focus on the extension of the methodology to broader classes of ZIF-4 polymorphs that are missing in this analysis, such as ZIF-1, ZIF-3, ZIF-6, and ZIF-10[13] and their associated phase transitions, or to the classification of adsorption-induced structural transitions of flexible ZIFs.[37,38] Overall, this study demonstrates that machine learning classifiers based on local descriptors provide a powerful and flexible framework for the automated analysis of polymorphism and phase behavior in metal−organic frameworks.

ASSOCIATED CONTENT

Data Availability Statement

Data and software availability: in order to make our work fully reproducible by others, representative input files for each type of simulation are available online in our data repositories at https://github.com/rosemino/MAGNIFY/tree/main/ZIFs_NN_classifiers.

Supporting Information

The Supporting Information is available free of charge at https://pubs.acs.org/doi/10.1021/acs.jcim.6c01175.

Full simulation details, analysis of force field dependence of the classifiers performances, assessment of influence of choice of mixing parameter in PCovC analyses, molecular snapshots, BPSF-based classifier results for ZIF-4-cp/ZIF-4-cp-II interconversions, and experimental phase diagram of ZIF-4 (PDF)

AUTHOR INFORMATION

Corresponding Authors

François-Xavier Coudert − *Chimie ParisTech, PSL University, CNRS, Institut de Recherche de Chimie Paris, 75005 Paris, France;* orcid.org/0000-0001-5318-3910; Email: fx.coudert@chimieparistech.psl.eu

Rocio Semino − *Sorbonne Université, CNRS, Physico-chimie des Electrolytes et Nanosystèmes Interfaciaux, PHENIX, F-75005 Paris, France;* orcid.org/0000-0003-3937-7414; Email: rocio.semino@sorbonne-universite.fr

Authors

Emilio Méndez − *Sorbonne Université, CNRS, Physico-chimie des Electrolytes et Nanosystèmes Interfaciaux, PHENIX, F-75005 Paris, France;* orcid.org/0000-0002-7577-3783

Léna Triestram − *Chimie ParisTech, PSL University, CNRS, Institut de Recherche de Chimie Paris, 75005 Paris, France*

Dune André − *Chimie ParisTech, PSL University, CNRS, Institut de Recherche de Chimie Paris, 75005 Paris, France*

Complete contact information is available at: https://pubs.acs.org/10.1021/acs.jcim.6c01175

Author Contributions

[§]E.M. and L.T. contributed equally to the work. F.-X.C. and R.S.: conceptualization, methodology, investigation, funding acquisition, resources, supervision, writing—original draft, and writing—review and editing. E.M., L.T., and D.A.: conceptualization, methodology, investigation, data curation, formal analysis, software, visualization, writing—original draft, and writing—review and editing.

Notes

The authors declare no competing financial interest.

https://doi.org/10.1021/acs.jcim.6c01175
J. Chem. Inf. Model. 2026, 66, 10046−10055

■ ACKNOWLEDGMENTS

Access to high-performance computing platforms was provided by GENCI grants A0190807069 and A0190915688. This work was partly funded under the Imperial College London—CNRS PhD joint programme. This work was partly funded by the European Union ERC Starting grant MAGNIFY, grant number 101042514.

■ REFERENCES

(1) Schneemann, A.; Bon, V.; Schwedler, I.; Senkovska, I.; Kaskel, S.; Fischer, R. A. Flexible metal−organic frameworks. *Chem. Soc. Rev.* **2014**, *43*, 6062−6096.
(2) Peh, S. B.; Karmakar, A.; Zhao, D. Multiscale Design of Flexible Metal−Organic Frameworks. *Trends Chem.* **2020**, *2*, 199−213.
(3) Coudert, F.-X. Recent advances in stimuli-responsive framework materials: Understanding their response and searching for materials with targeted behavior. *Coord. Chem. Rev.* **2025**, *539*, 216760.
(4) Aulakh, D.; Varghese, J. R.; Wriedt, M. The Importance of Polymorphism in Metal−Organic Framework Studies. *Inorg. Chem.* **2015**, *54*, 8679−8684.
(5) Cheetham, A. K.; Kieslich, G.; Yeung, H. H.-M. Thermodynamic and Kinetic Effects in the Crystallization of Metal−Organic Frameworks. *Acc. Chem. Res.* **2018**, *51*, 659−667.
(6) Castillo-Blas, C.; Chester, A. M.; Keen, D. A.; Bennett, T. D. Thermally activated structural phase transitions and processes in metal−organic frameworks. *Chem. Soc. Rev.* **2024**, *53*, 3606−3629.
(7) Hughes, J. T.; Bennett, T. D.; Cheetham, A. K.; Navrotsky, A. Thermochemistry of Zeolitic Imidazolate Frameworks of Varying Porosity. *J. Am. Chem. Soc.* **2013**, *135*, 598−601.
(8) Widmer, R. N.; Lampronti, G. I.; Chibani, S.; Wilson, C. W.; Anzellini, S.; Farsang, S.; Kleppe, A. K.; Casati, N. P. M.; MacLeod, S. G.; Redfern, S. A. T.; Coudert, F.-X.; Bennett, T. D. Rich Polymorphism of a Metal−Organic Framework in Pressure−Temperature Space. *J. Am. Chem. Soc.* **2019**, *141*, 9330−9337.
(9) Bumstead, A. M.; Ríos Gómez, M. L.; Thorne, M. F.; Sapnik, A. F.; Longley, L.; Tuffnell, J. M.; Keeble, D. S.; Keen, D. A.; Bennett, T. D. Investigating the melting behaviour of polymorphic zeolitic imidazolate frameworks. *CrystEngComm* **2020**, *22*, 3627−3637.
(10) Hou, J.; Sapnik, A. F.; Bennett, T. D. Metal−organic framework gels and monoliths. *Chem. Sci.* **2020**, *11*, 310−323.
(11) Smirnova, O.; Hwang, S.; Sajzew, R.; Ge, L.; Reupert, A.; Nozari, V.; Savani, S.; Chmelik, C.; Reithofer, M. R.; Wondraczek, L.; Kärger, J.; Knebel, A. Precise control over gas-transporting channels in zeolitic imidazolate framework glasses. *Nat. Mater.* **2024**, *23*, 262−270.
(12) Méndez, E.; Semino, R. Phase diagram of ZIF-4 from computer simulations. *J. Mater. Chem. A* **2024**, *12*, 31108−31115.
(13) Park, K. S.; Ni, Z.; Côté, A. P.; Choi, J. Y.; Huang, R.; Uribe-Romo, F. J.; Chae, H. K.; O'Keeffe, M.; Yaghi, O. M. Exceptional chemical and thermal stability of zeolitic imidazolate frameworks. *Proc. Natl. Acad. Sci. U. S. A.* **2006**, *103*, 10186−10191.
(14) Vervoorts, P.; Stebani, J.; Méndez, A. S. J.; Kieslich, G. Structural Chemistry of Metal−Organic Frameworks under Hydrostatic Pressures. *ACS Mater. Lett.* **2021**, *3*, 1635−1651.
(15) Song, J.; Pallach, R.; Frentzel-Beyme, L.; Kolodzeiski, P.; Kieslich, G.; Vervoorts, P.; Hobday, C. L.; Henke, S. Tuning the High-Pressure Phase Behaviour of Highly Compressible Zeolitic Imidazolate Frameworks: From Discontinuous to Continuous Pore Closure by Linker Substitution. *Angew. Chem., Int. Ed.* **2022**, *61*, No. e202117565.
(16) Robertson, G. P.; Anzellini, S.; Meneghini, C.; Herlihy, A.; Amboage, M.; Keen, D. A.; Irifune, T.; Bennett, T. D. Changes in the Long-Range Order and Local Atomic Structure of Zeolitic Imidazolate Frameworks under Extreme Conditions. *Inorg. Chem.* **2026**, *65*, 156−164.
(17) Bouëssel du Bourg, L.; Ortiz, A. U.; Boutin, A.; Coudert, F.-X. Thermal and mechanical stability of zeolitic imidazolate frameworks polymorphs. *APL Mater.* **2014**, *2*, 124110.
(18) Castel, N.; André, D.; Edwards, C.; Evans, J. D.; Coudert, F.-X. Machine learning interatomic potentials for amorphous zeolitic imidazolate frameworks. *Digital Discovery* **2024**, *3*, 355−368.
(19) Castel, N.; Coudert, F.-X. Challenges in Molecular Dynamics of Amorphous ZIFs Using Reactive Force Fields. *J. Phys. Chem. C* **2022**, *126*, 19532−19541.
(20) Méndez, E.; Semino, R. Microscopic mechanism of thermal amorphization of ZIF-4 and melting of ZIF-zni revealed via molecular dynamics and machine learning techniques. *J. Mater. Chem. A* **2024**, *12*, 4572−4582.
(21) Helfrecht, B. A.; Semino, R.; Pireddu, G.; Auerbach, S. M.; Ceriotti, M. A new kind of atlas of zeolite building blocks. *J. Chem. Phys.* **2019**, *151*, 154112.
(22) Larsen, P. M.; Schmidt, S.; Schiøtz, J. Robust structural identification via polyhedral template matching. *Modell. Simul. Mater. Sci. Eng.* **2016**, *24*, 055007.
(23) Rosset, L. A. M.; Drabold, D. A.; Deringer, V. L. Signatures of paracrystallinity in amorphous silicon from machine-learning-driven molecular dynamics. *Nat. Commun.* **2025**, *16*, 2360.
(24) Behler, J. Atom-centered symmetry functions for constructing high-dimensional neural network potentials. *J. Chem. Phys.* **2011**, *134*, 074106.
(25) Bartók, A. P.; Kondor, R.; Csányi, G. On representing chemical environments. *Phys. Rev. B* **2013**, *87*, 184115.
(26) Henke, S.; Wharmby, M. T.; Kieslich, G.; Hante, I.; Schneemann, A.; Wu, Y.; Daisenberger, D.; Cheetham, A. K. Pore closure in zeolitic imidazolate frameworks under mechanical pressure. *Chem. Sci.* **2018**, *9*, 1654−1660.
(27) Balestra, S. R. G.; Semino, R. Computer simulation of the early stages of self-assembly and thermal decomposition of ZIF-8. *J. Chem. Phys.* **2022**, *157*, 184502.
(28) Triestram, L.; Coudert, F.-X. Identifying phase transitions in zeolitic imidazolate frameworks: microscopic insight from molecular simulations. *Chem. Sci.* **2026**, *17*, 6734−6745.
(29) Andarzi Gargari, S.; Semino, R. Unveiling ZIF-8 Nucleation Mechanisms through Molecular Simulation: Role of Temperature, Solvent, and Reactant Concentration. *Chem. Mater.* **2025**, *37*, 9460−9470.
(30) Méndez, E.; Semino, R. Thermodynamic insights into the self-assembly of zeolitic imidazolate frameworks from computer simulations. *Chem. Sci.* **2025**, *16*, 11979−11988.
(31) Batatia, I.; Kovacs, D. P.; Simm, G.; Ortner, C.; Csányi, G. MACE: Higher order equivariant message passing neural networks for fast and accurate force fields. *Advances in neural information processing systems*; Curran Associates, Inc., 2022; Vol. *35*, pp 11423−11436.
(32) Batatia, I.; Kovács, D. P.; Simm, G. N. C.; Ortner, C.; Csányi, G. MACE: Higher Order Equivariant Message Passing Neural Networks for Fast and Accurate Force Fields. *arXiv* **2023**, arXiv.2206.07697.
(33) Agarap, A. F. Deep Learning using Rectified Linear Units (ReLU). *arXiv* **2019**, arXiv.1803.08375.
(34) Kingma, D. P.; Ba, J. Adam: A Method for Stochastic Optimization. *arXiv* **2017**, arXiv:1412.6980.
(35) Jorgensen, C.; Lin, A. Y.; Vasavada, R.; Cersonsky, R. K. Interpretable visualizations of data spaces for classification problems. *Mach. learn.: sci. technol.* **2026**, *7*, 025008.
(36) McInnes, L.; Healy, J.; Saul, N.; Großberger, L. UMAP: Uniform Manifold Approximation and Projection. *J. Open Source Softw.* **2018**, *3*, 861.
(37) Tanaka, H.; Ohsaki, S.; Hiraide, S.; Yamamoto, D.; Watanabe, S.; Miyahara, M. T. Adsorption-Induced Structural Transition of ZIF-8: A Combined Experimental and Simulation Study. *J. Phys. Chem. C* **2014**, *118*, 8445−8454.
(38) Chaplais, G.; Fraux, G.; Paillaud, J.-L.; Marichal, C.; Nouali, H.; Fuchs, A. H.; Coudert, F.-X.; Patarin, J. Impacts of the Imidazolate Linker Substitution (CH_3, Cl, or Br) on the Structural and Adsorptive Properties of ZIF-8. *J. Phys. Chem. C* **2018**, *122*, 26945−26955.