

Characterizing Defect Dynamics in Silicon Carbide Using Symmetry-Adapted Collective Variables and Machine Learning Interatomic Potentials

Soumajit Dutta,[†] Cunzhi Zhang,[†] Gustavo Perez Lemus,[†] Juan J. de Pablo,[‡]
Francois Gygi,[¶] Giulia Galli,^{†,§,||} and Andrew L. Ferguson^{*,†,§}

[†]*Pritzker School of Molecular Engineering, University of Chicago, Chicago, Illinois 60637,
United States of America*

[‡]*Tandon School of Engineering, New York University, New York, New York 11201, United
States of America*

[¶]*Department of Computer Science, University of California Davis, California 95616,
United States of America*

[§]*Department of Chemistry, University of Chicago, Chicago, Illinois 60637, United States of
America*

^{||}*Materials Science Division and Center for Molecular Engineering, Argonne National
Laboratory, Lemont, Illinois 60439, United States of America*

E-mail: andrewferguson@uchicago.edu

Abstract

Silicon carbide (SiC) divacancies are attractive candidates for spin defect qubits possessing long coherence times and optical addressability. The high activation barriers associated with SiC defect formation and motion pose challenges for their study by first-principles molecular dynamics. In this work, we develop and deploy machine learning interatomic potentials (MLIPs) to accelerate defect dynamics simulations while retaining *ab initio* accuracy. We employ an active learning strategy comprising symmetry-adapted collective variable discovery and enhanced sampling to compile configurationally diverse training data, calculation of energies and forces using density functional theory (DFT), and training of an E(3)-equivariant MLIP based on the Allegro model. The trained MLIP reproduces DFT-level accuracy in defect transition activation free energy barriers, enables the efficient and stable simulation of multi-defect 216-atom supercells, and permits an analysis of the temperature dependence of defect thermodynamic stability and formation/annihilation kinetics to propose an optimal annealing temperature to maximally stabilize VV divacancies.

1 Introduction

Solid-state spin defects play a crucial role in advancing various quantum technologies, including quantum sensing, communication, and computing.^{1,2} Silicon carbide (SiC), a wide-bandgap semiconductor, has emerged as a promising host for such defects owing to its well-established growth, doping, and fabrication techniques.³ Among these, the divacancy (VV) configuration, which forms by the interaction of a silicon vacancy (V_{Si}) and carbon vacancy (V_C) has garnered significant attention as a spin qubit due to its desirable characteristics, such as long coherence times, optical addressability, and high-fidelity readout.⁴

In SiC, VV defects are typically formed via ion implantation or electron irradiation followed by high-temperature annealing⁴ (Figure 1). Computational modeling offers a means to gain mechanistic understanding and insight into defect formation, transition, and migration, and can also predictively identify optimal thermal annealing protocols.⁵⁻¹¹ *Ab initio* simulations have, however, reported activation barriers for typical defect transitions of several electronvolts.^{9,10,12-16} Yan *et al.* performed *ab initio* density functional theory (DFT) calculations using the PBE functional and reported migration barriers of 4.3 eV for V_{Si} and 3.3 eV for V_C in the 4H-SiC polytype.¹⁵ Zhang *et al.* carried out calculations for the 3C-SiC polytype and obtained barriers of 4.05 eV for V_{Si} and 3.87 eV for V_C .¹⁶ These high energy barriers make defect transitions a rare event and render unbiased *ab initio* simulation of defect dynamics exceedingly computationally expensive. Simulation of defect dynamics using classical force fields is computationally efficient, but typically inaccurate due to the lack of explicit treatment of the important electronic effects.¹⁶

Machine learning interatomic potentials (MLIPs) have created an opportunity to break this trade-off between efficiency and accuracy.¹⁷ MLIPs predict atomic forces and energies by learning many-body interactions typically using either kernel-based methods¹⁸⁻²⁰ or neural networks.²¹ The field has evolved rapidly since seminal work by Behler and Parrinello²² that introduced a symmetric, invariant representation of atomic local environments to predict system energies using neural networks. Since then, numerous more sophisticated machine




Figure 1: Vacancy defects in silicon carbide (SiC) crystals. (A) Schematic illustration of the divacancy formation process in SiC. Vacancies are typically created by ion implantation or electron irradiation followed by temperature annealing to promote high defect mobility and the formation of stable divacancies. (B) A number of vacancy defects commonly appear and interconvert during the annealing process including divacancies (VV), antisite vacancies ($V_C V_{Si}$), carbon vacancies (V_C), and silicon vacancies (V_{Si}).

learning models and atomic environment descriptors have been developed to capture atomic interactions with greater fidelity. For example, the kernel-based FLARE MLIP,²³ built on the Atomic Cluster Expansion (ACE) descriptor,²⁴ was used to study pressure-induced phase transitions in SiC polytypes with accuracy comparable to *ab initio* calculations and with high sampling efficiency.²⁵ In recent years, equivariant neural network-based MLIPs have attracted growing interest due to their improved data efficiency and generalizability, enabled by the explicit encoding of physical symmetries in the hidden layers.^{26–30}

In this study, we employ the Allegro MLIP, which uses an E(3)-equivariant graph neural network architecture with atomic cluster expansion (ACE) descriptors.³¹ Allegro is selected among the variety of MLIP models for its attractive combination of high accuracy, data efficiency, and computational speed.^{32,33} To ensure stability in Allegro-driven molecular dynamics simulations, we employ a prior potential to accelerate learning, improve extrapolation and constrain the MLIP from sampling unrealistic conformations,³⁴ and adopt an active learning approach to expand the training data to cover diverse thermally relevant system configurations. Using this protocol, we train an Allegro-based MLIP possessing *ab initio* accuracy in just six rounds of active learning requiring DFT force and energy calculations

for 1800 training configurations. The resulting MLIP reproduces five key defect transition activation free energies within the thermal fluctuation error ($1 k_bT \approx 0.13$ eV at 1500 K) of the DFT calculations and enables stable long-term simulations of a 216-atom supercell of SiC at an execution speed of approximately 6 ns/day on an NVIDIA L40S GPU. We deploy the trained MLIP to study defect transitions in 3C-SiC crystal polytype and study the temperature variation effect on divacancy formation from the dynamic interactions of two monovacancies. Our analyses of these long unbiased simulation trajectories reveals a subtle interplay of the relative stability and kinetic interconversion rates between competing defect structures and can inform materials processing strategies to maximally stabilize VV spin defects in SiC.

2 Methods

2.1 SiC defect crystal

All calculations in this work were conducted on a 3C-SiC polytype with the same 4.416 Å at 1500 K lattice constant used in previous work.¹⁶ Unless otherwise specified, a simulation supercell containing 216 atoms was used. For the multi-defect study, a larger 512-atom supercell was employed. Defects were introduced by selectively removing carbon and silicon atoms from the lattice. The Atomic Simulation Environment (ASE) Python package³⁵ was used for system construction and for converting between different file formats.

2.2 Active learning training of MLIP for SiC defect dynamics

Recent studies have highlighted the instability of long-time molecular dynamics simulations under MLIPs.^{36–38} This drawback can often be attributed to the incorrect prediction of forces for conformations outside of training data due to limited extrapolation capabilities of the MLIP. Therefore, the accuracy and stability of the MLIP relies on the ability to effectively capture a diverse range of local environments in the training data to enable accurate inter-

polation of the local environments encountered during simulations. To overcome the dual goals of generating diverse conformations and constructing a reliable force field using those conformations, we adopt an active learning approach that interleaves rounds of enhanced sampling of configurational space using the MLIP and retraining of the MLIP using the expanded set of training configurations (Figure 2). We describe here an overview of the six steps in the active learning workflow, and provide full technical details in the Supplementary Methods in the Supporting Information.

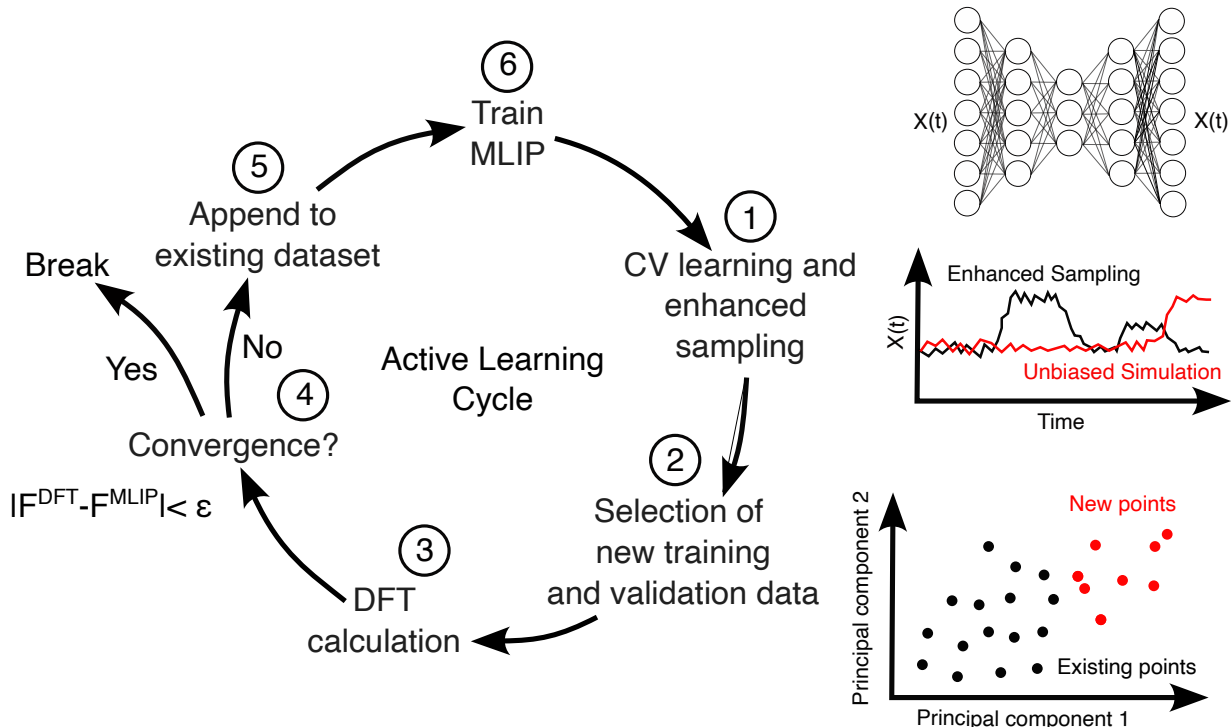


Figure 2: Active learning protocol for training an Allegro machine learning interatomic potentials (MLIP) for SiC. The cyclic process includes (1) enhanced sampling using an autoencoder-based collective variable (CV), (2) configuration selection, (3) density functional theory (DFT) evaluation of selected configurations, (4) convergence assessment, (5) training and validation data appending and (6) force field retraining with the updated dataset.

Step 1: CV learning and enhanced sampling. In the first step of the active learning cycle, we conducted enhanced sampling of configurational space using OPES-Explore, a variant of the On-the-fly Probability Enhanced Sampling (OPES) method, to efficiently explore diverse local defect configurations with the most recent Allegro MLIP trained in

the prior round of active learning (Figure S3). Simulations were driven in three collective variables (CVs) learned over the enhanced sampling trajectories conducted in the previous active learning cycle using Permutationally Invariant Networks for Enhanced Sampling (PINES), designed to respect the underlying translational, rotational, and permutational symmetries of the crystal lattice.³⁹ PINES learns symmetry-adapted CVs by transforming a pairwise distance vector into a permutationally invariant vector (PIV) through block-wise sorting, then feeds this vector into a variational autoencoder (VAE). The PIV can be conceived as an atom-type specific pair correlation function that serves as a translationally, rotationally, and permutationally invariant featurization of an atomic configuration,⁴⁰ and the low-dimensional CVs extracted from the bottleneck layer of the VAE can be conceived as a low-dimensional distillation of the PIV.⁴¹ Importantly, the learned CVs preserve the same translational, rotational, and permutational symmetries of the PIV vector, which is desirable in improving the data efficiency of learning, and are also differentiable with respect to atomic coordinates, thus permitting their use within enhanced sampling calculations.^{42–44}

Calculations were performed in a 216-atom supercell of the 3C-SiC polytype containing a V_C and a V_{Si} defect positioned within one coordination shell of each other. The atoms within the third coordination shell of the initial monovacancies were restrained to prevent the monovacancies from migrating too far apart and promote defect transitions and the exploration of diverse local defect environments. PIV vectors were constructed by considering all atoms residing within the two coordination shells around the defect, producing a length 435 vector. Atoms of the same element and coordination shell were treated as identical, yielding four permutationally indistinguishable atom-type groups that enabled the PINES CVs to distinguish defect reorientations while eliminating degeneracy from permutations of identical atoms. We employed a simple VAE architecture employing an encoder with hidden layers comprising 64, 32, and 16 neurons and a symmetric decoder. A bottleneck layer of 3 neurons was used to extract three learned CVs, which we found to be sufficient to resolve all of the important defect configurations and transitions. A custom PLUMED⁴⁵

module implementing PINES is available in our GitHub repository at https://github.com/Ferg-Lab/SiC_Allegro_MLIP.git.

Molecular dynamics simulations were performed using the LAMMPS simulation package (v. 27 June 2024),⁴⁶ employing GPU acceleration via the KOKKOS package.⁴⁷ All simulations were carried out in the canonical (NVT) ensemble using a Nosé-Hoover thermostat with a damping parameter of 0.1 ps to maintain the system temperature at 1500 K and employing a velocity Verlet integrator with a 1 fs time step.⁴⁸ Initial velocities were assigned from a Gaussian distribution with total momentum and angular momentum removed, and a 15 ps equilibration run conducted to relax the system. The neighbor list was updated every step with a bin size of 2.0 Å, and dynamic neighbor checking was enabled. In each round of active learning, 10 ns of enhanced sampling were performed and trajectories written with a period of 0.1 ps.

In the first round of active learning, we lack both an initial Allegro MLIP and learned PINES CVs. To ignite the active learning loop, we model the system using a pre-trained Fast Learning of Atomistic Rare Events (FLARE) MLIP developed for SiC crystals²⁵ and drive enhanced sampling using PINES CVs learned from unbiased simulation trajectories performed using the classical Environment-Dependent Interatomic Potential (EDIP).⁴⁹ The EDIP potential is known not to provide a satisfactory description of defect dynamics,¹⁶ but we use it to learn approximate initial CVs with which to initialize the first round of active learning (Figure S1-2). The pre-trained FLARE potential provided a reasonable starting model for exploring SiC crystal defect configurations (Figure S4), but due to the superior learning and data efficiency of equivariant graph neural network-based models,⁵⁰⁻⁵⁴ subsequent rounds of active learning are used to develop a SiC-specific Allegro-based MLIP. The FLARE and EDIP potentials were employed exclusively to initialize the first round of the active learning cycle and do not feature any further in the workflow.

Step 2: Selection of training and validation data. The enhanced sampling trajectories were analyzed to select new training and validation points for Allegro MLIP training

and performance evaluation. Dissimilarity of configurations in the current enhanced sampling trajectories relative to the current set of training points were evaluated using a smooth overlap of atomic positions (SOAP) kernel,⁵⁵ which computes similarity by the dot product of average power spectrum vectors⁵⁶ to present a translationally, rotationally and permutationally invariant representation of the local atomic environments obtained by expanding the local atomic density using orthogonal basis functions.⁵⁶ Descriptor-space similarity measures such as SOAP have been widely used in several MLIP frameworks, including GAP¹⁸ and FLARE,²⁰ where the prediction of energies and forces relies on kernel similarity between local atomic environments in the training set. As a local descriptor based on an expansion of the local neighbor density around a central atom within a cutoff, SOAP-based dissimilarity measures limit the dimensionality and computational complexity of distance calculations and scales well to large and complex systems. Calculations were performed using the Python package Dscribe.⁵⁷ The 300 most dissimilar configurations were added to the training set in each round of the active learning cycle. The validation set was defined by clustering all sampled configurations visited over the enhanced sampling trajectories collected to date into 100 cluster centers to provide a diverse sample of configurational space for evaluation of MLIP performance. We adopt 300 new training points and 100 new validation points based on the available parallel compute resources. With more resources available, one may choose to employ larger training batch sizes per round in an effort to reduce the number of active learning rounds required for convergence. The SOAP dissimilarity of points to the training data is positively correlated with the MLIP prediction error in the force relative to the DFT ground truth, providing support for a diversity maximizing strategy in curating the training data and demonstrating that structural dissimilarity is indeed a good proxy for MLIP accuracy that serves as an inexpensive surrogate metric for determining the set of new training and validation points to incorporate into the active learning campaign (Figure S5).⁵⁸

Step 3: DFT energy and force calculations. Energy and force calculations for the selected training and validation conformations were performed using DFT with Qbox.⁵⁹

Calculations were conducted using the PBE functional⁶⁰ with optimized norm-conserving Vanderbilt pseudopotentials⁶¹ and a plane-wave kinetic energy cutoff of 55 Ry. The simulation supercell was assumed to be neutrally charged, and the calculations were conducted in a triplet spin state. Previous studies have shown that the activation temperature of defect formation depends on effective barriers that in turn depend on the charge state of the defect.¹⁶ However the procedure adopted here to derive MLIPs requires the choice of a specific charge state. Our choice of a neutral state is an approximation which nevertheless allows us to understand trends on the stability of VVs as a function of temperature as well as kinetic effects, and to obtain valuable mechanistic insights. We demonstrate good agreement of our model predictions with experiment (*vide infra*), but propose that future work may consider the development of charge- and spin-adaptable MLIPs in pursuit of superior physical realism. All DFT calculations were performed on single Intel Xeon E5-2680 CPUs.

Step 4: Convergence assessment. The prediction accuracy of the Allegro MLIP relative to reference DFT forces was evaluated using a force-based convergence criterion. We seek to train a model with a force accuracy at the 1500 K simulation temperature of $1 k_B T / \text{\AA} \approx 3 \text{ kcal/mol.\AA} \approx 0.13 \text{ eV/\AA}$, which is in line with the typical performance expected from a MLIP for reliable molecular dynamics simulations and accurate property prediction.^{19,62} Guided by this benchmark, we defined convergence of the active learning cycle as the point at which the Allegro MLIP achieves a mean absolute error (MAE) below 0.1 eV/\AA on both the training and validation datasets. When this threshold is satisfied, the active learning loop terminates; otherwise, additional rounds of sampling, DFT labeling, and model retraining are performed.

Step 5: Training and validation data update. If the Allegro MLIP has not met the accuracy threshold on the validation data, we augment the existing training set with the 300 new training points and replace the existing validation set with the 100 new validation points spanning the configurational space.

Step 6: Allegro MLIP retraining. The Allegro MLIP was retrained over the aug-

mented training set. The Allegro model is implemented and trained using the NequIP v0.6.0²⁶ and Allegro v0.2.0³¹ Python packages and all calculations performed using a cutoff distance of 6 Å for neighbor list construction conducted in double-precision floating point. The model returns pairwise energies, which are then appropriately summed to provide the per-atom and overall system energy. Atomic forces are computed by automatic differentiation of the predicted energies through the computational graph to guarantee a conservative force field. Training against the ground truth DFT energies and forces was conducted using a mean squared error (MSE) loss function with equal weighting of the energy and force terms, and the model weights and biases updated using the Adam optimizer.⁶³ The model is exposed to all training points collected to date. The training set volume increases by 300 configurations in every pass through the active learning cycle, so later training rounds are exposed to comparatively more training data. To stabilize the force field in data-sparse regimes and prevent exploration of unphysical, high-energy configurations with atomic overlaps during enhanced sampling (Figure S6), we followed the approach adopted in several modern MLIP frameworks of employing a physically motivated prior potential.^{34,64,65} Specifically, we employed a Lennard-Jones (LJ) prior and used a delta-learning strategy in which Allegro learns only the residual between the LJ prior and the reference DFT energies and forces. The parameters of the LJ prior were learned by linear least squares fitting against the ground truth DFT energies and an Allegro MLIP trained to fit the residuals between the LJ and DFT energies and forces (Figure S7). This approach furnished a simpler learning problem for the Allegro MLIP, stabilized the model against exploring thermally irrelevant high-energy regions of configurational phase space in the aggressive OPES-Explore enhanced sampling calculations (Figure S8) without any degradation in predictive accuracy (Figure S9). A custom implementation of the Allegro MLIP under the LJ prior potential and delta learning strategy is available in our GitHub repository at https://github.com/Ferg-Lab/SiC_Allegro_MLIP.git. All Allegro training was performed with GPU acceleration using NVIDIA’s CUDA framework on an NVIDIA L40S GPU.

3 Results

3.1 Active learning training of Allegro MLIP for SiC defect dynamics

An active learning strategy was employed to train an Allegro MLIP for simulating defect dynamics in a 3C-SiC polytype with computational efficiencies approaching that of classical force fields (~ 6 ns/day on an NVIDIA L40S GPU) and *ab initio* accuracy in force predictions (< 0.1 eV/Å mean average error). The E(3)-equivariant nature of the Allegro MLIP makes training inherently data efficient and by employing an active learning strategy combining data-driven learning of symmetry-adapted collective variables, enhanced sampling of the thermally-accessible phase space, a physically-motivated prior potential, and a diversity maximizing selection of training configurations, we converged the Allegro MLIP using just 1800 training configurations over the course of six active learning rounds. A total of ~ 240 GPU-h were expended on conducting unbiased and enhanced sampling molecular simulations and ~ 43 GPU-h on MLIP training using a NVIDIA L40S GPU and $\sim 72,000$ CPU-h in computing ground truth energies and forces for the training and validation configurations on Intel Xeon E5-2680 CPUs.

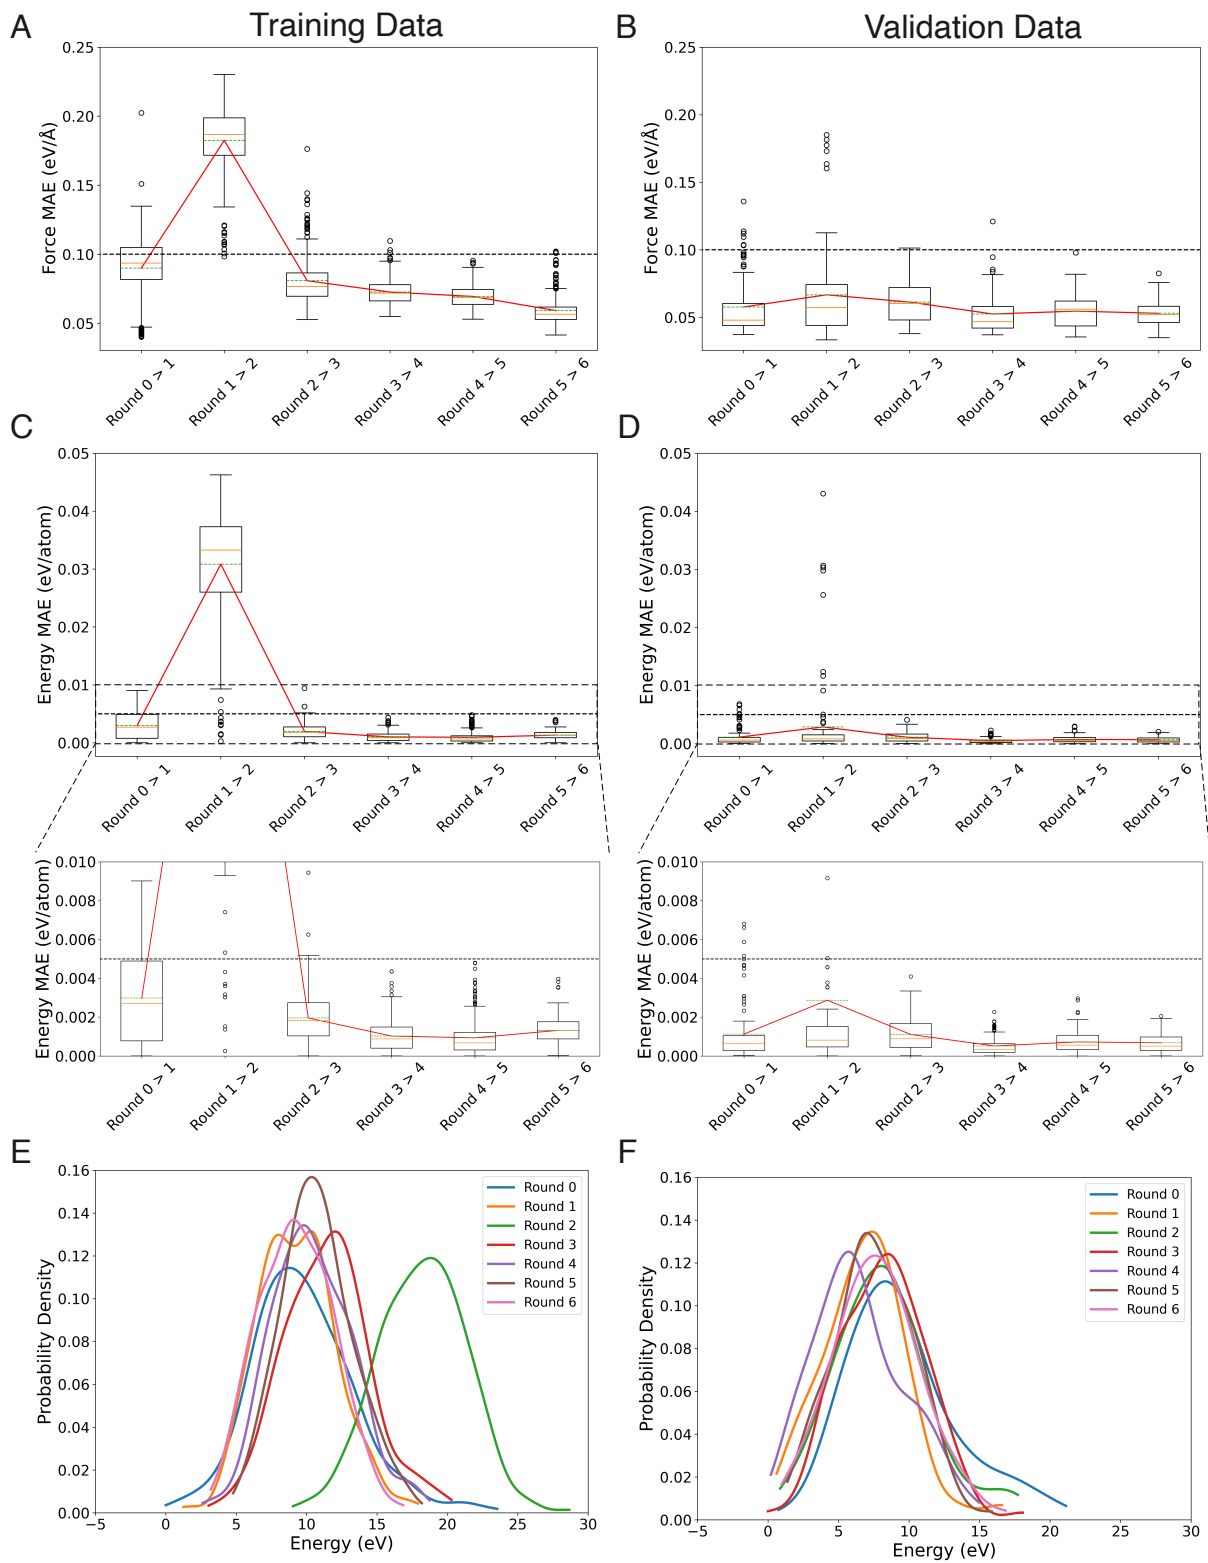


Figure 3: Accuracy of the trained Allegro MLIP after each round of active learning. Box plots of the mean average error (MAE) in the force predictions relative to the DFT ground truth over (A) the 300 training configurations and (B) 100 validation configurations in each round of the active learning cycle. The green horizontal line within the box represents the mean of the distribution, and the orange horizontal line represents the median. The limits of the box correspond to the first (Q1) and third (Q3) quartiles, and the whiskers extend to the minimum and maximum values within 1.5 times the interquartile range (IQR) from the quartiles. Outliers in the distribution are explicitly shown. The red line connecting the means indicates the trend. The dotted horizontal line represents the desired force MAE threshold of $0.1 \text{ eV}/\text{\AA}$ used to judge convergence of the model. (C,D) Analogous plots to panels (A,B), but illustrating the MAE in the per atom energy relative to the DFT ground truth. The dotted horizontal line represents an MAE in the energy of 0.005 eV/atom . Energy distribution of each round of (E) training and (F) validation configurations showing the diversity of configurations explored over each of the active learning rounds. The exploration of high-energy training configurations in Round 2 coincides with the spike in the MAE.

We present in Figure 3 the progression of the Allegro MLIP energy and force errors over the validation set at each round of the active learning cycle. The initial MLIP performs apparently quite well on the validation data collected in Round 1, with a mean average error (MAE) in the predicted forces relative to the DFT ground truth averaged over the training set of just $0.089 \text{ eV}/\text{\AA}$ and over the validation set of just $0.057 \text{ eV}/\text{\AA}$, lying below the $0.1 \text{ eV}/\text{\AA}$ threshold used to judge sufficiently accurate performance of the MLIP (Figure 3A,B). The apparently good performance of the Round 1 model is, however, deceptive, since it is largely a result of quite low configurational diversity in the training and validation sets (Figure 3E, F). The MLIP has not yet been challenged to train or predict on diverse defect configurations across the thermally accessible phase space. Furthermore, there is quite high variability in the force MAEs, with 35% of the training configurations and 2% of the validation configurations possessing force MAEs in excess of the $0.1 \text{ eV}/\text{\AA}$ target threshold. In Round 2, the enhanced sampling calculations generate much higher configurational diversity in the training set (Figure 3E), leading to an approximate doubling of the training set mean force MAE to $0.182 \text{ eV}/\text{\AA}$ and an uptick in the validation force MAE to $0.067 \text{ eV}/\text{\AA}$. The diversity maximizing strategy to accumulating training configurations quickly pays dividends in subsequent active learning rounds, with the training mean force MAE plummeting to $0.081 \text{ eV}/\text{\AA}$ and the validation force MAE to $0.061 \text{ eV}/\text{\AA}$ in Round 3. The training force MAE

monotonically decreases in all subsequent rounds while the validation force MAE plateaus in Round 4. We declare the MLIP converged in Round 6 where the training mean force MAEs meets the validation mean force MAE at a performance level of 0.055 eV/Å, the validation mean force MAE shows no further downward trend, and 1% of the training configurations and 0% of the validation configurations possess a force MAE above the 0.1 eV/Å threshold.

The MAE in the MLIP energy predictions relative to the DFT ground truth follow a similar trend to the forces (Figure 3C,D), with an uptick in the error as the configurational diversity increases in Round 2, followed by a precipitous decrease and plateau in the training and validation mean energy MAE around 0.001 eV/atom in Rounds 3-6. State-of-the-art MLIPs typically achieve energy MAEs in the range 0.001-0.005 eV/atom.⁶⁶ By the terminal Round 6 of the active learning campaign and after having been exposed to just 1800 training configurations, fully 100% of training configurations and 100% of validation configurations possess an energy MAE better than 0.005 eV/atom.

Taken together, these results demonstrate that the trained SiC Allegro MLIP has force and energy prediction accuracy at a level of *ab initio* DFT calculations and a diversity-maximizing active learning strategy enabled training of the model in a computationally and data efficient manner. The primary source of error in the MLIP in the present application was empirically observed to be associated with insufficient coverage of the thermally-relevant configurational ensemble (Figure S5), which was effectively addressed by the diversity-maximizing active learning strategy.^{58,67,68} We observe, however, that this may not be an adequate strategy for systems in which there are regions of configurational space that are intrinsically challenging to model and could benefit from additional training samples, or large, complex systems in which there are numerous diverse but thermally-irrelevant configurations that would be vanishingly unlikely to be observed in simulations of the system at thermodynamic conditions of interest. In such cases, strategies targeting training data accumulation to the regions of configurational space exhibiting large errors and/or limiting the exploration of high-energy configurations may be more appropriate. We also note that

MLIP fitting was conducted by matching energies and forces in the NVT ensemble and that it is expected that this will limit deployment of the fitted Allegro model to fixed volume systems at the density at which the model was trained. Indeed, an exploratory simulation in the NPT ensemble at 1500 K and 1 bar employing a Nosé–Hoover thermostat and a Nosé–Hoover barostat was found to be unstable. We anticipate that the Allegro fine tuning procedure could be augmented by direct fitting against DFT-calculated pressures and/or stresses at a range of different densities to explicitly train the model to accurately reproduce the pressure and enable stable NPT calculations, but we have not done so in the present work and this does currently restrict the model to deployments in the NVT ensemble at the density at which the model was trained.

3.2 Prediction of defect transition free energy barriers

We next evaluated the performance of the trained Allegro MLIP in predicting the activation free energies for a variety of known important defect transition processes in SiC.¹⁶ This is a quite stringent test for MLIP performance as it requires correct prediction of the dynamical transition mechanism, the transition state, and the associated energy. Five different defect transitions were considered: (i) reorientation of a VV divacancy via motion of a Si atom, (ii) migration of a V_C monovacancy, (iii) migration of a V_{Si} monovacancy, (iv) VV divacancy formation from a pair of monovacancies via C-atom hopping, and (v) VV divacancy formation from a pair of monovacancies via Si-atom hopping. The first three transitions are symmetric in the sense that the initial and final states represent equivalent configurations and so possess identical free energies, whereas the final two transitions are asymmetric, with the forward process representing divacancy formation and the reverse process divacancy annihilation. We present in Figure 4 a schematic illustration of each defect transition process along with the free energy landscapes for each process computed at 1500 K under the terminal trained Allegro MLIP using the the adaptive biasing force (ABF) method⁶⁹ implemented within the SSAGES software framework.⁷⁰ Details of the ABF implementation are provided in the

Supplementary Methods in the Supporting Information. In each case, the ABF calculations were carried out using a one-dimensional collective variable (CV), defined as the projection of the distance between the moving atom and the center of mass of a selected group of atoms onto the vector connecting the atom’s initial and final positions, where the choice of this CV and the corresponding atom selection for each process follow the methodology described by Zhang *et al.*¹⁶

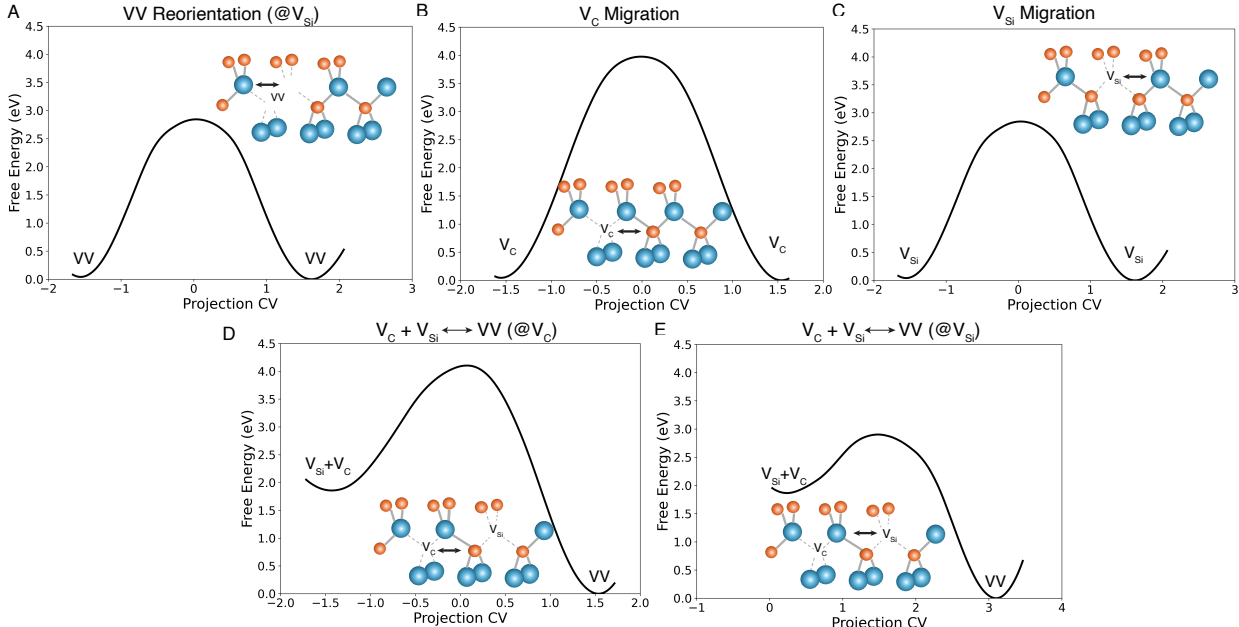


Figure 4: Free energy landscapes for five defect transition processes estimated under the trained Allegro MLIP using ABF: (A) reorientation of a VV divacancy via motion of a Si atom, (B) migration of a V_C monovacancy, (C) migration of a V_{Si} monovacancy, (D) VV divacancy formation from a pair of monovacancies via C-atom hopping, and (E) VV divacancy formation from a pair of monovacancies via Si-atom hopping. In each case, the free energy landscape is projected into one dimensional projection CV as described by Zhang *et al.*¹⁶

We present in Table 1 a quantitative comparison of the Allegro MLIP activation free energies to the corresponding values computed using *ab initio* molecular dynamics calculations with the ABF method. The *ab initio* molecular dynamics free energy barriers for VV reorientation and formation/annihilation were taken from prior work by Zhang *et al.*,¹⁶ and those for V_C and V_{Si} migration computed in this work using an identical protocol. All *ab initio* molecular dynamics calculations were performed in Qbox⁵⁹ at 1500 K using the SSAGES⁷⁰

implementation of ABF.⁶⁹ Full details of the calculations are provided in the Supplementary Methods. In all cases the Allegro MLIP free energy barriers are in agreement with the *ab initio* values within thermal fluctuations ($1 k_B T \approx 0.13$ eV at 1500 K). We note, however, that in the majority of cases the MLIP values tend to be slightly higher than the *ab initio* results, with the exception of divacancy formation via Si-atom hopping.

Table 1: Comparison of the free energy barriers for the five defect transition processes illustrated in Fig. 4 computed at 1500 K using *ab initio* molecular dynamics and molecular dynamics under the trained Allegro MLIP. Free energy landscapes were computed using the adaptive biasing force (ABF) method. For symmetric transitions, a single value for the barrier height is reported; for asymmetric transitions, the forward and backward barrier heights are reported. The final column reports the difference Δ between the MLIP and *ab initio* free energy barrier values. For reference, at 1500 K $1 k_B T \approx 0.13$ eV.

Defect Transformation (T = 1500 K)	<i>ab initio</i> (eV)	MLIP (eV)	Δ (eV)
	Forward (Backward)	Forward (Backward)	
VV reorientation (@V _{Si})	2.70	2.82	0.12
V _C migration	3.92	3.95	0.03
V _{Si} migration	2.95	2.98	0.03
V _C + V _{Si} → VV@V _C	2.24 (4.06)	2.25 (4.10)	0.01 (0.04)
V _C + V _{Si} → VV@V _{Si}	1.14 (2.98)	1.04 (2.90)	-0.10 (-0.08)

The excellent accuracy of the trained Allegro MLIP in predicting atomic energies and forces and defect transition free energy barriers relative to ground truth DFT and *ab initio* molecular dynamics calculations provides a strong validation of the accuracy and reliability of the trained MLIP. These results support the utility of the MLIP as a significantly more computationally efficient alternative to *ab initio* molecular dynamics for the study of defect formation and transitions in SiC crystals.

3.3 Transferability of trained MLIP to multi-defect systems

As a next test of the trained Allegro MLIP, we sought to assess its transferability to multi-defect systems despite having been trained exclusively on configurations containing divacancies and monovacancy pairs. This analysis tests the degree to which the local structural environments the model was exposed to in the divacancy and monovacancy pair training

configurations are sufficiently representative of those encountered in multi-defect systems to permit robust, accurate, and transferable force predictions.

First, we constructed two trivacancy systems within a 216-atom 3C-SiC simulation box – one containing two carbon and one silicon vacancy ($V_C V_{Si} V_C$) and one containing one carbon and two silicon vacancies ($V_{Si} V_C V_{Si}$) – and conducted 5 ns of unbiased molecular dynamics simulations under the trained Allegro MLIP in each of the two systems at 1500 K. To assess the accuracy of the MLIP force predictions, we extracted structures at 100 ps intervals and, using the same approach used to assign ground truth forces and energies to the MLIP training data, computed for comparison the forces using DFT under the PBE functional. As illustrated in Figure 5A and B, in each system the mean force MAE remains below approximately 0.10 eV/Å, corresponding to the accuracy threshold at which the models were trained. Moreover, this error remained flat and did not show any increase over the course of the simulation that may be indicative of an inherent instability in the MLIP.

Second, we constructed a larger 512-atom 3C-SiC simulation cell and introduced ten randomly distributed vacancy defects by removing ten atoms. Specifically, six carbon and four silicon atoms were removed, resulting in one divacancy, five carbon vacancies, and three silicon vacancies. We conducted 5 ns of molecular dynamics simulation at 1500 K under the Allegro MLIP, harvested configurations at 100 ps intervals, and computed the force MAE relative to DFT force calculations under the PBE functional. As illustrated in Figure 5C, the mean force MAE remained below approximately 0.10 eV/Å with the exception of a transient excursion to 0.15 eV/Å at around 2.7 ns.

These analyses demonstrate that the trained Allegro MLIP retains high predictive accuracy in applications to more complex multi-defect configurations to which it was not exposed during training. Not only does the MLIP enable the simulation of far larger systems for far longer time scales than would typically be accessible to *ab initio* calculations, but the simulations remain stable over the course of a 5 ns run and the predictive accuracy of the forces show little to no significant degradation relative to the 0.10 eV/Å accuracy threshold to

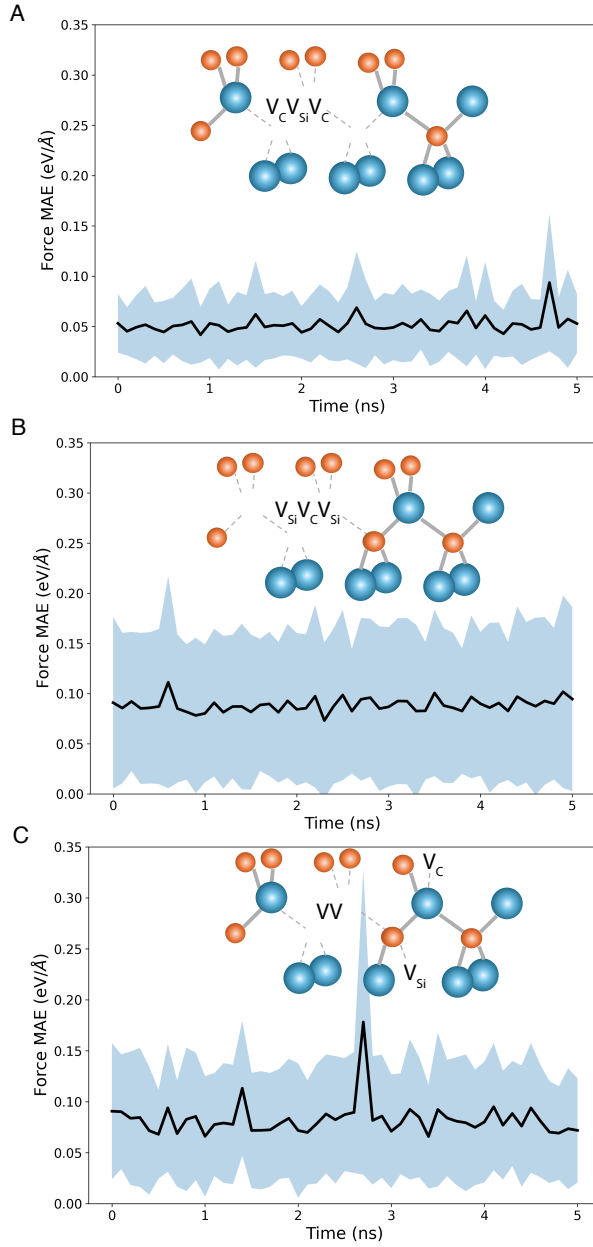


Figure 5: Force accuracy of trained Allegro MLIP in extrapolative predictions to (A) $V_C V_{Si} V_C$ and (B) $V_{Si} V_C V_{Si}$ trivacancy defect simulations and (C) multivacancy defect simulations. The MAE of the MLIP force predictions relative to the DFT calculations are illustrated over the simulation time series. The black lines and shaded regions represent the MAE and the standard deviation of the per-atom force error, respectively, computed at 100 ps time intervals.

which the model was trained.

3.4 Prediction of temperature effects

The MLIP was trained on DFT forces and energies collected from a variety of configurations generated by enhanced sampling. As such, the model was not parameterized at any particular temperature and, provided the configurations encountered at any temperature of interest were adequately represented within the enhanced sampling training data, the model is expected to be temperature transferable. Of particular interest in regards to temperature effects, is to predict how temperature influences the kinetics and thermodynamics of divacancy formation and to identify temperature annealing protocols to selectively promote and stabilize divacancies as putative qubits for quantum information technologies.

We performed 250 nanoseconds of unbiased simulation of a 216-atom SiC box containing a single divacancy at five different temperatures: 1000 K, 1250 K, 1500 K, 1750 K, and 2000 K. To characterize the kinetics and thermodynamics of defect transitions at each temperature, we employed a Markov State Modeling (MSM) framework⁷¹ using the PyEMMA software.⁷² In brief, we generated 250 ns of unbiased simulation trajectories at each temperature by employing 1000 independent 250 ps simulations initialized from diverse initial configurations, projected these data into a 3D latent space spanned by the three leading slowest dynamical modes of the 1000 K simulations, performed a microstate clustering using k-means into a microstate MSM with a lag time of 45 ps, followed by a macrostate clustering using the Perron Cluster Cluster Analysis (PCCA+) method⁷³ to generate a macrostate MSM for each temperature. The macrostate MSMs can be viewed as interpretable surrogate models with which to understand, predict, and control defect kinetics and thermodynamics. Full details of these calculations are provided in the Supplementary Methods (Figures S10-13).

Analysis of the macrostate MSMs at each temperature revealed each macrostate to clearly correspond to one of five defect environments: monovacancies, divacancy, divacancy intermediate, and two antisite defects. We illustrate in Figure 6 the structural interpretation of the five macrostates for the 1500 K model. Similar trends were observed for the MSMs at all five temperatures indicating that the thermally accessible configurations were well conserved

across the 1000-2000 K temperature window. We did not observe the appearance of any new defect configurations that were not observed in the training data, suggesting that the diversity maximizing strategy employed in the active learning cycle used to parameterize the model most likely explored all of the thermally-accessible defect states.

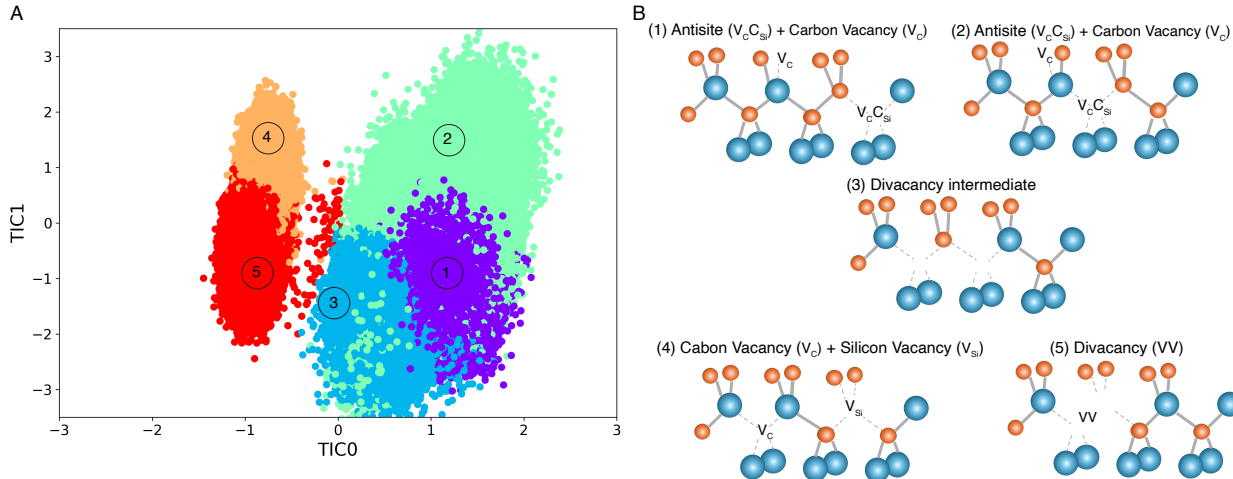


Figure 6: Structural interpretation of the macrostates within the 1500 K macrostate MSM. (A) Projection of the 1500 K simulation data into the leading two slowest dynamical modes of the system dynamics TIC0-TIC1 identified by applying TICA to the 1000 K simulation trajectories. Macrostates were defined by k-means clustering within a 3D TICA projection to form a microstate MSM, then PCCA+ clustering of the microstate MSM into a macrostate MSM. The projected configurations corresponding to each of the five macrostates is indicated by a different color and numerical label. (B) Inspection of the constituent configurations reveals each macrostate to correspond to a particular defect environment comprising a variety of structurally similar configurations. Schematic representations of the defect environment corresponding to each macrostate.

We next analyzed how the relative stabilities of, and transition rates between, the macrostates varied with temperature. First considering thermodynamic stability, we present in Figure 7A the stationary densities within each macrostate corresponding to the equilibrium probability distribution. We observe the VV divacancy to be the most stable defect configuration at all temperatures in the range, but that its stability varied in a non-monotonic fashion. In particular, the divacancy stationary density decreased between 1000 K and 1250 K, then increased again to a maximum at 1750 K. To understand this behavior, we estimated the rates of divacancy formation and annihilation by applying transition path theory (TPT) to the macrostate MSM to compute the net flux into (formation) and out of (annihilation)

the divacancy macrostate by all possible pathways.⁷⁴ TPT reveals the divacancy formation and annihilation rates to both increase with temperature, but that the between 1000 K and 1250 K, the annihilation rate grew relatively faster than the formation rate, leading to an overall destabilization of the divacancy macrostate at the higher temperature and a $\sim 14\%$ relative reduction in its equilibrium stationary density (Figure 7B). As required by overall conservation of probability, this drop coincided with corresponding increases in other defect populations: specifically the concentrations of the divacancy intermediates and antisite defects increased by approximately sevenfold and fourfold, respectively. As the temperature is increased beyond 1250 K, the divacancy stationary density rises again, peaking at 1750 K. Again, this trend can be understood from the temperature dependence of the formation and annihilation rates, wherein the former increases more rapidly than the latter over this range. This result is also consistent with the divacancy activation temperature previously predicted by first principles molecular dynamics.¹⁶ Finally, at the highest 2000 K temperature, the divacancy equilibrium probability falls off as the formation rate decreases slightly while the annihilation rate remains approximately constant.

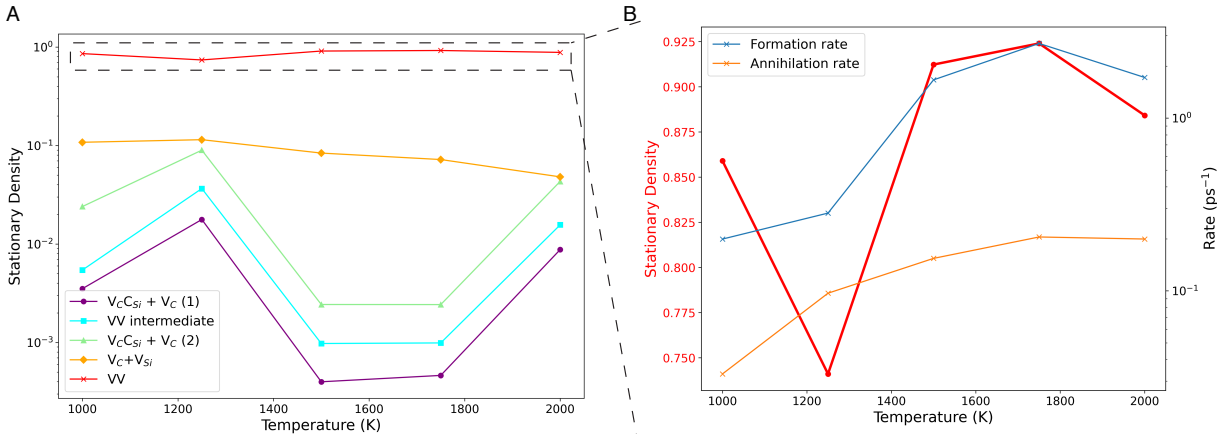


Figure 7: Thermodynamic stability of defect macrostates and kinetics of interconversion as a function of temperature exposed by the macrostate MSMs. (A) Stationary densities (i.e., equilibrium probability distribution) of the five defect macrostates at each temperature computed from the macrostate MSM. (B) Temperature dependence of the divacancy stationary density (red; left axis) and its formation and annihilation rates (blue and orange, respectively; right axis) computed by transition path theory (TPT) applied to the macrostate MSMs at each temperature.

The divacancy stability trends exposed by the macrostate MSMs predict the existence of a temperature optimum at which divacancy defects are maximally stable within the 1000-2000 K temperature window. Since MSMs were only constructed at five discrete temperatures in this range, we fitted a cubic spline to the temperature-dependent divacancy stability to predict a temperature of maximum stability of approximately 1625 K (Figure S14).

A number of prior experimental studies have considered the thermal stability and annealing behavior of VV defects, primarily in the 4H-SiC polytype. Experimental work by Carlsson *et al.* probed VV concentrations as a function of annealing temperature using electron paramagnetic resonance (EPR) spectroscopy to observe an approximately monotonic increase in VV spin defect concentration between 273 K and 1273 K.⁵ A study by Karsthof *et al.* employing photoluminescence (PL) spectroscopy came to a similar conclusion, with the VV-related PL signal monotonically increasing over the 573 K to 1273 K temperature range.⁷⁵ Son *et al.* conducted EPR studies of annealing temperatures over the range 873 K to 1873 K, and identified a peak in the VV concentration at around 1773 K.^{6,8} A higher-temperature EPR study by Carlos *et al.* measured a monotonic decrease in the VV concentration over a temperature range of 1673 K to 2073 K.⁷ A recent study by Kimura *et al.* employed PL spectroscopy to measure the VV concentration at annealing temperatures of 1123 K, 1273 K, and 1473 K, with the 1273 K annealing temperature producing the highest VV-related PL intensity.¹¹ By appealing to *ab initio* calculations,^{10,16} these observations were understood as the lower 1123 K annealing temperature as offering insufficient thermal energy to activate silicon vacancy (V_{Si}) migration necessary for aggregation with a carbon vacancy (V_C) to form the VV divacancy, whereas the higher 1473 K annealing temperature offered sufficient thermal energy to also mobilize the VV and V_C defects.^{75,76} The 1273 K annealing temperature is optimal in providing sufficient thermal energy for V_{Si} migration for VV formation, but insufficient thermal energy to activate VV migration and permit their migration into deeper regions of the crystal or towards the surface where they are annihilated. Since the VV divacancy is known to be stable up to much higher temperatures,⁷ the decrease

in the measured VV-related PL intensity is therefore predicted to be due to migration and/or annihilation of the VV divacancies as opposed to their intrinsic dissociation.

In comparing the annealing temperature predictions of our model to these experimental data, we must take care to view this through a number of caveats associated with the limitations of our model and simulation protocols. First, our simulations consider the 3C-SiC polytype, whereas the experiments pertain to the 4H-SiC polytype. Second, our simulations consider a 216-atom supercell with periodic boundaries containing a single divacancy. As such, our calculations effectively model a homogenous bulk crystal at a specified defect concentration. The defect can and does undergo transitions between the five defect configurations – monovacancies, divacancy, divacancy intermediate, and two antisite defects (Figure 6) – and the defects are also free to migrate. However, our simulation setup precludes migration-driven loss of experimentally observable VV divacancies through diffusion into the crystal interior where they may evade PL spectroscopy or migration to the crystal surface where they may undergo annihilation. Moreover, our calculations on a single divacancy also preclude the opportunity for more complex multi-defect configurations and dynamics. Third, our MLIP was parameterized for a single charge state (neutral) and spin state (triplet) that remain fixed over the course of the simulation. This is the most significant deficiency of the model and the most likely source of discrepancies with measurements of experimental systems that can, of course, undergo changes in charge and spin state over the course of their dynamical evolution.

With these caveats in mind, our predictions are consistent with the experimental work of Son *et al.* in that we predict an optimum temperature for maximization of the VV concentration, and our prediction of the maximum at around 1625 K is only 150 K lower than the 1773 K maximum in the measured VV-associated EPR signal.^{6,8} Our prediction is consistent with the EPR measurements of Carlos *et al.*, who report a monotonic decrease in the VV-associated signal from 1673 K to 2073 K, but this study does not report the lower temperature data required to test our prediction of a maximum at approximately 1625 K.⁷

Given the inherent insufficiency of our 216-atom simulation cell in modeling VV migration into the deep crystal or surface annihilation, we are wary to draw too close a comparison with the PL measurements of Kimura *et al.* given the centrality of these mechanisms in rationalizing the experimental observations.¹¹ Our calculations effectively isolate the intrinsic thermodynamic stability of a single divacancy, unaffected by extrinsic processes such as migration or surface annihilation. In this more restricted sense, however, we draw a point of consistency with Kimura *et al.* in that our calculations predict the VV divacancy to be intrinsically stable up to approximately 1650 K, supporting the mechanism put forward in that study that the reduction in the observed VV-associated PL intensity beyond 1273 K is not due to VV dissociation but rather migration into the crystal interior or surface.¹¹

4 Conclusions

In this study, we developed an E(3)-equivariant Allegro MLIP to model defect dynamics in the 3C-SiC crystal polytype. The trained model possesses *ab initio* accuracy in predicting atomic energies and forces, while possessing computational efficiencies approaching classical force fields that permit the efficient generation of long trajectories for large systems that are outside the reach of *ab initio* molecular dynamics. We devise an efficient active learning framework to accumulate training configurations for the model by integrating data-driven learning of symmetry-adapted collective variables, enhanced sampling of the thermally-accessible phase space, a physically-motivated prior potential, and a diversity maximizing selection of training configurations. Using this strategy, we converged the Allegro MLIP to a force prediction accuracy of 0.1 eV/Å using just 1800 training configurations over the course of six active learning rounds, and deploying a total of ~ 240 GPU-hours on MLIP molecular dynamics simulations, $\sim 72,000$ CPU-h in computing ground truth DFT energies and forces, and ~ 43 GPU-h in MLIP training. We demonstrated excellent agreement of the MLIP predictions of defect transition free energy barriers to *ab initio* calculations, verified

the transferability and stability of the trained MLIP to larger, multi-defect systems, and conducted an analysis of the temperature dependence of defect thermodynamic stability to propose an optimal annealing temperature to maximally stabilize VV divacancies within the SiC lattice and an analysis of the formation/annihilation kinetics to provide atomic-level dynamical insight into the stability mechanism.

This study makes two key contributions. First, it demonstrates a generalizable active learning framework for the data and compute efficient training of MLIPs with *ab initio* accuracy in atomic energy and force prediction. Second, it delivers an accurate and validated MLIP for 3C-SiC force field that can be extended to simulate, understand, and control the thermodynamics and dynamics of lattice defects. This study illustrates the value of MLIPs as an enabling tool to simulate defect dynamics at far larger time and length scales than are accessible to *ab initio* calculations to expose atomic-level understanding of the defect formation, migration, and stabilization mechanisms critical to defect engineering efforts. The simulation protocols presented in this work are also of utility for the efficient training of MLIPs in condensed matter systems, and their deployment in predicting important materials properties such as free energy barriers, optimal annealing temperatures, and defect transition mechanisms as emergent phenomena within unbiased molecular dynamics simulations, through the parameterization of MSMs as long-time kinetic models, and via coupling to efficient collective variable biasing and path-based enhanced sampling techniques.

The good agreement with experimental measurements of the temperature stability of VV divacancies and prior first principles predictions of activation energies is encouraging, but we also identify a number of avenues for additional model development. First, the present MLIP development assumes fixed charge and spin states for all configurations. In reality, these properties can change during defect transformations and strongly influence defect formation and migration energetics.^{10,16} While it would be straightforward to train an ensemble of MLIPs at different charge and spin states using the approach presented herein, it remains challenging to dynamically switch between charge and spin states as a function

of atomic configuration. Nevertheless, there have been recent advances in charge- and spin-aware potentials^{17,77-80} that we anticipate may open the door to future adaptations that would significantly enhance the physical realism of the model. To do so would necessitate a coupled evolution of the configurational and electronic dynamics in which the charge and spin state depend on the local geometry and the Fermi level. Second, while this work focuses on a single polytype of SiC, the material exists in multiple polytypes with distinct structural and electronic properties. Developing a transferable, universal MLIP that accurately describes defects across all major SiC polytypes would greatly broaden the applicability of the model. Third, we have verified that the trained MLIP produces stable and accurate simulation trajectories in applications to larger multi-defect systems, but deploying an analogous active learning loop employing larger and more diverse systems would be of value in further fine-tuning and improving the predictive accuracy of the model for more complex multi-defect configurations and transitions.

Author Contributions

A.L.F., G.G., and F.G. designed the research. S.D. designed and executed the active learning workflow for MLIP parameterization, including collective variable discovery and enhanced sampling, DFT energy and force calculations, training of the MLIP model, and convergence assessment. C.Z. and S.D. computed the defect transition free energy barriers. G.P.L. provided guidance in MLIP training and evaluation. A.L.F., G.G., F.G., and J.J.dP. supervised the research. S.D. and A.L.F. wrote the manuscript. A.L.F., G.G., F.G., and C.Z. reviewed and critically revised the manuscript. All authors read and approved the final manuscript.

Acknowledgements

This work was supported by MICCoM (Midwest Center for Computational Materials), as part of the Computational Materials Science Program funded by the U.S. Department of

Energy, Office of Science, Basic Energy Sciences, Materials Sciences and Engineering Division, through Argonne National Laboratory, under contract no. DE-AC02-06CH11357. This work was completed in part with resources provided by the University of Chicago Research Computing Center. We gratefully acknowledge the computing time at the University of Chicago high-performance GPU-based cyber infrastructure supported by the National Science Foundation under grant no. DMR-1828629.

Conflict of Interest Disclosure

A.L.F. is a co-founder and consultant of Evozyne, Inc. and a co-author of US Patent Applications 16/887,710 and 17/642,582, US Provisional Patent Applications 62/853,919, 62/900,420, 63/314,898, 63/479,378, 63/521,617, and 63/669,836, and International Patent Applications PCT/US2020/035206, PCT/US2020/050466, and PCT/US24/10805.

Code and Data Availability

The codes implementing the active learning workflow for training the Allegro MLIP – including PINES CV discovery, OPES-Explore enhanced sampling, SOAP dissimilarity training point selection, DFT calculation of forces and energies, convergence assessment, and Allegro MLIP training – along with the trained Allegro MLIP are made available via a GitHub repository at https://github.com/Ferg-Lab/SiC_Allegro_MLIP.git.

Supporting Information

Learning of symmetry-adapted CVs using PINES; enhanced sampling using OPES-Explore; selection of new training and validation data; training of the Allegro MLIP; ABF calculation of defect transition free energy barriers; Markov state modeling.

References

- (1) Weber, J. R.; Koehl, W. F.; Varley, J. B.; Janotti, A.; Buckley, B. B.; Van de Walle, C. G.; Awschalom, D. D. Quantum computing with defects. *Proceedings of the National Academy of Sciences* **2010**, *107*, 8513–8518.
- (2) Awschalom, D. D.; Hanson, R.; Wrachtrup, J.; Zhou, B. B. Quantum technologies with optically interfaced solid-state spins. *Nature Photonics* **2018**, *12*, 516–527.
- (3) Lohrmann, A.; Johnson, B. C.; McCallum, J. C.; Castelletto, S. A review on single photon sources in silicon carbide. *Reports on Progress in Physics* **2017**, *80*, 034502.
- (4) Wolfowicz, G.; Heremans, F. J.; Anderson, C. P.; Kanai, S.; Seo, H.; Gali, A.; Galli, G.; Awschalom, D. D. Quantum guidelines for solid-state spin defects. *Nature Reviews Materials* **2021**, *6*, 906–925.
- (5) Carlsson, P.; Son, N. T.; Gali, A.; Isoya, J.; Morishita, N.; Ohshima, T.; Magnusson, B.; Janzén, E. EPR and ab initio calculation study on the EI4 center in 4H- and 6H-SiC. *Physical Review B* **2010**, *82*, 235203.
- (6) Son, N. T.; Carlsson, P.; ul Hassan, J.; Janzén, E.; Umeda, T.; Isoya, J.; Gali, A.; Bockstedte, M.; Morishita, N.; Ohshima, T.; Itoh, H. Divacancy in 4H-SiC. *Physical Review Letters* **2006**, *96*, 055501.
- (7) Carlos, W. E.; Garces, N. Y.; Glaser, E. R.; Fanton, M. A. Annealing of multivacancy defects in 4H-SiC. *Physical Review B* **2006**, *74*, 235201.
- (8) Son, N. T.; Carlsson, P.; ul Hassan, J.; Magnusson, B.; Janzén, E. Defects and carrier compensation in semi-insulating 4H-SiC substrates. *Physical Review B* **2007**, *75*, 155204.
- (9) Lee, E. M. Y.; Yu, A.; de Pablo, J. J.; Galli, G. Stability and molecular pathways to the formation of spin defects in silicon carbide. *Nature Communications* **2021**, *12*, 6325.

- (10) Zhang, C.; Gygi, F.; Galli, G. Engineering the formation of spin-defects from first principles. *Nature Communications* **2023**, *14*, 5985.
- (11) Kimura, T.; Ahn, J.; Deegan, N.; Dibos, A.; Zhang, J.; Pingault, B.; Zhang, C.; Galli, G.; Awschalom, D.; Heremans, F. J. Computationally guided experimental validation of divacancy defect formation in 4H-SiC. *Applied Physics Letters* **2025**, *126*, 164001.
- (12) Bockstedte, M.; Mattausch, A.; Pankratov, O. Ab initio study of the migration of intrinsic defects in 3C-SiC. *Physical Review B* **2003**, *68*, 205201.
- (13) Rurali, R.; Hernández, E.; Godignon, P.; Rebollo, J.; Ordejón, P. First principles studies of neutral vacancies diffusion in SiC. *Computational Materials Science* **2003**, *27*, 36–42.
- (14) Wang, X.; Zhao, M.; Bu, H.; Zhang, H.; He, X.; Wang, A. Formation and annealing behaviors of qubit centers in 4H-SiC from first principles. *Journal of Applied Physics* **2013**, *114*, 194305.
- (15) Yan, X.; Li, P.; Kang, L.; Wei, S.-H.; Huang, B. First-principles study of electronic and diffusion properties of intrinsic defects in 4H-SiC. *Journal of Applied Physics* **2020**, *127*, 085702.
- (16) Zhang, C.; Gygi, F.; Galli, G. Charge state and entropic effects affecting the formation and dynamics of divacancies in 3C-SiC. *Physical Review Materials* **2024**, *8*, 046201.
- (17) Unke, O. T.; Chmiela, S.; Sauceda, H. E.; Gastegger, M.; Poltavsky, I.; Schütt, K. T.; Tkatchenko, A.; Müller, K.-R. Machine Learning Force Fields. *Chemical Reviews* **2021**, *121*, 10142–10186.
- (18) Bartók, A. P.; Payne, M. C.; Kondor, R.; Csányi, G. Gaussian Approximation Potentials: The Accuracy of Quantum Mechanics, without the Electrons. *Physical Review Letters* **2010**, *104*, 136403.

- (19) Chmiela, S.; Sauceda, H. E.; Müller, K.-R.; Tkatchenko, A. Towards exact molecular dynamics simulations with machine-learned force fields. *Nature Communications* **2018**, *9*, 3887.
- (20) Vandermause, J.; Torrisi, S. B.; Batzner, S.; Xie, Y.; Sun, L.; Kolpak, A. M.; Kozinsky, B. On-the-fly active learning of interpretable Bayesian force fields for atomistic rare events. *npj Computational Materials* **2020**, *6*, 20.
- (21) Smith, J. S.; Isayev, O.; Roitberg, A. E. ANI-1: an extensible neural network potential with DFT accuracy at force field computational cost. *Chemical Science* **2017**, *8*, 3192–3203.
- (22) Behler, J.; Parrinello, M. Generalized Neural-Network Representation of High-Dimensional Potential-Energy Surfaces. *Physical Review Letters* **2007**, *98*, 146401.
- (23) Vandermause, J.; Xie, Y.; Lim, J. S.; Owen, C. J.; Kozinsky, B. Active learning of reactive Bayesian force fields applied to heterogeneous catalysis dynamics of H/Pt. *Nature Communications* **2022**, *13*, 5183.
- (24) Drautz, R. Atomic cluster expansion for accurate and transferable interatomic potentials. *Physical Review B* **2019**, *99*, 014104.
- (25) Xie, Y.; Vandermause, J.; Ramakers, S.; Protik, N. H.; Johansson, A.; Kozinsky, B. Uncertainty-aware molecular dynamics from Bayesian active learning for phase transformations and thermal transport in SiC. *npj Computational Materials* **2023**, *9*, 36.
- (26) Batzner, S.; Musaelian, A.; Sun, L.; Geiger, M.; Mailoa, J. P.; Kornbluth, M.; Molinari, N.; Smidt, T. E.; Kozinsky, B. E(3)-equivariant graph neural networks for data-efficient and accurate interatomic potentials. *Nature Communications* **2022**, *13*, 2453.
- (27) Schütt, K. T.; Unke, O. T.; Gastegger, M. Equivariant message passing for the predic-

- tion of tensorial properties and molecular spectra. Proceedings of the 38th International Conference on Machine Learning (PMLR). 2021; pp 9377–9388.
- (28) Thölke, P.; Fabritiis, G. D. TorchMD-NET: Equivariant Transformers for Neural Network based Molecular Potentials. The Tenth International Conference on Learning Representations. 2022.
- (29) Klicpera, J.; Becker, F.; Günnemann, S. GemNet: Universal Directional Graph Neural Networks for Molecules. Advances in Neural Information Processing Systems. 2021.
- (30) Wang, Y.; Wang, T.; Li, S.; He, X.; Li, M.; Wang, Z.; Zheng, N.; Shao, B.; Liu, T.-Y. Enhancing geometric representations for molecules with equivariant vector-scalar interactive message passing. *Nature Communications* **2024**, *15*, 313.
- (31) Musaelian, A.; Batzner, S.; Johansson, A.; Sun, L.; Owen, C. J.; Kornbluth, M.; Kozinsky, B. Learning local equivariant representations for large-scale atomistic dynamics. *Nature Communications* **2023**, *14*, 579.
- (32) Nikidis, E.; Kyriakopoulos, N.; Tohid, R.; Kachrimanis, K.; Kioseoglou, J. Harnessing machine learning for efficient large-scale interatomic potential for sildenafil and pharmaceuticals containing H, C, N, O, and S. *Nanoscale* **2024**, *16*, 18014–18026.
- (33) Leimeroth, N.; Erhard, L.; Albe, K.; Rohrer, J. Machine-learning interatomic potentials from a user’s perspective: A comparison of accuracy, speed and data efficiency. *Modelling and Simulation in Materials Science and Engineering* **2025**, *33*, 065012.
- (34) Majewski, M.; Pérez, A.; Thölke, P.; Doerr, S.; Charron, N. E.; Giorgino, T.; Husic, B. E.; Clementi, C.; Noé, F.; De Fabritiis, G. Machine learning coarse-grained potentials of protein thermodynamics. *Nature Communications* **2023**, *14*, 5739.
- (35) Hjorth Larsen, A.; Jørgen Mortensen, J.; Blomqvist, J.; Castelli, I. E.; Christensen, R.; Dułak, M.; Friis, J.; Groves, M. N.; Hammer, B.; Hargus, C.; Hermes, E. D.;

- Jennings, P. C.; Bjerre Jensen, P.; Kermode, J.; Kitchin, J. R.; Leonhard Kolsbjerg, E.; Kubal, J.; Kaasbjerg, K.; Lysgaard, S.; Bergmann Maronsson, J.; Maxson, T.; Olsen, T.; Pastewka, L.; Peterson, A.; Rostgaard, C.; Schiøtz, J.; Schütt, O.; Strange, M.; Thygesen, K. S.; Vegge, T.; Vilhelmsen, L.; Walter, M.; Zeng, Z.; Jacobsen, K. W. The atomic simulation environment—a Python library for working with atoms. *Journal of Physics: Condensed Matter* **2017**, *29*, 273002.
- (36) Wang, Z.; Wu, H.; Sun, L.; He, X.; Liu, Z.; Shao, B.; Wang, T.; Liu, T.-Y. Improving machine learning force fields for molecular dynamics simulations with fine-grained force metrics. *The Journal of Chemical Physics* **2023**, *159*, 035101.
- (37) Fu, X.; Wu, Z.; Wang, W.; Xie, T.; Keten, S.; Gomez-Bombarelli, R.; Jaakkola, T. Forces are not Enough: Benchmark and Critical Evaluation for Machine Learning Force Fields with Molecular Simulations. *Transactions on Machine Learning Research* **2023**, 1–31.
- (38) Liu, Y.; He, X.; Mo, Y. Discrepancies and error evaluation metrics for machine learning interatomic potentials. *npj Computational Materials* **2023**, *9*, 174.
- (39) Herringer, N. S. M.; Dasetty, S.; Gandhi, D.; Lee, J.; Ferguson, A. L. Permutationally Invariant Networks for Enhanced Sampling (PINES): Discovery of Multimolecular and Solvent-Inclusive Collective Variables. *Journal of Chemical Theory and Computation* **2023**, *20*, 178–198.
- (40) Pietrucci, F. *Handbook of Materials Modeling*; Springer International Publishing, 2018; pp 1–23.
- (41) Kingma, D. P.; Welling, M. An Introduction to Variational Autoencoders. *Foundations and Trends in Machine Learning* **2019**, *12*, 307–392.
- (42) Sidky, H.; Chen, W.; Ferguson, A. L. Machine learning for collective variable discov-

- ery and enhanced sampling in biomolecular simulation. *Molecular Physics* **2020**, *118*, e1737742.
- (43) Chen, W.; Ferguson, A. L. Molecular enhanced sampling with autoencoders: On-the-fly collective variable discovery and accelerated free energy landscape exploration. *Journal of Computational Chemistry* **2018**, *39*, 2079–2102.
- (44) Chen, W.; Tan, A. R.; Ferguson, A. L. Collective variable discovery and enhanced sampling using autoencoders: Innovations in network architecture and error function design. *The Journal of Chemical Physics* **2018**, *149*, 072312.
- (45) Bonomi, M.; Bussi, G.; Camilloni, C.; Tribello, G. A.; Banáš, P.; Barducci, A.; Bernetti, M.; Bolhuis, P. G.; Bottaro, S.; Branduardi, D.; Capelli, R.; Carloni, P.; Ceriotti, M.; Cesari, A.; Chen, H.; Chen, W.; Colizzi, F.; De, S.; De La Pierre, M.; Donadio, D.; Drobot, V.; Ensing, B.; Ferguson, A. L.; Filizola, M.; Fraser, J. S.; Fu, H.; Gasparotto, P.; Gervasio, F. L.; Giberti, F.; Gil-Ley, A.; Giorgino, T.; Heller, G. T.; Hocky, G. M.; Iannuzzi, M.; Invernizzi, M.; Jelfs, K. E.; Jussupow, A.; Kirilin, E.; Laio, A.; Limongelli, V.; Lindorff-Larsen, K.; Löhr, T.; Marinelli, F.; Martin-Samos, L.; Masetti, M.; Meyer, R.; Michaelides, A.; Molteni, C.; Morishita, T.; Nava, M.; Paissoni, C.; Papaleo, E.; Parrinello, M.; Pfaendtner, J.; Piaggi, P.; Piccini, G.; Pietropaolo, A.; Pietrucci, F.; Pipolo, S.; Provasi, D.; Quigley, D.; Raiteri, P.; Raniolo, S.; Rydzewski, J.; Salvalaglio, M.; Sosso, G. C.; Spiwok, V.; Šponer, J.; Swenson, D. W. H.; Tiwary, P.; Valsson, O.; Vendruscolo, M.; Voth, G. A.; White, A. Promoting transparency and reproducibility in enhanced molecular simulations. *Nature Methods* **2019**, *16*, 670–673.
- (46) Thompson, A. P.; Aktulga, H. M.; Berger, R.; Bolintineanu, D. S.; Brown, W. M.; Crozier, P. S.; in 't Veld, P. J.; Kohlmeyer, A.; Moore, S. G.; Nguyen, T. D.; Shan, R.; Stevens, M. J.; Tranchida, J.; Trott, C.; Plimpton, S. J. LAMMPS - a flexible simulation

- tool for particle-based materials modeling at the atomic, meso, and continuum scales. *Computer Physics Communications* **2022**, *271*, 108171.
- (47) Johansson, A.; Weinberg, E.; Trott, C.; McCarthy, M.; Moore, S. LAMMPS-KOKKOS: Performance Portable Molecular Dynamics Across Exascale Architectures. Proceedings of the SC '25 Workshops of the International Conference for High Performance Computing, Networking, Storage and Analysis. 2025; pp 1217–1232.
- (48) Allen, M. P.; Tildesley, D. J. *Computer Simulation of Liquids*; Oxford University Press, 2017.
- (49) Justo, J. F.; Bazant, M. Z.; Kaxiras, E.; Bulatov, V. V.; Yip, S. Interatomic potential for silicon defects and disordered phases. *Physical Review B* **1998**, *58*, 2539–2550.
- (50) Perego, S.; Bonati, L. Data efficient machine learning potentials for modeling catalytic reactivity via active learning and enhanced sampling. *npj Computational Materials* **2024**, *10*, 291.
- (51) Owen, C. J.; Torrisi, S. B.; Xie, Y.; Batzner, S.; Bystrom, K.; Coulter, J.; Musaelian, A.; Sun, L.; Kozinsky, B. Complexity of many-body interactions in transition metals via machine-learned force fields from the TM23 data set. *npj Computational Materials* **2024**, *10*, 92.
- (52) Xu, Y.; Jin, Y.; García Sánchez, J. S.; Pérez-Lemus, G. R.; Zubieta Rico, P. F.; Delferro, M.; de Pablo, J. J. A molecular view of methane activation on Ni (111) through enhanced sampling and machine learning. *The Journal of Physical Chemistry Letters* **2024**, *15*, 9852–9862.
- (53) Perez-Lemus, G.; Xu, Y.; Jin, Y.; Zubieta Rico, P.; de Pablo, J. The importance of sampling the dynamical modes: Reevaluating benchmarks for invariant and equivariant features of machine learning potentials for simulation of free energy landscapes. *The Journal of Chemical Physics* **2024**, *161*, 244703.

- (54) Jin, Y.; Perez-Lemus, G. R.; Zubieta Rico, P. F.; de Pablo, J. J. Improving machine learned force fields for complex fluids through enhanced sampling: A liquid crystal case study. *The Journal of Physical Chemistry A* **2024**, *128*, 7257–7268.
- (55) Bartók, A. P.; Kondor, R.; Csányi, G. On representing chemical environments. *Physical Review B* **2013**, *87*, 184115.
- (56) De, S.; Bartók, A. P.; Csányi, G.; Ceriotti, M. Comparing molecules and solids across structural and alchemical space. *Physical Chemistry Chemical Physics* **2016**, *18*, 13754–13769.
- (57) Himanen, L.; Jäger, M. O.; Morooka, E. V.; Federici Canova, F.; Ranawat, Y. S.; Gao, D. Z.; Rinke, P.; Foster, A. S. DDescribe: Library of descriptors for machine learning in materials science. *Computer Physics Communications* **2020**, *247*, 106949.
- (58) Montes de Oca Zapiain, D.; Wood, M. A.; Lubbers, N.; Pereyra, C. Z.; Thompson, A. P.; Perez, D. Training data selection for accuracy and transferability of interatomic potentials. *npj Computational Materials* **2022**, *8*, 189.
- (59) Gygi, F. Architecture of Qbox: A scalable first-principles molecular dynamics code. *IBM Journal of Research and Development* **2008**, *52*, 137–144.
- (60) Perdew, J. P.; Burke, K.; Ernzerhof, M. Generalized Gradient Approximation Made Simple. *Physical Review Letters* **1996**, *77*, 3865–3868.
- (61) Schlipf, M.; Gygi, F. Optimization algorithm for the generation of ONCV pseudopotentials. *Computer Physics Communications* **2015**, *196*, 36–44.
- (62) Zuo, Y.; Chen, C.; Li, X.; Deng, Z.; Chen, Y.; Behler, J.; Csányi, G.; Shapeev, A. V.; Thompson, A. P.; Wood, M. A.; Ong, S. P. Performance and Cost Assessment of Machine Learning Interatomic Potentials. *The Journal of Physical Chemistry A* **2020**, *124*, 731–745.

- (63) Kingma, D. P.; Ba, J. Adam: A Method for Stochastic Optimization. The Third International Conference for Learning Representations. 2015.
- (64) Unke, O. T.; Chmiela, S.; Gastegger, M.; Schütt, K. T.; Sauceda, H. E.; Müller, K.-R. SpookyNet: Learning force fields with electronic degrees of freedom and nonlocal effects. *Nature Communications* **2021**, *12*, 7273.
- (65) Jacobs, R.; Morgan, D.; Attarian, S.; Meng, J.; Shen, C.; Wu, Z.; Xie, C. Y.; Yang, J. H.; Artrith, N.; Blaiszik, B.; Ceder, G.; Choudhary, K.; Csanyi, G.; Cubuk, E. D.; Deng, B.; Drautz, R.; Fu, X.; Godwin, J.; Honavar, V.; Isayev, O.; Johansson, A.; Kozinsky, B.; Martiniani, S.; Ong, S. P.; Poltavsky, I.; Schmidt, K.; Takamoto, S.; Thompson, A. P.; Westermayr, J.; Wood, B. M. A practical guide to machine learning interatomic potentials – Status and future. *Current Opinion in Solid State and Materials Science* **2025**, *35*, 101214.
- (66) Zhang, L.; Csányi, G.; van der Giessen, E.; Maresca, F. Efficiency, accuracy, and transferability of machine learning potentials: Application to dislocations and cracks in iron. *Acta Materialia* **2024**, *270*, 119788.
- (67) Baghishov, I.; Janssen, J.; Henkelman, G.; Perez, D. Application-specific machine-learned interatomic potentials: exploring the trade-off between DFT convergence, MLIP expressivity, and computational cost. *Digital Discovery* **2026**, *5*, 332–347.
- (68) Tan, A. R.; Dietschreit, J. C. B.; Gómez-Bombarelli, R. Enhanced sampling of robust molecular datasets with uncertainty-based collective variables. *The Journal of Chemical Physics* **2025**, *162*.
- (69) Darve, E.; Rodríguez-Gómez, D.; Pohorille, A. Adaptive biasing force method for scalar and vector free energy calculations. *The Journal of Chemical Physics* **2008**, *128*, 144120.

- (70) Sidky, H.; Colón, Y. J.; Helfferich, J.; Sikora, B. J.; Bezik, C.; Chu, W.; Giberti, F.; Guo, A. Z.; Jiang, X.; Lequieu, J.; Li, J.; Moller, J.; Quevillon, M. J.; Rahimi, M.; Ramezani-Dakhel, H.; Rathee, V. S.; Reid, D. R.; Sevgen, E.; Thapar, V.; Webb, M. A.; Whitmer, J. K.; de Pablo, J. J. SSAGES: Software Suite for Advanced General Ensemble Simulations. *The Journal of Chemical Physics* **2018**, *148*, 044104.
- (71) Prinz, J.-H.; Wu, H.; Sarich, M.; Keller, B.; Senne, M.; Held, M.; Chodera, J. D.; Schütte, C.; Noé, F. Markov models of molecular kinetics: Generation and validation. *The Journal of Chemical Physics* **2011**, *134*, 174105.
- (72) Scherer, M. K.; Trendelkamp-Schroer, B.; Paul, F.; Pérez-Hernández, G.; Hoffmann, M.; Plattner, N.; Wehmeyer, C.; Prinz, J.-H.; Noé, F. PyEMMA 2: A Software Package for Estimation, Validation, and Analysis of Markov Models. *Journal of Chemical Theory and Computation* **2015**, *11*, 5525–5542.
- (73) Röblitz, S.; Weber, M. Fuzzy spectral clustering by PCCA+: application to Markov state models and data classification. *Advances in Data Analysis and Classification* **2013**, *7*, 147–179.
- (74) Noé, F.; Schütte, C.; Vanden-Eijnden, E.; Reich, L.; Weikl, T. R. Constructing the equilibrium ensemble of folding pathways from short off-equilibrium simulations. *Proceedings of the National Academy of Sciences* **2009**, *106*, 19011–19016.
- (75) Karsthof, R.; Bathen, M. E.; Galeckas, A.; Vines, L. Conversion pathways of primary defects by annealing in proton-irradiated n-type 4 H-SiC. *Physical Review B* **2020**, *102*, 184111.
- (76) Bathen, M. E.; Coutinho, J.; Ayedh, H.; Ul Hassan, J.; Farkas, I.; Öberg, S.; Frodason, Y. K.; Svensson, B. G.; Vines, L. Anisotropic and plane-selective migration of the carbon vacancy in SiC: Theory and experiment. *Physical Review B* **2019**, *100*, 014103.

- (77) Yuan, Z.; Xu, Z.; Li, H.; Cheng, X.; Tao, H.; Tang, Z.; Zhou, Z.; Duan, W.; Xu, Y. Equivariant neural network force fields for magnetic materials. *Quantum Frontiers* **2024**, *3*, 1–10.
- (78) Ueno, K.; Ohuchi, S.; Ichikawa, K.; Amii, K.; Wakasugi, K. SpinMultiNet: Neural Network Potential Incorporating Spin Degrees of Freedom with Multi-Task Learning. 2024; <https://arxiv.org/abs/2409.03253>.
- (79) Deng, B.; Zhong, P.; Jun, K.; Riebesell, J.; Han, K.; Bartel, C. J.; Ceder, G. CHGNet as a pretrained universal neural network potential for charge-informed atomistic modelling. *Nature Machine Intelligence* **2023**, *5*, 1031–1041.
- (80) Kalita, B.; Zubatyuk, R.; Anstine, D. M.; Bergeler, M.; Settels, V.; Stork, C.; Spicher, S.; Isayev, O. AIMNet2-NSE: A Transferable Reactive Neural Network Potential for Open-Shell Chemistry. *Angewandte Chemie International Edition* **2025**, *65*, 307–392.

TOC Graphic

