

RESEARCH ARTICLE | JUNE 26 2023

Evaluating approaches for on-the-fly machine learning interatomic potentials for activated mechanisms sampling with the activation-relaxation technique nouveau **FREE**

Special Collection: [Machine Learning Hits Molecular Simulations](#)

Eugène Sanscartier ; Félix Saint-Denis ; Karl-Étienne Bolduc ; Normand Mousseau  



J. Chem. Phys. 158, 244110 (2023)

<https://doi.org/10.1063/5.0143211>



View
Online



Export
Citation

CrossMark



The Journal of Chemical Physics

Special Topic: Adhesion and Friction

Submit Today!

AIP
Publishing

AIP
Publishing

Evaluating approaches for on-the-fly machine learning interatomic potentials for activated mechanisms sampling with the activation-relaxation technique nouveau

Cite as: J. Chem. Phys. 158, 244110 (2023); doi: 10.1063/5.0143211

Submitted: 20 January 2023 • Accepted: 31 May 2023 •

Published Online: 26 June 2023



View Online



Export Citation



CrossMark

Eugène Sanscartier,^{a)}  Félix Saint-Denis,^{b)}  Karl-Étienne Bolduc,^{b)}  and Normand Mousseau^{b)} 

AFFILIATIONS

Département de physique and Regroupement québécois sur les matériaux de pointe, Université de Montréal, Case Postale 6128, Succursale Centre-ville, Montréal, Québec H3C 3J7, Canada

Note: This paper is part of the JCP Special Topic on Machine Learning Hits Molecular Simulations.

^{a)} **Electronic mail:** eugene.sanscartier@umontreal.ca

^{b)} **Author to whom correspondence should be addressed:** normand.mousseau@umontreal.ca.

URL: <https://normandmousseau.com>

ABSTRACT

In the last few years, much effort has gone into developing general machine-learning potentials capable of describing interactions for a wide range of structures and phases. Yet, as attention turns to more complex materials, including alloys and disordered and heterogeneous systems, the challenge of providing reliable descriptions for all possible environments becomes ever more costly. In this work, we evaluate the benefits of using specific vs general potentials for the study of activated mechanisms in solid-state materials. More specifically, we test three machine-learning fitting approaches using the moment-tensor potential to reproduce a reference potential when exploring the energy landscape around a vacancy in Stillinger–Weber silicon crystal and silicon–germanium zincblende structures using the activation-relaxation technique nouveau (ARTn). We find that a targeted on-the-fly approach specific to and integrated into ARTn generates the highest precision on the energetics and geometry of activated barriers while remaining cost-effective. This approach expands the types of problems that can be addressed with high-accuracy ML potential.

Published under an exclusive license by AIP Publishing. <https://doi.org/10.1063/5.0143211>

I. INTRODUCTION

As computational materials scientists turn their attention to ever more complex systems, they are faced with two major challenges: (i) how to correctly describe their physics and (ii) how to reach the appropriate size and time scale to capture the properties of interest. The first challenge is generally solved by turning to *ab initio* methods¹ that allow the solution of Schrödinger's equation with reasonably controlled approximations. These approaches, however, suffer from N^4 scaling, which limits their application to small system sizes and short time scales. The second challenge is met by a variety of methods that cover different scales. Molecular dynamics,² for example, which directly solves Newton's equation, accesses typical time scales between picoseconds and microseconds at the

very best. Other approaches, such as lattice^{3,4} and off-lattice kinetic Monte-Carlo,^{5,6} by focusing on physically relevant mechanisms, can extend this time scale to seconds and more, as long as the diffusion takes place through activated processes. Even though these methods are efficient, each trajectory can require hundreds of thousands to millions of force evaluations, becoming too costly with *ab initio* approaches and forcing modelers to use empirical potentials despite their incapacity at correctly describing complex environments.

Building on *ab initio* energy and forces, machine-learned potentials (MLP)^{7–10} open the door to lifting some of these difficulties by offering much more reliable physics at a small fraction of the cost of *ab initio* evaluations.

Since their introduction, ML potentials have been largely coupled with MD, focusing on the search for universal potentials

able to describe a full range of structures and phases for a given material.^{11–13} As we turn to more complex systems such as alloys and disordered and heterogeneous systems, it becomes more and more difficult to generate such universal potentials since the number of possible environments grows rapidly with this complexity. In this context, the development of specific potentials with on-the-fly learning that makes it possible to adapt to new environments becomes a strategy worth exploring.

In this work, we focus on the construction of machine-learned potentials adapted to the sampling of an energy landscape dominated by activated mechanisms, i.e., solid-state systems with local activated diffusion and evolution. This kinetics is associated with the aging and relaxation of disordered materials,¹⁴ sluggish diffusion in concentrated alloys,¹⁵ and defect diffusion in complex environments.¹⁶ Correct computational sampling, using open-ended methods such as the activation-relaxation technique (ART)¹⁷ and its revised version (ART nouveau or ARTn),^{18,19} requires a precise description of local minima and of the landscape surrounding the first-order saddle points that characterize diffusion according to the transition-state theory (TST).²⁰ These barriers can be high—reaching many electron-volts—and involve strained configurations that can be visited only very rarely with standard molecular dynamics. Yet, because of this need for long-time evolution, these problems have been largely out of reach of *ab initio* approaches and have been studied, until now, mostly with empirical potential.

More specifically, we compare three machine learning approaches to train a Moment Tensor Potential (MTP)^{10,21} for the diffusion of a vacancy in silicon and silicon–germanium alloy as sampled with ARTn. The first approach consists of pure MD learning, fitted at various temperatures, following steps that echo the work of Novoselov *et al.*;²² the second approach adds on-the-fly training during ARTn runs; and the third one focuses on purely on-the-fly training during ARTn runs. While on-the-fly learning is not a new approach and has been used before in various contexts, it has never been compared to the specific application of activated kinetics in solid state systems.

To generate the statistics necessary to offer solid conclusions, we selected to use the Stilliger–Weber empirical potential both for Si²³ and SiGe.²⁴ Clearly, the machine-learned potentials generated here have, therefore, limited physical relevance, but these are sufficient to allow us to assess the relative quality of our three approaches.

Results underline the efficiency gain in developing targeted ML potentials for specific applications by comparing the cost of fitting Si with SiGe. They also show the rapid increase in computation complexity associated with moving from element to alloy systems, which emphasizes the usefulness of a specific approach such as the one applied here to activated processes.

II. METHODOLOGY

A. ML potential

The Moment Tensor Potential (MTP)^{10,21} is a linear model of functions $B_\alpha(\mathbf{r}_i)$ built from contractions of moment tensor descriptors defined by the local neighborhood relative position \mathbf{r}_i of atom i within a sphere of influence of radius r_c respecting a set of invariances. This model has been shown to be fast while giving accuracy on

the order of \sim meV/atom and requiring a few hundred to thousands of reference potential calls²⁵ on-the-fly.

MTP has been used on a wide variety of problems, including on-the-fly MD simulation,^{21,22,26} search and minimization of new alloys,^{27,28} and diffusion processes²² on systems counting one or multiple species. In the following, we offer a summary of the method; more information on MTP is available in Ref. 21.

MTP approximates atomic configuration energy as the sum of local contributions. A local contribution is obtained through a sum over the included basis $\{B_\alpha(\mathbf{r}_i)\}$ as a linear combination of $B(\mathbf{r}_i)$ and ξ_α ,

$$V(\mathbf{r}_i) = \sum_{\alpha=1}^m \xi_\alpha B_\alpha(\mathbf{r}_i). \quad (1)$$

The “level” of a potential gives the number of different possible tensor $M_{\mu,\nu}(\mathbf{r}_i)$ descriptors. The $B_\alpha(\mathbf{r}_i)$ functions of Eq. (1) are constructed by a tensorial contraction of different $M_{\mu,\nu}(\mathbf{r}_i)$, and the number of different tensorial contractions sets m in Eq. (1). More information on MTP is available in Ref. 21.

The total energy of a N -atom configuration (\mathcal{R}) is then given by the sum of N local contributions,

$$E(\mathcal{R}) = \sum_{i=1}^N V(\mathbf{r}_i) = \sum_{i=1}^N \sum_{\alpha=1}^m \xi_\alpha B_\alpha(\mathbf{r}_i), \quad (2)$$

and the forces are obtained by taking the gradient of this quantity,

$$\mathbf{F}(\mathcal{R}) = -\nabla \sum_{i=1}^N \sum_{\alpha=1}^m \xi_\alpha B_\alpha(\mathbf{r}_i). \quad (3)$$

The parameters ξ_α are obtained by minimizing the loss function,

$$\sum_{\mathcal{R} \in A} \left[w_e (E(\mathcal{R}) - \hat{E}(\mathcal{R}))^2 + w_f \sum_i |\mathbf{f}_i(\mathcal{R}) - \hat{\mathbf{f}}_i(\mathcal{R})|^2 \right] \rightarrow \min_{\xi}. \quad (4)$$

Here, A is the training set made of configurations with known energy and forces. The goal is to minimize the difference between $E(\mathcal{R})$, $\mathbf{f}_i(\mathcal{R})$ (the real value) and $\hat{E}(\mathcal{R})$, $\hat{\mathbf{f}}_i(\mathcal{R})$ (predicted by the model), respectively, for all elements in A . Weights on contributions from energy and forces (w_e and w_f) are set to one.

B. Learning on-the-fly tools

On-the-fly atomic machine learning potential (OTF) involves the repeated training of the model potential as new atomic environments are generated through various procedures.

Following the work of Shapeev and collaborators,²¹ the reliability of the potential to evaluate a given configuration is performed using the D-optimality criterion, which states that the best training set is the one with maximal volume. The volume can be seen as a measure of the domain of the training set where the model allows reliable predictions. Assuming that the training set is of maximal volume, any configuration within or beyond this volume is inside or outside the known domain of the model, respectively, which we interpret as interpolation and extrapolation. We grade the D-optimality criterion of a given configuration by assessing the model’s reliability. For this, Shapeev *et al.* introduce a selection

algorithm (MaxVol) that tests whether or not this configuration should be added to the training set or replace a configuration already in it. While a detailed description can be found in Ref. 26, we provide here a brief summary of the retained approach.

The selection and extrapolation-grade algorithm can be applied using either a local-energy or a global-energy descriptor.

The local-energy descriptor is presented as a rectangular matrix $G_{N \times m}$ formed by the basis elements $B_\alpha(\mathbf{r}_i)$ associated with the neighborhood \mathbf{r}_i of all N atoms,

$$G = \begin{pmatrix} B_1(\mathbf{r}_1) & \dots & B_m(\mathbf{r}_1) \\ \vdots & \ddots & \vdots \\ B_1(\mathbf{r}_N) & \dots & B_m(\mathbf{r}_N) \end{pmatrix}. \quad (5)$$

For a given configuration, the global-energy description reduces this information to a vector \mathbf{g} ,

$$\mathbf{g} = (b_1(\mathfrak{R}) \quad \dots \quad b_m(\mathfrak{R})), \quad (6)$$

where each term, $\{b_\alpha(\mathfrak{R})\}$, is a sum over all neighborhoods for a specific basis element, $\{B_\alpha(\mathbf{r}_i)\}$,

$$\{b_\alpha(\mathfrak{R})\} = \sum_{i=0}^N \{B_\alpha(\mathbf{r}_i)\}.$$

For the global-energy descriptor of a given configuration with a training set A , solving for c_{ij} , in

$$\begin{pmatrix} c_{11} & \dots & c_{1m} \end{pmatrix} A = \mathbf{g} \quad (7)$$

Assuming that A is of maximal volume, $0 \leq c_{ij} \leq 1$. Conversely, if one finds $c_{ij} > 1$, then A is not of maximal volume. The latter case means that the model is extrapolating on this configuration and that updating A with this configuration will increase its volume. The extrapolation grade, γ , is then defined as the largest component of c_{ij} ,

$$\gamma(\mathfrak{R}) = \max |c_{ij}|. \quad (8)$$

The same approach is used for the local-energy description, applying Eq. (7) to the rows of matrix G rather than the vector \mathbf{g} . For non-linear and other ML potentials, in Eq. (5), one take the ML model's gradient with respect to its parameters.

In practice, for $\gamma(\mathfrak{R})$ below a certain threshold γ_0 , the model interpolation is reliable while for $\gamma_0 < \gamma(\mathfrak{R}) < \gamma_{\max}$, the model cannot be applied with confidence but can be adapted by adding this configuration to the training set. When $\gamma(\mathfrak{R}) > \gamma_{\max}$, the configuration is too far from the training set and is rejected as the model cannot be adapted with confidence. In this work, we set $\gamma_0 = 1.1$ and $\gamma_{\max} = 2.2$, taking the lower values of the parameters studied in Ref. 26 unless specified otherwise.

C. On-the-fly learning cycle workflow

Our workflow is similar to that of Ref. 21, with the main differences discussed in Sec. II F. We follow the same general machine-learning on-the-fly workflow for all sampling approaches tested here.

We split each simulation into one or multiple sequences of atomic configurations generated using either MD or ARTn. Each run unfolds as follows (see Fig. 1):

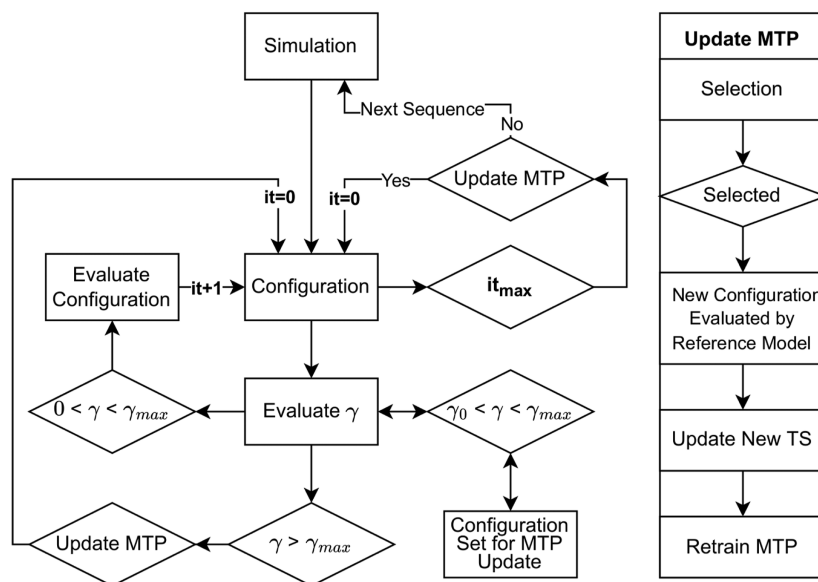


FIG. 1. On-the-fly machine learning workflow used with MD and ARTn (on the left). A potential update can take place at two points: when the sequence ends or when $\gamma > \gamma_{\max}$. The updating procedures are given in the box on the right.

1. Launch a sequence during which configurations are generated according to a sampling algorithm (MD or ARTn).

At each iteration step, the extrapolation-grade γ is evaluated.

- (a) If $0 < \gamma < \gamma_{\max}$, the energy and forces of the configuration are evaluated with MTP;
 - (b) if $\gamma_0 < \gamma < \gamma_{\max}$, the configuration is set aside for an update of MTP parameters;
 - (c) otherwise, if $\gamma > \gamma_{\max}$, energy and forces of the configuration are not evaluated with MTP, and the configuration is not kept for the update. The sequence is stopped, and we go directly to the update step (step 3).
2. Move on to the next iteration in the sequence (step 1).
 3. The model is updated if at least one configuration has been set aside for an update of MTP (i) at the end of a sequence or (ii) at any moment during the sequence if $\gamma > \gamma_{\max}$.
 4. If there is an update, restart a new sequence (go to step 1), or stop if no configuration with $\gamma > \gamma_0$ has been set aside during the predefined maximum length of the sequence.

The moment tensor potential model update is defined as follows (see Fig. 1, right-hand side):

1. A selection is made from the set aside configurations (with $\gamma > \gamma_0$) using MaxVol.²⁶
2. Each selected configuration is evaluated by the reference model.
3. The training set is updated with the new evaluated configurations.
4. The moment tensor potential is fitted on the new training set accordingly to Eq. (4).

More details of this procedure can be found in Ref. 26.

D. MD and ARTn

Two sampling approaches are used to generate a sequence of configurations: (1) molecular dynamics (MD) as implemented within LAMMPS²⁹ and (2) the activation-relaxation technique nouveau (ARTn) algorithm developed by Mousseau and collaborators.^{17–19} Since the MD is well known, we only give below a brief summary of ARTn.

ARTn is designed to explore the potential energy landscape of atomic systems through the identification of local transition states connecting nearby local minima. Its workflow can be summarized in three main steps (see, for a recent in-depth discussion of the ARTn version used in this work, Ref. 19) as follows:

1. *Leaving the harmonic well:* Starting from an energy minimum, an atom and its neighbors are moved iteratively in a direction selected at random until a direction of negative curvature on the potential energy surfaces, $\mathbf{d}(\lambda_{\min})$ with λ_{\min} , the lowest eigenvalue of the Hessian matrix, smaller than zero, emerges; this indicates the presence of a nearby first-order saddle point.
2. *Converging to a first-order saddle point:* The system is then pushed in the direction of negative curvature $\mathbf{d}(\lambda_{\min})$ while the force is minimized in the perpendicular plane until the

total force F passes below a threshold near F_0 , which indicates the saddle point has been reached.

3. *Relaxing into a new minimum:* The system is then pushed over the saddle point and relaxed into a connected new minimum.

At each step, λ_{\min} and $\mathbf{d}(\lambda_{\min})$ are found using an iterative Lanczos method.^{18,19,30} Perpendicular relaxation during activation and global minimization are performed using the Fast Inertial Relaxation Engine (FIRE) algorithm.³²

Generated events are accepted or rejected according to the Metropolis algorithm, where the acceptance probability p is given by

$$p = \min\left(1, e^{-\beta\Delta E}\right), \quad (9)$$

with $\Delta E = E_{\text{saddle}} - E_{\text{minimum}}$, the energy difference between the saddle and a connected minima, and $\beta = 1/k_B T$, where k_B is the Boltzmann factor and T is a fictitious temperature since thermal deformations are not taken into account. Potential energy landscape exploration consists of generating a number of events.

E. Systems studied

The fitting approaches are tested on two physical systems: (i) a Si diamond structure with Stillinger–Weber as a reference potential;²³ and (ii) a SiGe zincblende structure using the Stillinger–Weber potential with parameters from Ref. 24. Both models count 215 atoms and a vacancy. These potentials are selected as they are computationally light and facilitate the accumulation of statistics at the level needed here; to obtain physically-relevant results, one should train on density functional theory (DFT) and not on empirical potentials.

The Si system is fitted with an ML potential set at level 16, with 92 moment tensor functions [$B(\mathfrak{R})$, Eq. (1)]. For SiGe, a potential at this level (16) generates errors on the barrier of the order of 0.5 eV, which indicates that a richer set of parameters is needed to describe the chemical diversity, and a level 20 is chosen for this system, with 288 moment tensor functions. The relation between the number of moment tensor functions for Si and energy error is presented in the supplementary material, Fig. 1.

F. Fitting approaches

To evaluate the reliability of the various on-the-fly approaches to reproduce the reference potential on configurations of interest for complex materials, the training set is limited to structures visited during MD or ARTn simulations within the conditions described later. No additional information regarding alternative crystalline structures, defects, surfaces, pressure, etc., is provided.

For each of these two systems, we compare the following approaches:

1. ML-MD: The MTP potential is to train OTF on MD simulations. The potential is then evaluated, *without further update*, in the ARTn simulation.
2. OTF-MDART: Starting from the ML-MD generated potential, the MTP is re-trained following the OTF procedure during ARTn simulations.
3. OTF-ART: Training of the potential is performed uniquely during ARTn runs with OTF.

The ML-MD approach is in line with Ref. 22, where a potential is trained OTF during MD. However, while the potential is trained with MD, its accuracy is evaluated during an ARTn activated process search.

1. ML-MD: Simulations details

Nine sets of MTP ML-MD potentials are developed and trained independently during NVT MD simulations. Each set is trained at one specific simulation temperature ranging from 300 to 2700 K by steps of 300 K and starting from the same 215 atom crystalline structure with a vacancy. Each set consists of ten independently constructed MTP potentials for statistical purposes.

Training takes place on a series of sequences, each run for a maximum of 100 ps with steps of 1 fs, with an average of 75 ps per cycle. MTP potentials require about 34 ± 14 and 93 ± 43 learning cycles for Si and SiGe to be converged; the MTP potential is considered to have learned the potential when no configuration generated during a 100 ps second is found in the extrapolating zone of the potential (with $\gamma > \gamma_0$).

As long as this is not the case, the sequence is restarted from the same initial structure with different initial velocities. To facilitate convergence, ML-MD potentials are fitted over three sets of progressively more restricted reliability extrapolation parameter γ_0 . Moreover, because MD leads to global deformation, the extrapolation is computed using global descriptors (see Table I).

The final potential is then evaluated, in a fixed form, in ARTn simulations.

2. OTF ARTn simulations details

Each ARTn simulation is launched for 1500 events, with 24 parallel independent searches, for a total of 36 000 generated events. For ARTn, a sequence is either a search for a saddle point (successful or failed) or a minimization from the saddle to the minimum.

At each point, 24 sequences are generated in parallel, and the configuration selected for an update of the potential is based on the combined set of configurations to generate one training set. The sequence is restarted from the last accepted position or, in the case of the vacancy in Si, the ground state. When an activation step generates a configuration with $\gamma(\mathcal{R}) > \gamma_{\max}$, it is relaunched with the same initial deformation. As with MD, ten independent ARTn runs are launched for statistics.

In the bulk, diffusion of the vacancy in Si takes place through a symmetric mechanism bringing the vacancy from one state to an identical one, so all ARTn event searches are effectively started from the same state. Starting from a zincblende structure, SiGe evolves according to an accept-reject Metropolis with a fictitious temperature of 0.5 eV.³² Since the configurations explored by ARTn are

locally deformed, the extrapolation grade for ARTn generated configurations used for the OTF-MDART and OTF-ART approaches is evaluated with the local descriptors.

G. Analysis

Following the standard approach, the error is computed based on the energy and force differences between the MLP and reference potentials computed on the same structures. Here, however, this error is only measured on configurations generated during the ARTn procedure.

For the energy,

$$\Delta E = |E_{MLP}(X_{MLP}) - E_{ref}(X_{MLP})|, \quad (10)$$

and, for the forces,

$$\Delta F = \frac{1}{N} \sum_{i=0}^N \sqrt{\|\mathbf{f}_{MLP}^{(i)}(X_{MLP}) - \mathbf{f}_{ref}^{(i)}(X_{ref})\|^2}, \quad (11)$$

where the positions X_{MLP} are obtained from a simulation run with the machine-learned potential, and the energy on this exact configuration is computed with the reference and the machine-learned potentials. The same is performed for the error on forces.

Since this work is focused on the correct description of first-order transition states, we also compute the minimum and saddle barrier positions and energy convergence errors (ΔX_{conv} , ΔE_{conv}) as

$$\Delta X_{conv} = \sqrt{\sum_{i=0}^N \|\mathbf{x}_{MLP}^{(i)} - \mathbf{x}_{ref}^{(i)}\|^2}, \quad (12)$$

$$\Delta E_{conv} = |E_{MLP}(X_{MLP}) - E_{ref}(X_{ref})|, \quad (13)$$

where X_{MLP} and X_{ref} are the positions corresponding to the minimum or saddle point as defined by the MLP and the reference potentials, respectively, with $E_{MLP}(X_{MLP})$ and $E_{ref}(X_{ref})$ the corresponding energies; by definition, forces are zero at these points defined by the respective potentials.

While X_{MLP} and $E_{MLP}(X_{MLP})$ are obtained on the ARTn trajectories, X_{ref} and $E_{ref}(X_{ref})$ are obtained after reconverging the minima or the saddle point using the reference potential starting from X_{MLP} and following the ARTn procedure.

From an energy barrier $\delta E(X)$, the energy barrier error $\Delta \delta E_{barrier}$ is given by

$$\Delta \delta E_{barrier} = |\delta E_{MLP}(X_{MLP}) - \delta E_{ref}(X_{ref})|. \quad (14)$$

If no trend is observed between the different temperatures where potentials are trained, we calculate their average and deviation in order to effectively compare them with other approaches.

III. RESULTS

In this section, we first examine results for a vacancy in c-Si to establish the methods, then consider the same approaches on the more complex SiGe alloy.

A. ML-MD

The ML-MD approach serves as a benchmark to assess the efficiency of the various approaches to sampling energy barriers

TABLE I. Extrapolation and selection hyperparameter values used for the three on-the-fly approaches used in this work.

Approach	γ_0	γ_{\max}	Grade-mode
ML-MD	5.5/3.3/1.1	60/10/2.2	Global
OTF-MDART	1.1	2.2	Local
OTF-ART	1.1	2.2	Local




FIG. 2. Number of calls to the reference potential for each of the machine-learned potentials developed for Si as a function of the temperature referring to the one used during MD training. Since the configurations are relaxed to zero K in ARTn simulations, there is no associated temperature for this procedure. Vertical bars represent the standard deviation computed on ten independent realizations.

and diffusion mechanisms. Here, ten independent ML potentials are generated through on-the-fly MD simulations at nine different target temperatures ranging from 300 to 2700 K by step of 300 K and require between 253 ± 60 , at 300 K, and 369 ± 85 evaluations of the reference potential, at 2700 K, to complete learning cycles (see Fig. 2).

For the purpose of this work, the quality of the ML-MD potential is evaluated on configurations generated with ARTn as local activated events associated with a vacancy in a crystalline environment are generated. To avoid non-physical results, when an ARTn-generated configuration shows a $\gamma > 200$, the configuration is rejected, the event search is stopped, and a new event search is launched from the same initial minimum.

Figure 3 shows the standard validation error on energy and forces calculated over *all* configurations generated along pathways for the 36 000 successful events and the 10 080 failed saddle searches (a success rate of 78%). The error on energy increases almost exponentially with the sampling temperature, ranging from 0.44 ± 0.36 meV/atom at 300 K to 5.1 ± 1.7 meV/atom at 2700 K. The error on forces is essentially constant at 0.0123 eV/Å, on average, between 300 and 1800 K, and increases rapidly at high temperatures, reaching 0.0256 eV/Å at 2700 K.

Since the focus of this work is on transition states, Fig. 4 displays the error on the energy barriers as a function of MD-fitting temperature, computed with Eq. (12) and averaged over all generated barriers. This error is relatively uncorrelated with the MD temperature simulation with an average of 0.056 ± 0.022 eV, a minimum error of 0.024 ± 0.01 eV at 2400 K, and a maximum of 0.08 ± 0.03 eV at 1200 K. This error is lower than that for a general point on the energy landscape (Fig. 3) in part because it is computed as a difference between the saddle and initial minimum.

Errors in the position of the saddle point, associated with the capacity to reproduce correctly their geometry, are given in Fig. 5. The top panel indicates the average distance between saddle points




FIG. 3. Average energy (top) and mean absolute forces (bottom) errors per atom for Si measured over all configurations generated along pathways in ARTn for the three approaches. Temperature refers to the one used during MD training. Vertical bars represent the standard deviation computed on ten independent realizations.

converged with the reference and the ML potentials: it decreases from 0.16 ± 0.05 Å at 300 K to a minimum of 0.09 ± 0.02 Å between 1500 and 2100 K, going up at the two highest temperatures (2400 and 2700 K).

Overall, this straightforward fitting approach based on constant-temperature MD runs provides accurate diffusion barriers, ranging from 0.51 to more than 4 eV, for a vacancy in crystalline silicon at low computational costs (263 to 369 evaluations of the reference potential).

B. Revisiting ML-MD potential in ARTn: The OTF-MDART adjusting approach

To evaluate the possibility of improving ML-MD potentials for activated events, potentials are on-the-fly re-trained during ARTn learning cycles (OTF-MDART). Figure 2 gives the number of calls to the reference potential for this procedure during the ARTn runs (dashed orange line) as well as the total number of calls, including




FIG. 4. Average energy barrier error for Si as defined by Eq. (14) for all events generated in ARTn for the three approaches. Temperature refers to the one used during MD training. Vertical bars represent the standard deviation computed on ten independent realizations.

those made during ML-MD fitting (solid orange line). The number of calls during ARTn learning cycles ranges from 979 ± 153 at 300 K to 136 ± 38 at 2700 K for a total of 1232 ± 177 to 505 ± 109 , respectively, when including ML-MD calls.

The error in energy and forces remains correlated with the ML-MD temperature; it is higher when the error is higher at the ML-MD trained temperature. This correlation is particularly strong when retraining MD potentials fitted between 1500 and 2700 K (Fig. 3, solid orange line). The error on energy for OTF-MDART is almost constant between 300 and 2400 K, at 0.22 meV/atom, rising to 1.9 meV/atom at 2700 K, lower by 50% to 63% than ML-MD. A similar improvement is observed in the forces, which range from




FIG. 5. Mean position error on all saddle points for Si. Temperature refers to the one used during MD training. Vertical bars represent the standard deviation computed on ten independent realizations.

TABLE II. Average energy barrier error and mean position error on all saddle points for Si. The average error for ML-MD and OTF-MDART training is taken over all temperature sets. The standard deviation computed for all temperature sets.

Errors	ML-MD	OTF-MDART	OTF-ART
$\Delta\delta E_{\text{barrier}}$ (eV)	0.056 ± 0.022	0.040 ± 0.012	0.039 ± 0.008
ΔX_{conv} (Å)	0.114 ± 0.029	0.072 ± 0.010	0.072 ± 0.006

0.0103 eV/Å, on average, between 300 and 1800 K, increasing to 0.0173 eV/Å at 2700 K, representing a 16% to 32% decrease in error.

Between 300 and 1500 K, retrained potentials with OTF-MDART show more constant energy barrier errors than pure ML-MD models (Fig. 4), with an error of about 0.036 eV (OTF-MDART) vs an average of 0.064 eV (ML-MD), a 44% improvement. At the highest temperature, 1800 to 2700 K, however, as OTF-MDART calls for less learning cycles, errors and fluctuations are not reduced with respect to ML-MD. Interestingly, though, improvements in the saddle position are observed at all temperatures for OTF-MDART (Fig. 5) with an average error of 0.072 ± 0.010 Å.

Overall, by retraining ML-MD potential in ARTn, errors are reduced and results are more consistent, i.e., error distributions are narrower, irrespective of the temperature used in the initial MD training. This additional retraining leads to a 50% to 96% decrease in energy error (Fig. 3), a 29% improvement in average energy barrier errors (Table II), and a 37% reduction in mean saddle position errors, but with an additional number of calls to the reference potential increasing between 37% and 490%.

These results can be understood by looking at the fraction of MD-generated configurations that remain in the training set at the end of the simulation (Fig. 6). At temperatures between 300 and 1200 K, none of the ML-MD configurations remain in the final training set; this proportion goes from 1.3% to 38% between 1500 and 2700 K (left-hand axis, blue line). At these temperatures, the




FIG. 6. Fraction of original MD configurations (left scale) and total number of MD configurations (right scale) remaining in the final training set (TS) for Si. Temperature refers to the one used during MD training. Vertical bars represent the standard deviation computed on ten independent realizations.

system melts and generates a wider range of configurations. Since these configurations are far from ARTn-generated configurations, the selection algorithm keeps them in the set even though they do not help reduce errors for the configurational space of interest with ARTn.

C. The OTF-ART adjusting approach

Given the results for OTF-MDART, we now turn to an OTF approach entirely integrated into ARTn in an attempt to increase accuracy and reduce the cost and waste of evaluations of the reference potential.

Ten independent on-the-fly ML potentials are generated entirely in ARTn for a total of 36 000 events starting from the same initial minimum. Each potential is trained initially from the same configuration (the initial minimum) in the training set. Each parallel event search goes through a learning cycle if needed, and as the simulation progresses, the learning cycles become rarer. The values are averaged over the ten simulations and as they go through learning.

With an average total of 628 ± 283 reference potential evaluations, the cost of the OTF-ART is between that of the ML-MD and the OTF-MDART. Along pathways, the average energy error for these potentials is 0.22 ± 0.03 meV/atom, on par with the OTF-MDART potential based on low-temperature ML-MD fitting and 49% lower than the 300 K ML-MD potential. Errors on forces, at 0.011 ± 0.001 eV/Å, are in between ML-MD (0.012 eV/Å) and OTF-MDART (0.010 eV/Å) at low training temperatures. Comparing with the 2700 K potential fitting in MD, the OTF-ART error is 57% lower than ML-MD (0.026 eV/Å) and 36% lower than OTF-MDART (0.017 eV/Å).

Focusing on barrier energy, the average error is 0.039 ± 0.008 eV (see Fig. 4), about 2.5% lower than OTF-MDART and 30.3% better than ML-MD. The error of 0.072 ± 0.006 Å on the converged saddle position is similar to the 0.072 ± 0.010 Å obtained with OTF-MDART and 37% lower than with ML-MD (0.114 Å).

D. Reproducing the dominant diffusion mechanism

The exploration of the energy landscape around the vacancy leads to the generation of a wide range of activated mechanisms and associated barriers (both forward, associated with the diffusion of the vacancy, and backward, from the final minima back to the saddle point). Figure 7 presents the complete distribution of generated direct and inverse barriers connected to the ground state. The peak near 0 eV (around 10^{-2} to 10^{-1} eV) is associated with the inverse barrier to the direct saddle at 2.38, 2.70 eV, and higher (up to 5.5 eV), except for the inverse 0.45 eV barrier, which is linked to the 2.87 eV direct barrier. Direct barriers at 0.51 eV represent symmetric first neighbor vacancy diffusion, while barriers at 2.38 and 2.70 eV are associated with more complex vacancy diffusion mechanisms.³³ Events with barriers at 2.38 and 2.70 eV, for example, involve vacancy diffusion through complex bond-exchanges. Spectator events³⁴ where the diamond network around the vacancy is transformed by bond switching are also generated. This mechanism was proposed by Wooten, Winer, and Weaire (WWW) to describe the amorphization of silicon.³⁵ The main spectator event




FIG. 7. ARTn-generated energy barrier distributions for vacancy-diffusion events in Si, including direct barriers (from the ground state) and inverse barriers (from excited states), as generated with the MTP model (orange) and re-converged using the reference model (blue) from saddles and minima positions originally found with the MTP model.

occurs as two neighbors of the vacancy are pushed together, allowing the creation of a bond associated with the 2.87 eV barrier. Other mechanisms involve strong lattice distortion and bond formation not involving direct neighbors of the vacancy with very high energy barriers³³ of between 3.2 and 4.0 eV.

Since vacancy diffusion for this system is dominated by a 0.51 eV single barrier mechanism, with the next barrier at 2.35 eV, an accurate description of the dominant mechanism is essential to correctly capture defect kinetics in Si. Table III presents the error on this barrier for the three approaches described earlier. With an error of 0.019 ± 0.005 eV and a relative error of 3.7%, OTF-ART offers the closest reproduction of the reference barrier, followed by OTF-MDART and ML-MD, with respective errors of 0.022 ± 0.011 (relative error of 4.3%) and 0.026 ± 0.015 (5.1%). Overall, the error on the energy barrier is lower than that on the total energy presented earlier (0.046 ± 0.006 eV for OTF-ART, for example) due to a partial error cancellation associated with the energy difference taken to measure the barrier.

The validity of the barrier is also measured by the precision of the saddle geometry. For the 0.51 eV barrier, ML-MD converges with an error on the position of 0.088 ± 0.036 Å, with OTF-MDART and OTF-ART giving errors almost 50% lower, at 0.040 ± 0.017 and 0.047 ± 0.018 Å, respectively.

TABLE III. Average energy barrier errors and mean saddle position errors on the 0.51 eV vacancy diffusion for Si. The average error for ML-MD and OTF-MDART training is taken over all temperature sets. The standard deviation computed for all temperature sets.

Errors	ML-MD	OTF-MDART	OTF-ART
$\Delta\delta E_{\text{barrier}}$ (eV)	0.026 ± 0.015	0.022 ± 0.011	0.019 ± 0.005
ΔX_{conv} (Å)	0.088 ± 0.036	0.040 ± 0.017	0.047 ± 0.018




FIG. 8. SiGe barrier histogram, including direct barriers (from states accepted) and inverse barriers (from excited states), as found on-the-fly by the MTP model (orange) and re-converged by the reference model (blue) from saddles and minima positions originally given by MTP.

E. SiGe system

Having shown an interest in developing a specific potential by applying on-the-fly learning directly to activated events on a simple system such as *c*-Si with a vacancy, we tested this approach with a more complex alloy with the same overall reference potential to facilitate comparison. Starting from an ordered zincblende structure, the diffusion of a vacancy creates a chemical disorder that complexifies the landscape visited, as shown by the continuous distribution of activated barriers, including both direct and inverse barriers, found as the vacancy diffuses (Fig. 8). We note that the lowest barrier for a vacancy diffusing is around 0.6 eV, with lower barriers associated, as for Si, with reverse jumps from metastable states. The energy barrier distribution for vacancy diffusing in SiGe (Fig. 8) is much more complex than for Si due to the chemical disorder that builds as the vacancy diffuses.

As stated in the methodology, the additional complexity of the system imposes a richer machine-learning potential with a larger set of parameters to encompass the greater diversity in the components and the configurations due to chemical disorder. Combined, these two levels of complexity (set of parameters and configurational) result in an overall higher number of calls to the reference potential as compared to Si, irrespective of the approach used [see Fig. 9 (SiGe) vs Fig. 2 (Si)]: while ML-MD requires between 380 evaluations of the reference potential at 300 K and 1549 at 2700 K, OTF-MDART needs a total of around 3465 calculations of the reference potential, irrespective of the temperature, as original ML-MD configurations are progressively removed from the training set. This effort results in a number of calls to the reference potential for OTF-MDART that are 4% higher than with OTF-ART (3329 on average). To reduce computational costs, we omit the 1500 K run, as statistical behavior is smooth in this temperature region.

To disentangle the two contributions, we compare the cost of fitting a Si potential with the same level 20 potential as used for SiGe. Following the full OTF-ART procedure, creating a Si MLP requires 2926 calls to the reference potential. The intrinsic complexity of the landscape contributes, therefore, to about a 14% increase in the Si




FIG. 9. Number of calls to the reference potential for each of the OTF machine-learned potentials developed for SiGe as a function of the temperature referring to the one used during MD training. Since the configurations are relaxed to zero K in ARTn simulations, there is no associated temperature for this procedure. Vertical bars represent the standard deviation computed on ten independent realizations.

TABLE IV. Average energy barrier errors and mean saddle position errors on all barriers for SiGe. The average error for ML-MD and OTF-MDART training is taken over all temperature sets. The standard deviation computed for all temperature sets.

Errors	ML-MD	OTF-MDART	OTF-ART
$\Delta\delta E_{\text{barrier}}$ (eV)	0.082 ± 0.024	0.072 ± 0.014	0.066 ± 0.015
ΔX_{conv} (Å)	0.091 ± 0.020	0.076 ± 0.013	0.070 ± 0.014

baseline call count. In terms of accuracy, the Si MLP level 20 leads to an average error on the energy of 0.1 meV/atom, about 50% lower than with the level 16 potential described earlier (0.22 meV/atom). For SiGe, this error is (0.42 meV/atom), two times higher than for Si MLP level 16 and four times that of Si MLP level 20.

This can be understood by the number of different configurations visited: as opposed to the Si system, where each initial minimum is identical (as the vacancy moves in an otherwise perfect elemental crystal), the binary system is transformed as the vacancy diffuses and the chemical order is slowly destroyed: each of the 24 ARTn parallel trajectories used to define the potential over 1500 events evolves independently according to a probability given by the Metropolis algorithm with a fictitious temperature (since the network itself is structurally at 0 K) of 0.5 eV [Eq. (9)], providing a rich range of local environments.

Fitting a potential is clearly harder: with the parameters used—when a configuration graded at $\gamma > 200$ is encountered, the ARTn event search is stopped—no event could be generated using the ML-MD potential at 300 and 600 K, which explains the absence of data for these temperatures in Fig. 10 and Table IV. For SiGe, the error on energy (see Fig. 10) with the ML-MD at 900 K and the above-mentioned ranges from 0.5 to 1.4 meV/atom as a function of temperature. On average, these errors are between 14% and 69% lower with OTF-MDART or OTF-ART at around 0.43 meV/atom.




FIG. 10. Average energy (top) and mean absolute forces (bottom) errors for SiGe measured over all configurations generated along pathways in ARTn for the three approaches. Temperature refers to the one used during MD training. Vertical bars represent the standard deviation computed on ten independent realizations.

The OTF-ART approach gives an error in energy barrier of 0.066 ± 0.015 eV, which represents a 19.5% and 8.3% lower error than the ML-MD (0.082 ± 0.024 eV) and OTF-MDART (0.072 ± 0.014 eV), respectively (Table IV). The errors on the converged saddle position for OTF-ART and OTF-MDART are similar at 0.070 ± 0.014 and 0.076 ± 0.013 Å, respectively, and represent a 23% lower error than with ML-MD (0.092 Å). This accuracy is similar to that obtained with Si, in contrast to total energy and energy barrier errors.

We note that the advantage of ML-MD for SiGe is overstated, as shown by the proportion of events generated with ML-MD potential that are interrupted due to a too large extrapolation grade, $\gamma > 200$ for both SiGe and Si (Fig. 11): for SiGe, between 85% and 30% of events are aborted between 300 and 1200 K, respectively. This proportion falls to zero percent failure at 1800 K.




FIG. 11. Percentage of search interruptions during ML-MD potential evaluation in ARTn ($\gamma > 200$) for Si and SiGe as a function of ML-MD training temperature. Vertical bars represent the standard deviation computed on ten independent realizations.

IV. DISCUSSION

We compare three approaches aimed at the construction of potentials with machine learning on-the-fly for the exploration of activated mechanisms in the potential energy landscape. We evaluate these by computing their efficiency at reproducing the energy landscape around a vacancy in two systems: a relatively simple Si diamond system (Fig. 7) and a more complex SiGe zincblende system that disorders under vacancy diffusion (Fig. 8), both described with a reference empirical potential to allow for better statistical analysis.

The first approach, which sets the comparison level, constructs a more general machine learning potential with molecular dynamics (ML-MD); the second on-the-fly adjusts this generated potential during the search for activated events using ARTn; and the third approach constructs a specifically on-the-fly trained potential during the search for activated events (OTF-ART). The efficiency of these three procedures is measured by the quality of the reproduction of the reference potential during the search for activated events.

The baseline, defined by the ML-MD, is competitive with previously published work. Energy errors for the more standard ML MD approach with a level 16 potential range from 0.44 ± 0.36 meV/atom at 300 K to 5.1 ± 1.7 meV/atom at 2700 K (Fig. 3), an order of magnitude lower or similar to the 4 meV/atom on an MTP potential of level 24 for Si obtained by Zuo *et al.*,²⁵ with the difference explained by the fact that activated events involve local deformations from a zero-temperature crystal with a vacancy and that DFT potentials are more difficult to fit than empirical ones.²¹ That is, using DFT, we would expect that these conclusions hold as we are looking at the differences between our three approaches, not claiming specific accuracy of these.

Similarly, the relative energy error on the dominant 0.51 eV diffusion barrier for SW Si is 5.1% (0.026 eV) with the ML-MD approach and 3.7% (0.019 eV) with the OTF-ART. Using the same MTP potential trained using an OTF MD with an *ab initio* reference

potential, Novoselov *et al.* find a 0.20 eV barrier for vacancy diffusion in Si as compared with 0.18 eV with the reference potential, an error of 0.02 eV or a 10.0% relative error.

Overall, the ML-MD approach, especially when run at temperatures between 900 and 1800 K, can generate a generic ML potential with reasonable precision for describing activated mechanisms in Si and SiGe. Developing a more specific OTF potential, generated directly with ARTn on activated trajectories, however, offers a more accurate description of both the energy and geometry at the barriers.

It is possible to recover this precision by adjusting the original MD potential during ARTn runs; however, this increases the number of calls to the reference potential, raising the total costs beyond that of OTF-ART while largely erasing work made during the ML-MD training phase: for Si, between 300 and 1200 K, none of the ML-MD configurations are retained, while around 1.3% to 12.5% are retained for the potential trained in the range of 1500 and 2700 K (Fig. 6, right-hand axis, orange line), but at the cost of lowering the precision on barriers.

Moving to a more complex system, such as an evolving binary alloy, increases the overall cost of the procedure in terms of calls to the reference potential as more parameters need to be fit. Here also, the gain from using a specific potential constructed from ARTn trajectories is notable, both in the average errors and their fluctuations. Indeed, the ML-MD potential presents considerable instabilities while generating activated trajectories, as can be seen by the number of configurations considered outside of the potential's scope ($\gamma > 200$), see Fig. 11. This can also be seen by comparing the number of evaluations of the reference potential as a function of the temperature for OTF-MDART; for Si, this number decreases while, for SiGe, the total number of evaluations of the reference potential remains constant. This means that the potential requires significant re-adjustment, even from its high-temperature training, to adapt to the created disorder of the evolving SiGe system.

V. CONCLUSION

We compare the advantages of using a more general vs specific machine-learned potential (MLP) to describe activated mechanisms in solids. To do so, we first generate an MLP constructed with the Moment Tensor Potential formalism^{10,21} to replicate Stillinger–Weber potential for Si and SiGe crystals with a single vacancy using a standard molecular dynamics procedure (MD–ML).

Comparing the quality of the reproduction of activated mechanisms with an ML potential further refined during an activation-relaxation technique nouveau sampling of the energy landscape and a unique potential constructed on-the-fly within ARTn, we show that while a general potential can deliver high accuracy for both the barrier geometries and their related energies, error and fluctuations around the average value are significantly lowered by constructing a specific potential, with a number of calls to the reference potential that is lower than a combined approach (MD + ARTn) for a similar precision.

The advantage of using a specific potential remains when looking at more complex materials, such as the SiGe alloys considered here, even though the advantage in terms of calls to the reference is strongly reduced. The next steps will involve applying this strategy with DFT as a reference potential to attack problems that have

long been out of reach of computational materials sciences, allowing a much closer connection between modeling and experience.

SUPPLEMENTARY MATERIAL

See the supplementary material for a figure of the relation between accuracy and the level of the MTP potential for Si. In addition, there are figures displaying the accuracy of the saddle position for the 0.51 eV barrier in Si and all barriers in SiGe.

ACKNOWLEDGMENTS

This project is supported through a Discovery Grant from the Natural Science and Engineering Research Council of Canada (NSERC). Karl-Étienne Bolduc is grateful to NSERC and IVADO for summer scholarships. We are grateful to Calcul Québec and Compute Canada for their generous allocation of computational resources.

AUTHOR DECLARATIONS

Conflict of Interest

The authors have no conflicts to disclose.

Author Contributions

Eugène Sanscartier: Conceptualization (equal); Data curation (equal); Formal analysis (equal); Investigation (equal); Methodology (equal); Software (equal); Validation (equal); Visualization (equal); Writing – original draft (equal); Writing – review & editing (equal). **Félix Saint-Denis:** Conceptualization (supporting); Data curation (supporting); Investigation (supporting); Methodology (supporting). **Karl-Étienne Bolduc:** Conceptualization (supporting); Data curation (supporting); Investigation (supporting); Methodology (supporting). **Normand Mousseau:** Conceptualization (equal); Formal analysis (equal); Funding acquisition (equal); Methodology (equal); Project administration (equal); Resources (equal); Software (equal); Supervision (equal); Validation (equal); Writing – original draft (equal); Writing – review & editing (equal).

DATA AVAILABILITY

The ARTn packages as well as the data reported here are distributed freely. Please contact Normand Mousseau (normand.mousseau@umontreal.ca).

REFERENCES

- W. Kohn and L. J. Sham, “Self-consistent equations including exchange and correlation effects,” *Phys. Rev.* **140**, A1133–A1138 (1965).
- E. R. Lindahl, “Molecular dynamics simulations,” in *Molecular Modeling of Proteins* (Springer, 2008), pp. 3–23.
- A. F. Voter and J. D. Doll, “Transition state theory description of surface self-diffusion: Comparison with classical trajectory results,” *J. Chem. Phys.* **80**, 5832–5838 (1984).
- A. F. Voter, “Introduction to the kinetic Monte Carlo method,” in *Radiation Effects in Solids* (Springer, 2007), pp. 1–23.

- ⁵G. Henkelman and H. Jónsson, “Long time scale kinetic Monte Carlo simulations without lattice approximation and predefined event table,” *J. Chem. Phys.* **115**, 9657–9666 (2001).
- ⁶F. El-Mellouhi, N. Mousseau, and L. J. Lewis, “Kinetic activation–relaxation technique: An off-lattice self-learning kinetic Monte Carlo algorithm,” *Phys. Rev. B* **78**, 153202 (2008).
- ⁷J. Behler and M. Parrinello, “Generalized neural-network representation of high-dimensional potential-energy surfaces,” *Phys. Rev. Lett.* **98**, 146401 (2007).
- ⁸A. P. Bartók, R. Kondor, and G. Csányi, “On representing chemical environments,” *Phys. Rev. B* **87**, 184115 (2013).
- ⁹A. P. Thompson, L. P. Swiler, C. R. Trott, S. M. Foiles, and G. J. Tucker, “Spectral neighbor analysis method for automated generation of quantum-accurate interatomic potentials,” *J. Comput. Phys.* **285**, 316–330 (2015).
- ¹⁰A. V. Shapeev, “Moment tensor potentials: A class of systematically improvable interatomic potentials,” *Multiscale Model. Simul.* **14**, 1153–1173 (2016).
- ¹¹G. Sivaraman, J. Guo, L. Ward, N. Hoyt, M. Williamson, I. Foster, C. Benmore, and N. Jackson, “Automated development of molten salt machine learning potentials: Application to LiCl,” *J. Phys. Chem. Lett.* **12**, 4278–4285 (2021).
- ¹²P.-L. Kang, C. Shang, and Z.-P. Liu, “Large-scale atomic simulation via machine learning potentials constructed by global potential energy surface exploration,” *Acc. Chem. Res.* **53**, 2119–2129 (2020).
- ¹³G. Sivaraman, A. N. Krishnamoorthy, M. Baur, C. Holm, M. Stan, G. Csányi, C. Benmore, and Á. Vázquez-Mayagoitia, “Machine-learned interatomic potentials by active learning: Amorphous and liquid hafnium dioxide,” *npj Comput. Mater.* **6**, 104 (2020).
- ¹⁴L. K. Béland, Y. Anahory, D. Smeets, M. Guihard, P. Brommer, J.-F. Joly, J.-C. Pothier, L. J. Lewis, N. Mousseau, and F. Schiettekatte, “Replenish and relax: Explaining logarithmic annealing in ion-implanted *c*-Si,” *Phys. Rev. Lett.* **111**, 105502 (2013).
- ¹⁵Y. N. Osetsky, L. K. Béland, A. V. Barashev, and Y. Zhang, “On the existence and origin of sluggish diffusion in chemically disordered concentrated alloys,” *Curr. Opin. Solid State Mater. Sci.* **22**, 65–74 (2018).
- ¹⁶O. A. Restrepo, N. Mousseau, M. Trochet, F. El-Mellouhi, O. Bouhali, and C. S. Becquart, “Carbon diffusion paths and segregation at high-angle tilt grain boundaries in α -Fe studied by using a kinetic activation–relaxation technique,” *Phys. Rev. B* **97**, 054309 (2018).
- ¹⁷G. T. Barkema and N. Mousseau, “Event-based relaxation of continuous disordered systems,” *Phys. Rev. Lett.* **77**, 4358–4361 (1996).
- ¹⁸R. Malek and N. Mousseau, “Dynamics of Lennard-Jones clusters: A characterization of the activation–relaxation technique,” *Phys. Rev. E* **62**, 7723–7728 (2000).
- ¹⁹A. Jay, M. Gunde, N. Salles, M. Poberžnik, L. Martin-Samos, N. Richard, S. d. Gironcoli, N. Mousseau, and A. Hémerlyck, “Activation–relaxation technique: An efficient way to find minima and saddle points of potential energy surfaces,” *Comput. Mater. Sci.* **209**, 111363 (2022).
- ²⁰D. G. Truhlar, B. C. Garrett, and S. J. Klippenstein, “Current status of transition-state theory,” *J. Phys. Chem.* **100**, 12771–12800 (1996).
- ²¹I. S. Novikov, K. Gubaev, E. V. Podryabinkin, and A. V. Shapeev, “The MLIP package: Moment tensor potentials with MPI and active learning,” *Mach. Learn.: Sci. Technol.* **2**, 025002 (2020).
- ²²I. I. Novoselov, A. V. Yanilkin, A. V. Shapeev, and E. V. Podryabinkin, “Moment tensor potentials as a promising tool to study diffusion processes,” *Comput. Mater. Sci.* **164**, 46–56 (2019).
- ²³F. H. Stillinger and T. A. Weber, “Computer simulation of local order in condensed phases of silicon,” *Phys. Rev. B* **31**, 5262 (1985).
- ²⁴S. Ethier and L. J. Lewis, “Epitaxial growth of $\text{Si}_{1-x}\text{Ge}_x$ on $\text{Si}(100)2 \times 1$: A molecular-dynamics study,” *J. Mater. Res.* **7**, 2817–2827 (1992).
- ²⁵Y. Zuo, C. Chen, X. Li, Z. Deng, Y. Chen, J. Behler, G. Csányi, A. V. Shapeev, A. P. Thompson, M. A. Wood *et al.*, “Performance and cost assessment of machine learning interatomic potentials,” *J. Phys. Chem. A* **124**, 731–745 (2020).
- ²⁶E. V. Podryabinkin and A. V. Shapeev, “Active learning of linearly parametrized interatomic potentials,” *Comput. Mater. Sci.* **140**, 171–180 (2017).
- ²⁷E. V. Podryabinkin, E. V. Tikhonov, A. V. Shapeev, and A. R. Oganov, “Accelerating crystal structure prediction by machine-learning interatomic potentials with active learning,” *Phys. Rev. B* **99**, 064114 (2019).
- ²⁸K. Gubaev, E. V. Podryabinkin, G. L. W. Hart, and A. V. Shapeev, “Accelerating high-throughput searches for new alloys with active learning of interatomic potentials,” *Comput. Mater. Sci.* **156**, 148–156 (2019).
- ²⁹A. P. Thompson, H. M. Aktulga, R. Berger, D. S. Bolintineanu, W. M. Brown, P. S. Crozier, P. J. ‘t Veld, A. Kohlmeyer, S. G. Moore, T. D. Nguyen, R. Shan, M. J. Stevens, J. Tranchida, C. Trott, and S. J. Plimpton, “LAMMPS—A flexible simulation tool for particle-based materials modeling at the atomic, meso, and continuum scales,” *Comput. Phys. Commun.* **271**, 108171 (2022).
- ³⁰C. Lanczos, “An iteration method for the solution of the eigenvalue problem of linear differential and integral operators,” *J. Res. Natl. Bur. Stand., Sect. B* **45**, 255 (1950).
- ³¹E. Bitzek, P. Koskinen, F. Gähler, M. Moseler, and P. Gumbsch, “Structural relaxation made simple,” *Phys. Rev. Lett.* **97**, 170201 (2006).
- ³²N. Mousseau and G. T. Barkema, “Exploring high-dimensional energy landscapes,” *Comput. Sci. Eng.* **1**, 74–82 (1999).
- ³³F. El-Mellouhi, N. Mousseau, and P. Ordejón, “Sampling the diffusion paths of a neutral vacancy in silicon with quantum mechanical calculations,” *Phys. Rev. B* **70**, 205202 (2004).
- ³⁴Y. Kumeda, D. J. Wales, and L. J. Munro, “Transition states and rearrangement mechanisms from hybrid eigenvector-following and density functional theory: Application to $\text{C}_{10}\text{H}_{10}$ and defect migration in crystalline silicon,” *Chem. Phys. Lett.* **341**, 185–194 (2001).
- ³⁵F. Wooten, K. Winer, and D. Weaire, “Computer generation of structural models of amorphous Si and Ge,” *Phys. Rev. Lett.* **54**, 1392–1395 (1985).