

Machine learning interatomic potentials for solid-state precipitationLorenzo Piersante¹ and Anirudh Raju Natarajan^{1,2,*}¹Laboratory of Materials Design and Simulation (MADES), Institute of Materials,
École Polytechnique Fédérale de Lausanne, Lausanne, Switzerland²National Centre for Computational Design and Discovery of Novel Materials (MARVEL),
École Polytechnique Fédérale de Lausanne, Lausanne, Switzerland

(Received 19 January 2026; accepted 21 July 2026; published 4 September 2026)

Machine learning interatomic potentials (MLIPs) are routinely used to model diverse atomistic phenomena, yet parameterizing them to accurately capture solid-state phase transformations remains difficult. We present error metrics and data-generation schemes designed to streamline the parameterization of MLIPs for modeling precipitation in multicomponent alloys. We developed an algorithm that enumerates symmetrically distinct transformation pathways connecting chemical decorations of different parent crystal structures. Additionally, we introduce the weighted Kendall- τ coefficient and its semigrand-canonical generalization as metrics for quantifying MLIP accuracy in predicting low-temperature thermodynamics. We apply these approaches to parameterize an MLIP for a dilute Mg–Nd alloy. The resulting potential reproduces the complex early stage precipitation behavior observed experimentally. Large-scale atomistic simulations reveal competition between order-disorder and structural transformations. Furthermore, these results suggest a continuous transition between high-symmetry hcp and bcc crystal structures during aging heat treatments.

DOI: [10.1103/physrevmaterials.10.093802](https://doi.org/10.1103/physrevmaterials.10.093802)**I. INTRODUCTION**

Precipitation or age hardening is the primary strengthening mechanism in many commercial engineering alloys. Age hardening relies on the formation of secondary phases embedded within a disordered matrix phase. These precipitates increase material strength by impeding dislocation motion. To achieve peak strength, age-hardening heat treatments must be tuned to optimize precipitate size, shape, and number density. Insufficient aging times produce small precipitates that dislocations can easily overcome. Conversely, overaging leads to precipitate coarsening, which reduces their number density and decreases the strengthening effect. Optimizing the aging process requires quantitative knowledge of thermodynamic and kinetic parameters, such as free energies and diffusion constants, as well as an understanding of nucleation mechanisms and rates within the matrix phase. While several studies have focused on developing techniques and models to estimate the thermodynamic and kinetic parameters of materials [1–3], less is known about the nucleation mechanisms of precipitates.

Capturing the atomistic aspects of nucleation is important for modeling the early stages of precipitation. As the time and length scales required to study nucleation are inaccessible

to purely *ab initio* methods, machine learning interatomic potentials (MLIPs), trained on density functional theory (DFT) data, have emerged as reliable and computationally efficient alternatives. While the capabilities of MLIPs have advanced substantially in recent years, designing datasets and validation methods for specific atomistic processes is still challenging [4,5].

Precipitation in multicomponent alloys often proceeds through a series of phase transitions. The initial stage is typically an order-disorder transition, where precipitates form by the chemical ordering of elements on the same parent crystal structure as that of the disordered matrix phase. Longer aging promotes the formation of more stable precipitates, which can have crystal structures that are different from the matrix. This precipitation sequence has been proposed to arise from the lower nucleation barriers for precipitates formed through order-disorder phase transitions than for those formed through structural transformations. However, rigorous estimation of the barriers and driving forces for the competing transformations has remained elusive.

Multistage precipitation is common in many engineering alloys, including those based on aluminum, titanium, nickel, and magnesium. Rare-earth-containing magnesium alloys are a prototypical example, where experiments reveal a complex precipitation sequence involving ordering and structural transformations [6–11]. For instance, in dilute Mg–Nd alloys with $x_{\text{Nd}} \leq 0.5$ at.%, the decomposition of a supersaturated solid solution (SSSS) results in several intermediate phases during aging. The precipitation process begins with the formation of Guinier–Preston (GP) zones, followed by the β''' phase, which consists of rods of Nd atoms arranged in a zigzag pattern along the hcp [01 $\bar{1}$ 0] direction interwoven with hexagonal chains of

*Contact author: anirudh.natarajan@epfl.ch




FIG. 1. (a) Structural transformation from the ideal hcp decoration $h\text{-}\beta_1$ to the bcc structure β_1 . (b) Two possible sublattice decorations of a preexisting β' ordering: Nd substitution on the red sublattice yields β_1 , while substitution on the blue sublattice creates a local β'' ordering. Crystal structures are projected along the hcp [0001] direction. Orange and white circles represent Nd and Mg atoms, respectively. Atoms positioned between lattice sites reside in the layer above the basal plane. (c) DFT energy landscape along the transformation path between $h\text{-}\beta_1$ and β_1 .

solute atoms. The β''' precipitates comprise a family of ordered phases, with β' occurring at a composition of $x_{\text{Nd}} = 1/8$ and β'' at a composition of $x_{\text{Nd}} = 1/4$. The next precipitate, β_1 , requires a structural transformation from hcp to bcc. Finally, more complex structural transitions result in the formation of a metastable β and equilibrium β_e phases. Peak strength is achieved through a fine distribution of β''' and β_1 precipitates embedded within the disordered matrix [6,10].

The β_1 precipitate, one of the primary sources of strength in this alloy, forms through a structural phase transition from hcp to bcc. This hcp $\rightarrow$ bcc transformation occurs through the Burgers mechanism, which involves both a homogeneous deformation of the hcp lattice and a shuffle of atomic planes [12]. As illustrated in Fig. 1(a), the basal plane of hcp is homogeneously strained to form the {110} plane of bcc. Adjacent basal planes must then shuffle along the $\langle 01\bar{1}0 \rangle$ direction to attain the ideal bcc structure. The formation of the β_1 precipitate through the Burgers mechanism requires a specific arrangement of Mg and Nd atoms on the parent hcp structure, as shown in Fig. 1(b). DFT calculations in Fig. 1(c) predict a barrierless transformation from the ordered hcp precursor ($h\text{-}\beta_1$) to the stable bcc precipitate (β_1). The existence of a barrierless pathway indicates the possibility of a continuous homogeneous nucleation mechanism for the formation of β_1 .

A rigorous investigation of nucleation in this alloy requires an interatomic potential that accurately captures the hcp $\rightarrow$ bcc transformation, thermodynamics, defect energetics, and kinetic quantities such as vacancy migration barriers. An MLIP that faithfully models precipitation must predict the low-temperature phase stability. Models that yield zero-Kelvin ground states inconsistent with DFT or experiment, such as the modified embedded-atom-method potential for Mg–Nd by Kim *et al.* [13] (Sec. SII of the Supplemental

Material, SM [14]), have limited utility for studying precipitation. To capture multistage precipitation reactions, an MLIP must also describe both equilibrium and metastable thermodynamics. It must also reproduce salient features of structural transformations, such as the energy barriers that separate ordered phases on different parent lattices. Finally, estimating the effects of defects such as grain boundaries or dislocations requires training datasets with large structural and chemical diversity.

This article presents a strategy for parameterizing an MLIP capable of reliably modeling precipitation in the Mg–Nd alloy. We begin in Sec. II by introducing the MLIP formalism and summarizing the first-principles calculations. The following three sections present systematic methodologies for data generation, training dataset sampling, and validation, specifically tailored to assess phase stability and transformations. Section III describes the construction of a reference database of *ab initio* calculations and introduces an algorithm based on crystal symmetry to systematically map the potential energy surface of the hcp-to-bcc phase transition. Section IV outlines a procedure for sampling a training dataset from the reference database of DFT calculations, together with an iterative data-generation scheme that ensures all ground states predicted by the interatomic potential are included in the training set. Section V introduces post hoc validation metrics that go beyond standard statistical error measures. In particular, we propose the weighted Kendall- τ coefficient as an error metric for scoring phase stability in multicomponent systems. We then demonstrate in Sec. VI that these methodologies collectively enable the parameterization of an MLIP that accurately describes the atomistic behavior of the Mg–Nd system. Finally, Sec. VII applies the MLIP to model precipitation in dilute Mg–Nd alloys, revealing the atomistic nucleation mechanisms for both order-disorder and structural phase transformations.

II. INTERATOMIC POTENTIAL AND DFT CALCULATIONS

The MLIP for the Mg–Nd alloy was parameterized with the atomic cluster expansion (ACE) framework [15–17]. At present, the ACE represents a Pareto-optimal MLIP architecture with high computational efficiency [18] and linear scaling with the number of atoms, making it suitable for large-scale simulations. In the ACE formalism, the energy contribution of atom i is expressed as

$$E_i = (\rho_i^{(1)} + \sqrt{\rho_i^{(2)}}) \quad (1)$$

where $\rho_i^{(k)}$ is an atomic density expanded in a basis of cluster functions, B_{iv} , centered on site i . The atomic density is given by

$$\rho_i^{(k)} = \sum_v c_v^{(k)} B_{iv} \quad (2)$$

where $v = (\mu, n, l)$ is a multi-index that encodes the chemical species (μ), radial basis functions (n), and angular momenta (l) used to construct the rotationally invariant basis functions B_{iv} . The index v runs over all basis functions included in the ACE expansion up to the desired body order. The expansion

TABLE I. ACE basis set. The table records the maximum n and l values allowed in the vectors $\mathbf{n}$ and $\mathbf{l}$ in Eq. (2).

n/l	Pair	Triplet	Quadruplet	Quintuplet	Total
Mg	15/0	5/4	4/3	3/2	694
Nd	15/0	8/8	4/4	1/1	890
Mg-Nd	15/0	8/8	3/2		2662

coefficients, $c_v^{(k)}$, are the model fitting parameters determined during training.

The accuracy and speed of an MLIP are determined by the number of descriptors or basis functions used in its construction. Consequently, hyperparameters such as the cutoff radius, the largest many-body interaction, and the number of basis functions must be carefully chosen. As summarized in Table I, the Mg-Nd potential was constructed using 4246 cluster functions with an interaction distance of 7.8 Å. The basis set for pure Mg is identical to that used in [19]. A Ziegler-Biersack-Littmark (ZBL) potential was included to describe short-range ($r \lesssim 1-2$ Å) repulsive interactions between atoms, such as those encountered when computing dimer curves or during high-temperature molecular dynamics simulations. Section SIII of the SM [14] provides a comprehensive list of basis function cutoffs and other model hyperparameters.

Energies and forces for the structures in our database are computed with DFT as implemented in the Vienna *ab initio* Simulation Package (VASP, version 6.4.1) [20–23]. These calculations used the Perdew–Burke–Ernzerhof exchange-correlation functional within the generalized gradient approximation [24,25]. The projector-augmented-wave pseudopotentials treated the $2p^63s^2$ (Mg_pv 13Apr2007) and $5s^25p^64f^16s^2$ (Nd_3 06Sep2000) electrons as valence for Mg and Nd respectively. All calculations used a plane-wave cutoff of 500 eV, a reciprocal Γ -centered k-point mesh with a spacing of at most ≈ 0.02 Å⁻¹ for Brillouin-zone sampling and a second-order Methfessel-Paxton smearing of 0.2 eV. Formation energies of structures were computed relative to hcp-Mg and dhcp-Nd as

$$e_f = \frac{E(N, M) - Ne_{\text{hcp}} - Me_{\text{dhcp}}}{N + M} \quad (3)$$

where $E(N, M)$ is the total energy of the configuration, N and M are the numbers of Mg and Nd atoms, and e_{hcp} and e_{dhcp} are the per-atom energies of hcp-Mg, and dhcp-Nd, respectively.




FIG. 2. Structures in the reference database and their relationship to material properties.

III. CONSTRUCTING A REFERENCE DATABASE

Parameterizing an accurate interatomic potential requires a *reference database* encompassing the structural and chemical environments relevant to the properties of interest. A training dataset is sampled from this reference database to train an MLIP that reproduces the target properties accurately. We built our reference database from calculations relevant to the precipitation process in Mg–Nd alloys. Several classes of data were included following established protocols for constructing databases for multicomponent alloy potentials [26–29]. Figure 2 summarizes the types of structures in the database and their relationship to specific material properties. We describe additional strategies we developed to improve modeling accuracy for solid-state phase transformations in the Mg–Nd system.

The starting point for our reference database was a catalog of 267 structural prototypes of intermetallics, containing a range of diverse atomic environments from simple close-packed arrangements (e.g., fcc and hcp) to complex topologically close-packed phases [30]. This enabled us to sample a large degree of structural diversity using relatively small supercells. To span a wide range of densities, we included the volumetric distortions of several low-energy prototypes. Thermally disordered structures sampled from *ab initio* molecular dynamics simulations at several temperatures and densities provided finite-temperature information.

Accurately modeling precipitation in dilute alloys requires capturing the properties of the pure Mg matrix. The reference database included symmetry-distinct strain deformations of hcp-Mg [31]. It also contained a variety of structural defects, such as stable stacking faults (SSF), generalized stacking faults (GSF), grain boundaries (GB), vacancies, and surfaces. Some of the planar defects were sourced from the Crystallium database [32] and the work of Stricker *et al.* [29].

Constructing the reference database for the binary alloy was significantly more complex. Crystal structures of experimentally reported phases were obtained from the ICSD and OQMD databases [33,34]. Symmetrically distinct arrangements of the two elements were enumerated on the hcp, bcc, fcc, and C15 Laves parent structures using the CASM [35] code with the algorithm of Hart and Forcade [36,37]. We generated orderings in supercells containing up to 5, 8, 8, and 2 primitive unit cells for each parent structure. The atomic coordinates and lattice vectors of each resulting structure were then fully relaxed using DFT. The reference database also included structures that capture the mechanical response [31] of all stable and metastable phases in the binary alloy.

Chemically and structurally disordered configurations were created with high-temperature Monte Carlo molecular-dynamics simulations using the MEAM potential of Kim *et al.* [13,38]. The energies and forces for a representative selection of these disordered configurations were then recomputed with DFT for inclusion in the reference database. In addition, we applied normally distributed random displacements (with standard deviations matching the Debye mean-squared displacements of hcp-Mg or dhcp-Nd at representative temperatures) to supercell structures of all stable and metastable phases with $x_{\text{Nd}} \leq 1/2$. This approach conveniently incorporated composition-dependent finite-temperature vibrational effects into the database.




FIG. 3. (a) Crystal structures and symmetry-distinct transformation pathways of $h\text{-}\beta_1$ illustrated in the two-dimensional strain order parameter space (e_2, e_6) [12,31]. (b) DFT energy profiles along the three transformation paths.

Mg-rich compositions were oversampled, as the purpose of the MLIP is to model precipitation in Mg-rich alloys. A detailed description of the configurations and DFT data in the reference database is provided in Tables SII–SIV within the SM [14]. Additional orderings of Mg and Nd on the parent crystal structures were added to the reference database through an iterative scheme that identifies new stable and metastable ground states using preliminary MLIPs, as described in Sec. IV.

Enumerating transformation pathways

In addition to capturing the thermodynamic stability of stable and metastable phases, the MLIP must accurately reproduce the energy landscapes for structural phase transformations relevant to the precipitation process. The structural phase transitions connecting a pair of crystal structures can be enumerated in single-component materials through either domain knowledge or the automated algorithm of Thomas *et al.* [39]. The lower symmetry of an ordered phase compared to its undecorated parent structure complicates the computation of these energy landscapes in multicomponent alloys. Arranging two or more elements on a parent crystal structure breaks the symmetry of the parent structure, making transformation pathways that are equivalent in a pure element inequivalent in an ordered phase.

As an example, Fig. 3(a) illustrates the three possible variants of bcc that can be accessed through the Burgers transformation starting from $h\text{-}\beta_1$. Pathway 1 leads to the β_1 structure. Deforming the $h\text{-}\beta_1$ structure along pathway 2 results in an atomic arrangement on the bcc lattice distinct from the β_1 phase. The Nd atoms are arranged in pairs of columns distributed on the $(001)_{\text{bcc}}$ atomic planes. The bcc arrangement obtained from a deformation along pathway 3 is identical to that from pathway 2. These differences in the final atomic arrangements are reflected in their respective energy landscapes as shown in Fig. 3(b). For the transformation of $h\text{-}\beta_1$ to β_1 along pathway 1, the initial structure is unstable and transforms to the bcc phase without an energy barrier. In contrast, for pathways 2 and 3, the final bcc decorations are higher in energy than $h\text{-}\beta_1$, causing the energy to increase along these transformation paths. Parameterizing an MLIP

that reproduces the energy landscapes for structural phase transitions requires reference data containing transformations connecting distinct orderings on two or more parent crystal structures. The systematic generation of reference data for an MLIP must account for the underlying symmetry of the transformation pathway and structural decorations.

Any transformation pathway between two crystal structures can be decomposed into a homogeneous lattice distortion and atomic shuffles. An initial structure $\mathcal{P}_{\text{initial}}$ is defined by its lattice vectors $L_{\text{initial}} = [a, b, c]$, and the Cartesian coordinates of its n basis atoms, $C = [r_1, r_2, \dots, r_n]$. A supercell of the initial structure is generally required to describe the transformation pathway because of differing periodicities between the initial and final structures. This supercell lattice, $S_{\text{initial}} = L_{\text{initial}}T$, is computed using an integer transformation matrix T . The homogeneous distortion along the transformation pathway is described by a 3×3 deformation tensor F , which maps the initial supercell to the final lattice: $S_{\text{final}} = FS_{\text{initial}}$. The atomic shuffles are given by a displacement field D , which is a $3 \times d$ matrix containing the displacement vectors for each of the d atoms to attain the final structure $\mathcal{P}_{\text{final}}$. The complete transformation pathway can be defined by the tuple $\mathcal{T} = (S, F, D)$.

Our algorithm for generating transformation pathways between orderings on two different parent crystal structures consists of three steps. First, a single transformation pathway $\mathcal{T}_1$ is generated between the two undecorated parent structures. This is either computed with the structure mapping algorithm of [39] or provided as user input. All symmetrically distinct variants of the pathway are enumerated by applying the symmetry operations of the initial parent structure, $\mathcal{P}_{\text{initial}}$, to $\mathcal{T}_1$. This procedure yields a set of pathways, $\mathcal{S} = \{\mathcal{T}_1, \mathcal{T}_2, \dots\}$. For example, in the $\text{hcp} \rightarrow \text{bcc}$ transformation, one of the pathways identified by the algorithm of [39] will correspond to the Burgers transformation. The Burgers pathway has three distinct variants that connect the parent hcp structure to three different bcc orientations. These three variants are symmetrically equivalent if the starting point is an undecorated crystal of hcp.

In the second step, we enumerate all possible arrangements of the elements in the alloy on the parent crystal structure $\mathcal{P}_{\text{initial}}$. Symmetrically distinct orderings are enumerated with the algorithm of Hart and Forcade [36,37]. Each decoration, σ_i , represents a distinct arrangement of the elements on $\mathcal{P}_{\text{initial}}$. The set of orderings on the parent structure is collected in a set $\sigma = \{\sigma_1, \sigma_2, \dots\}$.

In the final step of the algorithm, the symmetry group of each ordering is used to partition the set of transformation pathways $\mathcal{S}$ into equivalent groups. The decoration of a crystal structure with several elements can lower the symmetry of the ordered phase relative to the unordered structure. The group of symmetry operations, $\mathcal{G}_{\sigma_i}$, that leaves an ordering σ_i invariant is a sub-group of the symmetry operations that leave the undecorated parent structure unchanged, $\mathcal{G}_{\sigma_i} \subseteq \mathcal{G}_{\mathcal{P}_{\text{initial}}}$. Applying the symmetry group of the ordered phase to the elements of $\mathcal{S}$ partitions the set into groups of transformation pathways that are related by symmetry. For instance, in the case of the $\text{hcp} \rightarrow \text{bcc}$ transformation of $h\text{-}\beta_1$ shown in Fig. 3, the three transformation pathways, $\mathcal{T}_1, \mathcal{T}_2, \mathcal{T}_3$, are partitioned into two symmetrically inequivalent sets: $\{\mathcal{T}_1\}$ and $\{\mathcal{T}_2, \mathcal{T}_3\}$. Pathways $\mathcal{T}_2$ and $\mathcal{T}_3$ are related by the symmetry operations of $h\text{-}\beta_1$,




FIG. 4. Data generation workflow for discovering new low-energy orderings in hcp, bcc, fcc, and C15 Laves phase lattices.

whereas no symmetry operation in $\mathcal{G}_{h-\beta_1}$ maps $\mathcal{T}_1$ to either of the other two. This equivalence is reflected in the transformation energy landscapes of Fig. 3. The energies for $\mathcal{T}_2$ and $\mathcal{T}_3$ are identical, while $\mathcal{T}_1$ is distinct, showing the initial hcp ordering to be unstable. Constructing a minimal set of reference structures requires generating structures along one representative pathway from each distinct set. Our algorithm implements this by selecting one pathway from each set (e.g., $\mathcal{T}_1$ and $\mathcal{T}_2$ for $h-\beta_1$) to generate reference structures.

IV. TRAINING DATASET CONSTRUCTION AND POTENTIAL PARAMETERIZATION

Parameterizing an accurate MLIP required two steps. We first ran an iterative active-learning loop to discover low-energy configurations not present in the initial reference database, augmenting the reference database with the corresponding DFT energies and forces. We then sampled a training dataset from this augmented reference database and fitted the potential.

A. Symmetry-distinct orderings in large supercells

In addition to capturing material properties of interest, MLIPs must also accurately reproduce the thermodynamic stability of stable and metastable phases. Capturing ground states has been a long-standing challenge for on-lattice cluster expansions and has been largely solved through the use of bespoke weighting schemes, loss functions, and cluster selection algorithms [40–42]. MLIPs require a different approach that necessitates iterative updates to the training and reference databases.

We utilized the iterative data generation scheme illustrated in Fig. 4. First, we trained an initial potential on a subset of structures from the reference database. We then used this potential to relax approximately 0.5 million symmetrically distinct orderings in supercells not present in the reference database. We assigned the MLIP-relaxed structures to parent lattices (hcp, bcc, fcc, and C15 Laves) based on a similarity metric computed from their ACE descriptors (see Sec. SVI of the SM [14]). We compared the relaxed formation energies of these new decorations against the convex hull of MLIP-predicted formation energies of the structures in the reference database, constructing separate convex hulls for orderings on each parent crystal structure. The structures in the reference database were classified using the structure matching algorithm of Thomas *et al.* [39]. Subsequently, we collected all configurations with convex hull distance $d_{ch} \leq 10$ meV/atom and compositions $x_{Nd} < 1/2$, along with all newly predicted




FIG. 5. Sampling workflow for assembling the training datasets.

ground states at higher Nd content. We restricted the selection to structures within a small convex hull distance on the Mg-rich side to limit the number of DFT calculations. We then recomputed the energies and forces with DFT and added them to both the reference database and the training subset used to parameterize the potential. We retrained the MLIP and repeated the loop. We considered the procedure complete once the loop produced no new ground states in the hcp, bcc, fcc, and C15 convex hulls below $x_{Nd} = 1/2$ and found no Nd-rich ground states that destabilized the Mg-rich ground states.

B. Dataset sampling and potential parameterization

The reference database described in Sec. III contains over 150 thousand configurations. Training an MLIP on such a large dataset is computationally expensive. Moreover, the potential has to simultaneously reproduce several different material properties. To ensure full coverage of the atomic environments important for modeling precipitation in a smaller dataset, we stratified the reference database into data classes related to bespoke properties of interest, as graphically summarized in Fig. 2. The *training dataset* was created by sampling a fixed number of data points from each class through an iterative process, as schematically illustrated in Fig. 5. First, a seed training dataset was constructed. Next, the Mahalanobis distance between the atomic environments of the configurations belonging to the selected class and the distribution of the seed dataset was computed [43] as

$$d_M(\mathbf{B}_i)^2 = (\mathbf{B}_i - \boldsymbol{\mu})^T \boldsymbol{\Sigma}^{-1} (\mathbf{B}_i - \boldsymbol{\mu}) \quad (4)$$

where $d_M(\mathbf{B}_i)$ is the Mahalanobis distance of an atomic environment centered around atom i with ACE descriptor $\mathbf{B}_i$. This is the vector obtained by concatenating the B -basis functions in Eq. (2), namely $\mathbf{B}_i = (B_{iv_1}, B_{iv_2}, \dots)^T$. $\boldsymbol{\mu}$ and $\boldsymbol{\Sigma}$ are the mean vector and covariance matrix computed from the ACE descriptors of all atomic environments in the current training dataset of size N_{dataset} . The structure containing the atom whose atomic environment has the largest Mahalanobis distance from the training dataset distribution is added to the training dataset and removed from the reference database. The mean vector and covariance matrix are recomputed with the update equations shown in Eqs. (S3) and (S4) based on the descriptors of the newly added structure. The loop was repeated until the target number of structures (N_{target}) was sampled from the selected data class.

We applied the iterative sampling procedure one data class at a time. This ensured a satisfactory coverage of the atomic environments important for modeling precipitation. The final training dataset for Mg, Nd, and Mg-Nd contained

7055, 4002, and 7985 configurations, respectively. A detailed description of the seed training datasets and the sampling methodology is provided in Sec. SVII of the SM [14].

The MLIP was trained using the *pacemaker* code [16]. We parameterized the unary interaction coefficients before fitting the binary interaction coefficients. First, an initial potential was parameterized, either through ladder training or by fitting all parameters to a subset of the training data. We then fine-tuned the resulting potential on the full training dataset. To prevent overfitting, we employed elastic-net regularization and regularization terms for the first and second derivatives of the radial basis functions, both available in the *pacemaker* loss function [17]. The L_1 and L_2 scaling factors were set to a common value of 10^{-6} , whereas the first and second derivative regularization coefficients were adjusted to $\approx 10^{-5}$ for Mg and Mg–Nd, and 10^{-8} for Nd.

Energies and forces were weighted to ensure that low-energy structures were accurately reproduced. The weighting scheme described in [17] was applied to most structures. For the Nd and Mg–Nd potentials, it was necessary to carefully rescale the weights to ensure an accurate reproduction of the ground and low-energy states. Sections VIA and VIB explain how weights were adjusted when training the final MLIPs.

V. VALIDATING INTERATOMIC POTENTIALS

The quality of an MLIP is typically quantified by error metrics that measure the difference between its predicted energies, forces, and stresses and the reference values in training and validation datasets. These global error metrics are supplemented by comparisons of material properties, such as elastic moduli and stacking fault energies, and by a qualitative assessment of defect structures like dislocation cores [28,44]. While global error metrics, such as the root-mean-square error (RMSE) or mean absolute error (MAE), provide a general measure of the potential's predictive accuracy, these property-specific calculations are necessary to ensure that the MLIP accurately captures important quantities for large-scale atomistic simulations. We use the metrics introduced in this section to benchmark the MLIP after training.

A. Assessing phase stability at constant composition

Beyond conventional error metrics, an interatomic potential for studying precipitation must correctly reproduce the thermodynamic stability of all relevant phases. For instance, consider the structures and energies of a pure element shown in Fig. 6, where the double-hcp (dhcp) structure, with an ABAC stacking of triangular layers, is the lowest-energy polymorph. An MLIP that incorrectly predicts another structure, such as fcc, to be more stable than the dhcp ground state would be unreliable for simulating the element's thermodynamic properties. This type of failure can be difficult to detect using only global error metrics like RMSE or MAE, as an incorrect energy ranking between two similar, low-energy polymorphs may only marginally increase the total error. Accurate reproduction of the relative stabilities of low-energy structures is important, as they may be thermally accessible during finite-temperature simulations. In contrast, the exact energy ordering of very high-energy polymorphs is less important, as




FIG. 6. Schematic energy spectrum of elemental polymorphs, depicting a scenario where multiple stacking-order phases such as fcc, hcp, and dhcp lie close to the global ground state.

these structures are unlikely to form under typical simulation conditions.

The weighted Kendall- τ [45] (τ_w) quantifies the similarities and differences between the energy spectra of a set of structural polymorphs computed with DFT and the MLIP. It is defined as

$$\tau_w = \frac{1}{N_\tau} \sum_{i,j>i} w(i,j) \text{sgn}(e_i^{\text{DFT}} - e_j^{\text{DFT}}) \text{sgn}(e_i^{\text{MLIP}} - e_j^{\text{MLIP}}) \quad (5)$$

where e_i^{DFT} and e_i^{MLIP} are the normalized formation energies of a structural polymorph i of an element relative to a reference state, $w(i,j)$ is a pairwise weighting function, N_τ is a normalization constant [$N_\tau = \sum_{i,j>i} w(i,j)$], and $\text{sgn}(x)$ is given by

$$\text{sgn}(x) = \begin{cases} +1 & x > 0 \\ 0 & x = 0 \\ -1 & x < 0 \end{cases} \quad (6)$$

The weighted Kendall- τ takes a value of +1 if the ordering of structural polymorphs between DFT and the MLIP is identical, and a value of -1 if the ordering of the structural polymorphs predicted by the MLIP is exactly opposite to the ordering of energies computed with DFT. Using a uniform weight of 1 for all pairs of structures, transforms the τ_w into the well-known standard unweighted Kendall- τ , simply denoted as τ [45]. The weighted Kendall- τ is particularly useful for analyzing large datasets with a wide energy spread, whereas the unweighted Kendall- τ is better suited for small datasets with a moderate energy spread.

The unweighted Kendall- τ considers all possible pairs of structural polymorphs and counts how often the sign of the energy difference predicted by the MLIP agrees with DFT. If the sign is the same, the pair is *concordant*, and the coefficient registers a point of similarity (+1). If the sign is different, the pair is *discordant*, and it is registered as a point of dissimilarity (-1). The prefactor in Eq. (5) is then the total number of pairs and normalizes the coefficient. The τ coefficient is similar to the concordance index used in [5].

The weighting function, $w(i,j)$, can be tuned to prioritize the importance of low-energy polymorphs over that of high-energy structural polymorphs. Two commonly used weighting functions are either additive, $w(i,j) = w(i) + w(j)$, or multiplicative, $w(i,j) = w(i) \times w(j)$. The weights for individual

structural polymorphs $w(i)$ can be any function that decays with the energy of the polymorph relative to the ground state crystal structure. Adjusting the weighting functions $w(i)$ and $w(i, j)$ enables one to resolve the energy ranking of specific subsets of polymorphs. As long as the coefficient is normalized by N_τ , the τ_w lies between -1 and 1 . Values closer to 1 indicate closer reproduction of the energy ordering of highly weighted structural polymorphs.

In this work, we make use of the hyperbolic weighting function proposed in [45]

$$w(i) = \frac{1}{r(e_i^{\text{DFT}}) + 1} \quad (7)$$

where $r(e_i^{\text{DFT}})$ is a function that converts the normalized DFT energy into the discrete energy-based rank of the structure, starting from 1 , within the list of polymorphs. A discussion of different weighting schemes in the context of ranking algorithms is provided in [45].

Multiplicative weights penalize mispredictions at low energy more strongly, whereas additive weights are more effective at identifying errors at high energy. Furthermore, the Kendall- τ has an intrinsic dependence on the number of structures in the list of polymorphs. This is evident in the unweighted Kendall- τ , where the normalization is simply the number of pairs, $N(N-1)/2$ for a list of N items. A misprediction reduces τ_w by less when the structure list is large; the effect is more pronounced for a small list. Because the multiplicative scheme provides improved low-energy phase stability resolution, we report results using this weighting scheme. A brief study of the influence of the weighting function is provided in Sec. SVIII of the SM [14].

The Kendall- τ described by Eq. (5) can be broadly applied to estimate the quality of the MLIP on any energy spectrum where the relative ordering of energies is important to reproduce. For example, MLIPs must reproduce not just the magnitude of surface energies of various surface orientations, but also capture the relative order of the surface energies correctly. Similarly, unstable stacking fault energies must also be reproduced to capture the properties of dislocations and cracks. The Kendall- τ is readily applied to this problem and can provide a quantitative metric that, in conjunction with error metrics such as the RMSE, could guide the parameterization of MLIPs.

B. Assessing phase stability in binary systems

Using the Kendall- τ to compare phase stability predictions in a multicomponent alloy is more challenging. A naive approach would be to evaluate the τ_w coefficient at individual compositions. Although this approach provides an indication of the accuracy of the potential, it neglects the relative stabilities between phases at different compositions. For example, Fig. 7 illustrates a comparison between the true and MLIP-predicted convex hulls for a hypothetical alloy containing elements A and B. Despite relative energies of orderings being reproduced exactly for all compositions, the equiatomic ordering is predicted to be unstable because the formation energies of all orderings are more positive, as predicted by the MLIP, compared with the ground truth formation energies.




FIG. 7. Schematic of formation energies for a hypothetical A-B alloy computed via DFT and MLIP. While the energy ordering of configurations is correctly predicted at each composition, the AB compound fails to lie on the convex hull in the MLIP prediction. This discrepancy manifests as a decrease in τ_w within the chemical potential range corresponding to AB stability.

The Kendall- τ needs to be modified to compare the relative orderings of energies across multiple compositions.

The stability of phases with varying compositions can be compared within the semigrand-canonical ensemble. The semigrand-canonical potential energy, $\omega(\tilde{\mu}) = e_f(x_B) - \tilde{\mu}x_B$, can be used to compare the stabilities of various degrees of order across crystal structures and compositions. At a fixed exchange chemical potential, $\tilde{\mu} = \mu_B - \mu_A$, the phase with the lowest value of ω will be the stable state. The relative values of the semigrand-canonical potential then indicate the degree of metastability of individual phases.

A metric of the predictive accuracy of the MLIP with respect to phase stability can be obtained by computing the weighted Kendall- τ as a function of the exchange chemical potential. At each value of the exchange chemical potential, the semigrand-canonical potential energy is computed for every phase, and the Kendall- τ of Eq. (5) is calculated by using the semigrand-canonical potential energy in place of the formation energies. This is then repeated for all relevant values of the exchange chemical potential. The output of this procedure is a $\tau_w - \tilde{\mu}$ trace that captures the quantitative accuracy of the MLIP in reproducing both the prediction of phase stability and the prediction of the coexistence chemical potentials between stable phases.

Figure 8(a) summarizes the algorithm to compute the τ_w with semigrand-canonical potential energies. First, the range of relevant exchange chemical potentials is estimated based on the formation energies computed with DFT. The interval of exchange chemical potentials is discretized, and at each $\tilde{\mu}$ grid point the semigrand-canonical potentials of the configurations are calculated from DFT and the MLIP. Next, weights are computed from Eq. (7) with $\omega(\tilde{\mu})$ obtained with DFT in place of the energy. Finally, the $\tau_w(\tilde{\mu})$ coefficient is evaluated. The trace of this quantity can be plotted, and the similarities and differences relative to a value of 1 can be used to infer the quality of the MLIP. For example, Fig. 7 illustrates the $\tau_w - \tilde{\mu}$ trace for the MLIP predictions. At low values of $\tilde{\mu}$, the τ_w is equal to 1 as the MLIP reproduces the stable structure of element A. As the exchange chemical potential increases, the MLIP does not predict the ground state with equiatomic composition to be stable. Consequently, the Kendall- τ degrades to a value below 1 . At larger chemical potentials, the MLIP




FIG. 8. (a) Workflow for computing the semigrand canonical τ_w . (b) Schematic of formation energies illustrating three common MLIP fitting errors and their corresponding τ_w - $\tilde{\mu}$ traces.

and DFT predict the same stable phase of element B and the Kendall- τ approaches 1 again.

The τ_w - $\tilde{\mu}$ trace can be more complex than the simple scenario illustrated in Fig. 7. Figure 8(b) shows three other possible scenarios. In scenario 1, the temporary degradation in Kendall- τ is because of a mismatch in the coexistence chemical potentials of the ground state between DFT and the MLIP. The MLIP predicts a slightly more negative coexistence chemical potential than DFT. As shown in the trace of the τ_w , the coefficient degrades between $\tilde{\mu}_{MLIP}$ and $\tilde{\mu}_{DFT}$. The τ_w recovers a value of 1 at chemical potentials above $\tilde{\mu}_{DFT}$. Figure 8(b) also illustrates a similar degradation at exchange chemical potential values that are away from coexistence chemical potentials. In case 2, the drop in accuracy is because of the incorrect prediction of phase stability away from the ground states. While such mismatches can result in quantitative discrepancies between the predictions of the electronic structure calculations and the MLIP, the MLIP is usually still useful in predicting the thermodynamics of the alloy. In the final scenario illustrated in Fig. 8(b), the MLIP predicts a ground state that is different from DFT. Though the MLIP and DFT predict a ground state at B-rich compositions, they differ in the precise structure. The τ_w at high values of the exchange chemical potential consequently degrades to a small value. The degradation of τ_w over an extended chemical potential range indicates a discrepancy in ground states. The MLIP will then need additional fine-tuning to reproduce the low-temperature thermodynamics of the material.

C. Assessing structural phase transitions

Beyond phase stability, we must also quantify the accuracy of the MLIP in reproducing transformation energy landscapes connecting pairs of crystal structures. A simple error metric such as the RMSE computed over structures corresponding to structural transformations provides a useful estimate of MLIP quality. However, ensuring qualitative agreement between MLIP predictions and DFT calculations is equally important.

Figure 9(a) illustrates six possible energy profiles for structural phase transformations relating an ordering on the hcp parent crystal structure to its bcc counterpart. Pathways 1 and 2 show landscapes where the ordering is unstable in one




FIG. 9. Schematic of the six energy profile classes for the Burgers transformation. These classes span physical scenarios ranging from barrierless transitions to pathways exhibiting barriers and stable intermediate states.

parent structure and undergoes a barrierless homogeneous transformation to an ordering on the other parent structure. Pathways 3 and 4 illustrate a different scenario with an energy barrier separating the higher-energy phase from the lower-energy structure. The last two pathway classes illustrate scenarios where the ordering is unstable in both parent structures and transforms to an intermediate state between them. The MLIP must capture the qualitative trends in these energy profiles to reliably simulate precipitation. Categorizing pathways by their qualitative behavior rather than their precise profile shape allows us to assess mismatches between DFT and MLIP using a confusion matrix. We first assign each pathway in the dataset to one of the classes in Fig. 9 based on its DFT-computed energy landscape. We then use the MLIP to predict energies along the transformation path and classify the pathway accordingly. A confusion matrix quantifies the MLIP's accuracy in reproducing each class.

VI. RESULTS

We used a large dataset of electronic structure calculations to parameterize an interatomic potential for Mg–Nd. We validated the interatomic potential with standard techniques and the error metrics introduced in this study. The ACE basis set and model hyperparameters were chosen as discussed in Sec. II. The MLIPs were then fitted on diverse training datasets following the procedure outlined in Secs. IV A and IV B. Without a careful selection of the weighting schemes, it was not possible to reproduce the stable and metastable ground states spanning the Mg–Nd composition range.

A. Unary potential

The Mg potential achieves training MAEs of 1.9 meV/atom and 1.3 meV/Å. Validation MAEs are 5.6 meV/atom for energies and 6.3 meV/Å for forces. Table SV within the SM [14] compares the predicted hcp-Mg properties with DFT and experimental values. The potential reproduces lattice parameters, the bulk modulus, elastic constants, and stable stacking fault energies [46–50]. It also captures vacancy migration barriers despite the absence of this class of data in the training set. Furthermore, Table II shows that the MLIP accurately predicts planar defect properties, including surface and grain boundary energies. The potential preserves the energetic ordering and magnitudes of planar defect energies, achieving a Kendall- $\tau \geq 0.9$. To assess the finite-temperature reliability

TABLE II. Comparison of energy rankings for various slab and grain boundary geometries. The Kendall- τ coefficient is unweighted.

Surfaces	Highest (mJ/m ²)			Lowest (mJ/m ²)			τ
	Plane	DFT	MLIP	Plane	DFT	MLIP	
hcp slabs	(20 $\bar{2}$ 1)	784	793	(0001)	549	573	0.91
fcc Σ GB	(110)	475	471	(111)	129	140	0.90

of the MLIP, we tested the stability of the potential over the temperature range 50–2000 K, as illustrated in Fig. S5(a). We also calculated the Gibbs free energy of hcp-Mg using the quasi-harmonic approximation as implemented in phonopy [51,52].

The MLIP also reproduces the energy landscape of the structural phase transition relating hcp-Mg to bcc-Mg through the Burgers transformation. Figure 10 illustrates the energy variation from the shuffle of alternate (110) bcc-Mg planes at different strain states. These shuffles are necessary to attain the alternating stacking of triangular layers in the hcp structure. Figure 10(a) compares MLIP and DFT shuffle energies




FIG. 10. Energy variation with the Burgers shuffle at various strain states relative to bcc. Panels show the effect of varying (a) tetragonal distortion e_3 and (b) volumetric strain e_1 .

along the e_3 strain-order parameter, which tracks the degree of Burgers transformation. Previous studies [12,31] rigorously define the strain-order parameters e_1 and e_3 and their relationship to the Burgers transformation. For the ideal bcc structure at zero shuffle magnitude and $e_1 = 0$, both DFT and the MLIP indicate that the (110) shuffle is only marginally more stable. As e_3 becomes negative, the bcc (110) planes distort toward a triangular arrangement. With increasing distortion, the ideal bcc structure becomes less stable, and local minima emerge at non-zero shuffle amplitudes. The hcp structure, the lowest energy state, is attained at $e_3 \approx -0.1$. Figure 10(b) illustrates the response to volumetric compression e_1 relative to the bcc reference. As the bcc structure is compressed, DFT predicts that the shuffle instability diminishes. The MLIP qualitatively reproduces this trend.

Parameterizing the Nd potential was more challenging as multiple crystal structures are close in energy to the 0 K ground state. DFT predicts the 0 K ground state is dhcp, with hcp and fcc structures within 20 meV/atom. A naive parameterization of the interatomic potential for Nd failed to reproduce the dhcp ground state. Figure 11(a) plots the τ_w metric for Mg and Nd MLIPs as a function of the weight fraction assigned to structures within 50 meV/atom of the minimum energy structure. The coefficient is calculated solely over the data of structural prototypes of intermetallics. A Mg potential parameterized on a comparable dataset (Table SVII within the SM [14]) maintains $\tau_w \approx 1$ regardless of the weight assigned to low-energy polymorphs. In contrast, the Nd τ_w drops to ≈ 0.4 when the cumulative weight of low-energy structures falls below 50%, as the MLIP incorrectly predicts the ground state to be fcc. Increasing the weight of low-energy structures improves the predictive quality of the MLIP. As shown by the case of Nd, standard MAEs are largely invariant to varying weights and fail to capture the discrepancies with DFT-predicted ground states. An analysis like that of Fig. 11(a) provides a quantitative guide toward adjusting weighting schemes such that MLIPs reproduce low-temperature phase stability.

Figure 11(a) demonstrates that assigning a total weight of 75% to the low-energy polymorphs reproduces the neodymium ground states, with τ_w approaching 1. Consequently, in fitting our Nd interatomic potential, we assigned a weight of 75% to the low-energy structures. The final Nd potential achieved training MAEs of 9.2 meV/atom and 9.4 meV/Å, and validation MAEs of 42.5 meV/atom and 31.5 meV/Å on held-out data. We confirmed the reproduction of the global dhcp-Nd ground state and the relative ordering of the low-energy polymorphs through the high τ_w value of 0.86. The volumetric distortions for dhcp, fcc, and hcp [Fig. 11(b)] indicate the potential reproduces relative phase stabilities across a range of volumes. An NPT MD heating-cooling cycle confirmed the potential's stability, Fig. S5(b) within the SM [14]. Finally, Table SVI within the SM [14] shows the MLIP reproduces the lattice parameters and bulk modulus of the dhcp ground state [53,54].

B. Binary potential

The Mg–Nd potential achieved training MAEs of 6.4 meV/atom and 22.6 meV/Å. Validation errors were




FIG. 11. (a) The upper panel shows τ_w as a function of the total weight assigned to data points in the 0–50 meV/atom range for Mg and Nd. Circles and squares denote MLIPs predicting incorrect and correct ground states, respectively. The lower panel illustrates the lack of correlation between MAE and τ_w for Nd MLIPs trained with varying low-energy weights. (b) Energy-volume curves for Nd dhcp, fcc, and hcp. DFT and MLIP energies are referenced to the dhcp ground state.

3.6 meV/atom and 16.1 meV/Å. One round of iterative sampling of symmetry-distinct orderings was sufficient to converge the hcp, bcc, fcc, and C15 Laves phase convex hulls below $x_{\text{Nd}} = 1/2$ and to stabilize the Nd-rich side. We compare the convex hulls predicted by the preliminary potential and the final MLIP in Fig S7 within the SM [14]. Table III breaks down the ≈ 0.5 million configurations relaxed with the preliminary MLIP across the four available parent lattices, and reports the number of structures with $d_{ch} < 10$ meV/atom. Most new low-energy orderings were found on the Nd-rich side.

The τ_w metric was used to gauge the accuracy of the ground states on the stable and metastable convex hulls. During training, additional weight was assigned to structures on or near the convex hull, as well as to those with Nd compositions $x_{\text{Nd}} \leq 1/4$, to ensure that τ_w was sufficiently close to

TABLE III. Breakdown of symmetry-distinct orderings enumerated and relaxed during the active learning loop. A subset of configurations did not correspond to any of the four parent lattices.

Enumeration		After relaxation				$d_{ch} < 10$ meV/atom			
Parents		hcp	bcc	fcc	C15	hcp	bcc	fcc	C15
hcp	332272	60037	201010	7635	8	1157	1159	1954	5
bcc	66905	165	41452	2758	0	11	800	599	0
fcc	66905	49	38782	11167	0	4	690	570	0
C15	31268	1630	7752	45	203	36	12	2	14

1 and that the major ground states were reproduced. Table IV summarizes the structural and mechanical properties of the relevant precipitate phases. The MLIP-predicted lattice parameters match DFT predictions. The bulk moduli and elastic constants for these phases agree with DFT. The potential also reproduces properties of stable phases that do not form during precipitation in magnesium alloys containing dilute Nd concentrations. For example, Table SVIII within the SM [14] compares the MLIP-predicted lattice parameters and bulk moduli of Mg_2Nd (C15 Laves) and MgNd (B2) phases with DFT values. Finally, the interatomic potential reproduces the transformation energy landscape for the Burgers transformation of $h\text{-}\beta_1$ to β_1 , as shown in Fig. 12. We assessed phase stability by examining 0 K convex hulls for Nd compositions $0 \leq x_{\text{Nd}} \leq 1/2$ and analyzing the weighted Kendall- τ trace as a function of the exchange chemical potential ($\tilde{\mu} = \mu_{\text{Nd}} - \mu_{\text{Mg}}$). Figure 13(a) to 13(c) compare MLIP formation energies with DFT values. For each convex hull, the DFT structures were relaxed with the MLIP prior to evaluating the formation energies. The $\tau_w\text{-}\tilde{\mu}$ trace in Fig. 13(a) remains close to 1 up to $\tilde{\mu} \approx -0.4$ eV/atom, the coexistence chemical potential between the β_1 and β_e phases. This τ_w value indicates that the potential accurately captures stable and metastable phases up to $x_{\text{Nd}} = 1/4$. Between this composition and $x_{\text{Nd}} = 1/2$, the MLIP predicts a spurious ground state at $x_{\text{Nd}} = 1/3$, an

TABLE IV. Comparison of structural and mechanical properties for phases with $x_{\text{Nd}} \leq 1/4$ from DFT, and MLIP.

Property	β'		β''		β_1		β_e	
	DFT	MLIP	DFT	MLIP	DFT	MLIP	DFT	MLIP
Lattice parameters ($\text{\AA}$)								
a	7.07	7.19	6.81	6.81	5.24	5.26	11.63	11.65
b	10.87	10.75	$\equiv a$	$\equiv a$	$\equiv a$	$\equiv a$	$\equiv a$	$\equiv a$
c	5.12	5.11	5.25	5.25	$\equiv a$	$\equiv a$	$\equiv a$	$\equiv a$
Mechanical properties (GPa)								
B	37.8	38.1	38.3	38.8	40.1	38.7	39.4	39.6
C_{11}	54.4	53.7	64.0	63.6	59.6	57.2	72.3	67.6
C_{12}	35.9	39.4	31.2	31.9	28.8	31.4	21.3	23.0
C_{13}	21.2	22.7	18.1	22.1	$\equiv C_{12}$	$\equiv C_{12}$	21.9	26.2
C_{33}	70.1	70.5	77.1	73.4	$\equiv C_{11}$	$\equiv C_{11}$	71.7	64.5
C_{44}	14.6	11.6	10.8	15.5	39.0	37.1	27.5	22.9
C_{55}	30.7	28.0	$\equiv C_{44}$	$\equiv C_{44}$	$\equiv C_{44}$	$\equiv C_{44}$	$\equiv C_{44}$	$\equiv C_{44}$
C_{66}	27.6	24.4	16.4	15.8	$\equiv C_{44}$	$\equiv C_{44}$	16.3	12.26




FIG. 12. Comparison of Burgers transformation pathways originating from $h\text{-}\beta_1$.

ordering of Mg and Nd over the ω parent structure. This produces a modest decline in the $\tau_w\text{-}\tilde{\mu}$ trace at chemical potentials above -0.4 eV/atom. Both the weighting scheme

and the large number of structures in the analysis keep the change in τ_w small. Within the chemical potential window where the MLIP predicts the ω phase to be stable, DFT ranks the spurious ground state 34th among the structures. With ≈ 2200 structures in the comparison, this single misprediction produces an expected drop of ≈ 0.01 in τ_w , consistent with the small τ_w fluctuations in Fig. 13(a). A larger drop appears beyond the $\beta_1\text{-B2}$ coexistence chemical potential, where the MLIP incorrectly ranks the energies of configurations with higher Nd composition. The incorrect ground states predicted by the MLIP likely stem from parameterization choices such as the basis set size. While this error precludes using the potential for $x_{\text{Nd}} > 1/4$, it does not affect the study of precipitation in Mg-rich alloys that contain less than 1 at.% Nd.

The MLIP also reproduces metastable convex hulls for chemical decorations of bcc and hcp parent crystal structures. In Fig. 13(b), the τ_w remains near 1, with deviations arising from mismatches in predicted coexistence chemical potentials. The metastable hcp convex hull in Fig. 13(c)




FIG. 13. Convex hulls and τ_w metrics for (a) the global system, (b) the bcc phase, and (c) the hcp phase. Vertical dashed lines indicate coexistence chemical potentials predicted by the MLIP (red) and DFT (gray). (d) Energies of β''' orderings [dotted box in (c)]. Energies are referenced to pure Mg and β'' ($x_{\text{Nd}} = 1/4$). (e) Global convex hull of all experimental phases and orderings enumerated on the bcc, hcp, fcc, and C15 Laves phase lattices relaxed with the MLIP.

exhibits a more complex $\tau_w - \tilde{\mu}$ trace. The metric remains near 1 for exchange chemical potentials up to ≈ -0.5 eV/atom, corresponding to the hcp- β' coexistence. At higher exchange chemical potentials, the τ_w declines. As shown in Fig. 13(d), the potential fails to predict the β''' orderings on the convex hull. Capturing these long-period orderings is challenging because of small energy differences (≈ 2 meV/atom), the influence of long-range strain interactions, and inherent DFT errors. However, the overall discrepancy is small, and the potential remains sufficiently accurate to model the precipitation process. Finally, for chemical potentials $\tilde{\mu} > -0.35$ eV/atom, the τ_w degrades because of incorrect phase stability predictions at $x_{\text{Nd}} > 1/4$.

Figure 13(e) compares the MLIP formation energies, computed over nearly 0.5 million structures, against the DFT formation energies for all ordered phases in the reference database. The MLIP convex hull is constructed over chemical decorations of Mg and Nd on the bcc, hcp, fcc, and C15 parent crystal structures, as well as all experimentally reported phases in this binary alloy. The MLIP predictions agree well with the DFT calculations, and the MLIP predicts no spurious ground states at compositions below $x_{\text{Nd}} = 1/4$, as the traces of Figs. 13(a)–13(c) indicate. A spurious ground state emerges at a Nd composition of $x_{\text{Nd}} = 1/3$. This spurious ground state does not displace any ground states at lower Nd compositions, which preserves the accuracy of the potential for studying precipitation in dilute Mg-rich alloys. The erroneous ground state precludes the use of this MLIP for modeling alloys with high Nd compositions.

Having achieved sufficient accuracy to model the low-temperature thermodynamics of Mg-rich Mg-Nd alloys with the MLIP, we assessed its reproduction of the hcp-to-bcc Burgers transformation using symmetrically distinct transformation pathways enumerated with the algorithm outlined in Sec. III A. Our algorithm enumerated 149 symmetrically distinct pathways. The dataset includes all pathways obtained from symmetrically distinct chemical decorations in supercells up to twice the size of the primitive hcp cell, orderings with Nd compositions of 1/8 and 1/4 in supercells containing six and eight atoms, and specific orderings in supercells containing 16, 24, and 96 atoms. The larger supercells capture structural phase transitions between precursor orderings on hcp and β_1 . As shown in Fig. 14(a), our systematic enumeration produces all six classes of transformation energy landscapes defined in Fig. 9. Figure 14(b) shows the confusion matrix comparing DFT and MLIP classifications. Energies of structures along transformation pathways connecting hcp and bcc orderings spanned 10–50 meV/atom. Misclassifications occur most frequently in pathways with barriers (classes 3–4). In both classes, the MLIP erroneously predicts $\approx 10\%$ of pathways to lack a barrier. End-state stability inversion is another common error, causing 24% of class 3 pathways to be mislabeled as class 4. Further analysis revealed that pathways with inverted end-state energies have an average endpoint energy difference of ≈ 4.5 meV/atom, well within the potential's validation error.

Modeling precipitate growth requires accurate relative energies for point and extended defects within precipitates. For instance, the β' phase transforms into $h\text{-}\beta_1$ by replacing some Mg atoms with Nd, as shown in Fig. 1. Figure 15 depicts two




FIG. 14. (a) Number of pathways assigned to each class defined in Fig. 9. Classification is based on the DFT energy profiles. (b) Confusion matrix comparing MLIP and DFT predictions for the qualitative shapes of transformation pathways. Pathway classes correspond to those defined in Fig. 9.

substitution mechanisms that progressively approach the $h\text{-}\beta_1$ ordering. Figure 15(a) plots the energy profile for the structural transformation from β' to its bcc counterpart, showing close agreement between MLIP and DFT. Substituting a Nd atom at a site that would transform the β' ordering toward $h\text{-}\beta_1$ shifts the energy minimum toward bcc, as illustrated in Fig. 15(b). We also computed transformation energies for orderings obtained by adding a second Nd atom along either the hcp c axis or the $[10\bar{1}0]$ direction. As shown in Fig. 15(c), both DFT and the MLIP indicate that addition along the $[10\bar{1}0]$ direction is energetically favorable.

In addition to capturing defect energies within precipitates, the MLIP must also reproduce defect energetics within the disordered matrix, including point defects such as dilute Nd in hcp-Mg and Nd segregation to planar defects like stacking faults and grain boundaries. DFT predicts the Nd dissolution energy in a $4 \times 4 \times 3$ hcp supercell to be -34 meV relative to pure reference states, while the MLIP predicts -41 meV. The MLIP also accurately describes Nd-Nd and Nd-defect interactions. Figure S8(a) within the SM [14] shows that Nd-Nd binding energies within hcp-Mg agree with




FIG. 15. Schematic of the initial steps of h - β_1 nucleation starting from β' . Energy profiles represent Burgers pathways from the depicted hcp ordering to the lowest-energy bcc state. Crystal structures shown in (a) (no Nd added) and (b) (one Nd addition), the basal and prismatic projections correspond to the same structure. In (c) (two Nd additions), the basal projection shows Nd addition along a zigzag path, while the prismatic projection illustrates addition along a column.

DFT. Similarly, binding energies to stable stacking faults [Fig. S8(b) within the SM [14]] and the Σ grain boundary [Fig. S8(c)] match DFT predictions. The MLIP also accurately captures vacancy defect energetics, reproducing vacancy formation energies in the relevant ground states (Fig. S9 within the SM [14]). Capturing the Nd-vacancy binding energy in the hcp-Mg matrix proved more difficult. As shown in Fig. S10 within the SM [14], DFT predicts a slight attraction between Nd and a vacancy at short distances, whereas the MLIP predicts a slight repulsion. This sign

error persisted despite increased weighting of short-distance Nd-vacancy configurations during training. The discrepancy likely reflects a limitation of the ACE model, though whether it originates from the basis set, the cutoff radius, or the embedding function remains unclear. To assess the reliability of the potential for kinetic properties, we computed the Nd tracer diffusion constant in hcp-Mg with the Green function method developed by Trinkle [55–57]. Hop barriers were computed with the nudged elastic band method [58], and frequency prefactors were obtained from Vineyard transition state theory




FIG. 16. (a) Tracer diffusion constant of a Nd solute in hcp Mg. (b) Gibbs free energy of hcp Mg, β' , and β_1 derived from the quasiharmonic approximation.

[59]. As shown in Fig. 16(a), the MLIP-predicted diffusion constant agrees well with the DFT-based estimate [60] when the same frequency prefactors are used, and with available experimental data [61]. The potential is therefore suitable for modeling vacancy-mediated diffusion in the dilute random alloy.

We next assessed the Gibbs free energies of pure Mg, β' , and β_1 over the experimentally relevant range of 0–500 K, using the quasiharmonic approximation as implemented in phonopy. Figure 16(b) shows the temperature dependence of the free energy for the three phases. The MLIP matches DFT closely for all three phases. A shift of a few meV/atom appears for β_1 at low temperature, which we attribute to a small difference between the formation energy predicted by the MLIP and the value computed with DFT. Overall, the MLIP reproduces finite-temperature thermodynamic and transport properties well.

VII. DISCUSSION

Precipitation hardening is widely used in engineering alloys to improve material strength. This study presents methods for parameterizing reliable MLIPs to simulate precipitation and structural phase transitions in metallic alloys. We introduce error metrics and data-generation schemes for MLIPs that model solid-state phase transformations during precipitation. We demonstrate these methods by parameterizing an MLIP for the Mg–Nd binary alloy that captures stable and metastable thermodynamics, elastic moduli, point defect energetics, and vacancy formation energies. The weighted Kendall- τ (τ_w) and its semigrand-canonical generalization, illustrated in Figs. 6 to 8, provide interpretable metrics for scoring MLIP predictions of low-temperature thermodynamics in multicomponent systems. We can compute τ_w using several weighting schemes, though the relative merits of different schemes require further exploration. The algorithm for enumerating symmetry-distinct transformation pathways complements the phase stability analysis by mapping the potential energy surface along pathways connecting chemical orderings on different parent crystal structures. Although we demonstrated this algorithm for the Mg–Nd system using the Burgers transformation linking hcp and bcc orderings, it applies to any well-defined pathway connecting different crystal structures. This methodology enables error quantification through a confusion matrix as shown in Fig 14(b).

Standard statistical error metrics that enter the loss function are insensitive to the correct energy ranking of structures, particularly when energy differences are small, as demonstrated in Fig. 11(a). The τ_w metric addresses this limitation. We use it to identify the weighting scheme yielding an MLIP that accurately reproduces the ground state of Nd [Fig. 11(a)], and to validate MLIP predictions of phase stability [Figs. 13(a)–13(c)]. We do not incorporate τ_w into the loss function. For on-lattice cluster expansions, a previous study [41] trained cluster expansion models to reproduce ground states by incorporating convex-hull distances as constraints in a quadratic programming framework. Incorporating a Kendall- τ term directly into the MLIP loss function could

similarly improve phase stability predictions and remains an avenue for future work.

The data generation and validation strategies introduced in this work are directly applicable to other binary systems. The τ_w metric and the confusion matrix for transformation pathways compress complex information into a convenient form, simplifying the fine-tuning of an MLIP. Multicomponent alloys containing several elements pose additional challenges for both data generation and validation. For instance, the enumeration of symmetrically distinct transformation pathways connecting different parent crystal structures extends to multicomponent alloys. However, the number of distinct decorations grows combinatorially with the number of elements, making exhaustive exploration impractical for systems with many components. Selective sampling through active learning or data-driven metrics of undersampled regions, such as the Mahalanobis distance [43] or the extrapolation grade [62], could be used to identify the most informative pathways. Developing systematic, efficient, and automated training and validation strategies for alloys containing several elements remains an open challenge.

Systematic enumeration of crystal structures, combined with careful fine-tuning of the MLIP based on the error metrics developed here, produces an interatomic potential that generalizes well to data not included during training. We tested the MLIP on several interface and kinetic properties. These included planar fault energies such as anti-phase boundary energies in β' , generalized stacking fault energies in β' , β'' , and β_1 , and the energies of several interfaces between hcp-Mg, β' and β_1 . We analyzed vacancy migration barriers for symmetrically distinct vacancy hops within β' and for a vacancy adjacent to a Nd atom in otherwise pure hcp-Mg. Figures S11, S12(a), S12(b), S13, S14, and Table SIX within the SM [14] show comparisons between DFT and the MLIP. The MLIP matches our DFT results and literature values [60,63–66] across all these properties, indicating that the potential is well parameterized and captures features not explicitly included during training.

Investigating precipitation at 0 K

Our MLIP also reproduces precipitate energies across varying sizes when benchmarked against DFT. We computed the energies of small nuclei within a 384-atom hcp supercell, corresponding to a $3 \times 2 \times 4$ supercell of the β' crystal structure. Simulations included the formation of β' within hcp-Mg, as well as β'' and β_1 embedded within β' . Precipitates of varying sizes were generated by systematically adding Nd atoms one at a time to the lowest-energy site while accounting for atomic relaxations. We used the MLIP to identify the site with the lowest Nd substitution energy, after which the energy was computed with DFT. Only the sublattice relevant to the precipitating phase was filled, with precipitate size restricted to at most 1/8 of the supercell.

Figure 17(c) compares nucleation energies predicted by the MLIP against DFT for the formation of three precipitates. We approximate the nucleation energy at 0 K as

$$\Delta\phi = E_f(N_{\text{Nd}}) - E_f(0) - N_{\text{Nd}}\tilde{\mu} \quad (8)$$




FIG. 17. (a) Schematic of the four precipitate types and their orientations within a disordered matrix. (b) Snapshots of β_1 formation within a disordered matrix phase. Only Nd atoms are shown. (c) Comparison of MLIP and DFT nucleation energies for small precipitates. (d) Nucleation energies for precipitate formation from the disordered matrix. (e) Average Q_4 -order parameter as a function of precipitate size.

where $E_f(N_{Nd})$ is the formation energy of the supercell after adding N_{Nd} solute atoms and $\bar{\mu}$ is the difference in chemical potential between Nd and Mg. We set the exchange chemical potential $\bar{\mu}$ to the dissolution energy of a single Nd atom within hcp-Mg.

The MLIP predictions are in excellent agreement with the DFT calculations shown in Fig. 17(c). At the chemical potential and the small simulation cell sizes used in Fig. 17(c), the formation of β_1 within β' is thermodynamically favored over that of β'' . This agrees with experimental results from several studies [8,67,68]. The MLIP predictions quantitatively match the DFT results even though the comparison is made in a large supercell. The energy differences between DFT and MLIP values are <2 meV/atom, suggesting that the accuracy of this MLIP for modeling precipitation is better than what the validation errors would indicate.

Having established that the MLIP captures precipitate formation energies, we used the atomistic model to perform a preliminary study of precipitation in a dilute Mg–Nd alloy. Figure 17(a) summarizes the precipitate geometries considered. Previous experiments [6,8,10,67] and phase field simulations [64,69] indicate that precipitates in Mg–Nd alloys form as flat plates with habit planes of $\{1\bar{2}10\}_{hcp}$ for β' and $\{1\bar{1}00\}_{hcp}$ for β_1 . GP zones typically form through Nd atoms arranged in a zigzag pattern with a habit plane of $\{01\bar{1}0\}_{hcp}$ [10]. We also studied the formation of precipitates with two rows of Nd atoms arranged as zigzags with a habit plane of $\{01\bar{1}0\}_{hcp}$ and spacing matching that of β' . This morphology has not been observed experimentally. In Fig. 17(a), these geometries are labeled precipitates 1 through 4, corresponding to β' , β_1 , Guinier–Preston zones, and the double zigzag structure, respectively.

In principle, the energies of precipitates of all possible shapes at a fixed size should be computed to assess nucleation energies. We use insights from experiments and previous simulations to select specific precipitate aspect ratios. Following the phase field simulations in [64] for β' and [69] for β_1 , we set the ratio between the height (along $[0001]_{\text{hcp}}$) and the length [along the growth direction indicated by the red arrow in Fig. 17(a)] to approximately 1. Experimental data at 15 minutes of aging [64] suggests a β' plate thickness (normal to the habit plane) of $\approx 10 \text{ \AA}$, corresponding to four Nd atoms in the zigzag. Owing to the lack of experimental data, we imposed the same thickness on the β_1 plate. We based the GP zone and the 2-zigzag precipitate geometries on crystallographic information described in [6,10]. The modeled precipitation mechanism involves simultaneous size increase along $[0001]_{\text{hcp}}$ and the growth direction [red arrow in Fig. 17(a)] while keeping the plate thickness fixed.

We assessed precipitate formation energetics by computing the energy difference required to form a coherent precipitate from a 0.5 at.% Nd solid solution. The initial state consisted of a large orthorhombic simulation box, $173 \times 199 \times 186 \text{ \AA}$, containing randomly distributed Nd atoms at the target composition. Precipitates were created by extracting solute atoms from the disordered matrix and rearranging them to match the shape, arrangement, and orientation of one of the precipitates shown in Fig. 17(a). The precipitate-containing cells were then relaxed to track the variation in system energy as a function of the number of Nd atoms in the precipitate. All precipitates were assumed to be stoichiometric and coherent. Figure 17(b) illustrates the assembly process for the β_1 precipitate.

Figure 17(d) shows the change in energy of the simulation cell as precipitate size increases. Nucleation of β_1 produces the largest energy reduction, as it is more stable than β' or other orderings on hcp. Growth of β' along the $[01\bar{1}0]_{\text{hcp}}$ direction is favored over the GP zone and the 2-zigzag precipitate, consistent with experimental observations. The energy reduction for the formation of β' , the GP zone, and the 2-zigzag structure are similar at small sizes, suggesting competition among these ordered precipitates during early aging. As precipitates grow, however, the energy of the experimentally observed coherent β' precipitate decreases more rapidly, making it more favorable than other ordered phases on hcp.

The predictions of Fig. 17(d) show no nucleation barrier. This is because of the large thermodynamic driving force for precipitation at 0 K. The absence of finite-temperature contributions significantly destabilizes the disordered phase, causing ordered precipitates to form without any barrier. Entropic contributions lower the chemical potential of the solid solution with increasing temperature, thereby stabilizing it against precipitation. For a sufficiently dilute alloy where solute interactions are negligible, the chemical potential of the random solid solution at 0 K can be approximated by the dissolution energy of the solute. Within this approximation, the chemical potential of the disordered Mg–Nd matrix, $\bar{\mu}_{SS}$, is -0.063 eV/atom . This value is much higher than the coexistence chemical potential between hcp-Mg and β' or between hcp-Mg and β_1 , which is predicted to be $\approx -0.49 \text{ eV/atom}$. As a result the driving force for nucleation is very large and

completely overcomes any interfacial or strain energy penalties. A future study will describe the effects of temperature on the nucleation barrier in this alloy system.

Steinhardt bond-order parameters [70] can quantify the extent of the structural phase transformation when precipitates such as β_1 form within a matrix phase of a different crystal structure. Figure 17(e) shows the Q_4 -order parameter, averaged over atoms in the precipitate, as a function of the size of β' and β_1 . Horizontal dashed lines indicate ideal order parameter values for perfect hcp and bcc.

For β' , the Q_4 value at small precipitate sizes is close to the ideal hcp value but progressively deviates from hcp as precipitate size increases. This deviation arises from the large stress-free transformation strains accompanying β' formation, which have been predicted to be as high as $\approx 12\%$ along the $[2\bar{1}\bar{1}0]_{\text{hcp}}$ direction [10,11]. As β' precipitates grow, the energy gained by relaxing atoms away from perfect hcp positions offsets the energy penalties from straining the surrounding Mg matrix.

The β_1 precipitates in Fig. 17(e) exhibit a different trend. Even at small sizes, atoms within the β_1 precipitate relax away from their ideal hcp positions, resulting in structural-order parameters that deviate from hcp more significantly than those of β' . As the precipitate grows, the average order parameter for β_1 gradually shifts toward bcc, indicating a continuous phase transition that transforms the underlying hcp crystal structure of the h - β_1 ordering into a bcc structure resembling β_1 . Achieving perfect bcc-like coordination requires larger precipitate sizes, as the energy gain from the structural phase transition must offset penalties from interface mismatch and the strain energy needed to maintain precipitate coherency.

VIII. CONCLUSIONS

In this work, we present a strategy for parameterizing MLIPs to study precipitation and structural phase transitions in metallic alloys. The training and validation database was large and diverse, encompassing the crystalline environments required to capture the features of the Mg–Nd system and to model nucleation. We supplemented conventional fitting procedures with tailored error metrics and data-generation schemes, including the weighted Kendall- τ coefficient and an algorithm to enumerate symmetry-distinct transformation pathways. The τ_w provides a measure of the ability of the MLIP to predict the relative stability of different polymorphs and to optimize weighting schemes. Its semigrand-canonical formulation quantifies the MLIP's reliability in describing phase stability across the full composition range. Through this careful parameterization and assessment, the Mg–Nd MLIP generalized well to crystal structures absent from the training set, particularly interfaces and vacancy hop configurations. A preliminary study of β' and β_1 precipitation showed that β_1 nucleation produces the largest energy reduction, with β' formation also lowering the total energy but to a lesser extent. These results agree well with previous experimental studies of precipitation in this alloy system. The methods developed here apply to parameterizing interatomic potentials and studying precipitation in other alloy systems, including multicomponent Al alloys [71–73], Ti alloys [74,75], and high-entropy alloys [76]. Our Mg–Nd MLIP reproduces solid-state

precipitation reactions in this binary alloy, and future work can use it to probe the atomistic mechanisms of precipitation and to quantify finite-temperature nucleation barriers.

ACKNOWLEDGMENTS

This research was supported by the Swiss National Science Foundation (Grant No. 215178). The authors also acknowledge access to Eiger at the Swiss National Super-

computing Centre with project ID mr30 under the NCCR MARVEL's share, a National Centre of Competence in Research funded by the Swiss National Science Foundation (Grant No. 205602). The authors acknowledge the use of pymatgen [77], ase [78], spglib [79], and LAMMPS [38].

DATA AVAILABILITY

The data associated with this study is available at [80].

- [1] S. Starikov, Y. Lysogorskiy, M. Qamar, A. Bochkarev, M. Mrovec, and R. Drautz, Atomic cluster expansion for the aluminum-magnesium-hydrogen system, *Phys. Rev. Mater.* **9**, 103606 (2025).
- [2] D. Marchand, Foundation models for metallurgy? *MRS Bull.* **50**, 805 (2025).
- [3] P. Srinivasan, S. Puri, K. Pacho Dominguez, A. P. Horsfield, M. R. Gilbert, and D. Nguyen-Manh, Atomic cluster expansion interatomic potentials for lithium: Investigating the solid and liquid phases, *Phys. Rev. B* **112**, 054108 (2025).
- [4] Y. Liu, X. He, and Y. Mo, Discrepancies and error evaluation metrics for machine learning interatomic potentials, *npj Comput. Mater.* **9**, 174 (2023).
- [5] Y. Liu and Y. Mo, Assessing the accuracy of machine learning interatomic potentials in predicting the elemental orderings: A case study of Li-Al alloys, *Acta Mater.* **268**, 119742 (2024).
- [6] J.-F. Nie, Precipitation and hardening in magnesium alloys, *Metall. Mater. Trans. A* **43**, 3891 (2012).
- [7] Z. Xu, M. Weyland, and J. Nie, Shear transformation of coupled β_1/β' precipitates in Mg-RE alloys: A quantitative study by aberration corrected STEM, *Acta Mater.* **81**, 58 (2014).
- [8] H. Liu, Y. Zhu, N. Wilson, and J. Nie, On the structure and role of β F' in beta1 precipitation in Mg-Nd alloys, *Acta Mater.* **133**, 408 (2017).
- [9] H. Xie, X. Zhao, J. Jiang, J. Bai, S. Li, H. Pan, X. Pang, H. Li, Y. Ren, and G. Qin, Diffusional-displacive transformation mechanism for the beta1 precipitate in a model Mg-rare-earth alloy, *Mater. Charact.* **174**, 111018 (2021).
- [10] A. R. Natarajan, E. L. Solomon, B. Puchala, E. A. Marquis, and A. Van Der Ven, On the early stages of precipitation in dilute Mg-Nd alloys, *Acta Mater.* **108**, 367 (2016).
- [11] A. R. Natarajan and A. Van der Ven, A unified description of ordering in HCP Mg-RE alloys, *Acta Mater.* **124**, 620 (2017).
- [12] A. R. Natarajan and A. Van der Ven, Toward an understanding of deformation mechanisms in metallic lithium and sodium from first-principles, *Chem. Mater.* **31**, 8222 (2019).
- [13] K.-H. Kim and B.-J. Lee, Modified embedded-atom method interatomic potentials for Mg-Nd and Mg-Pb binary systems, *Calphad* **57**, 55 (2017).
- [14] See Supplemental Material at <http://link.aps.org/supplemental/10.1103/xyb1-7j1p> for summary of crystallography, details concerning the ace mlip, in-depth description of the database and training dataset, and supplemental results on unary and binary potentials.
- [15] R. Drautz, Atomic cluster expansion for accurate and transferable interatomic potentials, *Phys. Rev. B* **99**, 014104 (2019).
- [16] Y. Lysogorskiy, C. V. D. Oord, A. Bochkarev, S. Menon, M. Rinaldi, T. Hammerschmidt, M. Mrovec, A. Thompson, G. Csányi, C. Ortner, and R. Drautz, Performant implementation of the atomic cluster expansion (PACE) and application to copper and silicon, *npj Comput. Mater.* **7**, 97 (2021).
- [17] A. Bochkarev, Y. Lysogorskiy, S. Menon, M. Qamar, M. Mrovec, and R. Drautz, Efficient parametrization of the atomic cluster expansion, *Phys. Rev. Mater.* **6**, 013804 (2022).
- [18] N. Leimeroth, L. C. Erhard, K. Albe, and J. Rohrer, Machine-learning interatomic potentials from a users perspective: A comparison of accuracy, speed and data efficiency, *Modell. Simul. Mater. Sci. Eng.* **33**, 065012 (2025).
- [19] E. Ibrahim, Y. Lysogorskiy, M. Mrovec, and R. Drautz, Atomic cluster expansion for a general-purpose interatomic potential of magnesium, *arXiv:2305.03577*.
- [20] G. Kresse and J. Hafner, *Ab initio* molecular dynamics for liquid metals, *Phys. Rev. B* **47**, 558(R) (1993).
- [21] G. Kresse and J. Hafner, *Ab initio* molecular-dynamics simulation of the liquid-metal-amorphous-semiconductor transition in germanium, *Phys. Rev. B* **49**, 14251 (1994).
- [22] G. Kresse and J. Furthmüller, Efficiency of ab-initio total energy calculations for metals and semiconductors using a plane-wave basis set, *Comput. Mater. Sci.* **6**, 15 (1996).
- [23] G. Kresse and D. Joubert, From ultrasoft pseudopotentials to the projector augmented-wave method, *Phys. Rev. B* **59**, 1758 (1999).
- [24] J. P. Perdew, K. Burke, and M. Ernzerhof, Generalized gradient approximation made simple, *Phys. Rev. Lett.* **77**, 3865 (1996).
- [25] J. P. Perdew, K. Burke, and M. Ernzerhof, Generalized gradient approximation made simple [Phys. Rev. Lett. 77, 3865 (1996)], *Phys. Rev. Lett.* **78**, 1396 (1997).
- [26] R. Kobayashi, D. Giofré, T. Junge, M. Ceriotti, and W. A. Curtin, Neural network potential for Al-Mg-Si alloys, *Phys. Rev. Mater.* **1**, 053604 (2017).
- [27] F. Maresca, D. Dragoni, G. Csányi, N. Marzari, and W. A. Curtin, Screw dislocation structure and mobility in body centered cubic Fe predicted by a Gaussian approximation potential, *npj Comput. Mater.* **4**, 69 (2018).
- [28] D. Marchand, A. Jain, A. Glensk, and W. A. Curtin, Machine learning for metallurgy I. A neural-network potential for Al-Cu, *Phys. Rev. Mater.* **4**, 103601 (2020).
- [29] M. Stricker, B. Yin, E. Mak, and W. A. Curtin, Machine learning for metallurgy II. A neural-network potential for magnesium, *Phys. Rev. Mater.* **4**, 103602 (2020).

- [30] S. K. Kolli, A. R. Natarajan, J. C. Thomas, T. M. Pollock, and A. Van Der Ven, Discovering hierarchies among intermetallic crystal structures, *Phys. Rev. Mater.* **4**, 113604 (2020).
- [31] J. C. Thomas and A. Van Der Ven, The exploration of nonlinear elasticity and its efficient parameterization for crystalline materials, *J. Mech. Phys. Solids* **107**, 76 (2017).
- [32] R. Tran, X.-G. Li, J. H. Montoya, D. Winston, K. A. Persson, and S. P. Ong, Anisotropic work function of elemental crystals, *Surf. Sci.* **687**, 48 (2019).
- [33] D. Zagorac, H. Müller, S. Ruehl, J. Zagorac, and S. Rehme, Recent developments in the Inorganic Crystal Structure Database: Theoretical crystal structure data and related features, *J. Appl. Crystallogr.* **52**, 918 (2019).
- [34] S. Kirklin, J. E. Saal, B. Meredig, A. Thompson, J. W. Doak, M. Aykol, S. Rühl, and C. Wolverton, The Open Quantum Materials Database (OQMD): Assessing the accuracy of DFT formation energies, *npj Comput. Mater.* **1**, 15010 (2015).
- [35] B. Puchala, J. C. Thomas, A. R. Natarajan, J. G. Goiri, S. S. Behara, J. L. Kaufman, and A. Van Der Ven, CASM — A software package for first-principles based study of multicomponent crystalline solids, *Comput. Mater. Sci.* **217**, 111897 (2023).
- [36] Gus L. W. Hart and R. W. Forcade, Algorithm for generating derivative structures, *Phys. Rev. B* **77**, 224115 (2008).
- [37] G. L. W. Hart and R. W. Forcade, Generating derivative structures from multilattices: Algorithm and application to hcp alloys, *Phys. Rev. B* **80**, 014120 (2009).
- [38] A. P. Thompson, H. M. Aktulga, R. Berger, D. S. Bolintineanu, W. M. Brown, P. S. Crozier, P. J. In 't Veld, A. Kohlmeyer, S. G. Moore, T. D. Nguyen, *et al.*, LAMMPS—A flexible simulation tool for particle-based materials modeling at the atomic, meso, and continuum scales, *Comput. Phys. Commun.* **271**, 108171 (2022).
- [39] J. C. Thomas, A. R. Natarajan, and A. Van Der Ven, Comparing crystal structures with symmetry and geometry, *npj Comput. Mater.* **7**, 164 (2021).
- [40] B. Puchala and A. Van der ven, Thermodynamics of the Zr-O system from first-principles calculations, *Phys. Rev. B* **88**, 094108 (2013).
- [41] W. Huang, A. Urban, Z. Rong, Z. Ding, C. Luo, and G. Ceder, Construction of ground-state preserving sparse lattice models for predictive materials simulations, *npj Comput. Mater.* **3**, 30 (2017).
- [42] J. G. Goiri and A. Van Der Ven, Recursive alloy Hamiltonian construction and its application to the Ni-Al-Cr system, *Acta Mater.* **159**, 257 (2018).
- [43] P. C. Mahalanobis, Reprint of: Mahalanobis, P.C. (1936) "On the generalised distance in statistics", *Sankhya A* **80**, 1 (2018).
- [44] D. Marchand and W. A. Curtin, Machine learning for metallurgy IV: A neural network potential for Al-Cu-Mg and Al-Cu-Mg-Zn, *Phys. Rev. Mater.* **6**, 053803 (2022).
- [45] S. Vigna, in *Proceedings of the 24th International Conference on World Wide Web* (International World Wide Web Conferences Steering Committee, Florence Italy, 2015), pp. 1166–1176.
- [46] G. Walker and M. Marezio, Lattice parameters and zone overlap in solid solutions of lead in magnesium, *Acta Metall.* **7**, 769 (1959).
- [47] L. J. Slutsky and C. W. Garland, Elastic constants of magnesium from 4.2°K to 300°K, *Phys. Rev.* **107**, 972 (1957).
- [48] M. Nishimura, K. Kinoshita, Y. Akahama, and H. Kawamura, Volume compression of Mg and Al to multimegabar pressure, in *Proceedings of the Joint 20th AIRAPT - 43rd EHPRG Conference on Science and Technology of High Pressure* (Forschungszentrum Karlsruhe GmbH, Karlsruhe, 2005).
- [49] R. Ahmad, B. Yin, Z. Wu, and W. Curtin, Designing high ductility in magnesium alloys, *Acta Mater.* **172**, 161 (2019).
- [50] P. Tzanetakis, J. Hillairet, and G. Revel, The formation energy of vacancies in aluminium and magnesium, *Physica Status Solidi (b)* **75**, 433 (1976).
- [51] A. Togo, First-principles phonon calculations with Phonopy and Phono3py, *J. Phys. Soc. Jpn.* **92**, 012001 (2023).
- [52] A. Togo, L. Chaput, T. Tadano, and I. Tanaka, Implementation strategies in phonopy and phono3py, *J. Phys.: Condens. Matter* **35**, 353001 (2023).
- [53] Nd Crystal Structure: Datasheet from "PAULING FILE Multinaries Edition – 2022" in SpringerMaterials.
- [54] K. A. Gschneidner, Jr. and L. Eyrin, *Handbook on the Physics and Chemistry of Rare Earths* (Elsevier/North-Holland Publishing, Amsterdam, 1978).
- [55] D. R. Trinkle, Diffusivity and derivatives for interstitial solutes: Activation energy, volume, and elastodiffusion tensors, *Philos. Mag.* **96**, 2714 (2016).
- [56] D. R. Trinkle, Automatic numerical evaluation of vacancy-mediated transport for arbitrary crystals: Onsager coefficients in the dilute limit using a Green function approach, *Philos. Mag.* **97**, 2514 (2017).
- [57] D. R. Trinkle, Variational principle for mass transport, *Phys. Rev. Lett.* **121**, 235901 (2018).
- [58] H. Jónsson, G. Mills, and K. W. Jacobsen, in *Classical and Quantum Dynamics in Condensed Phase Simulations* (World Scientific, Singapore, 1998), pp. 385–404.
- [59] G. H. Vineyard, Frequency factors and isotope effects in solid state rate processes, *J. Phys. Chem. Solids* **3**, 121 (1957).
- [60] R. Agarwal and D. R. Trinkle, Exact model of vacancy-mediated solute transport in magnesium, *Phys. Rev. Lett.* **118**, 105901 (2017).
- [61] M. Paliwal, S. K. Das, J. Kim, and I.-H. Jung, Diffusion of Nd in hcp Mg and interdiffusion coefficients in Mg–Nd system, *Scr. Mater.* **108**, 11 (2015).
- [62] Y. Lysogorskiy, A. Bochkarev, M. Mrovec, and R. Drautz, Active learning strategies for atomic cluster expansion models, *Phys. Rev. Mater.* **7**, 043801 (2023).
- [63] Y. Guo, B. Liu, W. Xie, Q. Luo, and Q. Li, Anti-phase boundary energy of β series precipitates in Mg-Y-Nd system, *Scr. Mater.* **193**, 127 (2021).
- [64] S. DeWitt, E. L. Solomon, A. R. Natarajan, V. Araullo-Peters, S. Rudraraju, L. K. Aagesen, B. Puchala, E. A. Marquis, A. Van Der Ven, K. Thornton, and J. E. Allison, Misfit-driven β'' precipitate composition and morphology in Mg-Nd alloys, *Acta Mater.* **136**, 378 (2017).
- [65] D. Choudhuri, R. Banerjee, and S. G. Srinivasan, Interfacial structures and energetics of the strengthening precipitate phase in creep-resistant Mg-Nd-based alloys, *Sci. Rep.* **7**, 40540 (2017).
- [66] H. Liu, Y. Gao, J. Liu, Y. Zhu, Y. Wang, and J. Nie, A simulation study of the shape of β' precipitates in Mg–Y and Mg–Gd alloys, *Acta Mater.* **61**, 453 (2013).

- [67] E. L. Solomon, V. Araullo-Peters, J. E. Allison, and E. A. Marquis, Early precipitate morphologies in Mg-Nd-(Zr) alloys, *Scr. Mater.* **128**, 14 (2017).
- [68] J. M. Meier, J. Miao, L. DeBeer-Schmitt, J. Ilavsky, and A. A. Luo, Stimulating β -series precipitation in Mg-Nd alloys via microalloying: A comparison of electron microscopy and small-angle scattering techniques, *Metall. Mater. Trans. A* **56**, 914 (2025).
- [69] H. Liu, Y. Gao, Y. Zhu, Y. Wang, and J. Nie, A simulation study of β 1 precipitation on dislocations in an Mg-rare earth alloy, *Acta Mater.* **77**, 133 (2014).
- [70] P. J. Steinhardt, D. R. Nelson, and M. Ronchetti, Bond-orientational order in liquids and glasses, *Phys. Rev. B* **28**, 784 (1983).
- [71] C. Sigli, F. De Geuser, A. Deschamps, J. Lépinoux, and M. Perez, Recent advances in the metallurgy of aluminum alloys. Part II: Age hardening, *Comptes Rendus. Physique* **19**, 688 (2018).
- [72] L. Bourgeois, Y. Zhang, Z. Zhang, Y. Chen, and N. V. Medhekar, Transforming solid-state precipitates via excess vacancies, *Nat. Commun.* **11**, 1248 (2020).
- [73] L. Ding, F. J. H. Ehlers, R. Hu, Z. Zhang, H. Nagaum, C. Hutchinson, Q. Liu, and Z. Jia, On the order-disorder transformation within a main hardening precipitate in Al-Mg-Si alloys, *Philos. Mag.* **105**, 169 (2025).
- [74] M. Li and X. Min, Origin of ω -phase formation in metastable β -type Ti-Mo alloys: Cluster structure and stacking fault, *Sci. Rep.* **10**, 8664 (2020).
- [75] X. Fu, X.-D. Wang, B. Zhao, Q. Zhang, S. Sun, J.-J. Wang, W. Zhang, L. Gu, Y. Zhang, *et al.*, Atomic-scale observation of non-classical nucleation-mediated phase transformation in a titanium alloy, *Nat. Mater.* **21**, 290 (2022).
- [76] Y. L. Müller and A. R. Natarajan, First-principles thermodynamics of precipitation in aluminum-containing refractory alloys, *Acta Mater.* **274**, 119995 (2024).
- [77] S. P. Ong, W. D. Richards, A. Jain, G. Hautier, M. Kocher, S. Cholia, D. Gunter, V. L. Chevrier, K. A. Persson, and G. Ceder, Python Materials Genomics (pymatgen): A robust, open-source Python library for materials analysis, *Comput. Mater. Sci.* **68**, 314 (2013).
- [78] A. H. Larsen, J. J. Mortensen, J. Blomqvist, I. E. Castelli, R. Christensen, M. Dułak, J. Friis, M. N. Groves, B. Hammer, C. Hargus, *et al.*, The atomic simulation environment—A Python library for working with atoms, *J. Phys.: Condens. Matter* **29**, 273002 (2017).
- [79] A. Togo, K. Shinohara, and I. Tanaka, Spglib: A software library for crystal symmetry search, *Sci. Technol. Adv. Mater.: Methods* **4**, 2384822 (2024).
- [80] The reference databases for Mg, Nd, and Mg-Nd used in the present work are available through the Materials Cloud, <https://doi.org/10.24435/materialscloud:1w-f4>.