



Accelerating training of MLIPs through small-cell training

Jason A. Meziere^{1,a)} , Yu Luo², Yi Xia³, Laurent Karim Béland², Mark R. Daymond², Gus L. W. Hart¹

¹Department of Physics, Brigham Young University, Provo, UT 84602, USA

²Department of Mechanical and Materials Engineering, Queen's University, Kingston, ON K7L 2N8, Canada

³Department of Mechanical and Materials Engineering, Portland State University, Portland, OR 97201, USA

^{a)}Address all correspondence to this author. e-mail: jmezier2@byu.edu

Received: 31 January 2023; accepted: 5 October 2023

While machine-learned interatomic potentials have become a mainstay for modeling materials, designing training sets that lead to robust potentials is challenging. Automated methods, such as active learning and on-the-fly learning, construct reliable training sets, but these processes can be resource-intensive. Current training approaches often use density functional theory calculations that have the same cell size as the simulations that the potential is explicitly trained to model. Here, we demonstrate an easy-to-implement small-cell training protocol and use it to model the Zr-H system. This training leads to a potential that accurately predicts known stable Zr-H phases and reproduces the α - β pure zirconium phase transition in molecular dynamics simulations. Compared to traditional active learning, small-cell training decreased the training time of the α - β zirconium phase transition by approximately 20 times. The potential describes the phase transition with a degree of accuracy similar to that of the large-cell training method.

Introduction

In the past decade, machine-learned interatomic potentials (MLIPs) have become an indispensable tool for modeling materials because MLIPs expand the length and time scales that can be modeled. These potentials fill the gap between density functional theory (DFT) and fast, physically motivated models such as embedded atom model (EAM) interatomic potentials [1, 2].

MLIPs are created by fitting flexible, parameter-rich models to DFT data [3–6]. These potentials can reproduce energies, forces, and stresses for structures nearly as accurately as DFT, but are orders of magnitude faster [7]. MLIPs simulate longer time scales than possible with DFT, successfully modeling material properties such as crack propagation and phase transitions or generating phase diagrams [8, 9]. On the other hand, EAMs have a simpler, physically motivated, mathematical form [10, 11], which makes them faster and more robust than MLIPs, but far less flexible. EAMs have tens of parameters while MLIPs can have thousands or tens of thousands. EAMs do not reproduce diverse chemical behavior because of their small number of parameters [12]. While EAM potentials have been able to model complex phenomena (e.g., in the Zr-H system [13–15]),

MLIPs can more accurately model material properties such as phonon bandstructures, elastic constants, point defects, and planar defects [12, 16].

Despite MLIPs' successes, creating accurate and robust potentials is still challenging [17]. A primary difficulty in creating an effective MLIP is determining a sufficient, but efficient, training set [8, 18]. In order for the MLIP to accurately predict a structure's energetics, the training set must contain atomic environments representative of that structure, but knowing which structures the MLIP will need to model before the training process begins is impossible. A *priori* construction of a training set relies on intuition rather than mathematical formalism and cannot guarantee that complex material properties will be modeled accurately.

There are two common approaches that leverage iterative, rather than *priori* construction, of the training set: active learning and on-the-fly learning [19–21]. Both methods use a metric that indicates the level of confidence in the MLIP's predictions. If the level of confidence in a prediction is insufficient, energies and forces of that structure are calculated with DFT and added to the training set. In this way, active learning and on-the-fly




Figure 1: The α , δ , ϵ , and γ phases are common pure Zr or Zr-H systems that have been experimentally measured. Green atoms are zirconium atoms and red atoms are hydrogen atoms.

TABLE 1: Relaxed structure properties for the primitive α , δ , ϵ , and γ phases.

	a (Å)	b (Å)	c (Å)	α (deg.)	β (deg.)	γ (deg.)
α (VASP)	3.230	3.230	5.167	90.0	90.0	120.0
α (MTP)	3.228	3.228	5.171	90.0	90.0	120.0
δ (VASP)	3.395	3.395	3.395	60.0	60.0	60.0
δ (MTP)	3.391	3.391	3.391	60.0	60.0	60.0
ϵ (VASP)	3.331	3.331	3.331	116.0	116.0	97.1
ϵ (MTP)	3.332	3.332	3.332	115.9	115.9	97.2
γ (VASP)	3.236	3.236	4.996	90.0	90.0	90.0
γ (MTP)	3.226	3.226	5.020	90.0	90.0	90.0

Graphically, these phases are shown in Fig. 1.

learning iteratively improve the training set until the resulting potential reproduces the material properties of interest [17, 22].

The philosophy of active learning differs from on-the-fly learning. On-the-fly learning continues trying to improve the potential whenever in use, so the potential is never “finished”. Active learning, on the other hand, uses a predefined set of simulations to conduct training. Once the potential successfully finishes all simulations in the predefined set without triggering more DFT calculations, the potential is considered reliable and training ends.

While active learning robustly provides training sets that sufficiently represent desired material properties, typical training approaches are inefficient, using large-cell structures to train potentials that model large-cell material properties (material properties that only manifest in simulations using large unit cells containing many atoms). We demonstrate that a protocol which prioritizes training on small-cell structures allows information gained from the small-cell structure training to be used when modeling the larger-cell structures. Other recent work has also suggested the small-cell training approach [23], and below we outline a systematic approach.

We demonstrate small-cell training with the moment tensor potential (MTP) as implemented in the MLIP-2 package, which includes active learning [3]. MTPs use an invariant polynomial basis to represent atomic densities. MLIP-2 employs the maximal volume (MV) grade [17, 22] to determine if a configuration

should be added to the training set or not. While we use the MTP/MLIP-2 framework, the methods we describe can be applied to other MLIPs that utilize active learning (Table 1).

Zr-H system

Many of zirconium’s industrial uses are related to nuclear power applications because of its low neutron cross-section. Most current commercial reactor designs include fuel cladding made of zirconium alloys. Furthermore, Canadian Deuterium Uranium (CANDU) reactors employ zirconium alloys as a structural material; CANDU pressure tubes and calandria tubes are made of zirconium alloys. There are two allotropic forms of pure zirconium, α -Zr (HCP) and β -Zr (BCC), which is the high temperature phase. Water chemistry on the primary side (reactor core) and indeed on the secondary side (steam generation) of nuclear power plants is typically controlled to provide highly reducing chemical conditions to favor the formation of protective metal oxide layers. Therefore, oxidation reaction in metals (anode) exposed to the water will be electrochemically balanced (cathode) by reduction of water and generation of hydrogen. Some fraction of this hydrogen is picked up by the zirconium alloys [24]. At typical reactor operating temperatures ($\sim 573\text{K}$) and hydrogen concentrations, the hydrogen will be in solid solution with the zirconium alloy. However, as reactors are powered down, the temperature decreases, which reduces the solubility

limit of hydrogen. This may lead to precipitation of zirconium hydrides, causing embrittlement and possible failure of the material, which is particularly a concern during storage of spent-fuel bundles [25]. Further, hydrogen diffuses ‘up’ hydrostatic tensile stress gradients [26] leading to elevated concentrations at notches or defects, and hence to the slow-cracking phenomenon of delayed hydride cracking [27], an issue of concern for nuclear power plant operators. Therefore, understanding how hydrogen pickup and hydride formation occur in zirconium are considered important for nuclear reactor safety [28, 29].

The crystal structure of the hydrides precipitated in zirconium is dependent on hydrogen concentration. At the relatively low hydrogen concentrations found in operational reactors, face-centered cubic δ -hydride is the most stable and most frequently observed hydride phase in most commercial zirconium alloys [30]. Meanwhile, face-centered tetragonal γ -hydrides are sometimes observed experimentally in commercial alloys [31] and are commonly seen in pure zirconium [32]. At higher hydrogen concentrations, face-centered tetragonal ϵ -hydrides are also reported [33]. These phases are shown in Fig. 1.

The specific case of phase transformation of zirconium hydrides is of particular interest. Despite decades of research, there is no scientific consensus as to the path of nucleation and growth mechanisms and the role that different hydride phases may play in the underlying precipitation of hydrides, with discussions dating to the 1970s [34] and continuing today [35, 36]. Further, experimental evidence strongly suggests that the pathway (whether thermodynamically or kinetically driven) for formation of γ versus δ hydride phases is finely balanced, with small changes sufficient to favor one over the other [32, 37, 38]. Yet, the relative likelihood as to which phase forms has practical application in predicting performance in-reactor due to the phases’ quite different volume expansions compared to zirconium [39] and hence impact on component behavior. From an atomistic simulation perspective, no current classical interatomic interaction potential is able to describe the structure and energetics of the multiple Zr-H phases, which limits atom-scale studies of this system to the small length and time scales accessible to DFT, such as in Refs. [40–48]. The development of a classical potential able to capture the main Zr-H phases is an important step towards a comprehensive understanding of the formation thermodynamics and kinetics of zirconium hydrides.

Small-cell structure training protocol

All active learning methods are iterative in nature. Training structures are systematically added, based on actual or predicted error, to improve the accuracy of potential. During the training process, the potential is tested for a variety of situations, errors assessed, and training structures that contribute the most new information are calculated with DFT and added to the training

set. This cycle continues until the potential passes all required robustness tests. In situations where the potential will be used for large-scale simulations (say more than 100 atoms), such as modeling structural phase transitions, large-cell structures are typically included in the training process. In contrast, we demonstrate that a reliable potential is possible with training sets that contain only small-cell structures.

In the small-cell approach, the first training set contains primitive unit cells of experimentally relevant phases. It is usually helpful to include a few other configurations to provide basic physics for the training. For example, primitive cells at several volumes provide basic equation of state information.

In the second step, slightly larger unit cells, modifications of the experimentally relevant phases with varying concentrations, are tested with the initial potential. These can be systematically enumerated by cell size [49, 50], providing a methodical approach to choosing the next set of structures to test the potential with. Structures that are not accurately handled by the potential are calculated in DFT, added to the training set, and the potential refitted. The potential is then tested on the enumerated structures of the next-largest cell size. Training stops when the potential effectively handles all the new structures in the added set.

In the case of MTP, a convenient test for the potential is to minimize the energy of a test structure with respect to its unit cell shape, size, and atomic positions. Training stops when the predicted error is sufficiently low for the entire structural “relaxation trajectory” for all of the enumerated structures in a test set. How to test (and improve) this potential (if necessary) for large-scale simulations will be demonstrated below.

We initially developed a Zr-H potential to search for thermodynamically stable structures over a large concentration range of hydrogen. Knowing the stable structures of the Zr-H system is a first step for understanding the energetic landscape of possible structures that occur in a nuclear environment. Currently, the structures considered to be thermodynamically stable at room temperature in the Zr-H system are the pure zirconium α phase, the γ (Zr-H) phase, the δ (Zr-H_{1.6–1.7}) phase and the Zr-H₂ ϵ phase. There is also a metastable β (pure zirconium) phase [37, 51]. While the ambient metastability of β -Zr is not in question, there continues to be extensive discussion in the literature as to the relative stability of the δ and γ phases when the H to Zr ratio is smaller than 1.0, with phase formation experimentally dependent on thermal history as well as the properties (e.g., dislocation density and minority chemistry) of the zirconium, as reviewed in [32]; we note in particular that for hydrides formed from high purity Zr when the H to Zr ratio is smaller than 1.0, γ is stable at room temperature, while from even somewhat impure Zr it is δ which appears stable [32]. The most widely accepted phase diagram for the pure ZrH binary [52] reports γ as a room temperature stable phase when the H to

Zr ratio is smaller than 1.0 and the temperature is below 286 °C. δ is considered stable at higher H to Zr ratios and at higher temperatures; this is in contrast to earlier phase diagrams [53].

Stability is assessed by comparing the formation enthalpy (at $T = 0$ K) of a candidate structure with the formation enthalpy of structures that are known experimentally. Structures that are observed in experiment will lie on or near the lower convex hull of formation enthalpy formed by known stable structures [54–56]. To increase the chances of finding new stable structures, the energy of candidate structures is minimized with respect to atom position and cell shape, a process called relaxation [57].

An efficient way to generate candidates for stability is to consider a few structures that are known to be stable and vary the hydrogen concentration ranges by creating vacancies. To create a set of candidate structures for training, we enumerate configurations [58] of the δ , γ , and ϵ phases while varying hydrogen concentration. Systematically enumerating possible structures is straightforward, but the number of these structures grows exponentially with cell size. Despite this large increase in the amount of information the potential must successfully model, small-cell training leads to a training set that efficiently represents these candidate structures and their relaxation trajectories.

Our iterative small-cell training began by applying active learning on both the relaxation trajectories of the primitive structures of the δ , γ , and ϵ phases with varying hydrogen concentration and the relaxation trajectories of structures needed to compute the elastic constants of the δ , γ , ϵ , and α structure. The elastic constant structures were added to the initial set to create a more practical potential and to give the potential information about the α phase. As shown in Fig. 2(a), training with active learning required around 300 structures to be added to the training set before the potential could successfully relax the set of primitive structures and elastic constant structures. Once the MLIP could relax these structures without the active learning protocol requesting additional DFT calculations, training began on structures with the next smallest supercells. DFT calculations were done in VASP [59, 60] using the parameters shown in Table 2. The process was repeated for successively larger cells until the MLIP could successfully relax all of the enumerated structures. After finishing training, the potential could successfully relax 75,032 structures while only having 584 structures in the training set [Fig. 2(a)]. Additionally, 95% of the structures in the training set had less than seven atoms.

As training continued with larger supercells, the number of required DFT calculations decreased rapidly. After three increases in supercell size, the number of DFT calculations needed at each subsequent increase became nearly zero. The size distribution of the training set [Fig. 2(b)] shows that by the fourth increase, 95% of the necessary DFT data had already

TABLE 2: Parameters used in all VASP calculations.

Parameter	Value
PREC	Accurate
ENCUT	400
NELMIN	2
NELM	30
ISMEAR	0
SIGMA	0.05
ISIF	3
EDIFF	1.0E-06
EDIFFG	-0.01
VOSKOWN	1
NBLOCK	1
NWRITE	1
NELM	60
ALGO	Normal
ISPIN	1
INIWAV	1
ISTART	0
ICHARG	2
LWAVE	.FALSE.
LCHARG	.FALSE.
ADDGRID	.FALSE.
LREAL	.FALSE.
NCORE	1

been collected. Even though only a small number of structures are added to the training set during training with large supercells, the number of candidate structures that the potential successfully predicts continues to increase exponentially. This shows that prioritizing small-cell structures allows for most of the necessary information to be gathered in cheap, small-cell DFT calculations that give relevant information about the larger-cell structures, significantly decreasing the total training time.

Additionally, using small-cell training still effectively represented the known convex hull well. The structures that sit on the convex hull are the α and ϵ phase, which matches DFT benchmarks. The γ phase is also well represented, lying just above the convex hull. Small-cell training also successfully found candidates for stable structures at higher temperatures. Figure 3 shows structures lying close to the convex hull at $Zr_1H_{1.5}$.

Small-cell training for large-scale simulations

When searching for stable structures of the Zr-H system, candidate structures for stability could be ordered by size to ensure that training starts with the small-cell structures, making small-cell training suitable. In contrast, a large-scale simulation can require hundreds or thousands of atoms for the physical

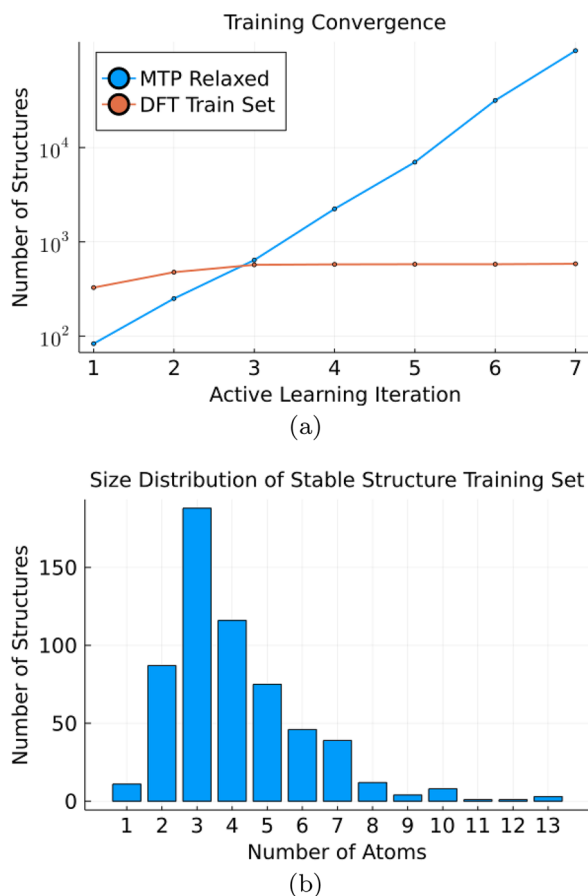


Figure 2: (a) Early in training, the number of structures needed for the training set is significantly higher than the number of structures the MLIP is predicting on. As training progresses to larger supercells, however, the need for additional structures becomes negligible. (b) Thus, the size of the structures in the training set are dominated by small-cell structures.




Figure 3: Small-cell training accurately reproduces the known convex hull of the Zr-H system. The δ and ϵ phases are stable with the γ phase lying just above the convex hull. Small-cell training also effectively explores candidate structures for stability, finding low enthalpy structures that approach the convex hull at $Zr_1H_{1.5}$.

phenomena to be apparent, making small-cell training seem inappropriate. However, small-cell training can be used to create potentials that model large-scale simulations if we train the potential with a sequence of simulations that increase in cell size and converge to the full simulation.

For example, a phase transition simulation is a large-scale simulation that often requires simulation cells with several hundred atoms. However, instead of starting training with a large-cell structure with hundreds of atoms, the primitive structure is used. When the potential can successfully finish this small-cell phase transition simulation without further DFT requests, the supercell size is increased to twice the primitive. Training progresses with larger and larger supercells until the potential can successfully capture the phase transition in a large-scale simulation.

From a physical perspective, this method is counter-intuitive. A phase transition simulation of only a few atoms cannot capture the phase transition and does not effectively model experimental conditions. From the perspective of taking a limit, however, the idea is reasonable. As training progresses with larger and larger supercells, the simulations used for training will more accurately approximate the large-scale simulation. Still, it is unclear how effective small-cell training is when some of the simulations used in training do not accurately model reality.

To assess the effectiveness of small-cell training for large-scale simulations, we consider the zirconium α - β phase transition. The α - β phase transformation of pure zirconium is well characterised and known experimentally, however the transformation is strongly affected by alloying elements [61]. Understanding the impact of alloying on this phase transformation is important both in terms of optimising alloying processing and to understand component integrity during reactor loss-of-coolant accidents (LOCA); future work could take the approach described here and apply it to determine the impact of alloying elements on phase transformation.

Because both the α - β phase transition and the Zr-H system are important for modeling zirconium in nuclear environments, the earlier potential developed to explore stable structures of the Zr-H system was further trained for modeling the α - β phase transition. In the first iteration of small-cell training, a two-atom heating simulation of α phase and a two-atom cooling simulation of β phase were performed in LAMMPS [62] with active learning. After these simulations ran successfully without requiring additional DFT calculations, small-cell training moved to four-atom simulations for both the heating and cooling simulations. The size of the supercells increased until the potential could model the full-scale phase transition simulation, finding that the potential captured the phase transition after training on molecular dynamics temperature ramps containing only eight atoms [Fig. 4(a)]. Figure 4(b) shows the effect

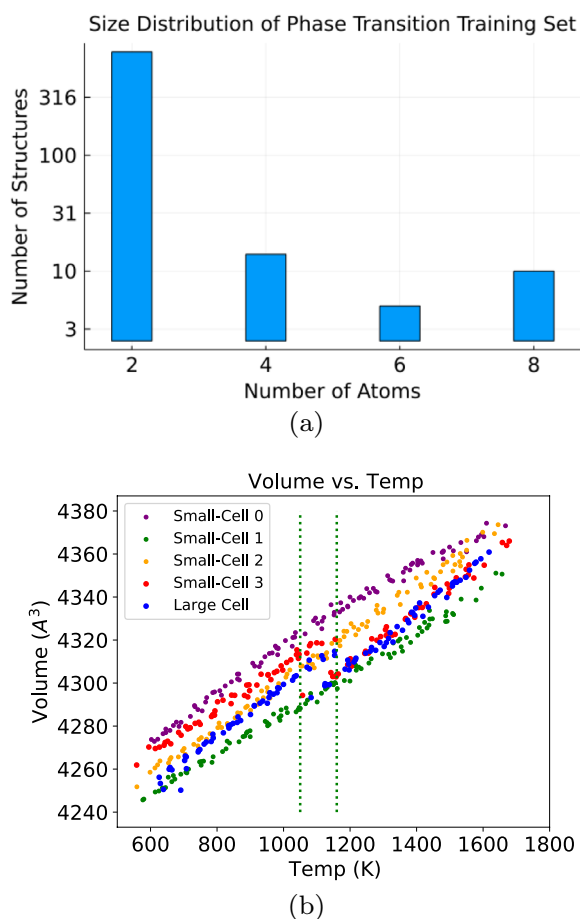


Figure 4: (a) The size distribution of the training set for phase transition training is dominated by two atom structures. While structures of up to eight atoms are used, these are a small subset of the total training set. (b) As Small-Cell 0 shows, the small-cell trained potential initially fails to reproduce the phase transition or the thermal expansion of the α or β phase (at temperatures above the correct phase transition temperature) that the large-cell trained potential (Large-Cell) successfully models. However, subsequent iterations of small-cell training first model the volume and thermal expansion correctly (Small-Cell 1 and Small-Cell 2) and finally are able to model the phase transition (Small-Cell 3). The vertical lines show that the phase transition temperature is in the same range for both the Small-Cell 3 potential and the Large-Cell potential.

of small-cell training on the potential. The potential initially fails to reproduce the phase transition and the thermal expansion of the α and β phases for temperatures above the phase transition temperatures (Small-Cell 0). However, as further iterations of small-cell training are completed, the thermal expansion of α (Small-Cell 1) and β (Small-Cell 2) are successfully modeled. After training is finished, the potential is able to model both the phase transition and the thermal expansion of the α and β phases (Small-Cell 3). DFT calculations were performed in VASP with the same parameters used earlier, as defined in Table 2.

It may be surprising that the potential can model the phase transition despite the training set containing only small

structures. A simulation with only eight atoms cannot model the α - β phase transition, even with a perfectly accurate interatomic potential, due to finite-size effects. However, these small-cell simulations provide a large and diverse sample of atomic environments similar to those that would be encountered by the large-scale phase transition simulation. After training, the local environments in the training set are diverse enough that the large-scale phase transition simulation is unlikely to encounter an atomic environment that is not well represented in the training set.

To evaluate the accuracy of the phase transition, we compare it to a second MTP that accurately modeled the phase transition but trained with large supercells [63]. The initial state of the training set of the second potential included large-cells with random atomic displacements. Then, the second potential was trained directly on the full-scale simulation with an active learning strategy. After three iterations of training, the MTP successfully modeled the phase transition between α and β phases of zirconium. Convergence was achieved within four iterations, giving rise to a training set containing a total number of 653 structures, which consist of 330 structures of α phase (54 atoms per structure) and 323 structures of β phase (48 atoms per structure), respectively.¹

When compared with the MTP trained with large supercells, the phase transition of our MTP has good agreement. The phase transition temperature is in the same range for both and the thermal expansion coefficient for the β phase also agrees. The thermal expansion coefficient for α phase is slightly lower with our MLIP, but for a given temperature, the volume is accurate to within 1% (Fig. 4). In addition to retaining the accuracy of the MLIP trained on large-cells, our small-cell trained potential also had a large reduction in training time.

The total cpu time used for DFT calculations during the training of our potential was 11.5 times less than used for the large-cell training. However, this speedup includes calculations from training for both the phase transition and when searching for stable Zr-H structures. When restricted to cpu time spent for training on the phase transition, the cpu time for DFT calculations is 17.5 times less than large-cell training.

¹ The parameters used for the VASP calculations of this training set were not the same as the parameters shown in Table 2. However, the only significant parameter change from the parameters in Table 2 was a change in ENCUT from 400 eV to 500 eV. Because 400 eV is a reasonable value of ENCUT for the Zr system, this should not affect the training results significantly.

Zr-H potential analysis

The Zr-H potential developed in this paper was trained to model a diverse set of candidates for thermodynamic stability, a few elemental structures, and the pure zirconium α - β phase transition. However, we wish to evaluate the potential's ability to model several other material properties such as the relaxed state, phonon bandstructure, and elastic properties of the α , δ , ϵ , and γ phases. All DFT predictions were calculated with the VASP parameters shown in Table 2.

The predictions of the relaxed state of the α , δ , ϵ , and γ phases is depicted in Table 1. These predictions are highly accurate, with less than a 1% relative error for all predictions. Given that the relaxation trajectories were included in the training when searching for thermodynamically stable structures, this may not be surprising. However, it does increase the confidence that the predictions of the convex hull are correct for these phases.

Using pymatgen [65], predictions of the elastic properties of the four phases were also reproduced with reasonable accuracy.

Tables 3 and 4 show that the elastic constants and other elastic properties of the α and γ phases are reproduced well. However, the potential fails to reproduce elastic properties for the δ and ϵ phases to high accuracy, often exceeding 30 GPa in absolute error.

As Fig. 5 shows, the potential is able to reproduce the general trend of the DFT predicted phonon bandstructures. However, for the ϵ phase, the acoustic branch fails to recreate an imaginary mode instability that DFT predicts. Additionally, the optical branches of the phonon bandstructures are underestimated by the MTP.

Ignoring configurational entropy, the potential also predicts the existence of the $Zr_1H_2\delta$ to $Zr_2H_2\gamma$ phase transition using the quasi-harmonic approximation for Gibbs free energy (Fig. 6). However, the predicted MTP phase transition temperature is inconsistent with DFT, estimating a phase transition temperature of ~ 150 K higher. Both the phonon bandstructures and estimated phase transition temperature were calculated with phonopy [66].

TABLE 3: Elastic constants for α , δ , ϵ , and γ phases.

	C_{11}	C_{22}	C_{33}	C_{44}	C_{55}	C_{66}	C_{12}	C_{13}	C_{23}
α (VASP)	147.6	146.9	163.3	23.5	23.5	29.4	62.8	65.8	66.2
α (MTP)	142.0	142.0	165.3	28.7	28.7	28.6	84.8	66.4	66.4
δ (VASP)	73.2	74.0	75.0	6.67	16.3	7.72	168.1	166.3	166.9
δ (MTP)	116.1	116.1	127.6	-11.7	-11.7	-5.97	116.5	105.0	105.0
ϵ (VASP)	190.2	186.0	168.1	29.0	41.0	29.4	104.3	120.6	110.7
ϵ (MTP)	207.1	197.7	266.8	39.9	41.4	40.4	133.4	132.4	131.3
γ (VASP)	193.0	193.0	179.1	44.8	44.8	0.43	61.8	88.2	88.2
γ (MTP)	167.8	167.8	187.8	48.0	48.0	-2.89	57.8	105.4	105.4

Elements of the elastic tensor are in units of GPa.

TABLE 4: Elastic properties for α , δ , ϵ , and γ phases.

	G_V	G_R	G_{VRH}	K_V	K_R	K_{VRH}	ν (unitless)	E
α (VASP)	32.8	30.3	31.6	94.5	94.2	94.3	0.35	85.2
α (MTP)	32.7	31.2	31.9	98.3	98.3	98.3	0.35	86.4
δ (VASP)	-12.5	2.24	-5.12	135.7	135.7	135.7	0.52	-15.5
δ (MTP)	-3.67	24.9	10.6	112.5	112.5	112.5	0.45	30.9
ϵ (VASP)	33.8	30.0	31.9	136.1	129.0	132.5	0.39	88.6
ϵ (MTP)	42.6	37.6	40.1	162.9	151.4	157.1	0.38	110.9
γ (VASP)	39.8	2.07	20.9	118.1	118.0	118.1	0.42	59.3
γ (MTP)	35.6	-19.9	7.83	117.9	112.2	115.0	0.47	23.0

Unless otherwise noted, reported elastic properties are in units of GPa. G and K represent bulk modulus and shear modulus, respectively, while V , R , and H represent Voigt, Reuss, and Hill, respectively. ν is Poisson's ratio and E is Young's modulus.

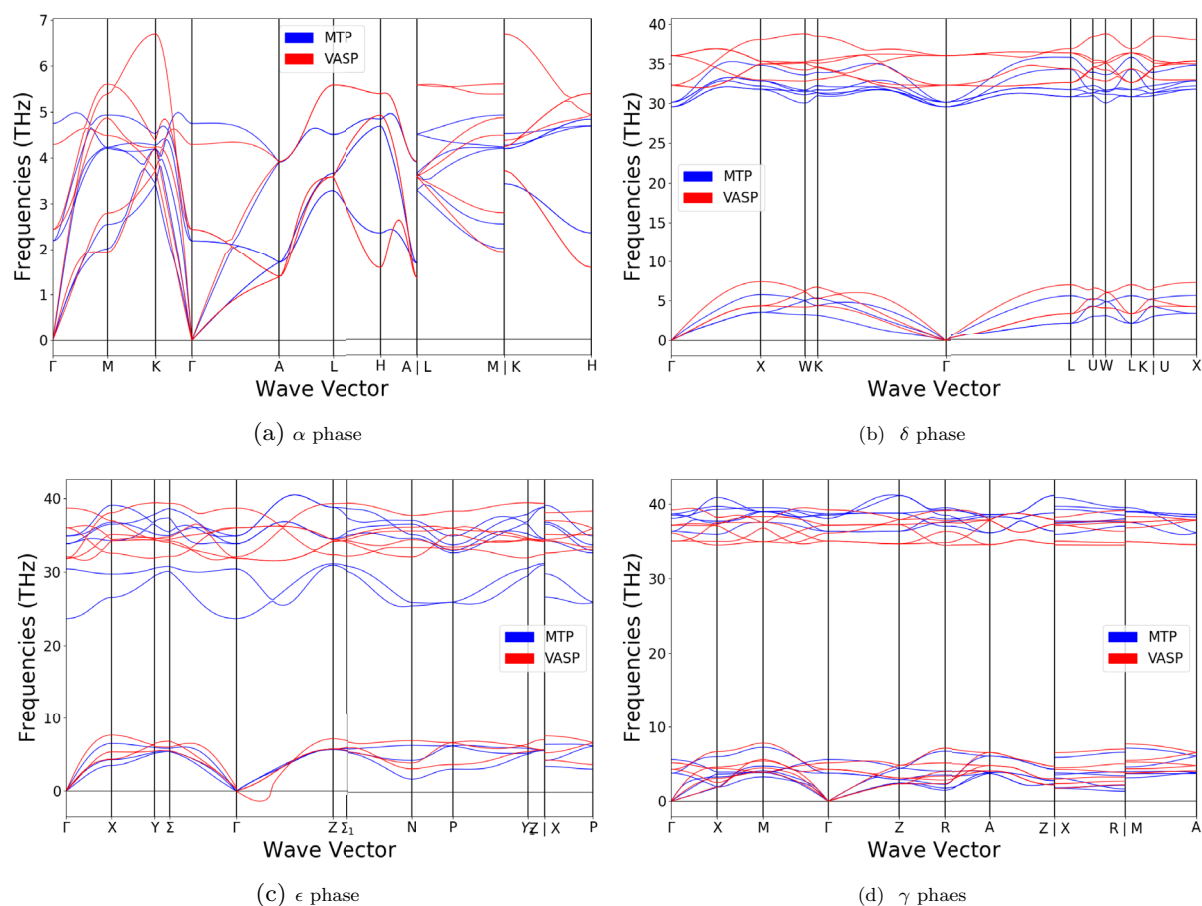


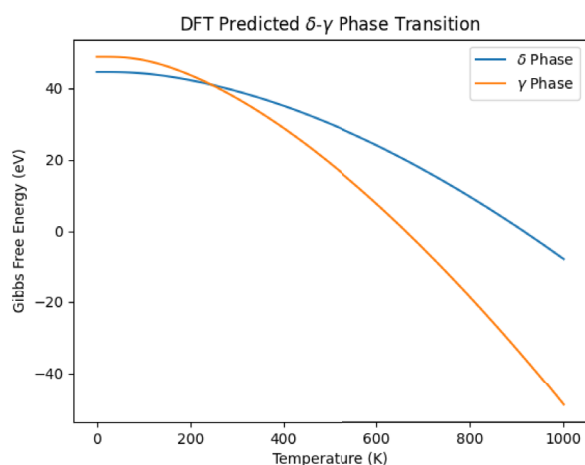
Figure 5: Phonon band structures for α , δ , ϵ , and γ phases.

Conclusion

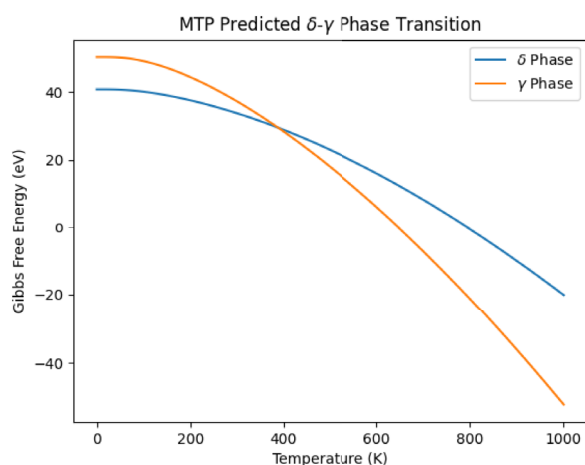
While active learning presents a methodical way to create a training set for complex material properties, it is often inefficient. Our small-cell training protocol allows inexpensive, small-cell DFT calculations to give information about large-cell structures with reasonable accuracy, avoiding computationally expensive calculations with larger supercells.

Small-cell training has the potential to improve other active learning approaches as well. Other active learning methods can incorporate small-cell training into their frameworks. Small-cell training only changes large simulations to small simulations when training, making it indifferent to the type of active learning used.

Small-cell training can effectively create a robust and comprehensive potential for modeling complex material properties in the Zr-H system. Our potential can model the relaxed state, elastic constants, and phonon bandstructures of the α phase and the hydride ϵ , δ , and γ phases as well as several candidates for stable structures at higher temperatures. Additionally, this potential can accurately model the α - β phase transition, reproducing the thermal expansion coefficients and the phase transition temperature. Small-cell training also models these material properties more efficiently than typical active learning methods alone. When compared with a potential trained on large supercells, training took up to 17.5 times longer.



(a) DFT Prediction



(b) MTP Prediction

Figure 6: Predictions of the δ - γ phase transition temperature using the quasi-harmonic approximation for Gibbs free energy.

Author contributions

JM: Conceptualization, Methodology, Formal analysis and investigation, Writing—original draft preparation, Writing—review and editing. YL: Formal analysis and investigation, Writing—review and editing. YX: Formal analysis and investigation, Writing—review and editing. LKB: Formal analysis and investigation, Writing—review and editing. MD: Formal Analysis and investigation, Writing—review and edition, Funding acquisition. GLWH: Conceptualization, Methodology, Formal analysis and investigation, Writing—review and edition, Funding acquisition, Resources, Supervision.

Funding

This work was funded by the Advanced Materials Simulation Engineering Tool (AMSET) project, sponsored by the US Naval Nuclear Laboratory (NNL) and directed by Materials Design,

Inc. YL, MRD, and LKB thank the Digital Research Alliance of Canada (formerly known as Compute Canada) for generous allocation of compute resources, and the Natural Sciences and Engineering Research Council of Canada (NSERC) and the NSERC/UNENE Industrial Research Chair in Nuclear Materials at Queen's for financial support.

Data availability

The moment tensor potential file (pot.mtp) and the training set file (train.cfg) are available at <https://github.com/jmeziere/Zr-H-Potential>.

Declarations

Conflict of interest On behalf of all authors, the corresponding author states that there is no conflict of interest.

References

1. V.L. Deringer, M.A. Caro, G. Csányi, *Adv. Mater.* **31**, 1902765 (2019)
2. Y. Lysogorskiy, C.V.D. Oord, A. Bochkarev, S. Menon, M. Rinaldi, T. Hammerschmidt, M. Mrovec, A. Thompson, G. Csányi, C. Ortner, R. Drautz, *npj Comput. Mater.* **7**, 97 (2021). <https://doi.org/10.1038/s41524-021-00559-9>
3. A.V. Shapeev, *Multisc. Model. Simul.* **14**, 1153 (2016). <https://doi.org/10.1137/15M1054183>
4. J. Behler, *Phys. Chem. Chem. Phys.* **13**, 17930 (2011). <https://doi.org/10.1039/C1CP21668F>
5. W.J. Szlachta, A.P. Bartók, G. Csányi, *Phys. Rev. B* **90**, 104108 (2014)
6. A.P. Bartók, M.C. Payne, R. Kondor, G. Csányi, *Phys. Rev. Lett.* **104**, 136403 (2010)
7. J.A. Vita, D.R. Trinkle, *Comput. Mater. Sci.* **200**, 110752 (2021)
8. A.P. Bartók, R. Kondor, G. Csányi, *Phys. Rev. B* **87**, 184115 (2013)
9. C.W. Rosenbrock, K. Gubaev, A.V. Shapeev, L.B. Pártay, N. Bernstein, G. Csányi, G.L.W. Hart, *npj Comput. Mater.* **7**, 1 (2021). <https://doi.org/10.1038/s41524-020-00477-2>
10. M.S. Daw, M.I. Baskes, *Phys. Rev. Lett.* **50**, 1285 (1983). <https://doi.org/10.1103/PhysRevLett.50.1285>
11. M.S. Daw, M.I. Baskes, *Phys. Rev. B* **29**, 6443 (1984). <https://doi.org/10.1103/PhysRevB.29.6443>
12. A. Takahashi, A. Seko, I. Tanaka, *Phys. Rev. Mater.* **1**, 063801 (2017)
13. E. Wimmer, M. Christensen, W. Wolf, W.H. Howland, B. Kammenzind, R.W. Smith, *J. Nucl. Mater.* **532**, 152055 (2020)
14. S. Starikov, D. Smirnova, *Comput. Mater. Sci.* **197**, 110581 (2021)
15. O.G. Nicholls, D.G. Frost, V. Tuli, J. Smutna, M.R. Wenman, P.A. Burr, *J. Nucl. Mater.* **584**, 154391 (2023)

16. A.P. Bartók, J. Kermode, N. Bernstein, G. Csányi, *Phys. Rev. X* **8**, 041048 (2018)
17. E.V. Podryabinkin, A.V. Shapeev, *Comput. Mater. Sci.* **140**, 171 (2017). <https://doi.org/10.1016/j.commatsci.2017.08.031>
18. M. Hodapp, A. Shapeev, *Phys. Rev. Mater.* **5**, 113802 (2021)
19. E.V. Podryabinkin, E.V. Tikhonov, A.V. Shapeev, A.R. Oganov, *Phys. Rev. B* **99**, 064114 (2019)
20. L. Zhang, D.Y. Lin, H. Wang, R. Car, *Phys. Rev. Mater.* **3**, 023804 (2019)
21. I.S. Novikov, K. Gubaev, E.V. Podryabinkin, A.V. Shapeev, *Mach. Learn. Sci. Technol.* **2**, 025002 (2021)
22. A. Shapeev, K. Gubaev, E. Tsymalov, E. Podryabinkin, *Active Learning and Uncertainty Estimation*, in *Machine Learning Meets Quantum Physics*. ed. by K.T. Schütt, S. Chmiela, O.A. von Lilienfeld, A. Tkatchenko, K. Tsuda, K.R. Müller (Springer, Cham, 2020), pp.309–329
23. C.J. Pickard, *Phys. Rev. B* **106**, 014102 (2022). <https://doi.org/10.1103/PhysRevB.106.014102>
24. S. Kass, *J. Electrochem. Soc.* **107**, 594 (1960)
25. C.R.F. Azevedo, *Eng. Fail. Anal.* **18**, 1943 (2011). <https://doi.org/10.1016/j.engfailanal.2011.06.010>
26. B. Kammenzind, B. Berquist, R. Bajaj, P. Kreyens, D. Franklin, *STP* **1354**, 196 (2000)
27. J. Kearns, *J. Nucl. Mater.* **22**, 292 (1967)
28. Z. Duan, H. Yang, Y. Satoh, K. Murakami, S. Kano, Z. Zhao, J. Shen, H. Abe, *Nucl. Eng. Des.* **316**, 131 (2017). <https://doi.org/10.1016/j.nucengdes.2017.02.031>
29. A.T. Motta, L. Capolungo, L.-Q. Chen, M.N. Cinbiz, M.R. Daymond, D.A. Koss, E. Lacroix, G. Pastore, P.C.A. Simon, M.R. Tonks, B.D. Wirth, M.A. Zikry, *J. Nucl. Mater.* **518**, 440 (2019). <https://doi.org/10.1016/j.jnucmat.2019.02.042>
30. C. Ells, *J. Nucl. Mater.* **28**, 129 (1968)
31. R. Daum, Y. Chu, A. Motta, *J. Nucl. Mater.* **392**, 453 (2009)
32. F. Long, N. Badr, Z. Yao, M. Daymond, *J. Nucl. Mater.* **543**, 152540 (2021)
33. R. Beck, *Am. Soc. Metals Trans. Quart.* **55**, 1–10 (1962)
34. G. Carpenter, *Acta Metall.* **26**, 1225 (1978)
35. G. McRae, C. Coleman, *J. Nucl. Mater.* **499**, 622 (2018)
36. N. Badr, F. Long, Y. Luo, M. Topping, L. Béland, Z. Yao, M. Daymond, *Acta Mater.* **247**, 118729 (2023)
37. L. Lanzani, M. Ruch, *J. Nucl. Mater.* **324**, 165 (2004). <https://doi.org/10.1016/j.jnucmat.2003.09.013>
38. E. Tulk, M. Kerr, M. Daymond, *J. Nucl. Mater.* **425**, 93 (2012)
39. G. Carpenter, *J. Nucl. Mater.* **48**, 264 (1973)
40. I.C. Njifon, E. Torres, *Acta Mater.* **202**, 222 (2021)
41. C. Domain, R. Besson, A. Legris, *Acta Mater.* **50**, 3513 (2002)
42. Y. Udagawa, M. Yamaguchi, H. Abe, N. Sekimura, T. Fuketa, *Acta Mater.* **58**, 3927 (2010)
43. W. Zhu, R. Wang, G. Shu, P. Wu, H. Xiao, *J. Phys. Chem. C* **114**, 22361 (2010)
44. F. Wang, H. Gong, *Int. J. Hydrog. Energy* **37**, 12393 (2012)
45. P. Olsson, A. Massih, J. Blomqvist, A.M.A. Holston, C. Bjerkén, *Comput. Mater. Sci.* **86**, 211 (2014)
46. X. Zhu, D.Y. Lin, J. Fang, X.Y. Gao, Y.F. Zhao, H.F. Song, *Comput. Mater. Sci.* **150**, 77 (2018)
47. M. Christensen, W. Wolf, C. Freeman, E. Wimmer, R. Adamson, L. Hallstadius, P. Cantonwine, E. Mader, *J. Phys. Condens. Matter* **27**, 025402 (2015)
48. G.Y. Huang, B.D. Wirth, *J. Appl. Phys.* **126**, 135105 (2019)
49. G.L.W. Hart, R.W. Forcade, *Phys. Rev. B* **77**, 224115 (2008). <https://doi.org/10.1103/PhysRevB.77.224115>
50. G.L.W. Hart, R.W. Forcade, *Phys. Rev. B* **80**, 014120 (2009)
51. L. Zhang, Y.H. Li, Y.Q. Gu, L.C. Cai, *Sci. Rep.* **9**, 16889 (2019). <https://doi.org/10.1038/s41598-019-53088-3>
52. E. Zuzek, J. Abriata, A. San-Martin, F. Manchester, *Phase Diagrams of Binary Hydrogen Alloys* (ASM International, Almere, 2000), p.309
53. E. Zuzek, J. Abriata, A. San-Martin, F. Manchester, *Phase Diagrams of Binary Hydrogen Alloys* (ASM International, Almere, 1990), p.385
54. C. Nyshadham, C. Oses, J.E. Hansen, I. Takeuchi, S. Curtarolo, G.L.W. Hart, *Acta Mater.* **122**, 438 (2017). <https://doi.org/10.1016/j.actamat.2016.09.017>
55. O. Levy, G.L.W. Hart, S. Curtarolo, *Acta Mater.* **58**, 2887 (2010). <https://doi.org/10.1016/j.actamat.2010.01.017>
56. O. Levy, G.L.W. Hart, S. Curtarolo, *J. Am. Chem. Soc.* **132**, 4830 (2010). <https://doi.org/10.1021/ja9105623>
57. K. Gubaev, E.V. Podryabinkin, G.L.W. Hart, A.V. Shapeev, *Comput. Mater. Sci.* **156**, 148 (2019). <https://doi.org/10.1016/j.commatsci.2018.09.031>
58. W.S. Morgan, G.L.W. Hart, R.W. Forcade, *Comput. Mater. Sci.* **136**, 144 (2017). <https://doi.org/10.1016/j.commatsci.2017.04.015>
59. G. Kresse, J. Furthmüller, *Comput. Mater. Sci.* **6**, 15 (1996). [https://doi.org/10.1016/0927-0256\(96\)00008-0](https://doi.org/10.1016/0927-0256(96)00008-0)
60. G. Kresse, J. Furthmüller, *Phys. Rev. B* **54**, 11169 (1996b). <https://doi.org/10.1103/PhysRevB.54.11169>
61. Y. Li, R. Rogge, R. Holt, *J. Nuclear Mater.* **437**, 10 (2006)
62. 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, R. Shan, M.J. Stevens, J. Tranchida, C. Trott, S.J. Plimpton, *Comput. Phys. Commun.* **271**, 108171 (2022). <https://doi.org/10.1016/j.cpc.2021.108171>
63. P. Liu, C. Verdi, F. Karsai, G. Kresse, *Phys. Rev. Mater.* **5**, 053804 (2021)
64. The parameters used for the VASP calculations of this training set were not the same as the parameters shown in Table 2. However, the only significant parameter change from the parameters in Table 2 was a change in ENCUT from 400 eV to 500 eV. Because 400 eV is a reasonable value of ENCUT for the Zr system, this should not affect the training results significantly

65. S.P. Ong, W.D. Richards, A. Jain, G. Hautier, M. Kocher, S. Cholia, D. Gunter, V.L. Chevrier, K.A. Persson, G. Ceder, *Comput. Mater. Sci.* **68**, 314 (2013). <https://doi.org/10.1016/j.commatsci.2012.10.028>
66. A. Togo, L. Chaput, T. Tadano, I. Tanaka, *J. Phys. Condens. Matter* **35**, 353001 (2023). <https://doi.org/10.1088/1361-648X/acd831>

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

Springer Nature or its licensor (e.g. a society or other partner) holds exclusive rights to this article under a publishing agreement with the author(s) or other rightsholder(s); author self-archiving of the accepted manuscript version of this article is solely governed by the terms of such publishing agreement and applicable law.