

---

# PFT: Phonon Fine-tuning for Machine Learned Interatomic Potentials

---

Teddy Koker<sup>1</sup> Abhijeet Gangan<sup>2,3</sup> Mit Kotak<sup>4</sup> Jaime Marian<sup>3</sup> Tess Smidt<sup>1</sup>

## Abstract

Many materials properties depend on higher-order derivatives of the potential energy surface, yet machine learned interatomic potentials (MLIPs) trained with standard a standard loss on energy, force, and stress errors can exhibit error in curvature, degrading the prediction of vibrational properties. We introduce phonon fine-tuning (PFT), which directly supervises second-order force constants of materials by matching MLIP energy Hessians to DFT-computed force constants from finite displacement phonon calculations. To scale to large supercells, PFT stochastically samples Hessian columns and computes the loss with a single Hessian-vector product. We also use a simple co-training scheme to incorporate upstream data to mitigate catastrophic forgetting. On the MDR Phonon benchmark, PFT improves Nequix MP (trained on Materials Project) by 55% on average across phonon thermodynamic properties and achieves state-of-the-art performance among models trained on Materials Project trajectories. PFT also generalizes to improve properties beyond second-derivatives, improving thermal conductivity predictions that rely on third-order derivatives of the potential energy.

## 1. Introduction

Many computational materials workflows use density functional theory (DFT) to compute materials properties from first principles. DFT's computational cost has motivated machine learned interatomic potentials (MLIPs) as fast sur-

rogates that can replace or accelerate DFT in large-scale screening and simulation (Yang et al., 2024; Merchant et al., 2023). Accurate predictions of phonon and vibrational properties require an MLIP to match the *curvature* of the DFT potential energy surface (PES), not just energies, forces and stress. However, existing “universal” MLIPs are typically trained with supervision over energy, force, and stress; this only indirectly constrains second derivatives, leading to errors in the curvature that degrade phonon dispersions and thus various thermodynamic properties (Deng et al., 2025). In this work, we show that Hessian error strongly correlates with error across multiple phonon thermodynamic metrics.

In materials, training directly on curvature is challenging because phonon force constants in periodic crystals are commonly obtained by finite-displacement calculations in supercells; the supercell must be large enough that displaced atoms do not interact with periodic images of themselves, and to capture interactions that extend beyond the primitive cell. These requirements often necessitate hundreds or thousands of atoms, where the  $3N \times 3N$  Hessian, or force-constant matrix scales quadratically in system size, making full Hessian training infeasible.

We introduce **phonon fine-tuning (PFT)**, a fine-tuning procedure that directly incorporates second-order force constants by matching energy Hessians from the MLIP to DFT-derived force constants (Fig. 1c(iv)). To scale to large supercells, PFT stochastically samples force-constant columns and computes the corresponding loss via a single Hessian-vector product, reducing the training step cost from quadratic to linear with respect to the number of atoms. Because of the relative diversity of existing phonon datasets and lack of non-equilibrium geometries, we introduce a simple co-training strategy to mitigate catastrophic forgetting by interleaving upstream pretraining data during fine-tuning.

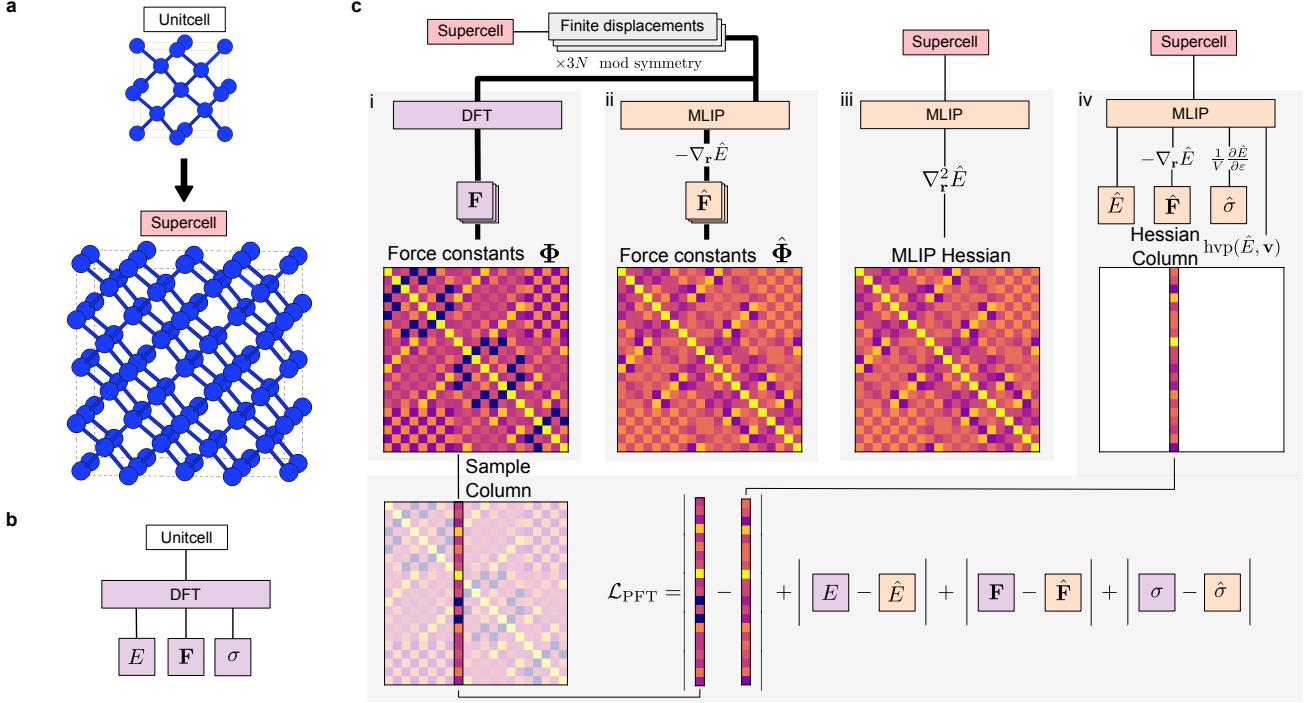
We evaluate PFT by fine-tuning the Nequix MP foundation model (Koker et al., 2025) and benchmarking on held out calculations from the PBE MDR Phonon benchmark (Togo, 2023a; Loew et al., 2025). At a fraction of the cost of the initial pretraining, PFT reduces property error by 55% on average across maximum phonon frequency, vibrational entropy, Helmholtz free energy, and heat capacity, achieving state-of-the-art performance among models trained on

<sup>1</sup>Department of Electrical Engineering and Computer Science, Massachusetts Institute of Technology, Boston, MA, USA

<sup>2</sup>Department of Civil and Environmental Engineering, University of California, Los Angeles, CA 90095, USA

<sup>3</sup>Department of Materials Science and Engineering, University of California, Los Angeles, CA 90095, USA

<sup>4</sup>Center for Computational Science and Engineering, Massachusetts Institute of Technology, Boston, MA, USA. Correspondence to: Teddy Koker <tekoker@mit.edu>.



**Figure 1. Overview of PFT framework.** **a** Finite-difference calculations rely on the construction of a supercell to obtain force constants from interactions beyond the unitcell. **b** MLIPs are pre-trained on standard unitcell DFT calculations. **c-i** Up to  $O(3N)$  atomic displacements are applied to the supercell, with the number reduced by crystal symmetries; forces are computed with DFT, and numerical derivatives yield the force constant matrix. **c-ii** The same workflow can be used for MLIPs, replacing DFT force calculations with force predictions from the model. **c-iii** Alternatively, the force constants can also be computed as the analytical Hessian of the predicted energy directly. **c-iv** In this work, we use the Hessian-vector product to efficiently compute columns of the Hessian, and align with sampled DFT force constant columns during training. Note the shown force constants are downsampled for clarity.

MPtrj (Deng et al., 2023; Jain et al., 2013). Furthermore, we demonstrate that PFT improves generalization to third-order derivatives, improving Matbench Discovery thermal conductivity predictions from 0.446 to 0.306  $\kappa_{\text{SRME}}$ , which is also state-of-the-art among MPtrj-trained models. Lastly, we demonstrate that through co-training, PFT degrades performance by less than 1% on the Matbench Discovery stability classification task, which is otherwise greatly impacted by training on phonon data alone.

Our contributions are:

- A fine-tuning objective that directly aligns the curvature of the MLIP PES by aligning energy Hessians with DFT-derived force-constants.
- A scalable strategy for training on phonons of large supercells by sampling columns of the Hessian, and training with Hessian-vector products.
- A simple co-training recipe that mitigates catastrophic forgetting by interleaving pre-training data into the finetuning procedure.
- Empirical results demonstrating that Hessian error strongly correlates with phonon property error, and

that the introduced training objective greatly improves performance on phonon properties while preserving performance on other tasks.

## 2. Background

### 2.1. Machine learned interatomic potentials

Machine learned interatomic potentials (MLIPs) seek to model the Born-Oppenheimer potential energy surface  $E(\mathbf{r}, \mathbf{z})$  for atoms with positions  $\mathbf{r} \in \mathbb{R}^{3N_a}$ , and species  $\mathbf{z} \in \mathbb{Z}_+^{N_a}$ , where  $N_a$  is the number of the atoms in the system. Neural network MLIPs typically build a model for predicted energy  $\hat{E}_\theta(\mathbf{r}, \mathbf{a})$  parameterized by weights  $\theta$ , based on local atomic environments (Unke et al., 2021). These models are trained on quantum chemistry calculations such as density functional theory (DFT) (Hohenberg & Kohn, 1964).

Recent work has demonstrated the success of so-called *universal* MLIPs, which seek to model the PES across a broad range of chemistries and geometries by training on vast DFT databases (Chen & Ong, 2022; Batatia et al., 2025a; Wood et al., 2025). These models offer the promise of acting as a surrogate to otherwise computationally expensive quantum

chemistry calculations, enabling higher-throughput computational materials workflows.

In this work we focus on *energy-conserving* MLIPs, where forces and stresses are obtained as derivatives of a scalar energy model, which ensures that the curl of the predicted forces are zero by design.

With atom positions  $\mathbf{r}$  in a periodic cell of volume  $V$ , predicted forces are computed as

$$\hat{\mathbf{F}}_a = -\nabla_{\mathbf{r}_a} \hat{E}(\mathbf{r}) \quad (1)$$

Predicted stress on the lattice is computed as the derivative of energy with respect to symmetric strain tensor  $\varepsilon$ :

$$\hat{\sigma}_{ij} = \frac{1}{V} \frac{\partial \hat{E}(\mathbf{r})}{\partial \varepsilon_{ij}} \Big|_{\varepsilon=0} \quad (2)$$

These quantities are convenient to compute with automatic differentiation (AD) from  $\hat{E}$ . Note that for the remainder of the paper we have dropped the dependence on atomic species  $\mathbf{z}$  for clarity.

MLIPs for periodic systems are typically trained using a loss over energies, forces, and stresses computed from DFT. This is constructed as a weighted sum of errors over each term:

$$\mathcal{L}_{\text{EFS}} = \lambda_E \mathcal{L}_E + \lambda_F \mathcal{L}_F + \lambda_\sigma \mathcal{L}_\sigma \quad (3)$$

with individual loss terms

$$\mathcal{L}_E = \left| \frac{\hat{E}}{N_a} - \frac{E}{N_a} \right| \quad \mathcal{L}_F = \frac{1}{N_a} \sum_{a=1}^{N_a} \left\| \hat{\mathbf{F}}_a - \mathbf{F}_a \right\|_2^2 \quad (4)$$

$$\mathcal{L}_\sigma = \frac{1}{9} \sum_{i=1}^3 \sum_{j=1}^3 |\hat{\sigma}_{ij} - \sigma_{ij}| \quad (5)$$

where  $E$ ,  $\mathbf{F}$ , and  $\sigma$  denote DFT-computed energy, force, and stress respectively, and coefficients  $\lambda$  are tunable hyperparameters to allow weighting of different quantities. Universal potentials are often trained on databases of relaxation trajectories (Deng et al., 2023; Jain et al., 2013; Schmidt et al., 2024) or by perturbing equilibrium structures to sample more of the PES (Barroso-Luque et al., 2024; Kaplan et al., 2025). Prior work has found that due to the bias of data towards equilibrium structures, there often exists a softening in the curvature of the PES (Deng et al., 2025).

## 2.2. Harmonic phonons and vibrational properties

Phonons arise from small lattice vibrations around a local minimum of the PES, and their spectra and scattering control

many materials properties of interest, including thermal conductivity, thermal expansion, and heat capacity (Ziman, 2001). They also dictate dynamic stability and, via electron-phonon coupling, can enable superconductivity (Giustino, 2017).

Phonon frequencies are calculated from the eigenvalues of the dynamical matrix, which is constructed by mass-weighting the real-space force constants  $\Phi$ , and applying a lattice Fourier transform that introduces the wavevector  $\mathbf{k}$  dependence. The force constants are defined as the second derivative of the PES with respect to atomic positions:

$$\Phi_{aibj} = \frac{\partial^2 E}{\partial r_{a,i} \partial r_{b,j}} \quad (6)$$

where  $i$  and  $j$  are Cartesian indices, and  $a$  and  $b$  are atom indices.

## 2.3. Phonon calculations

Phonon calculations are usually conducted with DFT in one of two ways: 1) using density functional perturbations theory (DFPT), or 2) DFT and with finite displacements. The latter is more common due to its generality across commonly used functionals (Miranda).

With finite displacement, second-order force constants are approximated with

$$\Phi_{aibj} \simeq -\frac{F_{b,j}(\Delta r_{a,i}) - F_{b,j}}{\Delta r_{a,i}} \quad (7)$$

$F_{b,j}(\Delta r_{a,i})$  are forces where atom  $a$  is displaced in direction  $i$ , and generally  $F_{b,j} = 0$  due to the atom positions being at equilibrium. This approximation works well due to the smoothness of DFT under small displacements, commonly 0.01 Å for these calculations. The number of finite-displacement calculations (3N) can typically be significantly reduced by leveraging the symmetry of the crystal (Togo et al., 2023)

It is important to note that phonon calculations using finite displacement require large supercells for calculations to avoid self-interaction with the displaced atom, and correctly model interatomic force constants that extend past the unit cell (Fig. 1a).

MLIPs can be used for phonon calculations in the same way as DFT, constructing the force constants by conducting force calculations under finite displacement (Fig. 1c(i-ii)), which has shown success in prediction of both second-order (Loew et al., 2025) and third-order (Póta et al., 2024) phonon properties.

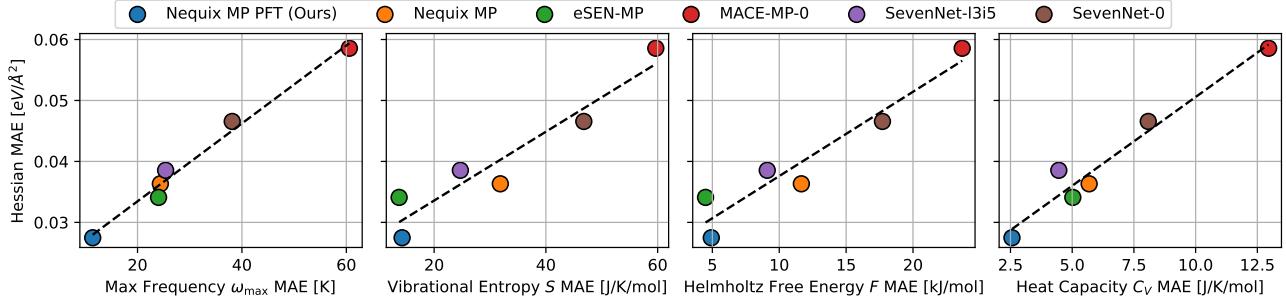


Figure 2. **Hessian error vs. phonon properties.** Error in the Hessians on the test subset of the MDR Phonon data are plotted against heat capacity errors for several foundation models trained on MPtrj. Hessian errors correlate with improved property prediction.

### 3. Methodology

#### 3.1. Hessian error and vibrational properties

Because phonon spectra and derived vibrational properties are functions of the second-order force constants, models must match the curvature of the PES to achieve accurate phonon predictions. Using the finite displacement method above, we calculate the error in the second-order force constants, or Hessians of several models trained on MPtrj (Deng et al., 2023; Jain et al., 2013), evaluating them on a portion of the force constants and phonon properties from the PBE MDR phonon dataset (Loew et al., 2025; Togo, 2023a). Figure 2 shows that lower hessian error correlates with lower error on multiple thermodynamic properties. This motivates that training directly on second-order force constants results in more accurate predictions of these properties.

#### 3.2. Phonon fine-tuning (PFT)

We propose phonon fine-tuning (PFT), a method for fine-tuning MLIPs directly on DFT-computed force constants.  $\mathcal{L}_{\text{PFT}}$  is a four term loss function, with terms that minimize error in the energy, force, stress, and second-order force constants:

$$\mathcal{L}_{\text{PFT}} = \lambda_E \mathcal{L}_E + \lambda_F \mathcal{L}_F + \lambda_\sigma \mathcal{L}_\sigma + \lambda_\Phi \mathcal{L}_\Phi \quad (8)$$

where we define

$$\mathcal{L}_\Phi = \frac{1}{3N_a} \sum_{a=1}^{N_a} \sum_{i=1}^3 \mathbb{E}_{b \sim \mathcal{U}[1, N_a]} \left| \frac{\partial^2 \hat{E}}{\partial r_{a,i} \partial r_{b,j}} - \Phi_{aibj} \right| \quad (9)$$

where  $N_a$  is the number of atoms in a system. Note that again the batch dimension is omitted for readability.  $\hat{E}$  is the predicted potential energy from the neural network; the first three terms follow the standard EFS loss from Eq. (3) with an MAE loss on energy and stress and a  $\ell_2$  loss on forces. Forces are predicted as the negative derivative of energy with respect to position  $\mathbf{r}$ , and stress is predicted as the derivative of energy with respect to a symmetric strain tensor. The

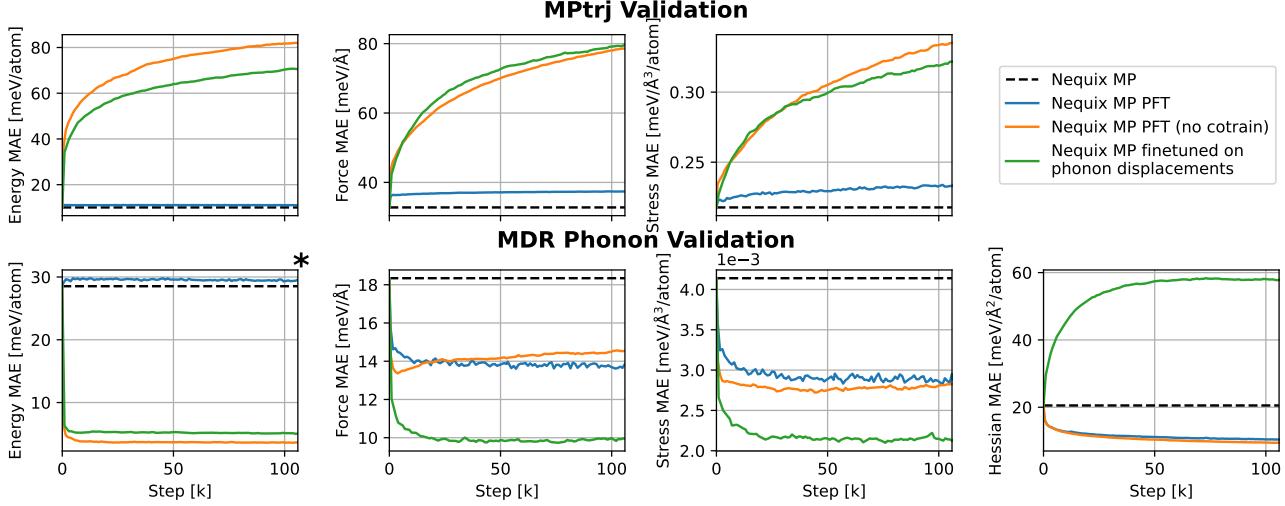
second-order force constants are predicted analytically as the hessian of the energy with respect to two atom positions.

To improve computational efficiency and enable training on large supercells, we uniformly sample one atom and Cartesian direction for each structure in the batch, effectively selecting one column of the hessian. This requires only a single Hessian-vector product (see next section) to compute the loss across the whole batch, while effectively still training on the full Hessian in expectation. This reduces the computational complexity of a training step from  $O(N^2)$  to  $O(N)$  for  $N$  number of atoms. Furthermore, by using a symmetry-aware  $E(3)$ -equivariant architecture, many Hessian elements are redundant due to the high symmetry of the force constants of a crystal structure (Califano et al., 1981). More consideration may be necessary with regards to sampling and data augmentation if non-equivariant architectures are used.

Since the DFT phonon calculations themselves consist of energy, force, and stress calculations under displacements of atoms, it is reasonable to assume that one could simply fine-tune models directly on this displacement data, as it contains sufficient information for constructing the full Hessian. Empirically we find this not to be the case, and observe that directly finetuning on the phonon displacements results in a *worse* hessian error. This suggests that the standard practice of EFS training on DFT calculations of rattled, or perturbed structures (Barroso-Luque et al., 2024; Kaplan et al., 2025) may not be sufficient for correctly modeling the PES curvature.

#### 3.3. Efficient computation via Hessian-vector product

PFT requires computing a subset of predicted second derivatives to compare to DFT-computed force constants for large supercells. For a structure with atomic positions  $\mathbf{r} = (\mathbf{r}_1, \dots, \mathbf{r}_{N_a}) \in \mathbb{R}^{3N}$  we form a selection vector  $\mathbf{v}$  that is composed of zeros everywhere except for the index corresponding to  $\mathbf{r}_{b,j}$ , the  $j$ -th Cartesian component of atom  $b$ , where it is one. The corresponding hessian column is then



**Figure 3. Training ablation.** The top row shows energy, force, and stress errors on the MPtrj validation set, while the bottom row shows energy, force, stress, and hessian errors on the MDR Phonon validation set during finetuning of Nequix MP. We compare phonon fine-tuning with and without co-training on MPtrj, as well directly finetuning energy, force, and stress on the phonon displacement calculations. Co-training mostly mitigates degradation of MPtrj performance at only a slight increase in hessian error, and that training directly on displacements worsens hessian performance over the base model. \*We note that larger energy errors on MDR Phonon data are likely due to a mismatch in energies between MPtrj and MDR Phonon, see Sec. A.1.

$$\nabla_{\mathbf{r}}^2 \hat{E}(\mathbf{r}) \mathbf{v} = \nabla_{\mathbf{r}} \left( \nabla_{\mathbf{r}} \hat{E}(\mathbf{r})^\top \mathbf{v} \right) \quad (10)$$

which can be computed using a Hessian-vector product (HVP) without materializing the full hessian using a forward-mode Jacobian-vector product (JVP) through a reverse-mode gradient (Pearlmutter, 1994). In practice, Hessian columns across a full batch of structures is computed by forming a single graph of atom positions where each structure is not connected, concatenating the sampled  $\mathbf{v}$  for each structure into a single vector, and computing the HVP with respect to the sum of energies across all structures. This facilitates calculation of the loss  $a$  with single HVP call, and enables training on GPUs that would otherwise not have enough memory for a single Hessian calculation, especially for the large supercells that are needed for accurate phonon calculations. Note that during training, this results in a “triple-backward” as gradient-based optimization requires a derivative of the HVP with respect to model weights.

### 3.4. Co-training

While fine-tuning models on phonon data will more closely align the curvature of the PES with that of DFT, this procedure may lead to catastrophic forgetting (French, 1999), where the fine-tuning procedure effects the performance of the model on its original upstream training data. This may be an issue as phonon calculations are always done at equilibrium, which can cause forgetting for out-of-distribution non-equilibrium configurations. Furthermore,

---

#### Algorithm 1 PFT training with co-training

---

**Require:** Phonon dataset  $\mathcal{D}_{\text{phonon}}$ , upstream dataset  $\mathcal{D}_{\text{up}}$ , co-training ratio  $K$ , model  $\hat{E}_\theta$

- 1: **for** batch in  $\mathcal{D}_{\text{phonon}}$  **do**
- 2:     Sample  $(b, j) \sim \mathcal{U}([1, N_a] \times [1, 3])$  per structure
- 3:     Update  $\theta$  using  $\mathcal{L}_{\text{PFT}}$  (Eq. 8)
- 4:     **for**  $k = 1$  to  $K$  **do**
- 5:         Sample batch from  $\mathcal{D}_{\text{up}}$
- 6:         Update  $\theta$  using  $\mathcal{L}_{\text{EFS}}$  (Eq. 3)
- 7:     **end for**
- 8: **end for**

---

if the relative quantity and diversity of the phonon fine-tuning dataset is less than that of the original pretraining or fine-tuning datasets (such as the case in our experimental setting)

We propose a simple solution to this by alternating training steps between PFT steps and a typical energy/force/stress (EFS) loss on the original upstream dataset, as outlined in Algorithm 1. The ratio of EFS on the upstream dataset to PFT steps on the phonon dataset  $K$  can be tuned by monitoring validation datasets for both the upstream and phonon datasets during training. As shown in Fig. 3, we demonstrate that without co-training, PFT causes energy, force, and stress errors to deviate quite significantly on the upstream validation dataset. The introduction of co-training, however, mostly eliminates this deviation with only a small reduction in Hessian MAE.

### 3.5. Analytical property prediction

As discussed, most existing methods for evaluating MLIPs on higher order derivatives of the PES do so with finite displacement (Loew et al., 2025; Póta et al., 2024). This is typically accurate if the smoothness and continuity of the PES is taken into account when designing the neural network architecture; however, it does introduce an additional hyper-parameter in the form of displacement distance, which can subtly effect results (Fu et al., 2025).

Leveraging the HVP, we can compute the full force constants, see Eq. (6), by iterating over Hessian columns, parallelizing HVP calls until GPU memory is saturated. While in theory this computation may need to be conducted on a graph that exceeds the receptive field of the neural network (Fang et al., 2024), simply using the same supercell as the DFT calculation will ensure we can sufficiently predict the same phonon modes as the ground truth. In practice we find that we achieve near identical results through analytical force constant prediction and finite displacement, which we show in Table 1.

## 4. Experiments

### 4.1. Training

For all experiments, we start with the Nequix MP (Koker et al., 2025), whose JAX (Bradbury et al., 2018) implementation enables convenient auto-differentiation through the hessian terms of the loss function. Nequix MP was trained on MPtrj (Deng et al., 2023), a dataset of relaxation trajectories from the v2022.10.28 release of Materials Project (Jain et al., 2013) consisting of 145,923 unique materials. MPtrj will serve as the upstream dataset  $\mathcal{D}_{\text{up}}$

For phonon calculation data  $\mathcal{D}_{\text{phonon}}$ , we use MDR Phonon calculation database (Togo, 2023a), which was recalculated by (Loew et al., 2025) using the PBE exchange correlation functional in order to match the settings of the Materials Project data. These calculations were conducted using finite displacement (see Sec. 2.3), and contain the original energy/force/stress calculations at each displacement. While the phonon data itself does not contain force, or stress labels needed for the PFT loss function (Eq. 8), we assume the force and stress labels are zero due to the strict structural relaxation procedure done prior to the phonon calculation. The energy label is computed by multiplying the provided unitcell energy by the number of repetitions within the supercell used for calculations. Of the 9,959 materials within the dataset, all but 91 exist in MPtrj based on `mp-id`. These 91 materials along with randomly selected calculations from the remaining data are used as a test set of 1,000 materials, with the rest of the data split into training/validation subsets with proportion 95/5. This ensures any benefit we see from PFT is due to the additional phonon information, and not

novel chemistry or geometries. We do note there is a discrepancy between the energies of materials in phonon data and MPtrj, which we detail in Sec. A.1.

We perform PFT both with ( $K = 4$ ) and without ( $K = 0$ ) co-training for 200 epochs on the MDR Phonon data. With co-training, this requires about 140 A100 hours, with about 60 A100 hours for PFT without co-training — significantly fewer than the 500 A100 hours needed to train the base Nequix MP model and a small fraction of other competitive MPtrj-trained MLIPs (Koker et al., 2025). For details on other hyper-parameters and their selection see Sec. A.2.

### 4.2. Phonon properties

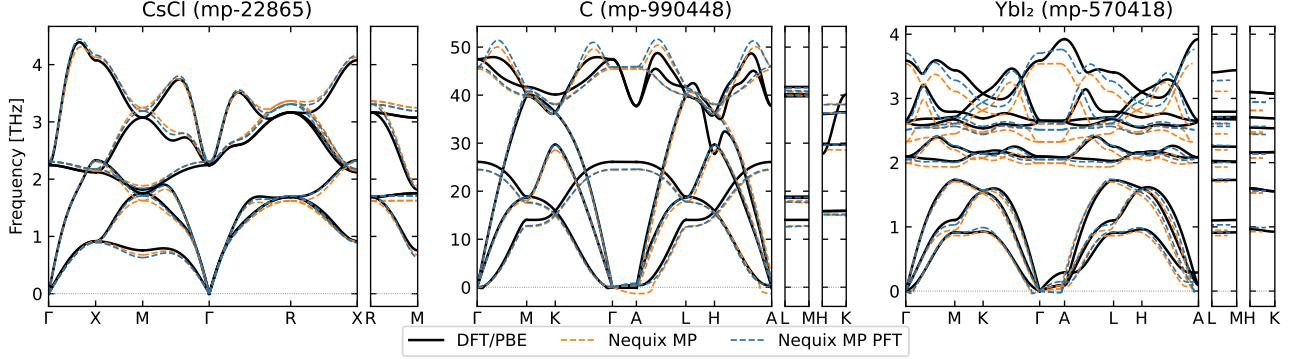
*Table 1.* Evaluation of models on held-out MDR Phonon data. Metrics are MAE of maximum phonon frequency  $\omega_{\text{max}}$  (K), vibrational entropy  $S$  (J/K/mol), Helmholtz free energy  $F$  (kJ/mol) and heat capacity at constant volume  $C_V$  (J/K/mol).

Model	$\omega_{\text{max}}$	$S$	$F$	$C_V$
MACE-MP-0	61	60	24	13
SevenNet-0	38	47	18	8
Nequix MP	24	32	12	6
SevenNet-l3i5	25	25	9	4
eSEN-MP	24	<u>14</u>	<b>4</b>	5
Nequix MP PFT	<u>12</u>	<u>14</u>	<u>5</u>	<u>3</u>
Nequix MP PFT (autodiff)	<u>12</u>	<u>14</u>	<u>5</u>	<u>3</u>
Nequix MP PFT (no cotrain)	<b>10</b>	<b>11</b>	<b>4</b>	<b>2</b>

Models are evaluated on the 1,000 material test subset of phonon calculations by following the same procedure as Loew et al. (2025), first performing a structural relaxation, and then running phonon calculations with finite displacement using phonopy (Togo et al., 2023). Thermodynamic properties are calculated at 300K. We re-run several competitive MPtrj-trained models (Batatia et al., 2025a; Park et al., 2024; Koker et al., 2025; Fu et al., 2025) on the subset. For Nequix MP PFT, we compute force constants with both finite displacement, and analytically via AD.

Table 1 shows the errors of each model on phonon related properties maximum phonon frequency, vibrational entropy, Helmholtz free energy, and heat capacity at constant volume. From the Nequix MP base model, PFT results in an average reduction in property prediction error of 55%, and achieves state-of-the-art error across all metrics for models trained on MPtrj materials despite being the smallest model. We find no significant differences between properties computed from forces constants obtained via finite displacement or AD, although the PFT model without co-training demonstrates a slight improvement in prediction error.

Lastly, Fig. 4 shows several computed phonon band structures before and after PFT. As is to be expected, we find the PFT leads to band structures that more closely align with the DFT computed band structures.



**Figure 4. Phonon band structures.** For ease of visualization, we display the phonon band structure for the three materials in the test split of MDR Phonon with the fewest number of atoms. In the case of Carbon structure (mp-990448), it results in a dynamically stable structure similar to DFT while the non-PFT model shows imaginary mode around the  $q$ -point  $A$ . Overall, we find that that PFT generally produces bands with closer alignment to those of the DFT ground truth. More phonon bandstructures are provided in Fig. A.2.

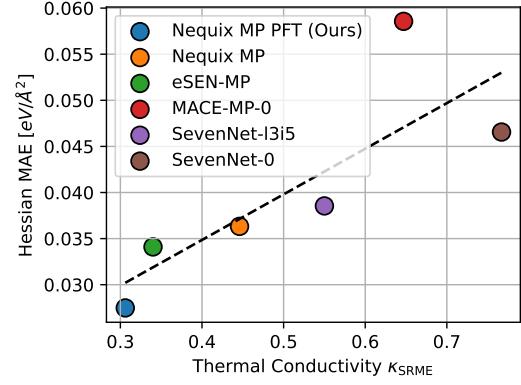
#### 4.3. Thermal conductivity

To demonstrate the effect of PFT on properties that are higher-order derivatives of the PES, we evaluate the model on thermal conductivity, a task from the Matbench Discovery (Riebesell et al., 2025) benchmark. Thermal conductivity  $\kappa$  is a function of third-order force constants (Póta et al., 2024), or the third derivative of the potential energy. Just as with second order force constants, third order force constants can be calculated with finite displacement, leveraging symmetry to reduce the number of computations (Togo, 2023b).

**Table 2.** Matbench Discovery “compliant” leaderboard for thermal conductivity, measured in symmetric relative mean error in predicted phonon mode contributions to thermal conductivity  $\kappa_{SRME}$ .

Model	Params	$\kappa_{SRME} \downarrow$
ORB v2 MPtrj	25.2M	1.725
eqV2 S DeNS	31.2M	1.676
MatRIS v0.5.0 MPtrj	5.83M	0.865
MACE-MP-0	4.69M	0.682
DPA-3.1-MPtrj	4.81M	0.650
HIENet	7.51M	0.642
SevenNet-l3i5	1.17M	0.550
GRACE-2L-MPtrj	15.3M	0.525
Nequip-MP-L	9.6M	0.452
Nequix MP	707K	0.446
Eqnorm MPtrj	1.31M	0.408
eSEN-30M-MP	30.1M	0.340
Nequix MP PFT	707K	0.306
Nequix MP PFT (no cotrain)	707K	<b>0.281</b>

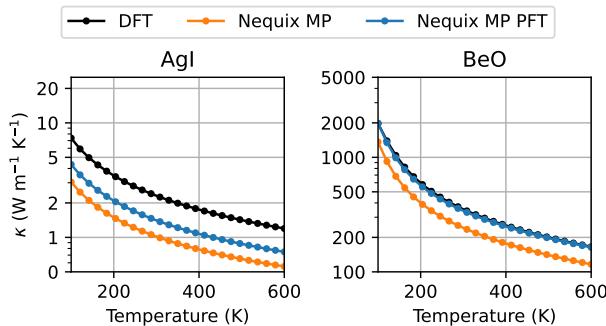
We follow the procedure used by Matbench Discovery (Riebesell et al., 2025; Póta et al., 2024) performing a structural relaxation and then calculation of the third-order force constants with a displacement of 0.03 Å. Results from all of the other MPtrj-trained, or “compliant” models (Neumann et al., 2024; Barroso-Luque et al., 2024; Zhou et al., 2025;



**Figure 5. Thermal conductivity vs. Hessian error.** Scatter plot of several models Hessian MAE on the MDR phonon test set vs. symmetric relative mean error in predicted phonon mode contributions to thermal conductivity  $\kappa_{SRME}$ . We find there is still a strong trend in between the two despite thermal conductivity using third-order force constants.

Batatia et al., 2025a; Zhang et al., 2025; Yan et al., 2025; Park et al., 2024; Bochkarev et al., 2024; Tan et al., 2025; Koker et al., 2025; Chen et al., 2025; Fu et al., 2025) sourced from the Matbench Discovery leaderboard are shown in Table 2.

Similarly to phonon properties we find a reduction in error of 31% from the base Nequix MP model, and state-of-the-art performance for MPtrj-trained models. In Fig. 6, we show macroscopic conductivity predictions for two zinc blends, AgI and BeO, showing again that PFT aligns predictions closer to the ground truth DFT calculation. This shows that PFT not only improves second order vibrational properties, but also generalizes to third order, or anharmonic, vibrational properties.



**Figure 6. Macroscopic conductivity predictions.** Macroscopic thermal conductivity calculations for zinc blend AgI (left), which exhibits low conductivity and zinc blend BeO (right), which exhibits high conductivity. DFT calculations are from (Togo, 2025). PFT produces conductivity predictions closer to the ground truth.

#### 4.4. Matbench Discovery

The main Matbench Discovery task (Riebesell et al., 2025) consists of a geometry optimization followed by energy prediction to determine material stability. Performance on this task is less influenced by curvature of the PES and largely dictated by energy error so it highlights any degradation of energy prediction caused by finetuning. Table 3 shows the evaluation of the geometry optimization task, measured in root mean squared distance (RMSD) from the DFT-optimized structures as well as the F1 score of the stability classification task.

**Table 3.** Evaluation of models on Matbench Discovery, consisting of geometry optimization, measured in RMSD ( $\text{\AA}$ ) and stable/unstable material classification F1 on unique prototypes.

Model	RMSD $\downarrow$	F1 $\uparrow$
Nequix MP	<b>0.0853</b>	<b>0.751</b>
Nequix MP PFT	0.0866	0.747
Nequix MP PFT (no cotrain)	0.0893	0.301

We find that neither PFT models improve performance over the base model. This may be expected, especially for F1, as the finetuning data has no additional energy data from extra chemistries or geometries, so performance of the architecture in this task will be saturated by the original MPtrj training data. However, the co-training is vital for preserving performance on this task; while the non co-trained PFT model worsens RMSD and F1 by 5% and 60% respectively, co-training reduces this difference to less than a 2% increase in geometry optimization RMSD and less than 1% reduction in stability classification F1.

#### 5. Related Work

Several works have proposed the use of Hessians for training interatomic potentials. Fang et al. (2024) first demonstrated

the ability to use AD for phonon prediction through the use of an extended graph, eliminating previously aforementioned issues introduced with self-interaction under finite displacement calculations. They also demonstrate the training of MLIP models on Hessians of small organic molecules. Gangan et al. (2025) demonstrate gradient-based optimization of the classical Stillinger-Weber and EDIP potentials to align with DFT-computed phonon and elastic constant calculations. Amin et al. (2025) introduced a method of model distillation by training smaller MLIP on the analytical Hessians of larger, more accurate models. They also randomly sample columns of the Hessian, but use the Jacobian-vector product of direct-force prediction (i.e. non-energy conserving) models for training.

Lastly, (Burger et al., 2025) proposed the direct prediction Hessians in small organic molecules, bypassing the need for AD or finite-difference calculations. While this may be beneficial in terms of computational cost, models would not be able to benefit from pre-training on the large quantities of available DFT calculations. Furthermore, tasks such as molecular dynamics, which require lower-order forces, or thermal conductivity which require higher third-order force constants, would be infeasible to conduct.

## 6. Discussion

Phonon fine-tuning (PFT) provides a simple, model agnostic method to improve vibrational and thermal property predictions by directly supervising PES curvature. Across MDR Phonon and Matbench Discovery thermal conductivity, we find that reducing Hessian error aligns phonon-derived observables and can transfer to downstream properties that depend on higher-order energy derivatives. PFT uses force-constants already produced by standard DFT phonon workflows, and remains scalable to large supercell via stochastic Hessian column sampling and Hessian-vector products.

We also highlight practical considerations for training on phonon calculations. Training necessitates access to higher-order derivatives through automatic differentiation, e.g. with JAX. Phonon datasets are biased towards relaxed equilibrium structures, and training on these alone may not resolve errors far from equilibrium, motivating continual learning strategies such as co-training.

Scaling PFT to larger and more diverse phonon datasets, or with more accurate base models with larger pre-training datasets could improve accuracy and generalization. Beyond force constants, this work could be extended to other DFT-computed higher-order energy derivatives such as elasticity or anharmonic force constants, as well as with properties obtained through experiment.

## Code and Data

We provide the source code used in this work, trained model weights, and preprocessed training data at <https://github.com/atomicarchitects/nequix/>.

## Acknowledgements

T. K. and T. S. were supported by the National Science Foundation through the AI Research Institutes program Award No. PHY-2019786 (The NSF AI Institute for Artificial Intelligence and Fundamental Interactions, <http://iaifi.org/>) and Award No. DMR-2433348 (The NSF AI Materials Institute, <https://aimi.cornell.edu/>), as well as by the Air Force Office of Scientific Research under Award No. FA9550-24-1-0067. A. G. and J. M. were supported by DOE Project No. DE-SC0024401. M. K. was supported by NSF Graduate Research Fellowship program under Grant No. DGE-1745302. This research used resources of the National Energy Research Scientific Computing Center (NERSC), a Department of Energy User Facility using NERSC awards ERCAP0033254 and ERCAP0036437.

## References

Amin, I., Raja, S., and Krishnapriyan, A. Towards fast, specialized machine learning force fields: Distilling foundation models via energy Hessians. *arXiv preprint arXiv:2501.09009*, 2025.

Barroso-Luque, L., Shuaibi, M., Fu, X., Wood, B. M., Dzamba, M., Gao, M., Rizvi, A., Zitnick, C. L., and Ulissi, Z. W. Open materials 2024 (omat24) inorganic materials dataset and models. *arXiv preprint arXiv:2410.12771*, 2024.

Batatia, I., Benner, P., Chiang, Y., Elena, A. M., Kovács, D. P., Riebesell, J., Advincula, X. R., Asta, M., Avaylon, M., Baldwin, W. J., et al. A foundation model for atomistic materials chemistry. *The Journal of chemical physics*, 163(18), 2025a.

Batatia, I., Lin, C., Hart, J., Kasoar, E., Elena, A. M., Norwood, S. W., Wolf, T., and Csányi, G. Cross learning between electronic structure theories for unifying molecular, surface, and inorganic crystal foundation force fields. *arXiv preprint arXiv:2510.25380*, 2025b.

Bochkarev, A., Lysogorskiy, Y., and Drautz, R. Graph atomic cluster expansion for semilocal interactions beyond equivariant message passing. *Physical Review X*, 14(2):021036, 2024.

Bradbury, J., Frostig, R., Hawkins, P., Johnson, M. J., Leary, C., Maclaurin, D., Necula, G., Paszke, A., VanderPlas, J., Wanderman-Milne, S., and Zhang, Q. JAX: composable transformations of Python+NumPy programs, 2018. URL <http://github.com/jax-ml/jax>.

Burger, A., Thiede, L., Rønne, N., Bernales, V., Vijaykumar, N., Vegge, T., Bhowmik, A., and Aspuru-Guzik, A. Shoot from the hip: Hessian interatomic potentials without derivatives. *arXiv preprint arXiv:2509.21624*, 2025.

Califano, S., Schettino, V., and Neto, N. Lattice dynamics. In *Lattice Dynamics of Molecular Crystals*, pp. 1–40. Springer, 1981.

Chen, C. and Ong, S. P. A universal graph deep learning interatomic potential for the periodic table. *Nature Computational Science*, 2(11):718–728, 2022.

Chen, Y., Fu, L., Yang, S., and Chen, L. Eqnorm MPtrj, 2025. URL <https://github.com/yzchen08/eqnorm>.

Deng, B., Zhong, P., Jun, K., Riebesell, J., Han, K., Bartel, C. J., and Ceder, G. Chgnet as a pretrained universal neural network potential for charge-informed atomistic modelling. *Nature Machine Intelligence*, 5(9):1031–1041, 2023.

Deng, B., Choi, Y., Zhong, P., Riebesell, J., Anand, S., Li, Z., Jun, K., Persson, K. A., and Ceder, G. Systematic softening in universal machine learning interatomic potentials. *npj Computational Materials*, 11(1):9, 2025.

Fang, S., Geiger, M., Checkelsky, J. G., and Smidt, T. Phonon predictions with e (3)-equivariant graph neural networks. *arXiv preprint arXiv:2403.11347*, 2024.

French, R. M. Catastrophic forgetting in connectionist networks. *Trends in cognitive sciences*, 3(4):128–135, 1999.

Fu, X., Wood, B. M., Barroso-Luque, L., Levine, D. S., Gao, M., Dzamba, M., and Zitnick, C. L. Learning smooth and expressive interatomic potentials for physical property prediction. *arXiv preprint arXiv:2502.12147*, 2025.

Gangan, A. S., Cubuk, E. D., Schoenholz, S. S., Bauchy, M., and Krishnan, N. A. Force-field optimization by end-to-end differentiable atomistic simulation. *Journal of Chemical Theory and Computation*, 21(12):5867–5879, 2025.

Giustino, F. Electron-phonon interactions from first principles. *Reviews of Modern Physics*, 89(1):015003, 2017.

Hohenberg, P. and Kohn, W. Inhomogeneous electron gas. *Physical review*, 136(3B):B864, 1964.

Jain, A., Ong, S. P., Hautier, G., Chen, W., Richards, W. D., Dacek, S., Cholia, S., Gunter, D., Skinner, D., Ceder, G., et al. Commentary: The materials project: A materials

genome approach to accelerating materials innovation. *APL materials*, 1(1), 2013.

Kaplan, A. D., Liu, R., Qi, J., Ko, T. W., Deng, B., Riebesell, J., Ceder, G., Persson, K. A., and Ong, S. P. A foundational potential energy surface dataset for materials. *arXiv preprint arXiv:2503.04070*, 2025.

Kasoar, E., Hart, J., Batatia, I., Elena, A., and Csányi, G. ml-peg, November 2025. URL <https://doi.org/10.5281/zenodo.17705093>.

Koker, T., Kotak, M., and Smidt, T. Training a foundation model for materials on a budget. *arXiv preprint arXiv:2508.16067*, 2025.

Liu, R., Liu, E., Riebesell, J., Qi, J., Ong, S. P., and Ko, T. W. MatCalc, January 2024. URL <https://github.com/materialsvirtuallab/matcalc>.

Loew, A., Sun, D., Wang, H.-C., Botti, S., and Marques, M. A. Universal machine learning interatomic potentials are ready for phonons. *npj Computational Materials*, 11(1):178, 2025.

Merchant, A., Batzner, S., Schoenholz, S. S., Aykol, M., Cheon, G., and Cubuk, E. D. Scaling deep learning for materials discovery. *Nature*, 624(7990):80–85, 2023.

Miranda, H. Phonons from density-functional perturbation theory. URL [https://vasp.at/wiki/Phonons\\_from\\_density-functional-perturbation-theory](https://vasp.at/wiki/Phonons_from_density-functional-perturbation-theory). Accessed: 2026-01-09.

Neumann, M., Gin, J., Rhodes, B., Bennett, S., Li, Z., Choubisa, H., Hussey, A., and Godwin, J. Orb: A fast, scalable neural network potential. *arXiv preprint arXiv:2410.22570*, 2024.

Park, Y., Kim, J., Hwang, S., and Han, S. Scalable parallel algorithm for graph neural network interatomic potentials in molecular dynamics simulations. *Journal of chemical theory and computation*, 20(11):4857–4868, 2024.

Pearlmutter, B. A. Fast exact multiplication by the hessian. *Neural computation*, 6(1):147–160, 1994.

Póta, B., Ahlawat, P., Csányi, G., and Simoncelli, M. Thermal conductivity predictions with foundation atomistic models. *arXiv preprint arXiv:2408.00755*, 2024.

Riebesell, J., Goodall, R. E., Benner, P., Chiang, Y., Deng, B., Ceder, G., Asta, M., Lee, A. A., Jain, A., and Persson, K. A. A framework to evaluate machine learning crystal stability predictions. *Nature Machine Intelligence*, 7(6):836–847, 2025.

Schmidt, J., Cerqueira, T. F., Romero, A. H., Loew, A., Jäger, F., Wang, H.-C., Botti, S., and Marques, M. A. Improving machine-learning models in materials science through large datasets. *Materials Today Physics*, 48:101560, 2024.

Tan, C. W., Descoteaux, M. L., Kotak, M., Nascimento, G. d. M., Kavanagh, S. R., Zichi, L., Wang, M., Saluja, A., Hu, Y. R., Smidt, T., et al. High-performance training and inference for deep equivariant interatomic potentials. *arXiv preprint arXiv:2504.16068*, 2025.

Togo, A. MDR phonon calculation database. <https://mdr.nims.go.jp/collections/8g84ms862?locale=en>, 2023a.

Togo, A. First-principles phonon calculations with phonopy and phono3py. *J. Phys. Soc. Jpn.*, 92(1):012001, 2023b. doi: 10.7566/JPSJ.92.012001.

Togo, A. Phono3py input data for 103 compounds calculated using the finite displacement method with PBE, 2025. URL <https://github.com/atztogo/phonondb>.

Togo, A., Chaput, L., Tadano, T., and Tanaka, I. Implementation strategies in phonopy and phono3py. *J. Phys. Condens. Matter*, 35(35):353001, 2023. doi: 10.1088/1361-648X/acd831.

Unke, O. T., Chmiela, S., Sauceda, H. E., Gastegger, M., Poltavsky, I., Schutt, K. T., Tkatchenko, A., and Muller, K.-R. Machine learning force fields. *Chemical Reviews*, 121(16):10142–10186, 2021.

Wood, B. M., Dzamba, M., Fu, X., Gao, M., Shuaibi, M., Barroso-Luque, L., Abdelmaqsoud, K., Gharakhanyan, V., Kitchin, J. R., Levine, D. S., et al. Uma: A family of universal models for atoms. *arXiv preprint arXiv:2506.23971*, 2025.

Yan, K., Bohde, M., Kryvenko, A., Xiang, Z., Zhao, K., Zhu, S., Kolachina, S., SarıTÜRK, D., Xie, J., Arróyave, R., et al. A materials foundation model via hybrid invariant-equivariant architectures. *arXiv preprint arXiv:2503.05771*, 2025.

Yang, H., Hu, C., Zhou, Y., Liu, X., Shi, Y., Li, J., Li, G., Chen, Z., Chen, S., Zeni, C., et al. Mattersim: A deep learning atomistic model across elements, temperatures and pressures. *arXiv preprint arXiv:2405.04967*, 2024.

Zhang, D., Peng, A., Cai, C., Li, W., Zhou, Y., Zeng, J., Guo, M., Zhang, C., Li, B., Jiang, H., et al. Graph neural network model for the era of large atomistic models. *arXiv preprint arXiv:2506.01686*, 2025.

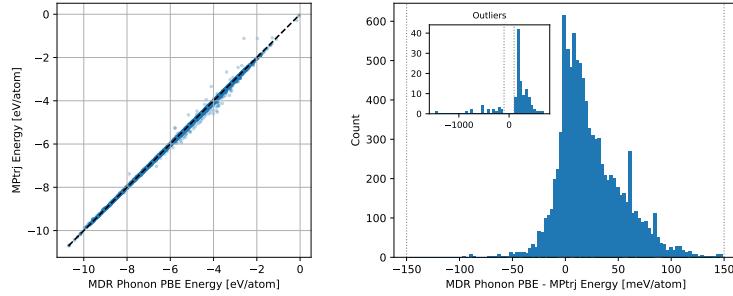
Zhou, Y., Hu, S., Wang, C., Tan, G., and Jia, W.  
MatRIS, 2025. URL <https://github.com/HPC-AI-Team/MatRIS>.

Ziman, J. M. *Electrons and phonons: the theory of transport phenomena in solids*. Oxford university press, 2001.

## A. Appendix

### A.1. Energy discrepancy between PBE MDR and MPtrj datasets

We compare the DFT energies within the PBE recalculation of the MDR Phonon database (Togo, 2023a) by Loew et al. (2025) to those within MPtrj (Deng et al., 2023; Jain et al., 2013). This is performed by selecting the 9959 matching materials across datasets by `mp_id`. For MPtrj we select the final relaxed structure energy, and for both datasets we normalize by number of atoms to offset differences caused by unitcell selection. The energies and the error between them are plotted in Fig. A.1. We find that, while the energies are generally in agreement, there exists a slight shift between the two datasets with a MAE of 31.60 meV/atom. Furthermore there are several outliers with energy differences up to 1.5 eV/atom. We suspect this systemic shift is the cause of the consistent  $\approx 30$  meV/atom MAE reported with all interatomic potentials used in Loew et al. (2025), and energy MAE seen in Fig. 3.



**Figure A.1. MDR Phonon and MPtrj energy comparison.** Left: Parity plot of energies from MPtrj (Deng et al., 2023; Jain et al., 2013) vs. PBE MDR Phonon (Togo, 2023a; Loew et al., 2025). Right: Histogram of differences in energy between materials in two datasets. Absolute differences less than 150 meV/atom are shown in the main plot with the remaining outliers in the inset.

### A.2. Hyper-parameters

Table A.1 shows the hyper-parameters and method for selection for the Nequix PFT models.

Table A.1. Hyper-parameters used and rationale behind selection

Hyper-parameter	Nequix MP PFT	Nequix MP PFT (no cotrain)	Notes/Rationale
Pretrained Model	Nequix MP	Nequix MP	See Koker et al. (2025) for architecture details.
Learning rate	0.0003	0.0003	Selected from {0.003, 0.001, 0.0003, 0.0001} based on validation performance early in training.
Optimizer	AdamW	AdamW	Standard optimizer.
Weight decay	0.001	0.001	From Nequix MP.
Phonon energy weight $\lambda_E$	0	20	From Nequix MP.
Phonon force weight $\lambda_F$	20	20	From Nequix MP.
Phonon stress weight $\lambda_\sigma$	5	5	From Nequix MP.
Phonon force constant weight $\lambda_\Phi$	100	100	Not tuned, selected to be larger than other weights.
Cotrain energy weight $\lambda_E$	500	n/a	Started with Nequix MP value $\times 10$ , increased until MPtrj validation energy does not diverge.
Cotrain force weight $\lambda_F$	200	n/a	Nequix MP value $\times 10$ .
Cotrain stress weight $\lambda_\sigma$	50	n/a	Nequix MP value $\times 10$ .
Cotrain ratio $K$	4	n/a	Trade off between training time and overfitting to phonon data; 4 cotraining steps for every phonon step is reasonable.
Batch size	16	16	Maximum for GPU memory.
Number of epochs	200	200	Not tuned, based on GPU budget. May benefit from longer training, but validation metrics are mostly converged at this duration.

### A.3. Elastic properties

We also benchmark the performance effects of PFT on materials elastic properties A.2. For the benchmark we use MatCalc’s Elasticity Calc(Liu et al., 2024) based the settings as described in and the ml-peg package (Kasoar et al., 2025; Batatia et al., 2025b) whereby the elastic tensor and derived properties like the bulk  $K$  and shear  $G$  modulus are obtained. Voigt-Reuss-Hill (VRH) average is used to report the final values for which we compute the metrics. We observe that the error on shear modulus is lowered for the PFT model while bulk modulus is worsened. Overall, performance on both modulus is improved for the PFT model. We also report the failure rate for the models for which MatCalc calculations fail to converge..

Table A.2. Evaluation of models on Materials Project elastic properties containing 12,122 different materials. Metrics are mean absolute error of bulk modulus  $K$  and shear modulus  $G$  in units of GPa, correlation score, as well as success rate of computations. To compare the models in a fair way we only consider the materials for which both models return a valid output.

Model	$K$ MAE	$G$ MAE	$R_K^2$	$R_G^2$	Failure
Nequix MP	<b>15.03</b>	18.16	<b>0.89</b>	0.44	18.37%
Nequix MP PFT	16.49	<b>16.41</b>	0.87	<b>0.54</b>	18.21%

### A.4. Phonon bandstructures

Figure A.2 shows additional examples of bandstructures calculated with Nequix MP PFT.

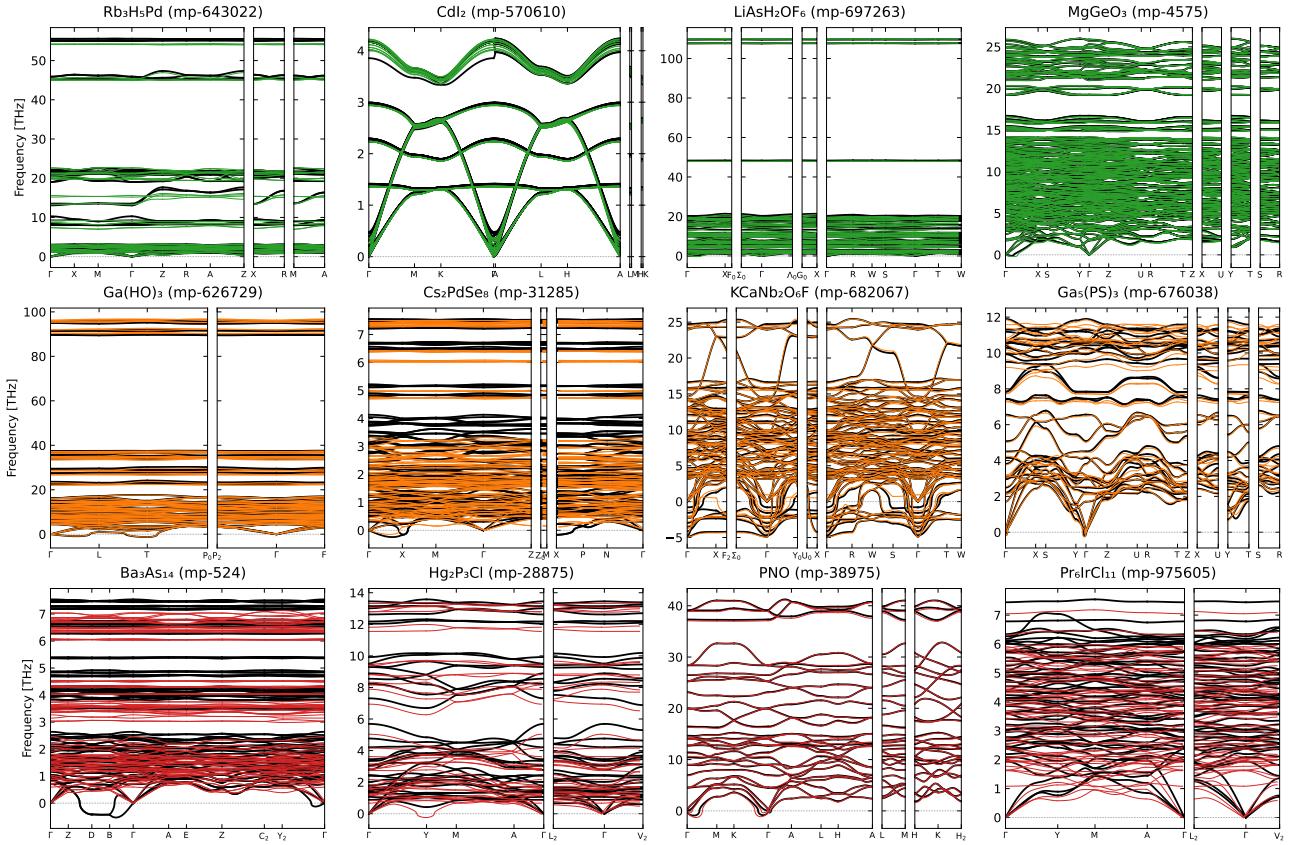


Figure A.2. Phonon bandstructures with Nequix MP PFT. Following (Fang et al., 2024), we display predicted bandstructures (with DFT in black) of four randomly selected materials from each tercile of Hessian error in the PBE MDR phonon test set. The top row shows lowest error, and the bottom row shows highest error materials. Bandstructures are calculated with phonopy (Togo et al., 2023). Note that the number bands is  $3 \times$  the number of atoms in the unit cell.