

Digital Discovery

Accepted Manuscript

This article can be cited before page numbers have been issued, to do this please use: D. Bidoggia, N. Manko, M. PERESSI and A. Marrazzo, *Digital Discovery*, 2025, DOI: 10.1039/D6DD00005C.



This is an Accepted Manuscript, which has been through the Royal Society of Chemistry peer review process and has been accepted for publication.

Accepted Manuscripts are published online shortly after acceptance, before technical editing, formatting and proof reading. Using this free service, authors can make their results available to the community, in citable form, before we publish the edited article. We will replace this Accepted Manuscript with the edited and formatted Advance Article as soon as it is available.

You can find more information about Accepted Manuscripts in the [Information for Authors](#).

Please note that technical editing may introduce minor changes to the text and/or graphics, which may alter content. The journal's standard [Terms & Conditions](#) and the [Ethical guidelines](#) still apply. In no event shall the Royal Society of Chemistry be held responsible for any errors or omissions in this Accepted Manuscript or any consequences arising from the use of any information it contains.

AiiDA-TrainsPot: Towards automated training of neural-network interatomic potentials

Davide Bidoggia,^{*a†} Nataliia Manko,^{*ab}, Maria Peressi^a and Antimo Marrazzo^{b‡}

Crafting neural-network interatomic potentials (NNIPs) remains a complex task, demanding specialized expertise in both machine learning and electronic-structure calculations. Here, we introduce AiiDA-TrainsPot, an automated, open-source, and user-friendly workflow that streamlines the creation of accurate NNIPs by orchestrating density-functional-theory calculations, data augmentation strategies, and classical molecular dynamics. Our active-learning strategy leverages on-the-fly calibration of committee disagreement against *ab initio* reference errors to ensure reliable uncertainty estimates. We use electronic-structure descriptors and dimensionality reduction to analyze the efficiency of this calibrated criterion, and show that it minimizes both false positives and false negatives when deciding what to compute from first principles. AiiDA-TrainsPot has a modular design that supports multiple NNIP backends, enabling both the training of NNIPs from scratch and the fine-tuning of foundation models. We demonstrate its capabilities through automated training campaigns targeting pristine and defective carbon allotropes, including amorphous carbon, as well as structural phase stability in monolayer $W_xMo_{1-x}Te_2$ alloys.

1 Introduction

In computational materials science, many key properties—such as phase transitions, diffusion, viscosity, thermal transport—arise from atomic and molecular motion over time. Chemical reactions are typically rare events, requiring long simulation times and enhanced sampling techniques. Because these properties emerge from dynamic processes rather than static configurations, molecular dynamics (MD) simulations are essential for capturing their time-dependent behavior and understanding material performance under different conditions. MD simulations enable the study of crucial phenomena in materials and molecular systems, such as heat conduction, ion transport in batteries, crack propagation in materials, just to name a few, and the behavior of soft matter and biomaterials under different conditions. Fully *ab initio* molecular dynamics (AIMD) simulations can be highly accurate but computationally expensive, as they require solving the Schrödinger equation every timestep to determine atomic forces. This makes AIMD impractical for large systems or long timescales. On the other hand, empirical interatomic potentials, which approximate atomic interactions with predefined functional forms, are computationally efficient but often lack the accuracy needed for quantum mechanical effects. Additionally, these potentials are typically parameterized for specific materials and struggle to generalize across diverse chemical environments, limiting their predictive power for complex or novel systems. For a long time, the Car-Parrinello method (CPMD)¹ offered a compromise between AIMD and empirical potentials by evolving both electronic and ionic degrees of freedom simultaneously with a Lagrangian formalism, reducing the need for explicit electronic structure calculations at every step. While CPMD improves efficiency compared to traditional AIMD, it still remains computationally demanding and is limited by the need for careful tuning of fictitious electron

mass parameters, which can affect accuracy and stability.

Machine learning (ML) and neural-network interatomic potentials (NNIPs) have revolutionized the field, by bridging the gap between *ab initio* accuracy and computational efficiency². These models are trained on high-fidelity density functional theory (DFT) data, allowing them to capture complex quantum mechanical interactions with far greater accuracy than traditional empirical potentials. Unlike fixed functional forms, ML potentials can flexibly generalize to diverse chemical environments while maintaining computational costs significantly lower than AIMD. This breakthrough has enabled large-scale and long-timescale simulations of materials with near-quantum accuracy, essentially extending the scope of what can be simulated from first principles.

Over the last years, accuracy and data efficiency of ML interatomic potentials have improved remarkably, often at the price of increased algorithmic complexity. In this context, NNIPs based on equivariance have emerged as particularly promising and several architectures have been proposed, including NequIP³, Allegro⁴, MACE⁵ and Point Edge Transformer (PET)⁶. More recently, foundation models leveraging large-scale pretraining on diverse chemical datasets, provided a path towards improved transferability, data efficiency, and accuracy across a wide range of materials and molecular systems^{7,8}.

As of today, the training of an accurate NNIP remains a complex and time-consuming task. First, high-quality NNIPs require training datasets of the order of thousands of supercell single-point *ab initio* calculations with hundreds of atoms in the unit cell; millions if foundation models for the entire periodic table are targeted. Most notably, the accuracy and extrapolation capabilities of NNIPs hinge on the careful choice of the training structures, which have to be sufficiently diverse and abundant to avoid overfitting. While foundation models promise transferability across chemical space, fine-tuning them to high accuracy for a given material family still requires curated datasets.

In this work, we introduce AiiDA-TrainsPot, an automated, open-source, and user-friendly framework for training neural-network interatomic potentials (NNIPs). It integrates automated

^a Dipartimento di Fisica, Università di Trieste, I-34151 Trieste, Italy

^b Scuola Internazionale Superiore di Studi Avanzati (SISSA), I-34136 Trieste, Italy

* These authors contributed equally to this work.

† davide.bidoggia@units.it

‡ amarrazz@sisssa.it



workflows for DFT calculations with neural-network training and classical MD to systematically explore the potential-energy surface (PES) via random distortions, strain, interfaces, neutral vacancies, and trajectories across a range of temperatures and pressures.

Widely used platforms such as DP-GEN⁹, SchNetPack¹⁰, FLARE¹¹, AL4GAP¹², ASPARAGUS¹³, and others implement or support active-learning strategies—including on-the-fly learning during molecular dynamics, uncertainty-driven sampling, and ensemble-based selection criteria—thereby demonstrating their effectiveness for ML interatomic potentials. Active-learning training is nowadays widely used, also thanks to key contributions such as the extrapolation-grade framework for moment tensor potentials¹⁴ and on-the-fly force-field learning during MD^{15,16}. More recent efforts have started incorporating these methods into standardized open-source workflow implementations, also improving exploration strategies to broaden configurational coverage, including uncertainty-driven dynamics for sampling¹⁷ and the integration of enhanced sampling (e.g., metadynamics) into active-learning loops¹⁸. As of today, an ecosystem of active-learning tools for ML interatomic potentials exists (including DP-GEN, SchNetPack and others), spanning different model families and exploration protocols. At the same time, most existing implementations remain closely coupled to specific potential families or training engines, while typically lacking automatic restart strategies and systematic provenance tracking across the full training pipeline—including dataset generation, model selection, and first-principles calculations. In addition, while committee-based methods are widespread in the community, they are not guaranteed to provide an accurate quantitative proxy for the true deviation from the reference quantum mechanical calculations^{19,20}. Here we focus on a complementary contribution that fills the gap: a provenance-tracked, restartable and automated workflow across all stages of training and simulation, and that *calibrates* committee disagreement against *ab initio* errors on-the-fly to provide quantitative selection thresholds and uncertainty estimates. We provide in Table S1 of the Supporting Information a detailed feature-by-feature comparison between AiiDA-TrainsPot and the widely used active-learning frameworks DP-GEN, SchNetPack, FLARE, and AL4GAP.

Indeed, AiiDA-TrainsPot is conceived as a modular and extensible automation layer rather than just as a active-learning protocol. It offers (i) code-agnostic modularity across quantum engines, ML architectures, and MD codes; (ii) an extensive suite of automated dataset-augmentation strategies (defects, slabs, clusters, substitutions, alloys); and (iii) a *calibrated* committee-disagreement scheme that provides quantitative uncertainty estimates, including in production runs. The architecture deliberately decouples sampling strategies from model training and first-principles backends, facilitating the future integration of additional quantum engines, ML frameworks, and MD protocols within the same automated workflow.

We validate the method through fully automated NNIP training and fine-tuning campaigns on a diverse set of carbon allotropes, including amorphous carbon, and on structural phase stability in monolayer $W_xMo_{1-x}Te_2$ alloys, achieving state-of-the-art ac-

curacy and data efficiency. Combined with AiiDA's reproducibility infrastructure, these features position AiiDA-TrainsPot as both an empowering tool for domain scientists and a robust platform for future foundation-model development.

2 Results and discussion

2.1 Automation strategy

AiiDA-TrainsPot is based on a two-stage augmentation process shown in Fig. 1. In a typical use case, users provide a handful of structures—from one up to tens—that are augmented to the order of thousands through structural manipulations. This happens in the first stage, where about a thousand of structures are calculated with *ab initio* methods (here DFT) and used to train the first generation of NNIPs. The second stage employs MD simulations at different thermodynamic conditions to accurately sample the basins of the PES, thus exploring regions that are particularly relevant for applications. All MD trajectories are obtained with the NNIPs trained in the previous workflow step. Some structural configurations are sampled from MD trajectories and calculated *ab initio* to train a second generation of NNIPs. The choice of the structures to label with *ab initio* results is based on the committee disagreement, i.e., the spread in the predictions of a committee of NNIPs initialized with different seeds and trained on the same data. At each iteration a new generation of potentials is trained and the committee disagreement is calibrated on actual deviations between NNIP predictions and the baseline level of theory, i.e., DFT.

The active learning loop continues until errors on energy, forces and stress tensor are below a user-defined threshold. We emphasize that AiiDA-TrainsPot supports multiple use cases depending on the available input data, which goes beyond the generation of a NNIP from scratch and includes the fine-tuning of foundation models (see Sec. 2.2). In the following, we discuss in detail each step of the workflow.

2.1.1 Input structures

AiiDA-TrainsPot can start from a small set of initial atomistic structures $\{\mathcal{X}\}^{(0)}$, determined by boundary conditions (periodic vs. open), cell parameters, atomic species and atomic positions. The number and diversity of input structures should reflect the target applications: for example, the study of temperature-dependent properties of diamond might require a single input structure, the development of NNIP for all carbon allotropes would probably include at least all known crystalline prototypes of carbon, while universal (a.k.a. foundation) models for the entire periodic table might require tens of thousands of input structures, which could be obtained from computational materials databases such as Materials Cloud²¹, Materials Project²², or crystallographic databases such as ICSD²³, COD^{24,25}, MPDS^{26–28}. While the user is responsible for providing these fundamental structures, the workflow progresses with automatic data augmentation to enhance dataset diversity without requiring exhaustive manual curation.

2.1.2 Dataset augmentation

In the dataset augmentation stage, additional structures are generated by manipulating the initial set $\{\mathcal{X}\}^{(0)}$. All manipulations



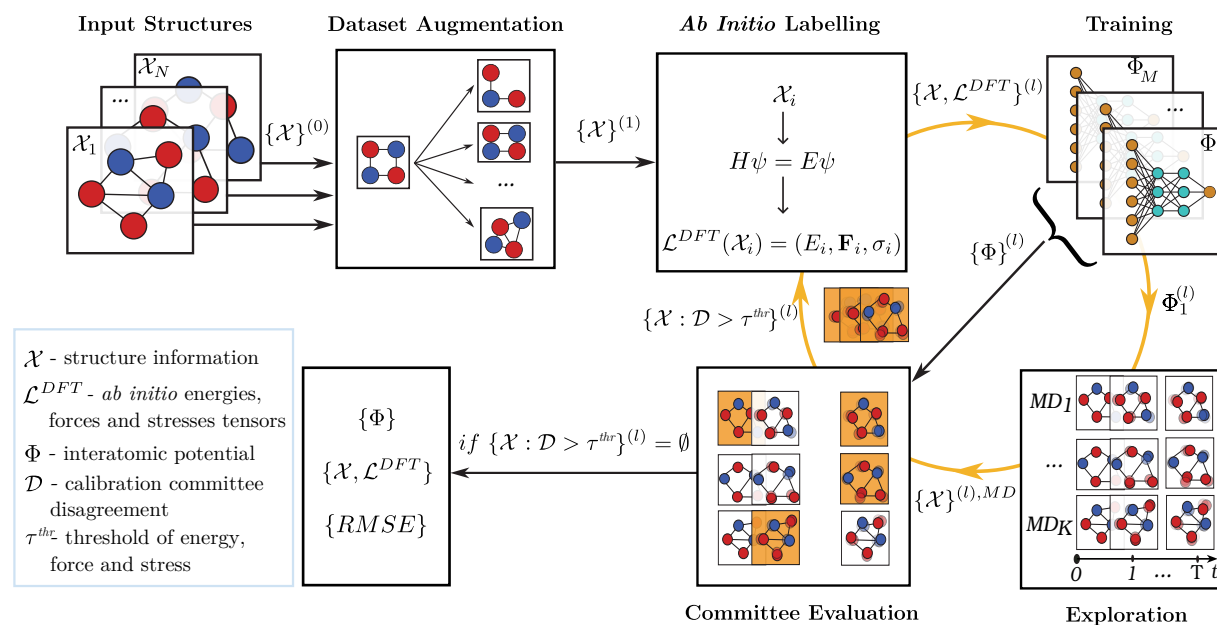


Fig. 1 Schematic representation of the AiIDA-TrainsPot automated workflow. Initial input structures ($\{\mathcal{X}\}^{(0)}$) can be augmented by creating configurations with random distortions, strain, vacancies, cluster and slab extraction. *Ab initio* calculations are performed on these configurations ($\{\mathcal{X}_i\}$) that are thus labeled by energies, forces, and stress tensors which constitute the *reference* data ($\{\mathcal{L}_i\}$). A committee of potentials ($\{\Phi_j\}$) is trained on these configurations and committee evaluation is used to compare their predictions ($\mathcal{L}_k(\Phi_j)$) on structures (\mathcal{X}_k) that are extracted from MD trajectories at different temperatures and pressures obtained through one of the potentials of the committee (Φ_1). If for a structure \mathcal{X}_k the disagreement (\mathcal{D}_k) averaged over the committee exceeds a threshold, the structure is added to the *reference* training dataset and the *ab initio* labelling is performed. This iterative process continues until convergence is achieved and workflow outputs labeled structures, trained potentials, and root mean square errors (RMSE) for energies, forces, and stress tensor components.

can be controlled through customizable parameters to tailor the augmentation process according to specific user needs; we group them in the following categories:

- **Supercells:** Initial structures are replicated aiming to ensure cells larger than a minimum threshold value (default: 18 Å, corresponding to twice the MACE default receptive field) while keeping the total atom count below a user-defined maximum limit (default: 450 atoms).
- **Random distortions:** Atomic positions are perturbed with random displacements, where the magnitude follows a uniform distribution up to a user-defined fraction (default: 30%) of the original nearest-neighbor distance. This introduces configurations away from equilibrium while preventing unphysically close atoms, which could lead to large forces that are difficult—and potentially uninformative—for NNIPs to learn.
- **Strain:** Strain can be applied to crystal structures by rescaling lattice parameters by a factor randomly sampled from uniform distribution between a user-defined range (default from -20% to +60%). This is key for predicting elastic properties.
- **Vacancies:** Vacancies are created by removing atoms at randomly selected sites, to learn about defect energetics and local relaxations around missing atoms. By default, vacancies

are introduced in 30% of the randomly distorted structures, with 2 atoms removed for each configuration.

- **Clusters:** Atomic clusters are constructed by assembling atoms with a next-neighbor distance between 1 and $\sqrt{3}$ times a user specified distance (default: 1.5 Å). This creates non-periodic environments that can become useful for training potentials capable of describing isolated molecules or clusters, and more in general surface or edge terminations.
- **Slabs:** Slabs are created by cutting bulk supercells along selected crystallographic directions (default: (111), (110) and (100)) ensuring a minimum slab thickness (default: 10 Å) unless a maximum number of atoms (default: 450) is reached. Those structures allow the NNIPs to learn and predict surface energetics, surface-specific forces, reconstructions and other relaxation phenomena.
- **Isolated atoms:** Single-atom configurations are included to establish reference energies for accurate calculations of the dissociation limit. These structures are computed using the same DFT settings as the rest of the dataset to ensure consistency, though we note that the procedure does not take into account the actual magnetic configuration for some atomic species.



- **Atomic substitutions:** For multi-species systems, randomly selected atoms of different elements are swapped to create chemical disorder and explore different local chemical environments. The user can define both the fraction of previously generated structures to undergo substitutions (default: 20%) and the number of swaps per structure as a fraction of total atom count (default: 20%). This helps to explore various chemical environments, substitutional defects, and atomic site preferences, ensuring robustness across different chemical compositions within the same structural motif.
- **Alloys:** Alloy configurations are generated by randomly mixing atomic species (specified via `alloy_species`), while optionally keeping some species fixed (specified via `fixed_species`). If no target compositions are specified, the workflow samples random alloy configurations with concentrations spanning the full range from 0 to 1.

The presented dataset augmentation techniques are applied by default; however, users can choose to apply only a subset of them or even skip the dataset augmentation stage altogether. After this stage the resulting dataset $\{\mathcal{X}\}^{(1)}$ integrates structures of different dimensions and boundary conditions: fully periodic (bulks), partially periodic (surfaces, nanowires, or 2D materials), and non-periodic (molecules, clusters) configurations. However, since DFT and MD calculations are performed in full periodic boundary conditions, for all structures that are non-periodic at least along one direction, the workflow ensures the presence of an appropriate vacuum buffer (default: 15 Å thick) along such directions in order to eliminate spurious interactions between periodic images.

2.1.3 *Ab Initio* Labelling

After the data augmentation stage, AiiDA-TrainsPot starts the active learning loop, which is represented by the orange circle in Fig. 1. Each structure \mathcal{X}_i in the augmented dataset is labeled through DFT calculations to obtain high-fidelity reference values for energies, forces, and stress tensors. We use the compact notation $\mathcal{L}^{DFT}(\mathcal{X}_i) = (E(\mathcal{X}_i), \mathbf{F}(\mathcal{X}_i), \boldsymbol{\sigma}(\mathcal{X}_i))$ to represent these computed properties. In subsequent sections, we denote specific quantities of interest as \mathcal{L}_α where $\alpha \in \{E, \mathbf{F}, \boldsymbol{\sigma}\}$. Ultimately, *ab initio* calculations directly determines an upper bound for the accuracy and precision of the trained NNIPs. While the overall accuracy is typically limited by considerations of computational efficiency and resources, precision can be substantially improved by enforcing well converged calculations and a consistent choice of key simulation parameters over cells of different sizes and dimensions. In this context, even if the workflow allows users to have full control over the DFT level of theory, by default AiiDA-TrainsPot enforces the use of well established simulations protocols originally introduced for high-throughput calculations. In particular, PBE pseudopotentials and cutoff parameters are given by version 1.3 of the branch of the Standard Solid State Pseudopotentials (SSSP) library optimized for precision^{29–31}, while the reciprocal space k-point density and smearing follow the *stringent* protocol defined by Nascimento et al.³².

By default, the workflow does not incorporate additional van der Waals (vdW) corrections at the DFT level, since NNIPs would

in principle require rather large radial cutoffs to accurately learn long-range interactions from the training data. However, for systems where dispersion forces are critical (e.g., layered materials or molecular crystals), users can enable empirical vdW corrections (such as Grimme-D2, D3) during subsequent MD simulations³³.

2.1.4 Training neural-network interatomic potentials

The labeled dataset $\{\mathcal{X}_i, \mathcal{L}^{DFT}(\mathcal{X}_i)\}^{(1)}$ is used to train a committee of M NNIPs $\{\Phi_j\}_{j=1}^M$, each with identical architecture but initialized with different random seeds. Prior to training, all structures are systematically partitioned into three subsets, ensuring representative sampling across different structural motifs while maintaining similar distributions of atomic environments:

- Training set (default: 80%): Used for model parameter optimization through gradient-based learning;
- Validation set (default: 10%): Used for hyperparameter tuning, early stopping decisions, and selection of optimal checkpoints during training;
- Test set (default: 10%): Reserved exclusively for final model evaluation, providing an unbiased assessment of generalization performance.

Throughout the active learning iterations, each structure remains in its initially assigned set, ensuring that the test set remains completely independent from the training and validation sets for reliable performance assessment. The training is performed using either MACE or Metatraining³⁴ with default hyperparameters, though these can be fully customized by the user of AiiDA-TrainsPot to suit specific requirements.

2.1.5 Exploration by molecular dynamics

After a committee of NNIPs is trained, the workflow employs MD simulations to systematically explore the potential energy landscape. This exploration phase is critical for identifying configurations where the NNIPs might have insufficient accuracy, thus guiding the selection of additional structures for *ab initio* calculations in subsequent training iterations.

At each iteration, one NNIP (Φ_1) is randomly selected from the committee to perform the MD simulations. By default, a set of 20 structures is randomly sampled from the initial augmented dataset $\{\mathcal{X}\}^{(1)}$ to serve as starting configurations. This selection strategy deliberately uses structures from the initial dataset rather than from the most recent active learning iteration dataset to ensure sampling of a configuration space that is not biased by previous explorations and remains representative of the user's target application domain. Users can customize this selection by specifying either the number of structures to sample or providing an explicit list of starting configurations.

All MD simulation parameters are fully customizable, including ensemble type, temperature, pressure, simulation time, and timestep. By default, simulations run for 10 ns with a 1 fs timestep, while temperatures and pressures range from 0 to 1000 K and -5 to 5 kbar respectively, systematically sampling diverse thermodynamic conditions. AiiDA-TrainsPot automatically selects



an isothermal-isobaric (NPT) ensemble with barostats acting only on periodic directions for bulk systems, or an isochoric-isothermal (NVT) ensemble for non-periodic systems. Additionally, even though dispersion corrections are disabled by default, users can activate Grimme's D2 van der Waals correction³⁵ during MD simulations, with coupling parameters automatically selected based on the atomic species present in the system.

To minimize correlations between sampled configurations, trajectory frames are extracted at regular intervals (default: 1 ns), ensuring the collection of statistically uncorrelated dataset that efficiently represent the accessible regions of the PES.

2.1.6 Committee Evaluation

This stage aims at identifying structures that are poorly predicted by the NNIPs; those are good candidates to be labeled with *ab initio* calculations and included in the training dataset. However, while Bayesian neural networks (NNs) come with a well-defined probabilistic uncertainty quantification, no such Bayesian error estimation can be defined for NNIPs³⁶. Here, similar to what has been done in recent works (e.g.,^{9,37,38}), we use the spread of predictions from a committee of NNIPs as a proxy for uncertainty quantification. Hence, at this step, each structure \mathcal{X}_k sampled from MD trajectories is evaluated by all NNIPs of the committee and we quantify the uncertainty of prediction through the committee disagreement metric $\mathcal{D}_\alpha(\mathcal{X}_k)$, which is calculated separately for each property α (energy, forces—averaged over all the atoms and components—and stress tensor):

$$\mathcal{D}_\alpha(\mathcal{X}_k) = \sqrt{\frac{1}{M} \sum_{j=1}^M \left(\mathcal{L}_\alpha^{\Phi_j}(\mathcal{X}_k) - \overline{\mathcal{L}_\alpha^\Phi}(\mathcal{X}_k) \right)^2} \quad (1)$$

where M is the committee size, $\mathcal{L}_\alpha^{\Phi_j}(\mathcal{X}_k)$ represents the prediction of property α from potential Φ_j for structure \mathcal{X}_k , and $\overline{\mathcal{L}_\alpha^\Phi}(\mathcal{X}_k)$ denotes the mean prediction across all committee members. Structures exhibiting disagreement beyond a specific threshold τ_α^{thr} are flagged as uncertain and prioritized for quantum mechanical labelling.

The error $\varepsilon_\alpha(\mathcal{X}_k)$ with respect to the *ab initio* labels is

$$\varepsilon_\alpha(\mathcal{X}_k) = \sqrt{\frac{1}{M} \sum_{j=1}^M \left(\mathcal{L}_\alpha^{\Phi_j}(\mathcal{X}_k) - \mathcal{L}_\alpha^{DFT}(\mathcal{X}_k) \right)^2} \quad (2)$$

and is not, at least in principle, strictly correlated to the committee disagreement, although empirical evidence suggests they might be linearly related^{19,20} (as we show and discuss more in detail in Sec. 2.3 through a validation study on carbon allotropes). After linear regression, the user-defined error tolerance threshold ε_α^{thr} (default 1 meV/atom, 100 meV/Å, 1 meV/Å³ for energy, forces and stress tensor respectively) is transformed into an equivalent disagreement threshold τ_α^{thr} that serves as the selection criterion for new structures by $\tau_\alpha^{thr} = a_\alpha \varepsilon_\alpha^{thr}$, where coefficients a_α are the slopes determined from fitting. The calibration, made at each iteration on the already *ab initio* labeled structures and compared with $\mathcal{D}_\alpha(\mathcal{X}_i)$ computed with the last generation potentials, ensures that the uncertainty quantification mechanism remains effective throughout the iterative learning process, as the NNIP

committee's overall accuracy improves with each active learning cycle. Notably, the calibrated committee disagreement $a_\alpha \mathcal{D}_\alpha(\mathcal{X}_k)$ can be used not only in the active learning scheme for the selection of worse predicted structure, but also in production runs, where it can provide an estimate of the uncertainty of the predictions of the NNIPs.

For large exploration sets $\{\mathcal{X}_k\}$, computational constraints often necessitate limiting the number of structures labeled in each active learning iteration. However, selecting only the structures with the highest disagreement values may not be optimal, as these could represent configurations far from the original dataset and the PES region of interest to the user. While such structures cannot be discarded a priori—since they may represent poorly predicted configurations that are nonetheless close to relevant regions of the PES—a more balanced selection strategy is needed. Therefore, we randomly select structures from those that exceed the threshold τ_α for any property α . This approach naturally prioritizes structures closer to the original dataset, as they are more likely to be represented in the configurations extracted from MD trajectories, while still capturing the most uncertain predictions. By default, at each active learning iteration a maximum number of 1000 new structures are selected. The active learning loop continues until one of two termination criteria is met: either all structures exhibit disagreement below the specified threshold τ_α^{thr} , indicating the achievement of the desired confidence across the configuration space of interest, or the maximum number of iterations L is reached. The final output of AiiDA-TrainsPot includes the latest NNIP committee $\{\Phi\}$, the full *ab initio* dataset $\{X, \mathcal{L}^{DFT}\}^{final}$ and quantitative RMSE performance metrics for each potential.

2.2 Code implementation

AiiDA-TrainsPot is built on the AiiDA infrastructure^{39,40}, providing a robust framework for managing complex computational workflows with full tracking of data provenance. The platform automatically persists all computational steps, input parameters, and generated data, ensuring complete reproducibility of results. The framework efficiently orchestrates job submissions across high-performance computing resources while leveraging existing AiiDA plugins for quantum mechanical and classical simulations.

The workflow architecture follows a hierarchical design of nested AiiDA WorkChains (Fig. 2), enabling both end-to-end automation and selective execution of individual components. The top-level `TrainsPotWorkChain` coordinates five specialized sub-processes that correspond to the major stages of a training campaign:

- `DatasetAugmentationWorkChain` for generation of highly-diverse training structures
- `AbInitioLabellingWorkChain` for quantum mechanical calculations
- `TrainingWorkChain` for training NNIP committees
- `ExplorationWorkChain` for exploration of the PES based on MD



- `EvaluationCalculation` for committee-based error estimation and identification of structures for which the NNIP yields predictions with low accuracy.

A key advantage of our implementation strategy and of the use of AiiDA as a workflow engine is the extensive reuse of existing AiiDA plugins, minimizing duplication of software and ensuring robustness through well-tested components. The `AbInitioLabellingWorkChain` calls `PwBaseWorkChain` from the AiiDA-Quantum ESPRESSO plugin^{39,40}, while the `ExplorationWorkChain` leverages `LammpsBaseWorkChain` from the AiiDA-LAMMPS plugin⁴¹. For specialized functionality not available in existing plugins, we develop custom components such as the `DatasetAugmentationWorkChain`, which implements various structure manipulation techniques through dedicated `calcfunctions` based on the Atomic Simulation Environment (ASE)⁴². Similarly, the `TrainingWorkChain` interfaces with NN and ML training codes and can be configured by the user to choose the training instrument, either `MaceTrainWorkChain` for MACE or `MetaTrainWorkChain` for Metatrain, handling all preprocessing, training, and postprocessing steps in a fully automated manner. We note that Metatrain currently implements both the PET architecture and other ML architectures not all based on neural-networks, such as Sparse Gaussian Approximation Potentials (GAP)⁴³ and Behler-Parrinello neural networks with SOAP features (SOAP BPNN)⁴⁴.

To optimize computational efficiency, AiiDA-TrainsPot implements parallel execution strategies within three main WorkChains (`AbInitioLabellingWorkChain`, `TrainingWorkChain`, and `ExplorationWorkChain`). Multiple DFT calculations, neural-network training sessions, and MD simulations are submitted concurrently, with results collected and analyzed collectively before proceeding to the next workflow stage.

This modular design provides several advantages beyond efficient execution. First, it offers maximum flexibility through multiple entry points, allowing users to bypass specific stages depending on their needs (see Fig. 3). For instance, users who would like to leverage access to existing datasets of labeled structures can proceed directly to training. The workflow also supports fine-tuning pretrained models, including foundation models.

Second, the architecture enables selective execution of individual components, permitting users to run specific stages independently. This is particularly valuable for scenarios such as structure labelling with DFT calculations without proceeding to the training phase, or for evaluating the accuracy of existing potentials on new trajectories using the committee disagreement.

Finally, the architecture maintains extensibility through AiiDA's plugin system, facilitating future integration with additional quantum engines, classical MD codes, or emerging ML frameworks. Individual components like `MaceTrainCalculation`, `MetaTrainCalculation` or `EvaluationCalculation` can be used as standalone tools outside the high-level WorkChain, making AiiDA-TrainsPot both a comprehensive platform and a flexible toolkit for specialized tasks in interatomic potential development.

To efficiently manage large datasets within the AiiDA framework, we introduce a custom AiiDA data type, `PESData`, a subclass

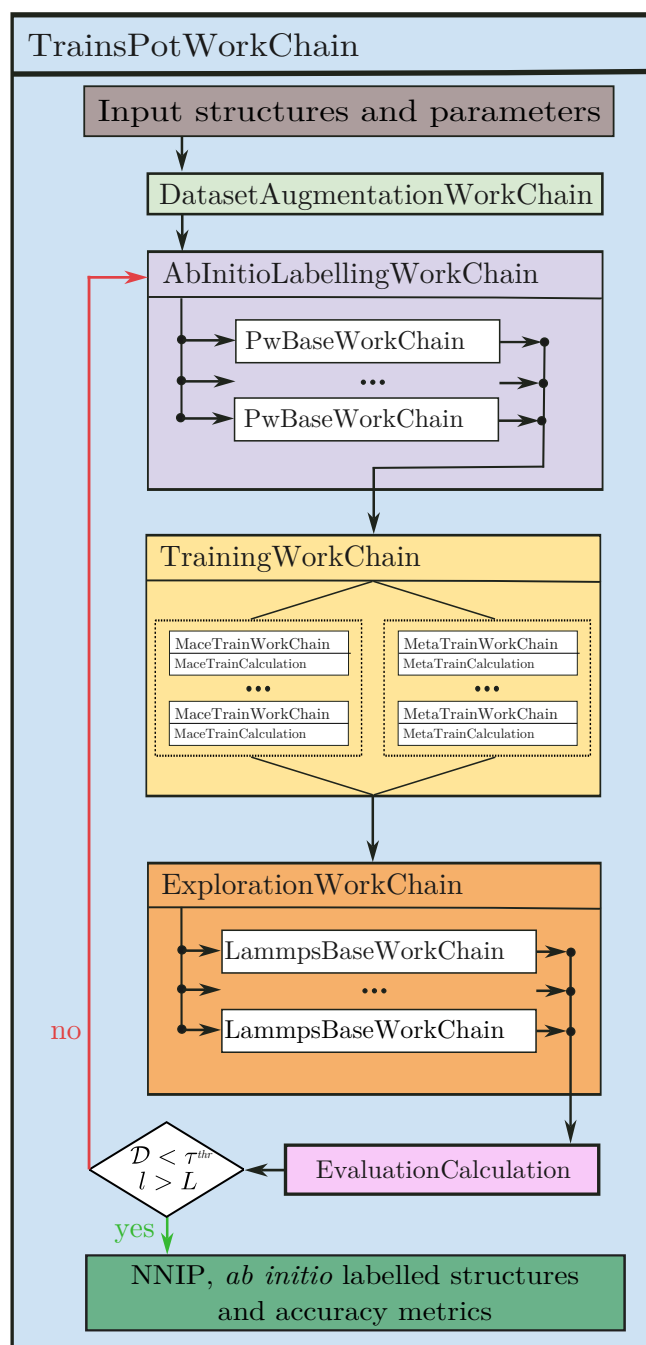


Fig. 2 Schematic representation of the TrainsPotWorkChain with its various computational tasks. The workflow begins with the initialization phase, where input structures, parameters, and control flags are set to determine which steps of the workflow will be executed. If dataset augmentation is enabled, the `DatasetAugmentationWorkChain` generates additional configurations. Next, the workflow enters the active learning loop. Within each iteration, the `AbInitioLabellingWorkChain` is called to label newly generated configurations by performing automated electronic structure calculations (`PwBaseWorkChain`). The `TrainingWorkChain` is then invoked to train interatomic potentials using MACE (`MaceTrainWorkChain`) or Metatrain (`MetaTrainWorkChain`). Subsequently, the `ExplorationWorkChain` executes MD simulations via LAMMPS (`LammpsBaseWorkChain`) to generate new configurations for further refinement. The `EvaluationCalculation` assesses the performance of the trained models using calibrated committee disagreement to determine whether additional iterations are required or not.



of `aiida.orm.Data`. This specialized data type enables the storage and manipulation of extensive sets of atomic structures along with their associated properties, while maintaining full compatibility with AiiDA's provenance tracking system.

PESData offers significant advantages over existing data types such as `TrajectoryData`, which is limited to structures containing identical numbers of atoms. Our implementation can seamlessly handle datasets with varying numbers of atoms per structure, making it ideal for diverse training sets that include bulk materials, surfaces, clusters, and defect structures. Beyond structural information and labeled properties, PESData can also store custom metadata including references to the original data sources, committee evaluation results, accuracy metrics, computational parameters used for labelling, and other application-specific information.

For performance optimization, the dataset is stored in the AiiDA repository as an HDF5 file using the `h5py` library⁴⁵. The class implements Python iterators to read data in chunks, enabling efficient handling of large datasets without overwhelming memory resources. This approach is crucial when working with training sets containing thousands of structures with hundreds of atoms each, and to maintain high performance and usability throughout the NNIP development workflow.

2.3 Validation

To validate our workflow and demonstrate its capabilities, we consider two representative case studies. First, we carry out a fully automated training campaign of a NNIP for a comprehensive set of carbon allotropes, showcasing the ability of the workflow to handle materials with widely different bonding topologies and dimensionalities. Second, we investigate 2D transition-metal dichalcogenides, focusing on monolayer $W_xMo_{1-x}Te_2$ alloys, to illustrate the flexibility of the approach for multicomponent systems and its accuracy in capturing structural phase stability as a function of composition.

Carbon allotropes

As a first validation study, we showcase AiiDA-TrainsPot with an active learning training of a NNIP for all carbon allotropes. Carbon represents an ideal test case due to its rich structural diversity, which includes 0D fullerenes, 1D nanotubes, 2D graphene, various 3D crystals, including diamond and layered forms such as graphite.

We conduct two independent validation runs, both initiated from the same set of 48 carbon structures primarily sourced from the Materials Cloud 2D and 3D databases²¹ (44 structures), supplemented with 4 selected low-dimensional structures (nanoribbons and fullerenes) from the dataset developed by Drautz *et al.*³³.

The first run, designated as *run A* or “fast exploration”, is designed to assess the workflow's capability to explore the PES and study the convergence of the active learning scheme for increasingly large datasets. The run starts by augmenting structures near equilibrium configurations (see Methods Sec. 4 for more details), that is followed by MD with a wide range of temperatures from 0 up to 5000 K. This also tests the model's ability to capture ex-

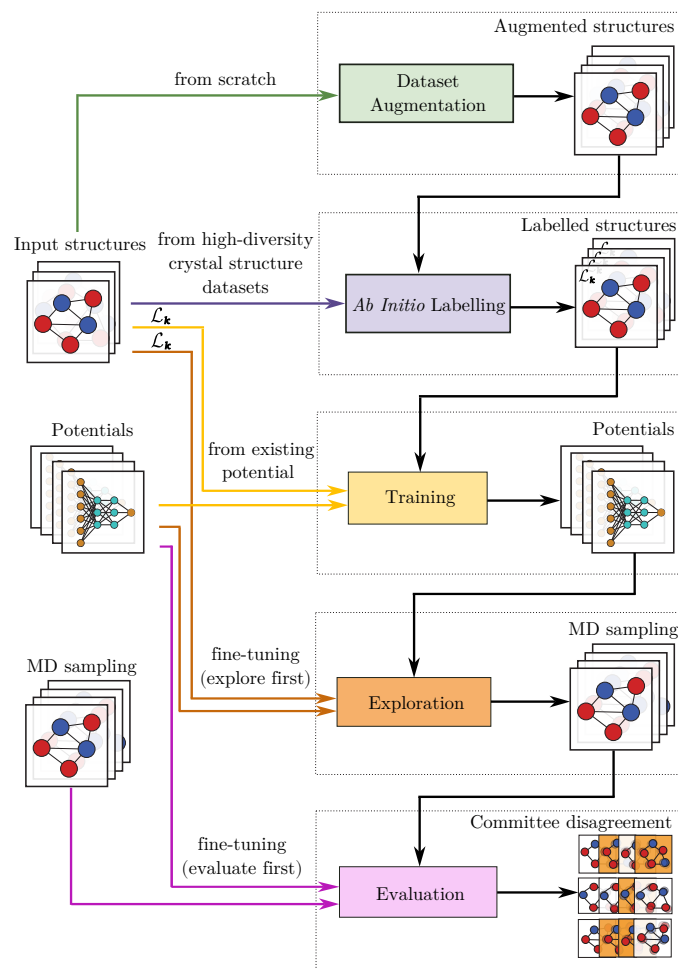


Fig. 3 Schematic representation of AiiDA-TrainsPot for different use cases. Users can execute individual subtasks or the entire workflow depending on their specific requirements. The workflow can start from Dataset Augmentation to expand data diversity, *Ab Initio* Labelling to perform DFT-based calculations of energy, force, and stress tensor, or Training to generate or improve existing machine-learned interatomic potentials. Users may also initiate Exploration for MD simulations or Evaluation to assess accuracy and performance.



treme conditions such as melting and formation of amorphous phases. As we want to explore convergence for rather large training datasets, which involve up to 10^4 single-point DFT calculations, we employ here the AiiDA-QuantumESPRESSO *fast* protocol³² for DFT calculations^{39,40}.

The second run, designated as *run B* or “accuracy and data-efficiency”, focuses, instead, on achieving rather accurate predictions for equilibrium and near-equilibrium conditions with small training datasets (around 2×10^3 calculations). The settings of data augmentation and the MD runs are refined for more efficient exploration, including smaller distortions of atomic positions, lower temperatures (up to 1000 K) and larger range of pressures (see Methods Sec. 4). *Run B* employs the AiiDA-QuantumESPRESSO *stringent* protocol, which uses more accurate pseudopotentials and stricter convergence criteria. While this differs from the *fast* protocol used in *run A*, it ensures higher numerical precision and reduced noise in the reference dataset.

The fast-exploration campaign runs for 10 iterations of active learning, Fig. 4 reports the evolution of the model performance in cold colors. The top panels track the RMSE for energy, forces, and stress tensor components across training, validation and test sets, as a function of the active learning step and, in parallel, the dataset size that range from the initial 1,177 structures (iteration 1) to 9,537 structures in the final iteration. For energy and stress tensor components, we observe a consistent decrease in prediction errors as the active learning progresses. Interestingly, errors on forces increase from the first to the second iteration, and then decrease monotonically for all the following iterations, suggesting that the active learning strategy first explores novel regions of the PES that require more data to learn. This is supported by the data analysis based on dimensionality reduction discussed later, which shows that early iterations sample an increasing number of distinct structural prototypes. The final RMSE values reach 4.3 meV/atom for energies, 293.0 meV/Å for forces, and 3.7 meV/Å³ for stress tensor components on the test set, comparable to those reported in Ref.³³ for a ML potential trained and tested on all carbon allotropes. The error bars in Fig. 4 represent the standard deviation across the model committee, which also decrease with iterations as the model becomes more consistent and robust. The bottom panels display parity plots comparing DFT reference values against NNIP predictions for the test set after the final iteration. The tight clustering of points along the diagonal line, together with the error distribution histograms (insets), demonstrates excellent agreement between the ML predictions and DFT calculations across all evaluated properties.

To better understand how our active learning strategy explores the potential energy landscape, we analyze data diversity in the training dataset by using the kinetic Spectral Operator Representation (SOREP) descriptor⁴⁶. The kinetic SOREP provides a compact electronic-structure fingerprint based on the density of states computed for the kinetic energy operator, which is evaluated on a basis set made of a customized version of the atomic natural orbitals (ANO) in terms of contracted Gaussian-type orbitals (cGTO)^{47–53}. We then use t-SNE (t-distributed Stochastic Neighbor Embedding) to visualize the high-dimensional SOREP descriptors in 2D space, with points colored according to the active

learning iteration in which they were generated (Fig. 5a). This depicts the evolution of the training set in electronic-structure space during the active-learning process. The clustering pattern reveals that early iterations (iterations 1-2) sample distinctly different and broad regions of the configuration space compared to the initial dataset (depicted as iteration 0 in Fig. 5a). This explains the initial increase in force errors, as the model encounters novel atomic environments with rather different electronic structures, requiring additional data to be learned accurately. In later iterations, after thorough sampling of these initial regions, the workflow begins exploring new domains. While the decrease in errors after the initial iterations can be attributed to good sampling of primary regions, the reduction is modest as exploration of new regions continues concurrently. Although the final model achieves rather good accuracy, further active learning iterations could be performed to explore a wider region of the PES and enhance even further model performance.

Beyond understanding the exploration strategy through SOREP analysis, we evaluate the effectiveness of using committee disagreement as an uncertainty metric for steering the growth of the training dataset. Figure 5b demonstrates the correlation between committee disagreement $\mathcal{D}_\alpha(\mathcal{X}_k)$ and true prediction errors $\varepsilon_\alpha(\mathcal{X}_k)$ at the final active learning iteration. The analysis reveals an approximately linear relationship between these quantities, confirming that committee disagreement can serve as a reliable proxy for accuracy¹⁹, although the proportionality constant is far from being unity. This aspect, anticipated in Sec. 2.1.6 and consistent with recent discussions on uncertainty quantification in ML atomistic models²⁰, is addressed here by introducing a calibration factor a_α determined via linear regression. This factor transforms user-defined error tolerances ε_α^{thr} into equivalent disagreement thresholds τ_α^{thr} ; the calibrated committee disagreement is then employed to select the structures to be labeled at the first-principles level. This calibration procedure ensures that the committee disagreement threshold appropriately reflects the true prediction errors, addressing cases where the uncalibrated disagreement metric might either overestimate (as observed for forces in Fig. 5b) or underestimate (as seen for energies and stress in our validation test) the actual deviations from DFT.

To quantitatively assess the reliability of this approach, we analyze two key metrics shown in the insets of Fig. 5b: the True Positive Rate (TPR) and Positive Predictive Value (PPV) as a function of the disagreement and the true error thresholds. The TPR is defined as:

$$TPR = \frac{TP}{TP + FN} \quad (3)$$

where TP is the number of true positives ($\mathcal{D}_\alpha(\mathcal{X}_k) > \tau_\alpha^{thr}$ and $\varepsilon_\alpha(\mathcal{X}_k) > \varepsilon_\alpha^{thr}$) and FN is the number of false negatives ($\mathcal{D}_\alpha(\mathcal{X}_k) < \tau_\alpha^{thr}$ and $\varepsilon_\alpha(\mathcal{X}_k) > \varepsilon_\alpha^{thr}$). The PPV is defined as:

$$PPV = \frac{TP}{TP + FP} \quad (4)$$

where FP is the number of false positives ($\mathcal{D}_\alpha(\mathcal{X}_k) > \tau_\alpha^{thr}$ and $\varepsilon_\alpha(\mathcal{X}_k) < \varepsilon_\alpha^{thr}$). These metrics quantify the reliability of using committee disagreement for structure selection. A high TPR indicates the approach successfully identifies structures with large



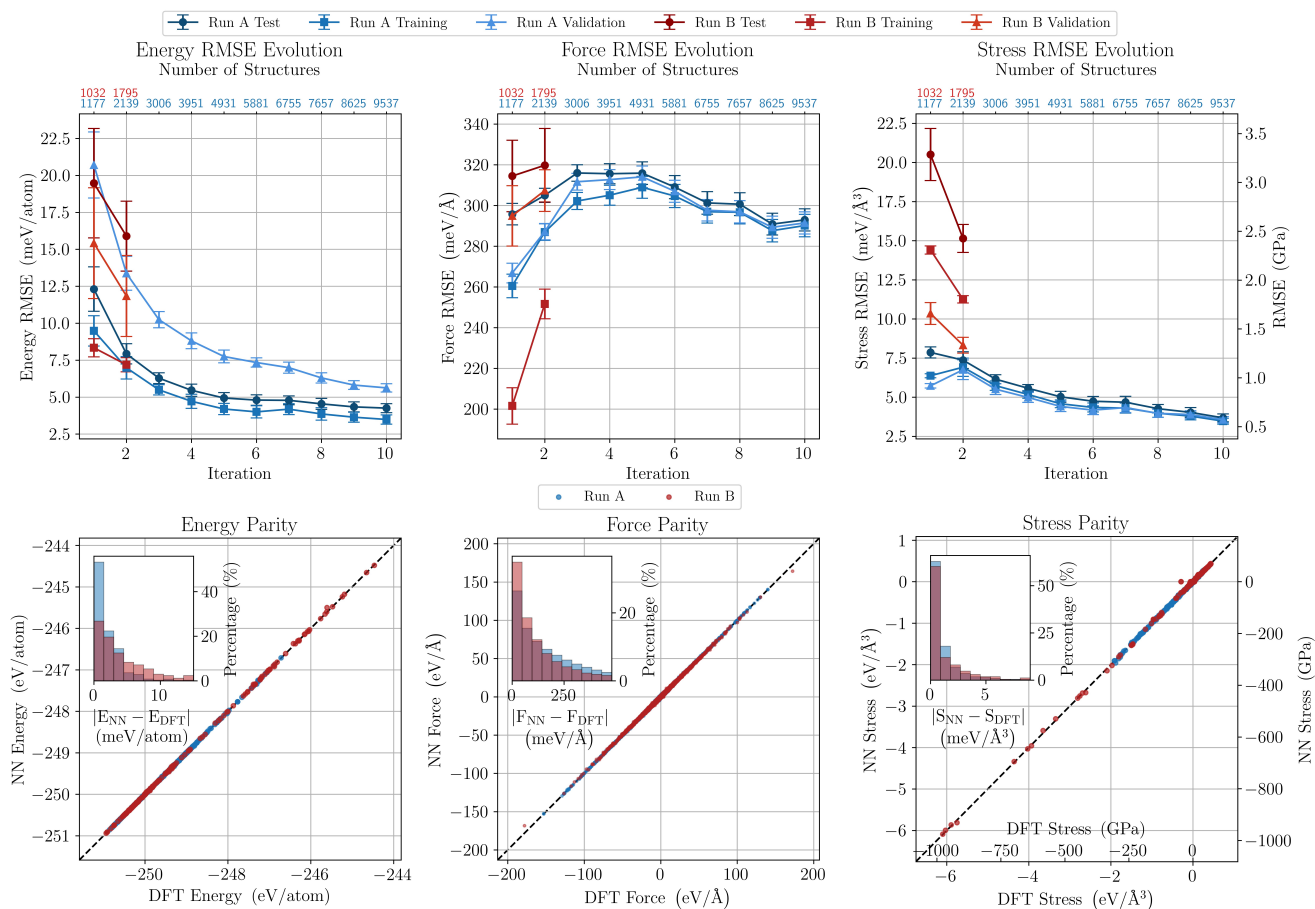


Fig. 4 Evolution of model accuracy over active learning. Top panels: RMSE for energies, forces, and stress tensor components are shown across training, validation, and test sets. Bottom panels: Parity plots comparing NNIP predictions (latest iteration of active learning for both runs) to DFT reference values on the final test set. Cold and warm colors identify the results of run A ("fast exploration") and run B ("accuracy and data-efficiency"), respectively. Error bars in the top panels represent the standard deviation across the model committee. Insets in the bottom panels show the error distribution histograms.



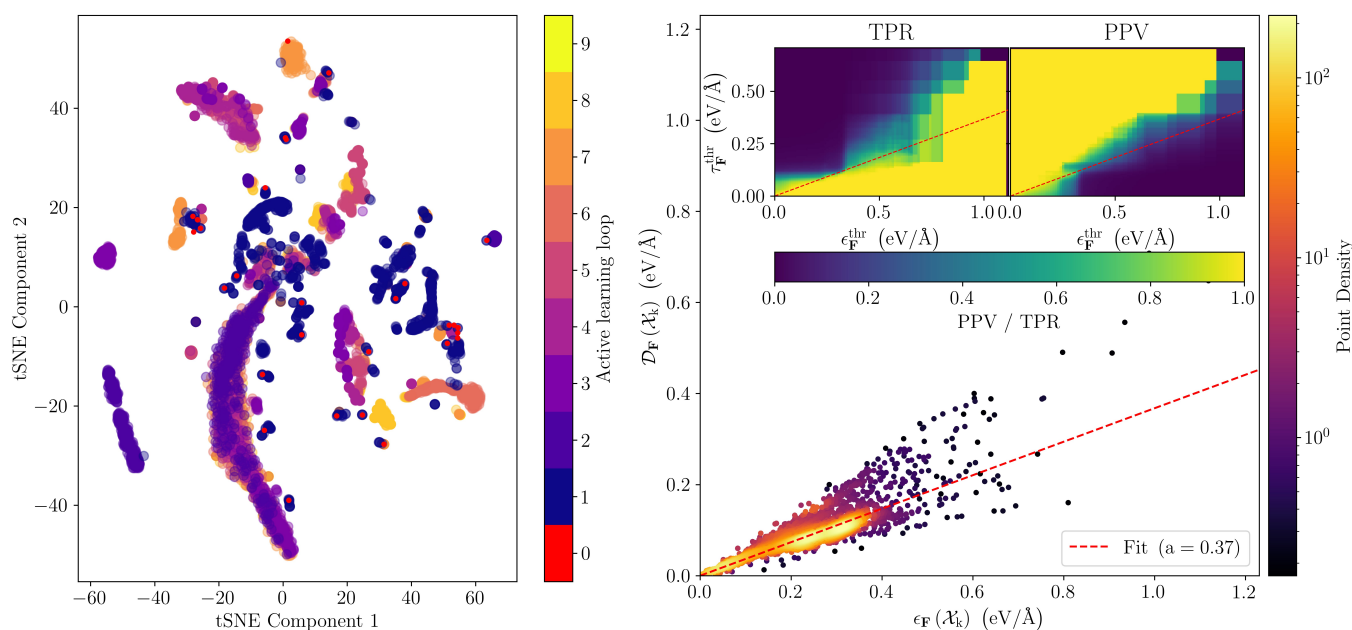


Fig. 5 Exploration of the potential energy surface (PES) and uncertainty quantification (run A). Left panel: t-SNE visualization of SOREP electronic-structure descriptors colored by active learning iteration, showing how the workflow systematically explores the PES through in two ways: by improving the sampling of known regions and by simultaneously exploring previously uncharted areas. Right panel: Committee disagreement versus true error (deviation from DFT) across different active learning iterations, showing strong correlation between model uncertainty and actual prediction uncertainty. We calibrate committee disagreement with linear regression against true errors, which enables quantitative uncertainty estimation in large-scale applications where reference DFT calculations are not feasible. Insets show the True Positive Rate (TPR) and the Positive Predictive Value (PPV) as functions of disagreement and true error thresholds. Both metrics approach unity along the fitted correlation line (dashed red), i.e., they are simultaneously maximized by a structure selection strategy based on calibrated committed disagreement.

true errors, while a high PPV confirms that selected structures genuinely require additional training. The analysis shows that along the fitted correlation line (dashed red line), i.e., using a calibrated committee disagreement, both TPR and PPV remain close to unity, particularly for force predictions up to several hundred $\text{meV}/\text{\AA}$ —which is the typical range for accuracy thresholds and where most data points lie. Furthermore, we find that the performance of the calibrated disagreement criterion is robust with respect to the progression of the active-learning procedure and to variations in the committee size, with no significant degradation of either the TPR or the PPV (see Supplementary Figure 1 in Supporting Information). This indicates not only that calibrated committee disagreement effectively identifies the structures requiring additional *ab initio* calculations, but it shows to be an optimal and robust strategy that simultaneously maximize both TPR and PPV, hence minimizing both the number of false positives and false negatives.

While the first validation study demonstrates the effectiveness of our active learning strategy, we observe diminishing returns in RMSE improvement as the number of iterations and the size of the dataset increase, due to the exploration of larger and larger regions of the PES. This suggests that strategic optimization of initial structures and data augmentation parameters can enhance the NNIP accuracy for the target application, while performing very few iterations of active learning. Therefore, for the second validation—focused on accuracy and data efficiency—we target the description of polymorphs near equilibrium with a refined

data augmentation strategy combined with constrained temperature ranges in MD simulations: strain range is increased with respect to the previous run, while atomic rattling is reduced and MD temperatures range up to 1000 K (see Methods 4 for details). This approach delivers high-quality potentials in just two active learning iterations and about 1800 training configurations, hence making more affordable and sustainable the use of the *stringent* protocol for DFT calculations, which is computationally more expensive.

The accuracy metrics are reported as warm colors in Fig. 4: the NNIPs score (test set RMSE) 15.9 meV/atom for energies, 319.8 $\text{meV}/\text{\AA}$ for forces, and 15.1 $\text{meV}/\text{\AA}^3$ for stress tensor components. While the overall errors on the test set seem comparable to the first run, we target here accurate energetic and vibrational properties, which are shown later to be in good agreement with DFT.

Although our training data is obtained with the semi-local Perdew-Burke-Ernzerhof (PBE) functional⁵⁴, we efficiently include long-range van der Waals interactions by adding Grimme's D2 dispersion corrections³⁵ on top of the NNIPs. We check that the approach works in practice by calculating the energy profile as a function of interlayer distance in graphite (see Fig. 6), comparing the NNIP with PBE—both with and without D2 corrections, where the in-plane lattice parameters are fixed with DFT structural optimization. Except for defect formation energies, the following validation tests are all performed with Grimme's D2 dispersion correction applied on top of the NNIP and compared with DFT calculations that include D2 corrections as implemented in



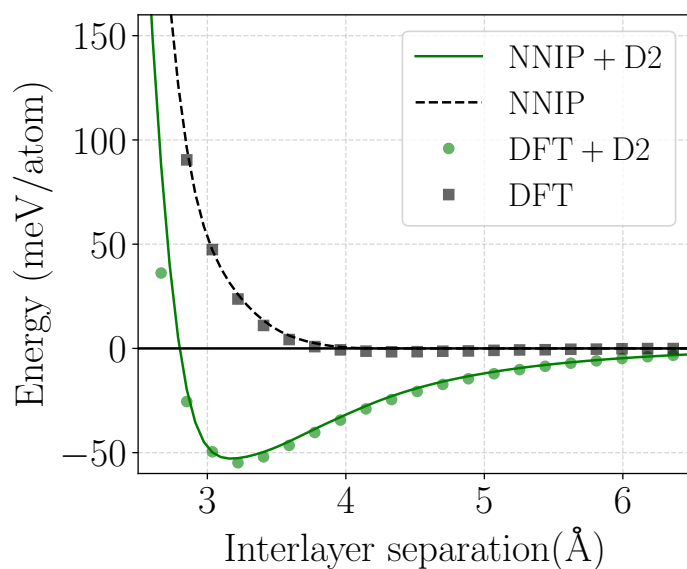


Fig. 6 Energy as a function of interlayer separation in graphite. The NNIP-predicted binding curve closely follows the DFT reference, both with and without Grimme's D2 dispersion corrections. Decoupling van der Waals interactions from the NN helps to keep the descriptor more local and reduce the computational cost of training and using NNIPs.

QUANTUM ESPRESSO.

Figure 7 shows the equation of states for various carbon allotropes, including graphite, graphene, diamond, dimer, simple cubic (sc), face-centered cubic (fcc), and body-centered cubic (bcc) structures. Binding energies are evaluated with reference to the energy of isolated atoms. This benchmarking approach is widely adopted in the literature^{33,55}, as it offers a compact yet informative way to assess how accurately a potential reproduces bonding behavior across diverse local geometries.

The left panel of Fig. 7 shows the excellent agreement between NNIP predictions and DFT references in around equilibrium; discrepancies become more noticeable at extreme bond compressions or expansions, which correspond to configurations that are underrepresented in the training data. Notably, the EOS for the dimer is reasonably accurate, even if no dimer configurations were included in the initial training set. The right-hand panel compares the EOS for graphite, graphene, and diamond, focusing on their relative energetic ordering. This inset is especially informative because graphite and diamond exhibit nearly degenerate formation energies in DFT. The NNIP correctly reproduces the energy hierarchy, demonstrating its ability to capture subtle thermodynamic trends. Similar benchmarking practices have been applied in the development of the ACE interatomic potentials³³.

To assess the transferability of our NNIP to defective structures, we compute the formation energies of three representative point defects in graphene: the monovacancy, divacancy, and Stone-Wales defect. The formation energy is defined as:

$$E_{\text{form}} = E_{\text{defected}} - E_{\text{pristine}} + n \cdot E^0,$$

where E_{defected} and E_{pristine} are the total energies of the relaxed defective and pristine graphene supercells, respectively; n is the

Relax-Energy	Monovacancy	Divacancy	Stone-Wales
NNIP-NNIP	7.30	7.91	4.74
NNIP-DFT	8.03	7.47	4.67
DFT-NNIP	7.72	8.00	4.74
DFT-DFT	7.72	7.39	4.64

Table 1 Formation energies (eV) of representative defects in graphene: monovacancy, divacancy, and Stone-Wales. The first two rows correspond to structures relaxed with the NNIP, with subsequent energy evaluation using either the NNIP (row 1) or DFT (row 2). The third and fourth rows correspond to structures relaxed with DFT, with subsequent energy evaluation using the NNIP (row 3) or DFT (row 4).

number of carbon atoms removed ($n = 0$ for Stone-Wales); and E^0 is the chemical potential of a carbon atom, estimated from bulk graphene.

As summarized in Table 1, we evaluate defect formation energies using four complementary approaches: structures relaxed with the NNIP and subsequently evaluated with either the NNIP or DFT, and structures relaxed with DFT and evaluated with either the NNIP or DFT. This protocol allows us to separately assess the contributions of structural relaxation versus energy evaluation to the overall accuracy.

For all the cases considered, the calculated formation energies agree well with literature values^{33,56} and our DFT references. For the divacancy and Stone-Wales defects, residual NNIP-DFT differences stem mainly from energy evaluation and not from the relaxation method. For the monovacancy, however, differences in the relaxed structures matter more: spin-unpolarized DFT yields a Jahn-Teller-like reconstruction with a slight out-of-plane displacement of one dangling atom. This subtle rearrangement is not fully reproduced by the NNIP, which returns indeed highly accurate energetics on the DFT-relaxed structure, but describes less accurately the local PES curvature. Instead, the NNIP predicts a completely flat configuration to be the energetically more stable for the monovacancy. However, we should note that a physically accurate description of vacancies in graphene—particularly monovacancies—would anyway require spin-polarized DFT calculations to properly account for the unpaired π orbitals⁵⁷, which were not used for our reference dataset.

We present in Fig. 8 the phonon dispersion and density of states for graphene, graphite, and diamond. The comparison between NNIP predictions and density-functional perturbation theory (DFPT)⁵⁸ calculations demonstrates good agreement across all three carbon allotropes, confirming the potential's ability to accurately capture vibrational properties. The observed small imaginary ZA phonons near Γ are a well-known numerical issue for 2D materials, which can be solved by adopting prohibitively tight parameters and in particular very high plane-wave cutoffs⁵⁹. Notably, despite being trained on DFT data exhibiting this, the NNIP does not display such unphysical behavior.

In order to further assess the transferability of our NNIP, we consider amorphous carbon, which features a disordered network of mixed sp^2 and sp^3 bonds that was not explicitly included in the rattled and strained crystalline configurations included in the training set. We compute the radial distribution function (RDF), $g(r)$, which characterizes short-range order in disordered sys-



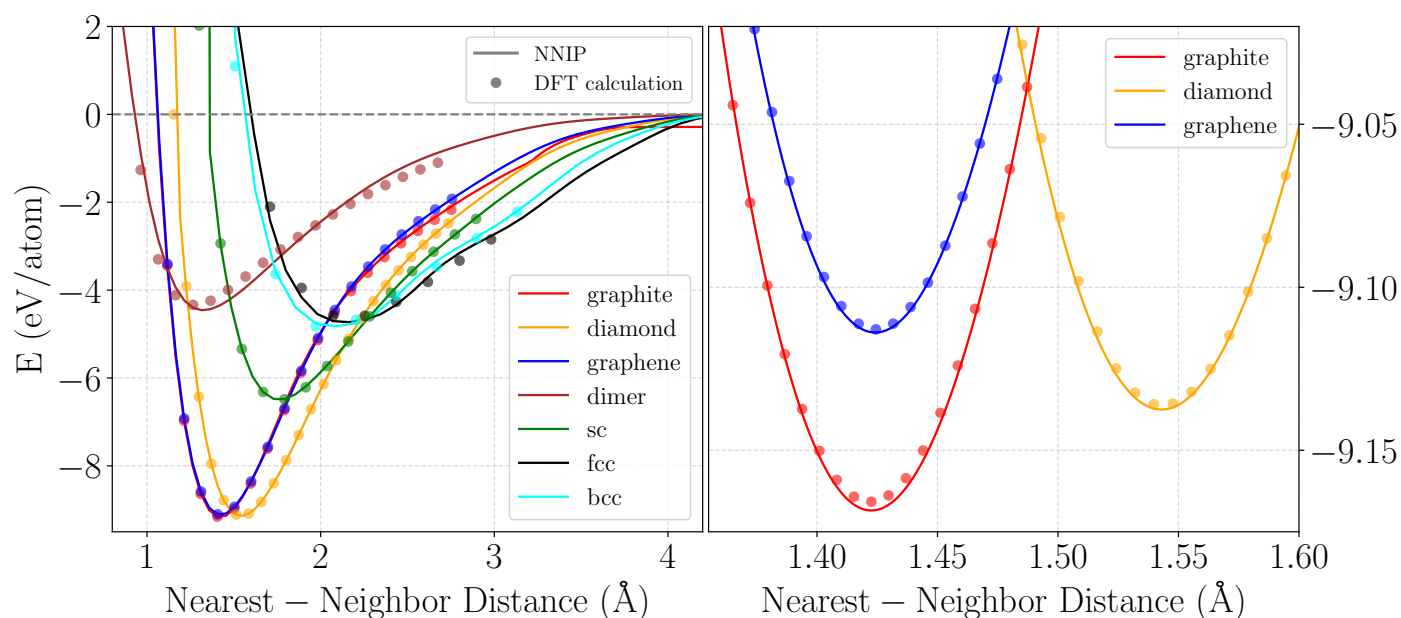


Fig. 7 Equations of state. Left panel: Binding energy per atom as a function of nearest-neighbor bond distance for various carbon structures, comparing predictions based on the NNIP with DFT calculations. Right panel: Zoomed-in view for graphite, graphene, and diamond; the model preserves the correct stability ordering at the meV level.

tems, on independent amorphous configurations at a density of 3.5 g/cm^3 that have been obtained using the melt-and-quench procedure described in Ref.⁶⁰. Figure 9 compares the RDFs for the committee of potentials and the committee average with reference *ab initio* molecular dynamics (AIMD) data from Ref.⁶¹. The NNIPs closely reproduce the AIMD reference, accurately capturing the position and height of the first and second peaks, confirming the ability of our potentials to reliably describe the local structure of amorphous carbon.

As a further transferability test, we evaluate one NNIP on the SACADA dataset, which contains 1635 carbon allotropes collected from the scientific literature⁶². Unlike the standard test sets, which are generated using the same protocols as the training sets (*run A* and *run B*) and therefore include structures obtained from data augmentation and MD sampling, the SACADA structures can differ substantially from the training set and span a wide variety of densities and space groups. This makes the SACADA evaluation a genuine test of the model's ability to generalize beyond its training domain. Out of the 1635, we exclude 13 structures due to un-converged DFT labeling. Then, we discard 19 structures where at least one of the potentials yields a committee-based disagreement larger than the following loose thresholds: 0.2 eV for energy, 50 eV/Å for forces, and 1 eV/Å³ for the stress tensor. The remaining 1603 structures (98% of the SACADA dataset) represent the overlap set between converged DFT simulations and successful NNIPs predictions, as required for a meaningful comparison. Here, we also evaluate the universal PET-MAD potential (v1.0.2)⁸, both in its original form and after fine-tuning on the carbon dataset generated in *run B*. Since *run B* was obtained with PBE functionals, while PET-MAD was trained with PBEsol⁶³, we relabel the SACADA dataset with PBEsol for evaluating the accuracy of PET-MAD potentials. In passing, we note how the modular structure of

AiiDA-TrainsPot permits a seamless exchange of training engines (MACE or Metatrain) and architectures, the reuse of existing labeled datasets, and uniform evaluation procedures across models. We report the parity plots and RMSE in Fig. 10: the MACE potential trained from scratch with AiiDA-TrainsPot on 1795 structures performs similarly to the PET-MAD foundation model, with comparable error on forces and stress tensors while doing slightly worse on energies (89 vs. 54 meV). Notably, fine-tuning PET-MAD on the PBEsol-labeled *run B* datasets lowers the error on all metrics (energies, forces and stress) both compared to the MACE potential trained from scratch and to PET-MAD.

Phase stability in monolayer $\text{W}_x\text{Mo}_{1-x}\text{Te}_2$ alloys

As a further validation, we apply AiiDA-TrainsPot to a more challenging problem: the relative phase stability of $\text{W}_x\text{Mo}_{1-x}\text{Te}_2$ monolayer alloys at zero temperature. This system is well known to exhibit a composition-dependent transition between the H phase and T' phase, which can be driven by temperature, strain, doping, or alloying⁶⁴. In this benchmark, we specifically assess the ability of different models to reproduce the evolution of the formation-energy difference between the two phases as a function of alloy composition.

For each W fraction x , the formation energy of a given configuration is defined as:

$$E_f(x) = E(\text{W}_x\text{Mo}_{1-x}\text{Te}_2) - xE(\text{WTe}_2) - (1-x)E(\text{MoTe}_2), \quad (5)$$

where all energies refer to total energies computed in the most stable polymorph at zero temperature (H for MoTe_2 and T' for WTe_2). The quantity reported in Fig. 11 is the formation energy difference $E_f^{\text{H}}(x) - E_f^{\text{T}'}(x)$, which directly determines the relative structural phase stability at zero temperature.



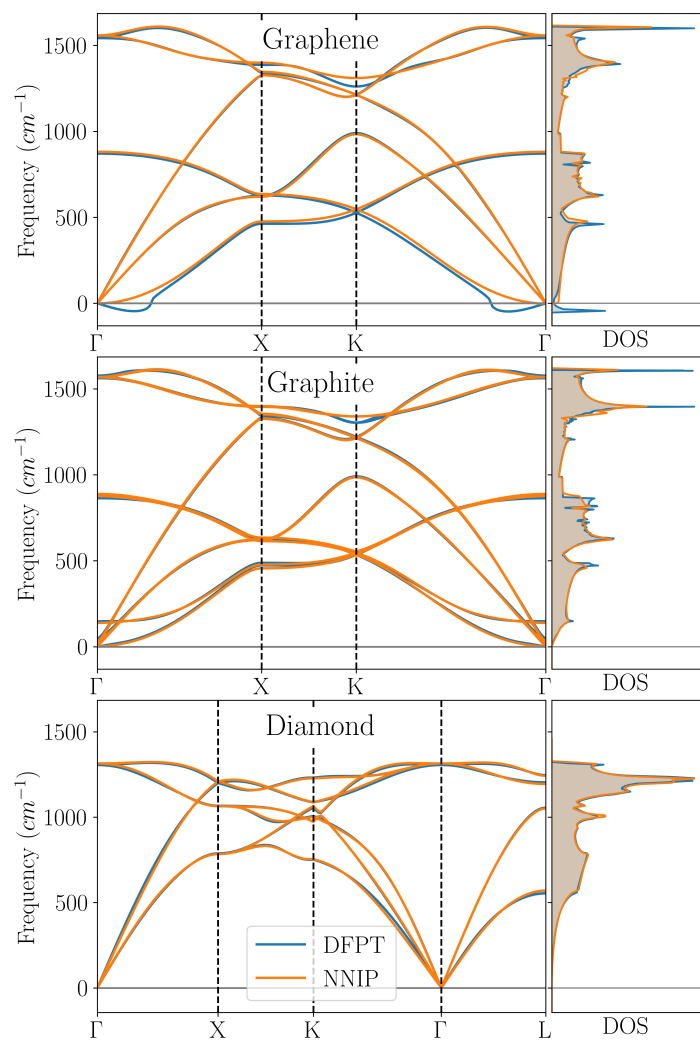


Fig. 8 Phonon dispersion and density of states for graphene, graphite, and diamond. The comparison between NNIP predictions and DFPT calculations shows good agreement for all three carbon allotropes.

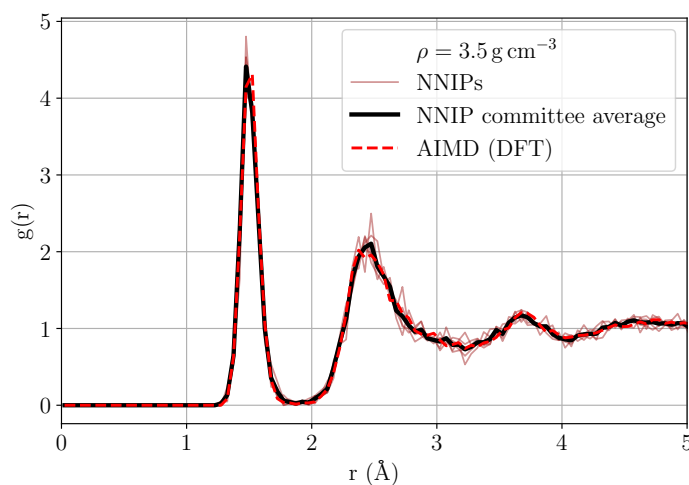


Fig. 9 Radial distribution functions $g(r)$ of amorphous carbon structures with a density of 3.5 g/cm^3 , generated using a melt-and-quench protocol and analyzed with our five NNIPs (thin brown lines) and their committee average (thick black line). The results demonstrate that our potentials accurately reproduce the AIMD, including the positions of the first and second peaks, and the characteristic nonzero minimum at $\sim 1.9 \text{ \AA}$.

To generate validation configurations, we employ special quasirandom structures (SQS)⁶⁵ to approximate random W/Mo distributions on the metal sublattice at a given composition x . For each intermediate composition shown in Fig. 11, we construct and evaluate 10 independent SQS realizations to ensure statistical robustness of the estimated formation-energy differences.

As shown in Fig. 11, existing foundation models struggle with this task. MATTERSIM (v1.0.0–1M) and MACE (MATPES–PBE–0) incorrectly predict the T' phase to be the ground state for all compositions, while PET–MAD (v1.0.2), although capturing a structural phase stability and quite well the formation energy difference for end-members, places the critical concentration around $x \sim 0.7$ —far from the DFT reference value at about $x \sim 0.3$.

Then, we employ AiiDA-TrainsPot to train three additional MACE models, all starting from the same four pristine input structures (H- and T' -phase MoTe_2 and WTe_2) taken from the Materials Cloud 2D database⁵⁹:

1. Alloy model trained from scratch. A dataset of 1687 alloy configurations is generated by randomly sampling the W/Mo distribution in both phases during the dataset-augmentation stage, and a committee of models is trained from scratch on this dataset.
2. Fine-tuned alloy model. The same alloy dataset is used to fine-tune the MACE MATPES–PBE–0 foundation model.
3. End-member model with active learning. A fully independent model was trained from scratch using only configurations of the pure end-members in both phases (no alloys, clusters, surfaces, or substitutional configurations were generated during dataset augmentation). Twelve active-learning iterations are performed, yielding a final dataset of 6965 structures.



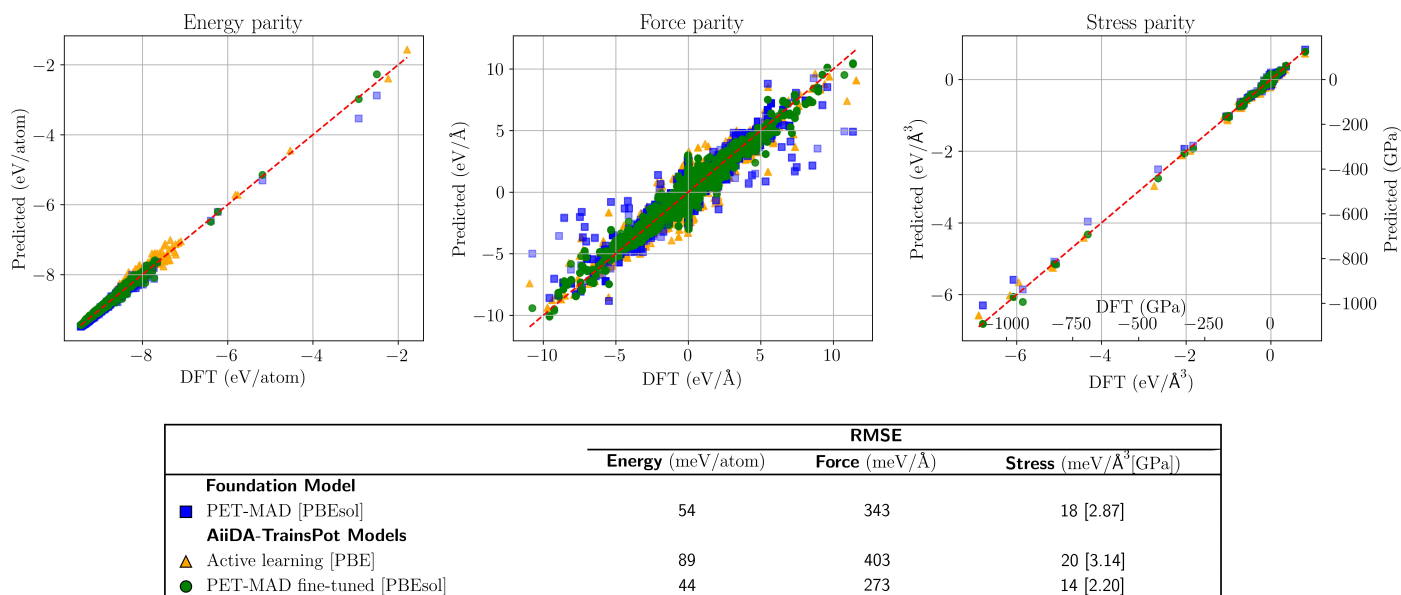


Fig. 10 Parity plots for energy, forces, and stress on 98% of the SACADA dataset, comparing three models: (i) our MACE-based potentials generated in *run B* (orange), (ii) the universal PET-MAD potential (blue), and (iii) the PET-MAD potential fine-tuned using the carbon dataset produced in *run B* (green). The dashed line indicates perfect agreement with the reference DFT data.

All three models markedly outperform the existing foundation models, in particular accurately capturing both the composition dependence of the formation-energy difference and the H-T' phase-stability crossover. Remarkably, the best-performing model is the one trained only on end-member configurations: despite never having seen alloy structures, it correctly reproduces the formation-energy trend and predicts the phase-stability crossover close to the DFT reference value. This result highlights the effectiveness of active learning with calibrated committee disagreement—and, more generally, of the AiiDA-TrainsPot protocol—in sampling the relevant regions of configuration space and constructing a compact dataset that yields a highly transferable NNIP.

3 Conclusions

We have demonstrated that a fully-automated strategy based on data augmentation and active learning—steered by a calibrated committee-disagreement for energy, forces and stress tensor components—provides an effective way to explore the PES and to train accurate NNIPs in a data-efficient manner with minimal human intervention. For instance, a potential describing carbon allotropes has been obtained with as few as 48 initial input structures that were pulled from publicly available crystal structure databases; these prototypes have been transformed into thousands of uncorrelated, diverse and relevant configurations—all generated and calculated with no human supervision. The data augmentation is performed right at the beginning of the process to obtain about 1,000 diverse and uncorrelated training configurations: this is enough to produce sufficiently accurate and stable NNIPs that can be used for classical MD simulations, which are computationally much cheaper than AIMD and can produce additional uncorrelated configurations for further refinement of the NNIP at low cost.

The SOREP-based dimensionality reduction (tSNE) has shown

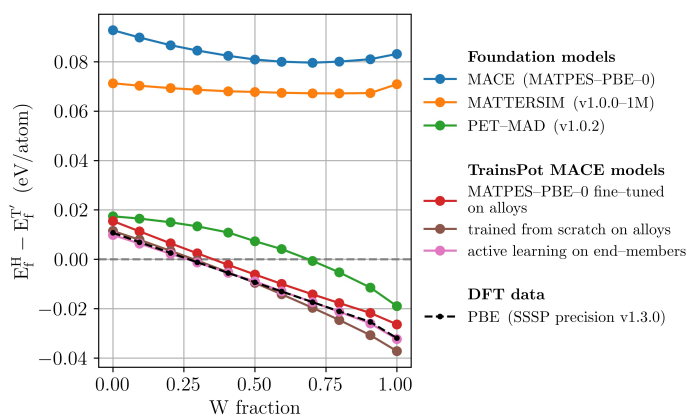


Fig. 11 Comparison of the formation-energy difference $E_f^H(x) - E_f^{T'}(x)$ as a function of W concentration x for several NNIPs, comparing foundation models with models trained or fine-tuned through AiiDA-TrainsPot. The MACE potentials generated with AiiDA-TrainsPot are trained on datasets produced via data augmentation starting from pure MoTe_2 and WTe_2 in both H and T' phases. The models explicitly trained or fine-tuned on alloys share the same dataset of 1687 alloy structures generated in the dataset-augmentation stage. In contrast, the active-learning runs on end-members alloys only, i.e., no clusters, surfaces, nor substitutional configurations are generated during dataset augmentation and hence are not present in the training dataset. All DFT simulations are computed with the same level of theory, except for the training sets originally employed to build the foundation models by their respective authors.



that subsequent MD-based active learning steps explore the PES through a dual mechanism: dense sampling of already-known regions and, simultaneously, exploration of entirely new basins. A compelling example is the spontaneous formation of carbon nanotubes during the MD simulations in the active learning cycle: carbon nanotubes were absent from the original dataset but become then automatically incorporated by AiiDA-TrainsPot into subsequent training iterations, suggesting the ability of the automated workflow to find novel and quite different metastable or stable structures without prior knowledge. Indeed, the automated training strategy delivered potentials that were also able to reproduce the correct RDF of amorphous carbon.

A key aspect is the use of calibrated committee disagreement to guide the selection of new training structures. This strategy improves the efficiency and reliability of active learning, while ensuring that the model is exposed just to configurations that enhance its predictive power. Notably, the relationship between committee disagreement and actual error appears to be linear over all active learning iterations and across different properties (e.g., energies, forces, stress tensors): that support the reliable use of calibrated committee disagreement also in production simulations, i.e., when reference *ab initio* simulations typically cannot be performed.

Furthermore, as demonstrated for both the carbon allotropes and the $W_xMo_{1-x}Te_2$ alloy benchmarks, AiiDA-TrainsPot can be effectively employed to fine-tune existing foundation models to specific applications. Fine-tuning can be performed either in a *single shot* on a fixed dataset or embedded into subsequent active-learning iterations, further improving model accuracy and robustness.

It would be interesting to investigate other data-augmentation approaches—such as random structure search⁶⁶, non-diagonal supercells⁶⁷, or generative models⁶⁸—as well as advanced exploration strategies beyond NPT and NVT MD—such introducing metadynamics⁶⁹ by interfacing AiiDA-TrainsPot to PLUMED^{70–72}. More in general, we hope that the AiiDA's plugin system and the modular structure of AiiDA-TrainsPot will encourage and facilitate future upgrades, as well as the integration of new tools. An example would be supporting multiple NNIP backends (beyond MACE and Metatrain) and electronic structure codes (beyond QUANTUM ESPRESSO), in the spirit of previous efforts on code-agnostic common workflows for EOS and dissociation curves⁷³. Powerful upgrades would be enabled by interfacing AiiDA-TrainsPot with existing specialized AiiDA workflows, for instance using DFT+U calculations where the Hubbard U can be automatically calculated for each configuration either with DFPT^{74–76} through the AiiDA-Hubbard workflow⁷⁷ or even more efficiently through ML methods⁷⁸.

As a side note, AiiDA-TrainsPot inherit from the AiiDA infrastructure the tracking of data-provenance graphs, enabling external validation and assessment of the published training data. This capability is crucial for public foundational models and, more broadly, for the reuse of training datasets and their corresponding NNIPs by the community. In the Supporting Information, we include a representative provenance graph from a small test run to illustrate how AiiDA automatically records the complete data

lineage throughout a training campaign. Supplementary Fig. 2 shows the full provenance graph, whereas Supplementary Fig. 3 presents a simplified view that retains only the most relevant nodes and connections.

While AiiDA-TrainsPot can operate autonomously for general-purpose NNIP development, domain experts retain full flexibility to incorporate their physical and chemical intuition in the automation strategy. The modular architecture is designed to enable full customization of all key components: initial structure selection, dataset augmentation parameters, MD simulation conditions, and computational settings for integrated codes—QUANTUM ESPRESSO, MACE, Metatrain and LAMMPS. Such level of customization allows users to balance the power of automation with the flexibility that is needed to support a wide range of applications, which require tailoring the active learning process to specific research objectives. Indeed, while AiiDA-TrainsPot automates the entire process—traditionally long, tedious and prone to human errors—of developing NNIPs, optimal results still benefit from careful consideration of the system of interest. In other words, the selection of initial structures, augmentation strategies, and MD conditions, can—and often should—be tailored to reflect the target application and desired properties: AiiDA-TrainsPot makes that effort straightforward. On top of that, the enforcement of standardized protocols (either already established, e.g., SSSP pseudopotentials²⁹ or “fast”/“stringent” QE protocols³², or introduced in this work) contribute to precision, reproducibility and seamless integration with future efforts in training larger models.

AiiDA-TrainsPot democratizes the access to high-quality NNIPs tailored to the application of interest, hopefully encouraging computational scientists with limited expertise in electronic structure and ML to tackle challenging phenomena and materials, pushing the frontier of what can be simulated, understood and designed with *ab initio* accuracy.

4 Methods

All calculations were performed using AiiDA-TrainsPot, QUANTUM ESPRESSO^{79–81} v7.3.1, MACE⁵ v0.3.12, metatrain³⁴ v2025.10, and LAMMPS⁸² v8Feb2023.

Carbon allotropes

The initial dataset for carbon allotropes included 48 structures, primarily sourced from the Materials Cloud 2D and 3D databases^{59,83}, with additional nanoribbons and fullerenes from the Drautz dataset³³.

For both runs, structures were replicated up to a maximum of 600 atoms and a minimum cell length of 24 Å. A total of 80 cluster structures (up to 30 atoms each with minimum interatomic distance 1.5 Å) were generated. Slab configurations were created with a minimum thickness of 10 Å and a maximum of 600 atoms, along the (100), (110), (111), (001), (011), (010), and (101) directions. Non-periodic directions were padded with 15 Å of vacuum. Vacancies (2 per structure) were created in 30% of the structures.

For *run A*, random distortions and strains were introduced with a *rattle_fraction* of 0.4, a *max_compressive_strain* of 0.2 and a



max_tensile_strain of 0.2. DFT calculations used the *fast* protocol³² for k-point grid ($\lambda = 0.30 \text{ \AA}^{-1}$) and smearing ($\sigma_{cold} = 0.0275 \text{ Ry}$). MD simulations explored temperatures ranging from 0 to 5000 K and pressures from -5 to 5 kbar.

For *run B*, dataset augmentation parameters were optimized for near-equilibrium conditions: *rattle_fraction* was reduced to 0.3, while strain ranges were increased to *max_compressive_strain* of 0.3 and *max_tensile_strain* of 0.6 to better sample elastic deformations. DFT calculations employed the *stringent* protocol for enhanced accuracy ($\lambda = 0.1 \text{ \AA}^{-1}$, $\sigma_{cold} = 0.0125 \text{ Ry}$). MD exploration was constrained to temperatures from 0 to 1000 K and pressures from -20 to 20 kbar.

Both runs utilized the SSSP PBE precision library v1.3 pseudopotentials^{29–31}, total energy convergence threshold of 10^{-8} Ry , MACE training with radial cutoff of 4.5 \AA , two message-passing layers, batch size of 1, and up to 500 epochs. MD simulations were performed in NPT (for fully or partially periodic systems) or NVT (for non-periodic systems) ensembles using a 1 fs timestep and extracting trajectory frames every 1 ns. Since van der Waals interactions were not included at the DFT level, Grimme's D2 dispersion correction³⁵ was enabled in MD simulations via the *momb* pair style⁸⁴ in LAMMPS. Active learning thresholds on energy, forces, and stress tensor were set to 2 meV, 50 meV/\AA , 10 meV/\AA^3 , respectively, with a maximum of 1000 structures selected per iteration.

SACADA dataset

DFT calculations employed both the SSSP PBE and PBEsol precision pseudopotential libraries (v1.3)^{29–31}, using the same computational parameters as those adopted in the active learning. For fine-tuning, the PET-MAD potential (v1.0.2)⁸ was trained using the default hyperparameters, with a batch size of 3 and 100 training epochs.

$W_x\text{Mo}_{1-x}\text{Te}_2$ alloys

The initial dataset for $W_x\text{Mo}_{1-x}\text{Te}_2$ alloys included 4 structures comprising the H and T' polymorphs of pristine MoTe_2 and WTe_2 monolayers, sourced from the Materials Cloud 2D database⁵⁹.

During dataset augmentation, structures were replicated up to a minimum cell length of 18 \AA . The non-periodic direction was padded with 15 \AA of vacuum. Vacancies (2 per structure) were created in 30% of the structures. No substitutional configurations, clusters, or surfaces were generated during dataset augmentation. Only for the alloys dataset, random W/Mo distributions on the metal sublattice were sampled by setting in dataset augmentation stage Te as *fixed_species* and Mo and W as *alloy_species*.

The same DFT computational settings as in *run B* for carbon allotropes were used also in this case, while MD simulations performed in NPT ensemble (with barostat applied only to in-plane directions) explored temperatures ranging from 0 to 1800 K and pressures from -10 to 10 kbar, without additional VdW corrections. MACE trainings were performed with radial cutoff of 6.0 \AA , two message-passing layers, batch size of 5, and up to 500 epochs, except for fine-tuning run for which same parameter of the original MACE MATPES–PBE–0 model were used.

All the calculations in this work were performed on the CINECA Leonardo supercomputer, using nodes equipped with 4 NVIDIA A100 GPUs. Taking as reference the active-learning run used to train the MACE potential for MoTe_2 and WTe_2 (alloy benchmark in Fig. 11), where each structure contains about 108 atoms, we estimate an average computational cost of ~ 1.2 GPU hours per structure for DFT calculations, ~ 25 GPU hours for each MACE training run, and ~ 0.1 GPU hours for each LAMMPS MD simulation. After 12 active-learning iterations, resulting in a final dataset of 6965 structures, the total computational cost for training this committee of MACE potentials amounts to about 11 000 GPU hours, 84% of which was spent on DFT labelling.

Author Contributions

D.B. and N.M. contributed equally. D.B.: Investigation, Data curation, Formal analysis, Software, Methodology, Visualization, Writing - original draft, Writing - review & editing. N.M.: Investigation, Data curation, Formal analysis, Software, Methodology, Visualization, Writing - original draft, Writing - review & editing. M.P.: Supervision, Formal analysis, Funding acquisition, Project administration, Methodology, Visualization, Writing - original draft, Writing - review & editing. A.M.: Conceptualization, Supervision, Formal analysis, Funding acquisition, Methodology, Visualization, Resources, Project administration, Writing - original draft, Writing - review & editing.

Conflicts of interests

There are no conflicts to declare.

Data availability

The training datasets and trained models are available on the Materials Cloud Archive⁸⁵ with license Creative Commons Attribution 4.0 International. An archived copy of the AiiDA-TrainsPot software is also deposited in the same Materials Cloud Archive. The AiiDA-TrainsPot source code is publicly available on GitHub at <https://github.com/aiida-trieste-developers/aiida-trains-pot> under the MIT license. The repository additionally includes example input files, configuration templates, and user documentation covering installation procedures, an end-to-end tutorial, and the workflow API. The documentation is also available at <https://aiida-trieste-developers.github.io/aiida-trains-pot/>.

Acknowledgements

All authors acknowledge support from the ICSC – Centro Nazionale di Ricerca in High Performance Computing, Big Data and Quantum Computing, funded by European Union - NextGenerationEU (CUP Grant. No. J93C22000540006, PNRR Investimento M4.C2.1.4), and ENI through the Innovation Grant APECAR. The authors acknowledge CINECA, under CINECA-SISSA agreements, for the availability of high-performance computing resources and support. A.M. and N.M. acknowledge partial support from the European Commission through the Centre of Excellence “MaX - Materials Design at the Exascale” (HORIZON-EUROHPC, Grant No. 101093374). The views and opinions expressed are solely those of the authors and do not necessarily



reflect those of the European Union, nor can the European Union be held responsible for them.

Notes and references

- R. Car and M. Parrinello, *Phys. Rev. Lett.*, 1985, **55**, 2471–2474.
- V. L. Deringer, M. A. Caro and G. Csányi, *Advanced Materials*, 2019, **31**, 1902765.
- S. Batzner, A. Musaelian, L. Sun, M. Geiger, J. P. Mailoa, M. Kornbluth, N. Molinari, T. E. Smidt and B. Kozinsky, *Nature Communications*, 2022, **13**, 2453.
- A. Musaelian, S. Batzner, A. Johansson, L. Sun, C. J. Owen, M. Kornbluth and B. Kozinsky, *Nature Communications*, 2023, **14**, 579.
- I. Batatia, D. P. Kovacs, G. N. C. Simm, C. Ortner and G. Csanyi, *Advances in Neural Information Processing Systems*, 2022, **35**, 11423–11436.
- S. Pozdnyakov and M. Ceriotti, *Advances in Neural Information Processing Systems*, 2023, pp. 79469–79501.
- D. P. Kovács, J. H. Moore, N. J. Browning, I. Batatia, J. T. Horton, Y. Pu, V. Kapil, W. C. Witt, I.-B. Magdäu, D. J. Cole and G. Csányi, *MACE-OFF: Transferable Short Range Machine Learning Force Fields for Organic Molecules*, 2025.
- A. Mazitov, F. Bigi, M. Kellner, P. Pegolo, D. Tisi, G. Fraux, S. Pozdnyakov, P. Loche and M. Ceriotti, *PET-MAD, a lightweight universal interatomic potential for advanced materials modeling*, 2025.
- Y. Zhang, H. Wang, W. Chen, J. Zeng, L. Zhang, H. Wang and W. E., *Computer Physics Communications*, 2020, **253**, 107206.
- K. T. Schütt, P. Kessel, M. Gastegger, K. A. Nicoli, A. Tkatchenko and K.-R. Müller, *Journal of Chemical Theory and Computation*, 2018, **15**, 448–455.
- J. Vandermause, S. B. Torrisi, S. Batzner, Y. Xie, L. Sun, A. M. Kolpak and B. Kozinsky, *npj Computational Materials*, 2020, **6**, 20.
- J. Guo, V. Woo, D. A. Andersson, N. Hoyt, M. Williamson, I. Foster, C. Benmore, N. E. Jackson and G. Sivaraman, *The Journal of Chemical Physics*, 2023, **159**, 024802.
- K. Töpfer, L. I. Vazquez-Salazar and M. Meuwly, *Computer Physics Communications*, 2025, **308**, 109446.
- E. V. Podryabinkin and A. V. Shapeev, *Computational Materials Science*, 2017, **140**, 171–180.
- Z. Li, J. R. Kermode and A. De Vita, *Phys. Rev. Lett.*, 2015, **114**, 096405.
- R. Jinnouchi, F. Karsai and G. Kresse, *Physical Review B*, 2019, **100**, 014105.
- M. Kulichenko, K. Barros, N. Lubbers, Y. W. Li, R. Messerly, S. Tretiak, J. S. Smith and B. Nebgen, *Nature Computational Science*, 2023, **3**, 230–239.
- V. Vitartas, H. Zhang, V. Juraskova, T. Johnston-Wood and F. Duarte, *Digital Discovery*, 2026, **5**, 108–122.
- L. Kahle and F. Zipoli, *Physical Review E*, 2022, **105**, 015311.
- F. Grasselli, S. Chong, V. Kapil, S. Bonfanti and K. Rossi, *Digital Discovery*, 2025, **4**, 2654–2675.
- L. Talirz, S. Kumbhar, E. Passaro, A. V. Yakutovich, V. Granata, F. Gargiulo, M. Borelli, M. Uhrin, S. P. Huber, S. Zoupanos, C. S. Adorf, C. W. Andersen, O. Schütt, C. A. Pignedoli, D. Passerone, J. VandeVondele, T. C. Schulthess, B. Smit, G. Pizzi and N. Marzari, *Scientific Data*, 2020, **7**, 299.
- A. Jain, S. P. Ong, G. Hautier, W. Chen, W. D. Richards, S. Dacek, S. Cholia, D. Gunter, D. Skinner, G. Ceder and K. A. Persson, *APL Materials*, 2013, **1**, 011002.
- G. Bergerhoff, R. Hundt, R. Sievers and I. D. Brown, *Journal of Chemical Information and Computer Sciences*, 1983, **23**, 66–69.
- S. Gražulis, D. Chateigner, R. T. Downs, A. F. T. Yokochi, M. Quirós, L. Lutterotti, E. Manakova, J. Butkus, P. Moeck and A. Le Bail, *Journal of Applied Crystallography*, 2009, **42**, 726–729.
- S. Gražulis, A. Daškevič, A. Merkys, D. Chateigner, L. Lutterotti, M. Quirós, N. R. Serebryanaya, P. Moeck, R. T. Downs and A. Le Bail, *Nucleic Acids Research*, 2011, **40**, D420–D427.
- P. Villars, N. Onodera and S. Iwata, *Journal of Alloys and Compounds*, 1998, **279**, 1–7.
- P. Villars, M. Berndt, K. Brandenburg, K. Cenxual, J. Daams, F. Hulliger, T. Massalski, H. Okamoto, K. Osaki, A. Prince, H. Putz and S. Iwata, *Journal of Alloys and Compounds*, 2004, **367**, 293–297.
- E. Blokhin and P. Villars, in *The PAULING FILE Project and Materials Platform for Data Science: From Big Data Toward Materials Genome*, ed. W. Andreoni and S. Yip, Springer International Publishing, Cham, 2018, pp. 1–26.
- G. Prandini, A. Marrazzo, I. E. Castelli, N. Mounet and N. Marzari, *npj Computational Materials*, 2018, **4**, 72.
- P. E. Blöchl, *Phys. Rev. B*, 1994, **50**, 17953–17979.
- A. Dal Corso, *Computational Materials Science*, 2014, **95**, 337–350.
- G. d. M. Nascimento, F. J. d. Santos, M. Bercx, D. Grassano, G. Pizzi and N. Marzari, *Accurate and efficient protocols for high-throughput first-principles materials simulations*, 2025.
- M. Qamar, M. Mrovec, Y. Lysogorskiy, A. Bochkarev and R. Drautz, *Journal of Chemical Theory and Computation*, 2023, **19**, 5151–5167.
- F. Bigi, J. W. Abbott, P. Loche, A. Mazitov, D. Tisi, M. F. Langer, A. Goscinski, P. Pegolo, S. Chong, R. Goswami, S. Chorna, M. Kellner, M. Ceriotti and G. Fraux, *Metatensor and Metatomic: Foundational Libraries for Interoperable Atomistic Machine Learning*, 2025.
- S. Grimme, *Journal of Computational Chemistry*, 2006, **27**, 1787–1799.
- M. Abdar, F. Pourpanah, S. Hussain, D. Rezazadegan, L. Liu, M. Ghavamzadeh, P. Fieguth, X. Cao, A. Khosravi, U. R. Acharya, V. Makarenkov and S. Nahavandi, *Information Fusion*, 2021, **76**, 243–297.
- J. Behler, *Journal of Physics: Condensed Matter*, 2014, **26**, 183001.
- L. Chen, I. Sukuba, M. Probst and A. Kaiser, *RSC Advances*, 2020, **10**, 4293–4299.



- 39 S. P. Huber, S. Zoupanos, M. Uhrin, L. Talirz, L. Kahle, R. Häuselmann, D. Gresch, T. Müller, A. V. Yakutovich, C. W. Andersen, F. F. Ramirez, C. S. Adorf, F. Gargiulo, S. Kumbhar, E. Passaro, C. Johnston, A. Merkys, A. Cepellotti, N. Mounet, N. Marzari, B. Kozinsky and G. Pizzi, *Scientific Data*, 2020, **7**, 300.
- 40 M. Uhrin, S. P. Huber, J. Yu, N. Marzari and G. Pizzi, *Computational Materials Science*, 2021, **187**, 110086.
- 41 The AIIDA-LAMMPS plugin is available at <https://github.com/aiidaplugins/aiida-lammps>.
- 42 A. H. Larsen, J. J. Mortensen, J. Blomqvist, I. E. Castelli, R. Christensen, M. Dułak, J. Friis, M. N. Groves, B. Hammer, C. Hargus, E. D. Hermes, P. C. Jennings, P. B. Jensen, J. Kermode, J. R. Kitchin, E. L. Kolsbjerg, J. Kubal, K. Kaasbjerg, S. Lysgaard, J. B. Maronsson, T. Maxson, T. Olsen, L. Pastewka, A. Peterson, C. Rostgaard, J. Schiøtz, O. Schütt, M. Strange, K. S. Thygesen, T. Vegge, L. Vilhelmsen, M. Walter, Z. Zeng and K. W. Jacobsen, *Journal of Physics: Condensed Matter*, 2017, **29**, 273002.
- 43 A. P. Bartók, M. C. Payne, R. Kondor and G. Csányi, *Phys. Rev. Lett.*, 2010, **104**, 136403.
- 44 J. Behler and M. Parrinello, *Phys. Rev. Lett.*, 2007, **98**, 146401.
- 45 A. Collette, *Python and HDF5*, O'Reilly, 2013.
- 46 A. Zadoks, A. Marrazzo and N. Marzari, *npj Computational Materials*, 2024, **10**, 278.
- 47 B. P. Pritchard, D. Altarawy, B. Didier, T. D. Gibson and T. L. Windus, *Journal of Chemical Information and Modeling*, 2019, **59**, 4814–4820.
- 48 D. Feller, *Journal of Computational Chemistry*, 1996, **17**, 1571–1586.
- 49 K. L. Schuchardt, B. T. Didier, T. Elsethagen, L. Sun, V. Gurmooorthi, J. Chase, J. Li and T. L. Windus, *Journal of Chemical Information and Modeling*, 2007, **47**, 1045–1052.
- 50 B. O. Roos, R. Lindh, P. Malmqvist, V. Veryazov and P. Widmark, *The Journal of Physical Chemistry A*, 2004, **108**, 2851–2858.
- 51 B. O. Roos, R. Lindh, P. Malmqvist, V. Veryazov and P. Widmark, *The Journal of Physical Chemistry A*, 2005, **109**, 6575–6579.
- 52 B. O. Roos, R. Lindh, P. Malmqvist, V. Veryazov and P. Widmark, *Chemical Physics Letters*, 2005, **409**, 295–299.
- 53 B. O. Roos, R. Lindh, P. Malmqvist, V. Veryazov, P. Widmark and A. C. Borin, *The Journal of Physical Chemistry A*, 2008, **112**, 11431–11435.
- 54 J. P. Perdew, K. Burke and M. Ernzerhof, *Phys. Rev. Lett.*, 1996, **77**, 3865–3868.
- 55 T. W. Ko, J. A. Finkler, S. Goedecker and J. Behler, *Nature Communications*, 2021, **12**, 398.
- 56 F. Banhart, J. Kotakoski and A. V. Krasheninnikov, *ACS Nano*, 2011, **5**, 26–41.
- 57 J. J. Palacios, J. Fernández-Rossier and L. Brey, *Physical Review B*, 2008, **77**, 195428.
- 58 S. Baroni, S. de Gironcoli, A. Dal Corso and P. Giannozzi, *Rev. Mod. Phys.*, 2001, **73**, 515–562.
- 59 N. Mounet, M. Gibertini, P. Schwaller, D. Campi, A. Merkys, A. Marrazzo, T. Sohler, I. E. Castelli, A. Cepellotti, G. Pizzi and N. Marzari, *Two-dimensional materials from high-throughput computational exfoliation of experimentally known compounds*, 2020.
- 60 Y. Shaidu, E. Küçükbenli, R. Lot, F. Pellegrini, E. Kaxiras and S. de Gironcoli, *npj Computational Materials*, 2021, **7**, 52.
- 61 V. L. Deringer and G. Csányi, *Phys. Rev. B*, 2017, **95**, 094203.
- 62 R. Hoffmann, A. A. Kabanov, A. A. Golov and D. M. Proserpio, *Angewandte Chemie International Edition*, 2016, **55**, 10962–10976.
- 63 J. P. Perdew, A. Ruzsinszky, G. I. Csonka, O. A. Vydrov, G. E. Scuseria, L. A. Constantin, X. Zhou and K. Burke, *Phys. Rev. Lett.*, 2008, **100**, 136406.
- 64 C. Zhang, S. KC, Y. Nie, C. Liang, W. G. Vandenberghe, R. C. Longo, Y. Zheng, F. Kong, S. Hong, R. M. Wallace and K. Cho, *ACS Nano*, 2016, **10**, 7370–7375.
- 65 A. Zunger, S.-H. Wei, L. G. Ferreira and J. E. Bernard, *Physical Review Letters*, 1990, **65**, 353–356.
- 66 C. J. Pickard and R. J. Needs, *Journal of Physics: Condensed Matter*, 2011, **23**, 053201.
- 67 J. H. Lloyd-Williams and B. Monserrat, *Phys. Rev. B*, 2015, **92**, 184301.
- 68 X. Luo, Z. Wang, P. Gao, J. Lv, Y. Wang, C. Chen and Y. Ma, *npj Computational Materials*, 2024, **10**, 254.
- 69 A. Laio and M. Parrinello, *Proceedings of the National Academy of Sciences*, 2002, **99**, 12562–12566.
- 70 M. Bonomi, D. Branduardi, G. Bussi, C. Camilloni, D. Provasi, P. Raiteri, D. Donadio, F. Marinelli, F. Pietrucci, R. A. Broglia and M. Parrinello, *Computer Physics Communications*, 2009, **180**, 1961–1972.
- 71 G. A. Tribello, M. Bonomi, D. Branduardi, C. Camilloni and G. Bussi, *Computer Physics Communications*, 2014, **185**, 604–613.
- 72 The PLUMED consortium, *Nature Methods*, 2019, **16**, 670–673.
- 73 S. P. Huber, E. Bosoni, M. Bercx, J. Bröder, A. Degomme, V. Dikan, K. Eimre, E. Flage-Larsen, A. Garcia, L. Genovese, D. Gresch, C. Johnston, G. Petretto, S. Poncé, G.-M. Rignanese, C. J. Sewell, B. Smit, V. Tseplyaev, M. Uhrin, D. Wortmann, A. V. Yakutovich, A. Zadoks, P. Zarabadi-Poor, B. Zhu, N. Marzari and G. Pizzi, *npj Computational Materials*, 2021, **7**, 136.
- 74 I. Timrov, N. Marzari and M. Cococcioni, *Phys. Rev. B*, 2018, **98**, 085127.
- 75 I. Timrov, N. Marzari and M. Cococcioni, *Phys. Rev. B*, 2021, **103**, 045141.
- 76 I. Timrov, N. Marzari and M. Cococcioni, *Computer Physics Communications*, 2022, **279**, 108455.
- 77 L. Bastonero, C. Malica, E. Macke, M. Bercx, S. Huber, I. Timrov and N. Marzari, *npj Computational Materials*, 2025, **11**, 183.
- 78 M. Uhrin, A. Zadoks, L. Binci, N. Marzari and I. Timrov, *npj Computational Materials*, 2025, **11**, 19.



- 79 P. Giannozzi, S. Baroni, N. Bonini, M. Calandra, R. Car, C. Cavazzoni, D. Ceresoli, G. L. Chiarotti, M. Cococcioni, I. Dabo, A. Dal Corso, S. de Gironcoli, S. Fabris, G. Fratesi, R. Gebauer, U. Gerstmann, C. Gougoussis, A. Kokalj, M. Lazzeri, L. Martin-Samos, N. Marzari, F. Mauri, R. Mazzarello, S. Paolini, A. Pasquarello, L. Paulatto, C. Sbraccia, S. Scandolo, G. Sclauzero, A. P. Seitsonen, A. Smogunov, P. Umari and R. M. Wentzcovitch, *Journal of Physics: Condensed Matter*, 2009, **21**, 395502.
- 80 P. Giannozzi, O. Andreussi, T. Brumme, O. Bunau, M. Buongiorno Nardelli, M. Calandra, R. Car, C. Cavazzoni, D. Ceresoli, M. Cococcioni, N. Colonna, I. Carnimeo, A. Dal Corso, S. de Gironcoli, P. Delugas, R. A. DiStasio, A. Ferretti, A. Floris, G. Fratesi, G. Fugallo, R. Gebauer, U. Gerstmann, F. Giustino, T. Gorni, J. Jia, M. Kawamura, H.-Y. Ko, A. Kokalj, E. Küçükbenli, M. Lazzeri, M. Marsili, N. Marzari, F. Mauri, N. L. Nguyen, H.-V. Nguyen, A. Otero-de-la Roza, L. Paulatto, S. Poncé, D. Rocca, R. Sabatini, B. Santra, M. Schlipf, A. P. Seitsonen, A. Smogunov, I. Timrov, T. Thonhauser, P. Umari, N. Vast, X. Wu and S. Baroni, *Journal of Physics: Condensed Matter*, 2017, **29**, 465901.
- 81 P. Giannozzi, O. Baseggio, P. Bonfà, D. Brunato, R. Car, I. Carnimeo, C. Cavazzoni, S. de Gironcoli, P. Delugas, F. Ferrari Ruffino, A. Ferretti, N. Marzari, I. Timrov, A. Urru and S. Baroni, *The Journal of Chemical Physics*, 2020, **152**, 154105.
- 82 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 and S. J. Plimpton, *Computer Physics Communications*, 2022, **271**, 108171.
- 83 S. Huber, M. Bercx, N. Hörmann, M. Uhrin, G. Pizzi and N. Marzari, *Materials Cloud three-dimensional crystals database (MC3D)*, 2022.
- 84 Y. Zhou, W. A. Saidi and K. A. Fichthorn, *The Journal of Physical Chemistry C*, 2014, **118**, 3366–3374.
- 85 D. Bidoggia, N. Manko, M. Peressi and A. Marrazzo, *AiiDA-TrainsPot: Towards automated training of neural-network interatomic potentials*, 2026, <https://doi.org/10.24435/materialscloud:ws-qe>.



Data Availability Statement for the manuscriptView Article Online
DOI: 10.1039/D6DD00005C*AiiDA-TrainsPot: Towards automated training of neural-network interatomic potentials*

By D. Bidoggia, N. Manko, M. Peressi and A. Marrazzo

The training datasets and trained models are available on the Materials Cloud Archive <https://doi.org/10.24435/materialscloud:ws-qe> with license Creative Commons Attribution 4.0 International. An archived copy of the AiiDA-TrainsPot software is also deposited in the same Materials Cloud Archive. The AiiDA-TrainsPot source code is publicly available on GitHub at <https://github.com/aiida-trieste-developers/aiida-trains-pot> under the MIT license. The repository additionally includes example input files, configuration templates, and user documentation covering installation procedures, an end-to-end tutorial, and the workflow API. The documentation is also available at <https://aiida-trieste-developers.github.io/aiida-trains-pot/>.

