

MLIP Studio: An Open Platform for Interactive Benchmarking and Atomistic Simulations Using Machine Learning Interatomic Potentials

Manas Sharma,^{*} Sudeep Punnathanam, and Ananth Govind Rajan^{*}

Department of Chemical Engineering, Indian Institute of Science, Bengaluru, Karnataka 560012, India

E-mail: manassharma@iisc.ac.in; ananthgr@iisc.ac.in

Abstract

Universal machine learning interatomic potentials (MLIPs) are foundation AI models transforming atomistic simulations, but their practical use remains hindered by fragmented software ecosystems, dependency conflicts, and the lack of accessible benchmarking tools. These models approach first-principles density functional theory (DFT) accuracy at a fraction of the computational cost. We introduce MLIP Studio (available at <https://mlipstudio.iisc.ac.in>), an open and free platform that brings more than 60 universal MLIPs into a unified interactive interface for molecules and materials. The platform enables end-to-end MLIP-driven workflows, including property prediction, geometry optimization, vibrational and equation-of-state analysis, spin-state determination, custom model deployment, and high-throughput benchmarking against reference data. Automated parity plots and sortable error tables facilitate rapid identification of element-wise outliers and problematic data points. We demonstrate that MLIP-based pre-optimization can reduce subsequent DFT optimization effort by $\sim 33\times$.

Additionally, the application enables benchmarking of computational performance. Through a comprehensive case study involving the 2D magnetic material CrCl_3 on a sapphire substrate, we show how cross-model comparisons of various properties and potential-energy landscapes can guide task-specific MLIP selection. Overall, MLIP Studio lowers the barrier to the reliable use of foundation models in end-to-end research workflows, benchmarking, and education in computational chemistry and materials science.

1 Introduction

Machine learning interatomic potentials (MLIPs) are data-driven models that approximate the quantum-mechanical potential energy surface.^{1–3} In contrast to empirical force fields, which rely on fixed functional forms with limited flexibility, MLIPs employ flexible neural network architectures that can systematically approach the accuracy of the target quantum-mechanical method on which the model is trained. Quantum-mechanical density functional theory (DFT) calculations are the workhorse of modern-day computational materials science.^{4–8} Though highly accurate, the computationally expensive nature of DFT^{9–13} has motivated the development of MLIPs that combine near ab-initio accuracy with orders-of-magnitude improvements in computational speed. Over recent years, a variety of universal or foundation MLIPs have emerged that can be applied across a broad range of chemical systems, from simple molecules to complex solids.^{14–22} Universal MLIPs are large-scale interatomic models pretrained on chemically diverse DFT datasets with millions of datapoints to achieve transferable, zero-shot accuracy across molecules, solids, and interfaces without system-specific retraining. The burgeoning field of universal MLIPs is rapidly transforming atomistic simulations in chemistry and materials science.¹⁴

Despite rapid architectural innovation, the MLIP ecosystem remains fragmented across independent software stacks, incompatible library dependencies, and heterogeneous interfaces. As a result, systematic benchmarking and reproducible comparison of models developed by different groups remain technically challenging. Furthermore, while there are several interactive applications for quantum-mechanical or classical molecular calculations (e.g., TMoleX,²³ GaussView,²⁴

BURAI,²⁵ Quantum ATK,²⁶ and Avogadro²⁷), comparable platforms for MLIPs remain largely absent, limiting the accessibility of universal MLIPs for educational training and rapid prototyping. MLIP Studio addresses these challenges by providing a free, source-available, and fully interactive interface that permits users to upload or import any chemical structure, select from among more than sixty universal MLIP models, and perform various calculations, including energy, force, stress, bandgap, and dipole moment prediction, geometry optimization, vibrational analysis, and the evaluation of bulk modulus and cohesive and atomization energies, all within a single, seamlessly integrated environment.

The foundation of MLIP Studio lies in the philosophy of accessible scientific software, a philosophy that is gaining momentum in the broader materials informatics community.^{28–30} Several recent initiatives, such as MLIP Arena³¹ and MatBench Discovery,³² have demonstrated the importance of fairness and transparency in MLIP benchmarking by considering evaluations that incorporate physical consistency, transferability, and real-world applicability. At the same time, a host of interactive web applications, built with modern frameworks such as Streamlit,³³ have emerged across domains ranging from drug discovery to materials property prediction. For example, Materials Project’s^{29,30,34} web-based application provides a searchable database of crystal structures and their computed thermodynamic and electronic properties, along with APIs and visualization tools. Similarly, C2DB³⁵ provides a web-based platform to explore and visualize data for two-dimensional materials, while Materials Cloud³⁶ offers curated datasets, interactive structure viewers, and reproducible computational workflows. Together, these web-based platforms make materials data and related tools easily accessible for exploration and analysis. Despite these significant developments, there remain very few freely accessible platforms that provide comprehensive MLIP-oriented workflows while supporting models developed by multiple research groups. Most existing solutions are proprietary commercial services, such as RowanSci,³⁷ Quantistry,³⁸ Matalantis,³⁹ and QpiAI,⁴⁰ which rely on closed-source implementations and have high subscription costs.

MLIP Studio is designed to fill this gap by leveraging freely available tools to provide a

universal and extendable web-based framework for MLIP evaluation and MLIP-based atomistic simulations. MLIP Studio provides access to sixty-two universal MLIPs from five major model families: multi-atomic cluster expansion (MACE),²⁰ FairChem,¹⁷ MatterSim,¹⁸ orbital materials (ORB),^{15,16} SevenNet,²¹ and point-edge transformer (PET),^{41,42} within the same computational environment. The web application is built using Streamlit,³³ a Python-based framework known for its rapid prototyping capabilities and ease of integration with powerful back-end libraries such as the Atomic Simulation Environment (ASE)⁴³ and Py3Dmol,⁴⁴ which allows for interactive three-dimensional visualization of molecular structures. This easy-to-use graphical user interface (GUI) ensures that MLIP-powered atomistic simulations are accessible to researchers with minimal coding expertise. The use of Python as the backend language ensures that the underlying architecture is modular and easily extendable, and works well with the existing ecosystem of diverse datasets, materials informatics libraries, and MLIP frameworks. As a result, users can seamlessly switch between MLIP models and directly compare their performance on their own datasets, without writing any code.

Unlike commercial alternatives, which impose licensing fees and offer a limited choice of MLIPs, MLIP Studio is provided entirely free under the Academic Software License (ASL). The full source code is available on our GitHub⁴⁵ repository. Additionally, we host a dedicated instance at <https://mlipstudio.iisc.ac.in>, running on our own server (with GPU acceleration), to provide a reliable and persistent access point for the community for demonstration purposes (with limits on system sizes). However, users can clone the repository and run it on their local machine to circumvent the limits placed on the free online version. This commitment to openness fosters reproducibility, community-driven development, and transparent benchmarking across the growing ecosystem of universal MLIPs.

Beyond licensing and accessibility concerns, a significant practical barrier to MLIP adoption is the difficulty of installing and managing multiple MLIP frameworks within a single computational environment. Popular MLIP packages from FAIRChem,^{17,19} MACE,²⁰ and ORB,^{15,16} often depend on specific and mutually incompatible versions of core libraries like PyTorch, Py-

Torch Geometric, NumPy, or e3nn.⁴⁶ In practice, this means that a researcher wishing to compare predictions across several MLIPs must maintain separate virtual environments for each model, activating and switching between them for every new calculation, which is tedious and hampers the setup of a seamless workflow for systematic benchmarking. MLIP Studio eliminates this friction by providing a carefully curated environment in which all supported MLIPs are installed with mutually compatible library versions. This was achieved through extensive compatibility testing, in which we systematically identified library versions that, while not always matching those specified in each package’s default requirements file, nonetheless produce correct and validated results across all supported models. Users can therefore switch between over 60 models instantaneously within a single session, without concern for dependency conflicts or environment management, enabling seamless and fair side-by-side comparisons that would otherwise require substantial setup effort.

The platform’s modular, plugin-based architecture is designed to accommodate the rapid pace of MLIP development. New models, simulation protocols, and post-processing tools can be integrated with minimal effort, ensuring that MLIP Studio remains current as new methods and datasets such as MAD,^{47–49} OMAT24,⁵⁰ OMOL25,⁵¹ MaPES,⁵² and MPTraj,^{29,30,34} continue to emerge. By supporting over sixty MLIP models spanning diverse architectural families in one place, the platform enables side-by-side comparisons that go beyond standard error metrics to include assessments of physical consistency, stability, and transferability based on users’ own datasets.

Furthermore, user experience and accessibility have been key design criteria in the development of MLIP Studio. Recognizing that robust computational tools are only as effective as their usability, the web app has been engineered to support a minimal learning curve for end users. The integration of a drag-and-drop interface for file uploads, combined with intuitive selection menus for choosing simulation protocols and MLIP models, ensures that even users with limited computational backgrounds can harness the power of universal MLIPs to explore complex chemical phenomena. Therefore, beyond research workflows, MLIP Studio also serves as a pedagogical platform. By providing real-time visualization of geometry optimization trajectories, vibrational

spectra, equation-of-state fitting, and spin-state scans, the application allows users to interactively explore core atomistic simulation concepts without writing code. This makes the platform suitable for classroom demonstrations and graduate-level computational chemistry and materials science courses.

In addition to improving accessibility and unifying diverse MLIP frameworks, MLIP Studio is designed with a strong emphasis on practical workflows and model evaluation. It facilitates MLIP-based geometry pre-optimization, which yields structures close to the DFT minimum, significantly reducing subsequent first-principles optimization cost. The platform also enables rapid screening of configuration ensembles, where energy profiling and downloadable data allow straightforward identification of lowest-energy structures for further simulations. When reference data are available, automatically generated parity plots and sortable error tables facilitate identification of high-error configurations and efficient diagnosis of problematic training frames. All calculations report wall times, enabling systematic benchmarking of model performance across hardware. The platform also supports spin-state determination, providing useful initial guesses for DFT calculations. Additionally, it facilitates cross-model comparisons for potential energy surfaces generated through different MLIPs.

The remainder of this manuscript is organized as follows. Section 2 introduces the conceptual foundations of universal MLIPs and briefly summarizes the architectural families supported in the platform. Section 3 describes the general workflow of the web application, including structure input, model selection, and task configuration. Section 4 details the supported simulation tasks and illustrates representative applications. Additionally, Section 4 presents computational performance benchmarks on CPU and GPU hardware. Section 5 demonstrates an end-to-end case study involving the precursor for the 2D magnetic material CrCl_3 on a sapphire substrate, which combines these capabilities for model evaluation and selection on a realistic molecule–surface system. Finally, Section 6 concludes with an outlook on future developments and extensions.

2 Universal Machine Learning Interatomic Potentials Supported by MLIP Studio

An MLIP learns a mapping from atomic numbers and Cartesian coordinates to quantum-mechanical observables of interest. Usually, the primary target is the potential energy surface. In conservative formulations, forces are obtained as the negative gradient of the predicted potential energy; in direct-force formulations, forces are predicted directly (without differentiation) for improved throughput and reduced memory footprint. Beyond energies and forces, MLIPs can also be trained to predict properties such as dipole moments, band gaps, partial charges, density of states (DOS), and so on.⁴⁹ Most MLIPs developed by research groups for a particular problem are bespoke: they are trained for a specific material, phase, or reaction pathway and optimized for accuracy within that restricted chemical space. Their predictive reliability typically degrades under distribution shifts, requiring retraining or active learning when new environments are encountered. Universal or foundation MLIPs pursue a different objective. They are pretrained on chemically diverse datasets comprising millions of quantum-mechanical DFT calculations. The goal is transferable accuracy across broad regions of the chemical space without system-specific retraining. In this sense, a universal MLIP functions as a pretrained foundation model that can be deployed in a zero-shot manner or finetuned for downstream tasks.

Architecturally, most modern MLIPs are geometric graph neural networks (GNNs).¹⁴ In this formulation, an atomic configuration is represented as a graph whose nodes correspond to atoms and whose edges connect neighboring atoms within a finite radial cutoff. Message passing layers iteratively update atomic feature vectors by aggregating information from neighboring atoms, thereby constructing many-body representations of the local chemical environment. The radial cutoff defines the spatial extent of interactions included in the model and ensures linear scaling with system size by restricting computation to local neighborhoods.⁵³

Physical symmetries impose additional constraints on the predictions made by MLIPs. The predicted total potential energy must be invariant to global translations, rotations, and permutations of

atoms, while forces must transform equivariantly under rotations.⁵³ State-of-the-art MLIPs enforce these symmetries either explicitly through equivariant tensor representations^{20,22} using spherical harmonics and Clebsch-Gordan coefficients to represent and update atomic environments (e.g., MACE²⁰ and NequIP⁴⁶) or implicitly through data augmentation and regularization (equigrad).¹⁵ For a systematic explanation of geometric GNN architectures, symmetry handling, tensor representations, and their mathematical foundations, we refer the reader to the detailed review in ref. 53.

Within this landscape, several architectural families of universal MLIPs have emerged that differ in how they enforce symmetry, scale model capacity, and optimize computational efficiency. Below, we briefly summarize the major model families supported by MLIP Studio and highlight the design choices that distinguish them.

- **MACE:** MACE introduced higher-order equivariant message passing in which many-body correlations can be captured within each interaction block, building upon the atomic cluster expansion (ACE) formulation.⁵⁴ By increasing body order without increasing network depth, MACE achieves high expressivity with only a few layers, improving data efficiency and parallel scalability. Its use of irreducible spherical tensor features ensures strict $E(3)$ equivariance while keeping parameter growth controlled relative to earlier tensor field networks.²⁰
- **ORB:** The ORBv3 family charts the performance-speed-memory Pareto frontier by offering both conservative models (that enforce rotational invariance through equigrad regularization) and lightweight direct-force models. The latter provides a $\approx 10\times$ reduction in latency and $\approx 8\times$ reduction in memory compared to traditional MLIPs, enabling stable molecular dynamics for systems exceeding 100,000 atoms.¹⁵
- **UMA (FAIRChem):** The universal model for atoms (UMA)¹⁷ family is based on the equivariant smooth energy network (eSEN) architecture¹⁹ and leverages a mixture of linear experts (MoLE) architecture to scale model capacity (up to 1.4B parameters). Trained on

a massive multi-domain dataset of 459 million atomic systems, UMA supports multi-task learning. Users can specify the target domain, such as `omat` (materials),⁵⁰ `omo1` (molecules),⁵¹ `oc20` (catalysis), `omc` (molecular crystals), or `odac` (MOFs), via a task embedding that routes representations through specialized experts. A UMA model with the `omo1` task head can take in spin and charge as inputs in addition to the atoms and their positions.

- **MatterSim (Microsoft):** Designed for generalization across extreme thermodynamic conditions, MatterSim¹⁸ is trained on the MPF-TP dataset. This dataset covers temperatures from 0 to 5000 K and pressures up to 1000 GPa, and training is performed using an uncertainty-aware sampling and active learning framework. This makes MatterSim particularly useful for various application areas, such as, lattice dynamics, mechanical properties, and thermodynamics.
- **SevenNet:** This equivariant GNN potential is specifically optimized for large-scale molecular dynamics through a scalable parallel algorithm compatible with spatial decomposition. By restricting communication to the original cutoff radius while exchanging node features and energy gradients, SevenNet achieves over 80% parallel efficiency in weak-scaling scenarios on multi-GPU clusters.
- **PET:** PET⁴² adopts a fundamentally different design from equivariant tensor GNNs by using a transformer-based message passing architecture operating on directed edges. Instead of enforcing $E(3)$ symmetry explicitly through spherical tensor algebra (as in MACE or NequIP), PET is rotationally unconstrained and learns approximate invariance/equivariance through data augmentation. At each message-passing step, edge features corresponding to bonds are updated via attention-based transformer blocks, allowing flexible and effectively high-order interactions with a small number of layers. This yields a model that is a universal approximator even at shallow depth, with virtually unrestricted body order and angular resolution. In practice, this leads to lower inference cost and favorable scaling, placing PET-based models on a competitive accuracy–efficiency Pareto frontier along with models such as

MACE, SevenNet, and ORB. Additionally, PET models support both conservative energy-based training and direct-force prediction heads, similar in spirit to ORB, offering further flexibility in performance–accuracy trade-offs. PET developers provide models trained on popular datasets such as MAD,^{47–49} OMAT,⁵⁰ and SPICE.⁵⁵

2.1 Available MLIP Models

Table 1 provides a comprehensive list of the universal MLIPs supported by the platform, organized by their respective model families:

Table 1: Universal MLIP models supported in MLIP Studio, organized by family.

MACE ^{20,56–58}	FAIRChem (UMA/ESEN) ^{17,19}	ORB v3 ¹⁵	MatterSim ¹⁸	SevenNet ^{21,59,60}	PET ^{41,42,47,49}
MACE MPA Medium	UMA Small 1.1	V3 OMOL Conservative	V1 SMALL	7net-0	PET-MAD-XS (v1; v1.5)
MACE OMAT Medium	UMA Small 1.2	V3 OMOL Direct	V1 LARGE	7net-l3i5	PET-MAD-S (v1; v1.5)
MACE OMAT Small	ESEN MD Direct All OMOL	V3 OMAT Conservative (inf)		7net-omat	PET-OAM-L
MACE MATPES r2SCAN Medium	ESEN SM Conserving All OMOL	V3 OMAT Conservative (20)		7net-mf-ompa	PET-OMAT-XS
MACE MATPES PBE Medium	ESEN SM Direct All OMOL	V3 OMAT Direct (inf)			PET-OMAT-S
MACE MP 0a Small		V3 OMAT Direct (20)			PET-OMAT-M
MACE MP 0a Medium		V3 MPA Conservative (inf)			PET-OMAT-L
MACE MP 0a Large		V3 MPA Conservative (20)			PET-OMATPES-L
MACE MP 0b Small		V3 MPA Direct (inf)			PET-SPICE-S
MACE MP 0b Medium		V3 MPA Direct (20)			PET-SPICE-L
MACE MP 0b2 Small					PET-MAD-DOS
MACE MP 0b2 Medium					PET-OMAD-XS
MACE MP 0b2 Large					PET-OMAD-S
MACE MP 0b3 Medium					PET-OMAD-L
MACE ANI-CC Large (500k)					
MACE OMOL-0 XL 4M					
MACE OMOL-0 XL 1024					
MACE OFF 23 Large					
MACE OFF 23 Medium					
MACE OFF 23 Small					
MACE OFF 24 Medium					
MACE POLAR 1 Small					
MACE POLAR 1 Medium					
MACE POLAR 1 Large					

3 General Workflow

The MLIP Studio web application is designed around a streamlined, four-step workflow: (i) structure input, (ii) model selection, (iii) task configuration, and (iv) execution and output visualization. Figure 1 provides an annotated overview of the home page, highlighting each of these stages.

3.1 Structure Input

The application supports multiple input methods, as shown in Figure 1a. Users may choose from the following options:

- **Select Example:** A set of predefined example structures is provided, allowing users to quickly explore the application without preparing any input files.
- **Upload File:** Users can upload a structure file from their local device. Supported file formats include CIF, extXYZ, XYZ, POSCAR, TURBOMOLE, and MOL.
- **Paste Content:** Instead of uploading a file, users may paste the contents of a structure file in any of the supported formats directly into a text box.
- **Materials Project ID:** A crystal structure can be imported directly from the Materials Project database³⁰ by specifying its Materials Project identifier (mp-id).
- **PubChem:** Molecular structures can be imported from the PubChem database^{61–64} by providing either the molecule name or its chemical formula.
- **Batch Upload:** Multiple structure files in any combination of the supported formats can be uploaded simultaneously for batch processing.
- **extXYZ Trajectory Upload:** An extended XYZ trajectory file containing multiple frames or configurations can be uploaded for sequential evaluation.

Once a structure is provided through any of these methods, it is parsed and rendered in an interactive three-dimensional viewer using Py3Dmol,⁴⁴ as shown in Figure 1c. Users can rotate, zoom, and inspect the structure, and select from various visualization styles, such as, ball-and-stick, stick, or space-filling representations.

3.2 Model Selection

The model selection panel, shown in Figure 1b, allows users to choose from five major MLIP families (already described above) as well as semi-empirical (GFN2-xTB)⁶⁵ and classical (UFF)⁶⁶ calculators and a standalone D3 dispersion correction calculator.⁶⁷ Having a separate dispersion calculator is useful because not all models support a van der Waals dispersion correction option. Using this option, the dispersion correction can be applied later on for unsupported models as well. Upon selecting a model family, the user is presented with the available models within that family. Additionally, for MACE-architecture models, users have the option to provide a custom model, either by uploading a model file or by specifying a URL from which the model can be downloaded. This enables the use of finetuned or in-development models within the same interface.

3.3 Calculation Setup and Execution

The calculation setup panel, shown in Figure 1d, summarizes the current configuration before execution. It displays the selected model type (e.g., MACE), the specific model name (e.g., MACE MPA Medium), whether dispersion correction is enabled (applicable to MACE models that support D3 corrections⁶⁸), and the computational device (CPU or CUDA, if a compatible GPU is available). The panel also shows the selected calculation task. Once all parameters are configured, the user initiates the computation by clicking the “Run Calculation” button. Upon completion, results are displayed directly in the browser, as described in Section 4.

4 Features and Applications

MLIP Studio supports a diverse set of atomistic calculation tasks, each accessible through the task selection menu in the calculation setup panel (Figure 1d). In this section, we describe each available task, the corresponding output, and illustrative applications.




Figure 1: Annotated screenshot of the MLIP Studio home page. (a) Input options panel showing the available methods for providing a chemical structure: selecting a built-in example, uploading a file, pasting file contents, importing from Materials Project or PubChem, batch uploading multiple files, or uploading an extXYZ trajectory. (b) Model selection panel listing the supported MLIP families, semi-empirical and classical calculators, and options for using predefined models, uploading a custom model, or specifying a model URL. (c) Interactive three-dimensional structure visualization of the selected or imported structure (here, a 4×4 hBN monolayer), rendered using Py3Dmol. (d) Calculation setup panel displaying the selected model type and name, dispersion correction status, computational device, selected task, and the “Run Calculation” button.

4.1 Energy, Force, and Stress Evaluation

The most fundamental task is the single-point evaluation of the total energy, atomic forces, and stress tensor for a given structure. Upon completion, the application reports the total potential energy (in eV), the maximum atomic force magnitude (in $\text{eV}/\text{\AA}$), the individual force components on each atom in a downloadable table, and the full 3×3 stress tensor. The wall time for the evaluation is also displayed. Figure 2a shows representative output for a water molecule, including the tabulated per-atom force components.

4.2 Atomization and Cohesive Energy

The atomization/cohesive energy task computes the energy required to dissociate a system into its constituent isolated atoms. For molecular (nonperiodic) systems, the quantity reported is the atomization energy; for periodic systems, it is the cohesive energy. In both cases, the definition is

$$E_{\text{at/coh}} = \sum_i E_{\text{atom},i} - E_{\text{system}}, \quad (1)$$

where E_{system} is the total energy of the system and $E_{\text{atom},i}$ is the energy of each isolated constituent atom computed with the same MLIP. The application reports $E_{\text{at/coh}}$, along with E_{system} and $\sum_i E_{\text{atom},i}$ individually. This provides a rapid means of comparing the relative stability of different compounds or polymorphs without requiring a full DFT calculation. Figure 2b shows the result of a cohesive energy calculation performed using the MACE MPA Medium model.




Figure 2: Screenshots of MLIP Studio output for representative calculations. (a) Energy, force, and stress evaluation results for a water molecule, showing the total energy, maximum force, time taken, and the tabulated per-atom force components. (b) Atomization/cohesive energy calculation results using the MACE MPA Medium model, reporting the atomization energy, system energy, and total isolated atom energy. (c) Geometry optimization input settings showing the task selection, maximum number of steps, force convergence threshold, and optimizer selection (here, LBFGS).

4.3 Geometry Optimization

Geometry optimization relaxes the atomic coordinates (and, optionally, the unit cell for periodic systems) to minimize the total potential energy. MLIP Studio offers a selection of optimizers available through ASE,⁴³ including LBFGS,⁶⁹ BFGS,^{70–73} LBFGSLineSearch,⁷⁴ BFGSLineSearch,⁷⁴ FIRE,⁷⁵ GPMIn,⁷⁶ and MDMIn. Users can specify the maximum number of optimization steps and the force convergence criterion (in eV/Å) through the interface, as shown in Figure 2c.

During the optimization, the application displays a live-updated table of the energy and maximum force at each iteration. Upon convergence, the following outputs are provided: (i) the final optimized energy and maximum residual force, (ii) the total number of steps taken and whether convergence was achieved, (iii) the optimized geometry, visualized interactively, (iv) a downloadable table of per-step energies and forces, (v) a plot of the energy as a function of the optimization cycle, and (vi) the full optimization trajectory, which can be viewed and downloaded. Figure 3a shows the tabulated per-step optimization data, and Figure 3b shows the energy convergence plot for an ibuprofen molecule optimized using the `UMAs1.1` model with the LBFGS optimizer. These comprehensive outputs also make MLIP Studio a valuable pedagogical tool for introducing geometry optimization concepts and comparing the behavior of different optimization algorithms in lecture settings.

In addition to the standard ASE optimizers, we have implemented a custom multi-stage optimizer called Fast-MSO (Fast Multi-Stage Optimizer), which chains three optimizers in a deterministic, monotonic sequence: FIRE $\rightarrow$ MDMIn $\rightarrow$ LBFGS. The rationale is that FIRE is robust for configurations with large residual forces, MDMIn provides efficient downhill relaxation at intermediate force magnitudes, and LBFGS achieves rapid final convergence near the minimum. Stage transitions occur when the maximum force (absolute) drops below user-configurable thresholds (f_{FIRE} and f_{MDMIn} , with default values of 0.8 and 0.25 eV/Å, respectively). Upon each transition, the new optimizer is freshly initialized to avoid contamination from any stale internal states (e.g., Hessian approximations or velocity vectors) accumulated by the previous optimizer. Once the transition happens, it is not reversible even if the maximum force increases at some future step. This

is done to ensure reproducible behavior.




Figure 3: Geometry optimization output for ibuprofen using the UMA OMOL s1.1 model and the LBFGS optimizer with a force convergence threshold of 0.01 eV/Å. (a) Tabulated energy and maximum force ($f_{\max}$) for the final optimization steps, showing convergence in 195 iterations. The final energy, maximum force, number of steps, and convergence status are reported below the table. (b) Energy convergence profile plotted as a function of the optimization cycle, with the final converged energy indicated by a dashed red line.

4.3.1 Application: MLIP Pre-optimization for Accelerating DFT Calculations

A particularly impactful application of the geometry optimization feature is the pre-optimization of structures prior to DFT calculations. In conventional computational workflows, initial geometries constructed using software such as Materials Studio⁷⁷ or Avogadro²⁷ are often pre-relaxed with UFF⁶⁶ before being submitted to DFT optimization. However, the accuracy of UFF is inferior to that of modern universal MLIPs, resulting in pre-optimized geometries that still require substantial DFT effort to converge. By contrast, MLIP pre-optimization yields structures that are already close to the DFT potential energy minimum, dramatically reducing the number of subsequent DFT optimization steps.

To demonstrate this, we pre-optimized three representative systems using MLIP Studio and subsequently performed full DFT geometry optimizations, comparing the results with DFT optimizations starting from the same initial (non-pre-optimized) geometries. The three systems: caf-

caffeine, ibuprofen, and a periodic box of 64 water molecules are shown in Figure 4. The MLIP pre-optimization convergence criterion was set to 0.01 eV/Å for the maximum force in all cases.




Figure 4: Structures used for the MLIP pre-optimization study. (a) Caffeine molecule. (b) Ibuprofen molecule. (c) Periodic simulation box containing 64 water molecules, (H₂O)₆₄, generated using Packmol.⁷⁸ Visualizations generated using CrysX-3D Viewer.⁷⁹

For caffeine and ibuprofen, pre-optimization was performed using the `UMA OMOL s1.1` model, which is trained on the OMOL dataset generated with the ω B97M-V functional.⁸⁰ The subsequent DFT optimizations were carried out using TURBOMOLE v7.9^{81,82} with the def2-TZVP basis set,⁸³ universal auxiliary basis set,⁸³ and the ω B97M-V functional (convergence criterion: 10^{-3} a.u./Bohr for forces and 10^{-3} a.u. for energy), ensuring consistency between the MLIP training data and the DFT reference. For comparison with a conventional classical pre-relaxation strategy, the same molecular systems were also pre-optimized using UFF⁶⁶ before the DFT optimization. For ibuprofen, the MLIP pre-optimization converged in 195 LBFGS iterations. The subsequent DFT optimization required only 7 cycles to converge. In contrast, the UFF-pre-optimized geometry required 162 DFT optimization cycles, while DFT optimization starting from the same initial geometry without pre-optimization required 185 cycles. Thus, UMA pre-optimization reduces the DFT optimization effort by over $26\times$ relative to the unoptimized starting geometry and by over $23\times$ relative to UFF pre-optimization. For caffeine, DFT optimization after UMA pre-optimization converged in just 6 cycles, compared to 30 cycles without pre-optimization and 33

cycles after UFF pre-optimization, yielding a speedup of approximately $5\times$ over the from-scratch DFT optimization.

For the water box, pre-optimization was performed using the MACE OMAT-0 Medium model, which is trained on DFT data obtained with the Perdew-Burke-Ernzerhof (PBE) exchange-correlation functional^{84,85} at the generalized gradient approximation level of theory. The subsequent DFT optimization was performed using the Vienna Ab initio Simulation Package (VASP)^{86,87} (functional: PBE; force convergence criterion: $0.05\text{ eV/\AA}$; plane-wave cutoff: 450 eV ; Γ point) using the projector augmented wave (PAW)^{88,89} potentials. With MLIP pre-optimization, the VASP geometry optimization converged in only 24 ionic steps, whereas starting from the unoptimized Packmol geometry required approximately 649 ionic steps, showcasing a reduction of over $27\times$. The benefit is even more striking when measured in terms of total electronic self-consistent field (SCF) steps: the pre-optimized run required only 180 total SCF iterations compared to 5,939 for the from-scratch optimization, corresponding to a speedup of approximately $33\times$. This dramatic improvement arises because the randomly packed Packmol geometries contain substantial atomic overlaps and high initial forces. Without pre-optimization, the maximum absolute force at the first VASP ionic step was approximately $1.3\text{ eV/\AA}$, whereas the MLIP pre-optimized structure started with a maximum force of just $0.18\text{ eV/\AA}$, placing it immediately much closer to the local minimum, enabling efficient convergence of the DFT optimizer. These results are summarized in Table 2.

Table 2: Comparison of the DFT geometry optimization effort with and without MLIP pre-optimization. For caffeine and ibuprofen, DFT calculations were performed with TURBOMOLE v7.9 (ω B97M-V/def2-TZVP; convergence: 10^{-3} a.u./Bohr), and pre-optimization used the UMA OMOL s1.1 model (convergence: $0.01\text{ eV/\AA}$). UFF-pre-optimized starting geometries were also tested for the two molecular systems. For $(\text{H}_2\text{O})_{64}$, DFT calculations were performed with VASP (PBE; convergence: $0.05\text{ eV/\AA}$) and pre-optimization used the MACE OMAT-0 Medium model (convergence: $0.01\text{ eV/\AA}$).

System	MLIP Steps	DFT Optimization Cycles/Ionic Steps			Speedup
		With MLIP pre-opt.	With UFF pre-opt.	Without pre-opt.	
Ibuprofen	195	7	162	185	$> 26\times$
Caffeine	26	6	33	30	$\sim 5\times$
$(\text{H}_2\text{O})_{64}$	720	24	–	649	$> 27\times$

4.4 Vibrational Mode Analysis

The vibrational mode analysis task computes the normal-mode frequencies of a structure using finite-difference approximations to the forces. Upon completion, the application reports a table of all vibrational modes with their frequencies (in cm^{-1}) and classification (physical mode or low-frequency translational/rotational mode), a histogram of the frequency distribution, the zero-point vibrational energy (ZPE) at a user-specified temperature, and a downloadable CSV file of the vibrational data. This feature is particularly useful for confirming whether a stationary point obtained, for example, from a nudged elastic band (NEB) calculation, is a true transition state (characterized by exactly one imaginary frequency) or a local minimum. Figure 5 shows the vibrational analysis output for a water molecule computed using the MACE OFF24 Medium model. The predicted vibrational modes at 1626, 3839, and 3945 cm^{-1} are in excellent agreement with CCSD(T) values of 1650, 3835, and 3945 cm^{-1} , respectively.⁹⁰ Additionally, the predicted ZPE of 0.587 eV is in excellent agreement with the experimental value of 0.584 eV.⁹¹

4.5 Equation of State

The equation of state (EOS) task allows users to compute the energy–volume relationship and extract mechanical properties for periodic systems. The user specifies the number of volume points and the volume range as a percentage of the equilibrium volume ($V_0 \pm x\%$), and selects an equation of state for fitting: Birch–Murnaghan,⁹² Murnaghan,⁹³ or Vinet.⁹⁴ The application uniformly scales the unit cell over the specified volume range, evaluates the energy at each volume using the selected MLIP, and fits the chosen EOS to the resulting $E(V)$ data.

For example, the third-order Birch–Murnaghan equation of state is given by

$$E(V) = E_0 + \frac{9V_0B_0}{16} \left\{ \left[\left(\frac{V_0}{V} \right)^{2/3} - 1 \right]^3 B'_0 + \left[\left(\frac{V_0}{V} \right)^{2/3} - 1 \right]^2 \left[6 - 4 \left(\frac{V_0}{V} \right)^{2/3} \right] \right\}, \quad (2)$$

where E_0 is the equilibrium energy (minimum of the $E(V)$ curve), V_0 is the equilibrium volume,




Figure 5: Vibrational mode analysis results for an H₂O molecule using the MACE OFF24 Medium model. (a) Vibrational mode summary table listing all nine modes (including translational and rotational modes), their frequencies in cm⁻¹, and their classification. Three physical vibrational modes are identified at 1626, 3839, and 3945 cm⁻¹. (b) Frequency distribution histogram and the computed zero-point vibrational energy (ZPE) of 0.587 eV.

B_0 is the bulk modulus at equilibrium defined as

$$B_0 = V \left. \frac{\partial^2 E}{\partial V^2} \right|_{V=V_0}, \quad (3)$$

and B'_0 is the first pressure derivative of the bulk modulus,

$$B'_0 = \left. \frac{\partial B}{\partial P} \right|_{P=0}. \quad (4)$$

Here, B_0 quantifies the resistance of the material to volume compression, while B'_0 describes how this resistance changes under applied pressure.

Similarly, the Murnaghan equation of state is expressed as

$$E(V) = E_0 + \frac{B_0 V}{B'_0(B'_0 - 1)} \left[\left(\frac{V_0}{V} \right)^{B'_0} (B'_0 - 1) + 1 \right] - \frac{B_0 V_0}{B'_0 - 1}, \quad (5)$$

and the Vinet equation of state is given by

$$E(V) = E_0 + \frac{2B_0 V_0}{(B'_0 - 1)^2} \left\{ 2 - \left[5 + 3(B'_0 - 1) \left(\frac{V}{V_0} \right)^{1/3} - 3 \frac{V}{V_0} \right] \exp \left[-\frac{3}{2}(B'_0 - 1) \left(\left(\frac{V}{V_0} \right)^{1/3} - 1 \right) \right] \right\}. \quad (6)$$

The MLIP Studio outputs include the bulk modulus B_0 (in GPa), its pressure derivative B'_0 , the equilibrium volume V_0 , and the equilibrium energy E_0 , all with associated fitting uncertainties. A table of energy versus volume data, including the lattice constants at each volume, is provided for downloading. The fitted equation of state is also plotted along with the computed data points, with the equilibrium point marked. Figure 6 shows both the input settings and the output for a silicon crystal computed using the MACE OMAT-0 Medium model with the Birch–Murnaghan EOS.

4.6 Spin State Determination

The spin determination task leverages the ability of certain MLIPs, notably the UMA OMOL family, trained on the OMOL dataset that includes spin and charge information, to accept spin




Figure 6: Equation of state calculation for silicon using the MACE OMAT-0 Medium model. (a) Input settings showing the task selection, number of volume points (11), volume range ($\pm 10\%$), and equation of state (Birch–Murnaghan). (b) Output showing the bulk modulus ($B_0 = 94.36$ GPa), pressure derivative ($B'_0 = 5.349$), equilibrium volume ($V_0 = 159.52$ Å³), and equilibrium energy ($E_0 = -43.42$ eV). The energy–volume data table and the fitted Birch–Murnaghan curve (red line) with computed data points (blue circles) and the equilibrium point (green star) are also shown.

multiplicity and charge as input parameters. By scanning over a range of spin multiplicities (i.e., numbers of unpaired electrons) for a given molecule, the application identifies the ground-state spin as the configuration with the lowest total energy. This is directly applicable to providing a good guess value of the magnetic moment (MAGMOM in VASP) or the fixed number of unpaired electrons (NUPDOWN in VASP) to specify in subsequent VASP DFT calculations.

We validated this feature on five molecules with known ground-state spin multiplicities: NO (doublet, 1 unpaired electron), CH₃ (doublet, 1 unpaired electron), MoCl₅ (doublet, 1 unpaired electron), MoCl₄ (triplet, 2 unpaired electrons), and CrCl₃ (quartet, 3 unpaired electrons). In all cases, the `UMA OMOL s1.1` model correctly identified the ground-state spin, as shown by the energy vs. number of unpaired electrons plots in Figure 7. The energy minimum in each plot corresponds to the known spin ground state, confirming the physical reliability of the spin-dependent energy predictions.

4.7 Batch Processing and Trajectory Analysis

MLIP Studio provides two closely related bulk-processing modes that enable efficient evaluation of large numbers of configurations: the batch uploader and the extXYZ trajectory uploader.

4.7.1 Batch Uploader

The batch uploader allows users to upload multiple structure files simultaneously in any combination of the supported formats (CIF, extXYZ, XYZ, POSCAR, TURBOMOLE, MOL). The application evaluates the energy, forces, and stress tensor for each structure using the selected MLIP. Upon completion, it reports a summary table listing the energy, maximum force, and stress of each configuration, histograms showing the distributions of energies, forces, and stresses across all configurations, and a plot of the potential energy as a function of the configuration index.




Figure 7: Spin-state determination using the UMA OMOL s1.1 model. Energy as a function of the number of unpaired electrons for (a) NO, (b) CH₃, (c) MoCl₅, (d) MoCl₄, and (e) CrCl₃. In each panel, the red star marks the energy minimum, corresponding to the predicted ground-state spin. The UMA OMOL s1.1 model correctly identifies the ground state as a doublet for NO, CH₃, and MoCl₅; a triplet for MoCl₄; and a quartet for CrCl₃, in agreement with expected values.

4.7.2 extXYZ Trajectory Uploader

The extXYZ trajectory uploader accepts a single extXYZ file containing multiple frames and evaluates each frame sequentially. This mode has two principal applications:

Potential Energy Surface Screening. In many computational workflows, simulation cells must be initialized with randomly packed molecules. For example, a box of water molecules generated by Packmol.⁷⁸ Since Packmol generates configurations stochastically, the resulting structures can vary significantly in energy, and selecting a low-energy initial configuration is desirable for subsequent molecular dynamics simulations. MLIP Studio enables rapid screening of such configuration ensembles. As a demonstration, we uploaded 1,000 configurations of a periodic box containing 64 water molecules and evaluated the potential energy of each using the MACE OMAT-0

Medium model. The energies ranged from -913 eV (most stable) to -904 eV (least stable), a spread of ~ 9 eV. Using the energy profile and the downloadable tabulated data, the lowest-energy configuration can be readily identified and extracted for further simulations.

This trajectory analysis also provides an opportunity to assess the consistency of predictions across different universal MLIPs, a critical consideration for establishing confidence in MLIP-based workflows. If models trained on similar data with the same level of theory (e.g., the PBE functional) disagree substantially on the relative energetics, the reliability of their predictions would be called into question. We therefore evaluated the 1,000 water-box configurations using eight different MLIPs: UMA OMAT s1.1, MACE OMAT-0 Medium, MACE MPA-0 Medium, MACE OMAT Small, MatterSim V1 Large, MatterSim V1 Small, ORBv3 OMAT Conservative (∞), and ORBv3 OMAT Direct (∞). Figure 8 shows the PES for the first 100 of the 1,000 total configurations so as to be able to see a clear distinction between the models. All models produce qualitatively consistent potential energy surfaces, agreeing on the relative ordering of configurations and identifying the same configurations as the highest- and lowest-energy structures. Similarly, MACE OMOL and UMA OMOL models, both trained on the OMOL dataset, also show excellent mutual agreement. This cross-model consistency is encouraging and supports the use of universal MLIPs as reliable tools for configuration screening.

Benchmarking Against Reference Data. If the input extXYZ file contains reference energies, forces, and/or stresses (e.g., from DFT calculations), the trajectory uploader automatically generates parity plots comparing the MLIP predictions with the reference data. For energies and stresses, scalar parity plots are produced; for forces, an atom-decomposed parity plot colored by element is generated, enabling identification of which atomic species exhibit the largest prediction errors. Error metrics including the mean absolute error (MAE), root-mean-square error (RMSE), and coefficient of determination (R^2) are reported.

As a demonstration, we uploaded a trajectory of 500 frames from an ab initio molecular dynamics (AIMD) simulation of a CrCl_3 molecule on a sapphire (Al_2O_3) substrate, performed in ref




Figure 8: Potential energy profiles for the first 100 of 1,000 randomly packed $(\text{H}_2\text{O})_{64}$ configurations, evaluated using eight different universal MLIPs trained on PBE-level data: UMA OMAT s1.1, MACE OMAT-0 Medium, MACE MPA Medium, MACE OMAT Small, MatterSim V1 Large, MatterSim V1 Small, ORBv3 OMAT Conservative (∞), and ORBv3 OMAT Direct (∞). All models produce qualitatively consistent energy landscapes and agree on the identification of the highest- and lowest-energy configurations.

95 by the authors using VASP at 800 K. The MACE OMAT-0 Medium model was used to predict the energies and forces for all 500 frames. The resulting force parity plot, shown in Figure 9, demonstrates good agreement between the MLIP predictions and the DFT reference, with an overall force RMSE of $0.097 \text{ eV/\AA}$ and $R^2 = 0.992$. The element-resolved coloring reveals that Cr atoms exhibit slightly larger force errors compared to Al, O, and Cl atoms, providing actionable diagnostic information for potential model refinement. Additionally, MLIP Studio also displays a downloadable table with reference and predicted values as well as the errors for each frame in the uploaded trajectory. It is also possible to sort the table by descending values of errors. This is extremely crucial to identify the frames with larger deviations from the reference values and enables quick diagnostics.




Figure 9: Benchmarking of the MACE OMAT-0 Medium model against DFT reference data for 500 AIMD frames of CrCl_3 on a sapphire substrate at 800 K. (Left) Energy parity plot. (Right) Atom-decomposed force component parity plot, colored by element (Al, O, Cr, Cl). The overall MAE is 0.065 eV/Å, RMSE is 0.097 eV/Å, and $R^2 = 0.992$ over $N = 546,000$ force components. The dashed red line indicates perfect agreement.

4.8 Custom Model Upload

The rapidly evolving MLIP landscape means that new MACE-based models are being released at an accelerating pace by diverse research groups. To provide maximum flexibility, MLIP Studio allows users to supply their own MLIP with MACE architecture, either by uploading a model file directly to the web application or by providing a URL from which the model can be downloaded. The uploaded custom model can then be used for all available inference tasks: energy evaluation, geometry optimization, vibrational analysis, equation of state, and so on, within the same unified interface.

This capability is particularly powerful when combined with the trajectory uploader with reference data (Section 4.7.2). Users can upload a custom finetuned model and immediately benchmark it against their DFT reference dataset, obtaining parity plots and error metrics without writing any code. We anticipate that this feature will be highly useful in model development workflows, where

rapid diagnostic evaluation and visual comparison of model performance are essential for guiding iterative refinement.

4.9 Computational Performance Benchmarking

All calculations in MLIP Studio report the wall-clock time upon task completion, enabling systematic benchmarking of MLIP inference speed. This is an important practical consideration, since model selection involves a trade-off between accuracy and computational cost, and the relative performance of different models can vary substantially depending on the hardware.

To provide indicative performance comparisons, we benchmarked the energy, force, and stress inference speed of representative MLIPs by evaluating 1,000 configurations of a periodic box containing 64 water molecules (the same configurations discussed in Section 4.7.2; structure files are provided in the Supporting Information). All simulations used `float32` numerical precision. For ORBv3, we used the `float32-high` precision setting, which was the fastest available option.

4.9.1 CPU Benchmarks

CPU benchmarks were performed on an Intel Core i9-12900K processor (12th generation) with 32 GB memory running Windows 11. The results, shown in Figure 10 (top panel), reveal a wide range of inference speeds. PET-MAD XS v1.5 Direct is the fastest model taking only 72 s, which is impressively fast compared to other models. MatterSim V1 Small and MACE OMAT-0 Small are the next-fastest models, requiring around 329 s and 409 s, respectively, for 1,000 evaluations. Despite its larger parameter count (5M parameters), MatterSim V1 Large also performs efficiently at 731 s. The medium-sized MACE models: OMAT, MPA, OFF24, and MatPES r2SCAN, all exhibit similar performance, completing 1,000 evaluations in approximately 1,000 s (i.e., ~ 1 configuration per second). ORBv3 OMAT Direct (∞) performs slightly faster at 916 s, while ORBv3 OMAT Conservative (∞) is notably slower at 1,905 s, taking roughly twice as long as the direct variant. The slowest models on CPU are the UMA s1.1 models (OMAT and OMOL task heads), each requiring roughly 5,400 s, and the MACE OMOL XL model, which required approxi-

mately 6,800 s for 1,000 evaluations.

4.9.2 GPU Benchmarks

GPU benchmarks were performed on an NVIDIA GeForce RTX 5070 (12 GB VRAM; rated FP32 throughput of 27.6 TFLOPS, compared to approximately 0.8 TFLOPS for the i9-12900K CPU). The GPU temperature remained below 60°C throughout all simulations. The GPU results are shown in Figure 10 (bottom panel).

On the GPU, the performance landscape changed considerably compared to the CPU. While the PET-MAD XS v1.5 Direct model still remained the fastest on the GPU taking merely ~ 23 s to evaluate 1000 configurations, the ORBv3 OMAT Direct (∞) was not far behind, emerging as the second-fastest model, completing 1,000 evaluations in just 34.4 s (over $26\times$ faster than on CPU). MACE OMAT-0 Small was the third fastest at 40.3 s. The medium-sized MACE models (OMAT, MPA, OFF24, MatPES r2SCAN) and ORBv3 OMAT Conservative (∞) clustered around 66–69 s. Notably, MatterSim V1 Small, which was the second-fastest model on the CPU, lost its advantage on the GPU (65 s), performing comparably to the medium-sized MACE models. MatterSim V1 Large similarly did not benefit as much from GPU acceleration (86.0 s) relative to its CPU performance. The MACE OMOL XL 1024 and UMA models remained the slowest, with the MACE model requiring 307 s and the UMA models taking 225–233 s on GPU, though this still represents a $\sim 23\times$ speedup over their CPU performance.

Based on these results, we recommend PET-MAD XS v1.5, MACE OMAT-0 Small and MatterSim V1 Small for CPU devices, and PET-MAD XS v1.5, ORBv3 OMAT Direct (∞), and MACE OMAT-0 Small for GPU devices, when inference speed is the primary consideration. However, for more accurate calculations, the MACE medium models are expected to be superior. These models can be further finetuned on domain-specific data to improve accuracy without sacrificing their speed advantage.

We emphasize that benchmarking the performance of computational codes is inherently sensitive to specific configurations of software libraries, compiler optimizations, and hardware. The




Figure 10: Computational performance benchmarks for universal MLIPs, measured as wall-clock time for 1,000 energy, force, and stress evaluations on (H₂O)₆₄ configurations. (Top) CPU benchmarks on an Intel Core i9-12900K (32 GB RAM, Windows 11). (Bottom) GPU benchmarks on an NVIDIA GeForce RTX 5070 (12 GB VRAM). All simulations used float32 precision. Note the different horizontal axis scales between the two panels.

results presented here are intended to provide an indicative comparison and should be interpreted accordingly. Reproducing the exact timings reported by the original developers of each MLIP would require matching their specific software and hardware environments. Nevertheless, the relative performance ratios observed in our benchmarks are physically reasonable and internally consistent, and we believe they provide a qualitatively reliable guide for model selection.

4.10 Band (HOMO-LUMO) Gap and Density of States

The electronic band gap is the energy difference between the valence band maximum and conduction band minimum, and is a key quantity governing the electrical and optical properties of materials. For molecules, the electronic energy gap arises from the difference between the energies of the lowest unoccupied molecular orbital (LUMO) and the highest occupied molecular orbital (HOMO). The DOS describes the distribution of available electronic states as a function of energy and provides a detailed picture of the electronic structure. The DOS, Fermi level, and band gap can be predicted on MLIP Studio using the PET-MAD-DOS model,⁴⁹ a universal model for electronic properties of molecules and materials. Fig. 11(b) shows how MLIP Studio can be used to predict these electronic properties for the silicon crystal. Efficient DOS prediction enables rapid screening of materials for electronic and optoelectronic applications, and can be combined with atomistic simulations to estimate finite-temperature electronic properties such as electronic heat capacity.

Additionally, we provide a lightweight, in-house developed model (MLIP Studio QM9 Gap) for predicting HOMO–LUMO gaps of organic molecules. The model is a four-layer message-passing graph neural network trained on the QM9 dataset.⁹⁶ It achieves a test MAE of 0.062 eV, which is comparable to the performance of established architectures such as SchNet,⁹⁷ PPGN,⁹⁸ Cormorant,⁹⁹ and MGCN.¹⁰⁰




Figure 11: Prediction of dipole moment, atom-resolved partial charges, band gap, and DOS. (a) Output of the dipole moment and partial charges prediction for a water molecule using the MACE-POLAR-1 M model. The model accounts for long-range electrostatics and polarization effects, enabling a consistent description of charge distribution across the system. (b) Output of the electronic DOS, Fermi level, and band gap for crystalline silicon using the PET-MAD-DOS model.

4.11 Dipole Moment and Partial Charges

The dipole moment is a measure of the separation of positive and negative charges in a system and provides insight into molecular polarity and electrostatic interactions. Atom-resolved partial charges further quantify the distribution of electronic charge across atoms, enabling a detailed description of charge transfer and local electrostatic environments. It is possible to compute the dipole moments (along x, y, z) and atom-resolved partial charges on MLIP Studio using the recently released MACE-POLAR-1 models,⁵⁶ which extend the MACE framework to include long-range electrostatics and polarization via a non-local, charge-density-based formulation. An example for a water molecule is shown in Fig. 11a. This allows for a physically consistent description of charge redistribution and induction effects across diverse molecular systems. The predicted partial charges can be used to analyze electrostatic interactions, hydrogen bonding, and charge transfer processes. They may also serve as an improved initial guess for electronic structure calculations, for example, by providing environment-dependent atomic charges in superposition of atomic densities (SAD) schemes.

5 Case Study: End-to-End Workflow for CrCl_3 Adsorption on Sapphire

To illustrate the practical use of MLIP Studio as an end-to-end simulation and benchmarking platform, we consider the CrCl_3 molecule and dimer interacting with sapphire ($\alpha\text{-Al}_2\text{O}_3$) as a representative molecule-surface interaction case study. This system is important from the perspective of modeling the chemical vapor deposition synthesis of CrCl_3 , which is a 2D magnetic material. The authors in ref. 95 showed that MACE OMAT Medium performs reasonably well for this system. Here, we compare the DFT results from that work with other universal MLIPs (trained on the PBE level DFT data) as facilitated by MLIP Studio. This system combines several levels of atomistic complexity: a bulk oxide substrate, an isolated transition-metal halide molecule with

unpaired electrons, an Al_2O_3 surface, an adsorbed interface, and a finite-temperature AIMD trajectory. It therefore provides a chemically meaningful test case for demonstrating multiple MLIP Studio capabilities within a single workflow.

The case study proceeds in four stages. We first validate the bulk Al_2O_3 substrate by comparing MLIP-predicted lattice parameters, cohesive energies, bulk moduli, and band gaps with DFT or literature reference values. We then check the capabilities of universal MLIPs, trained on spin data, to predict the correct spin state of CrCl_3 . Note that for spin-state determination, we use OMOL-based MLIPs, for which the training data is not at the PBE level of theory but rather at the $\omega\text{B97M-V}$ level. Next, we evaluate the Al_2O_3 surface energy and the potential energy surface of the CrCl_3 dimer on sapphire using several MLIPs and compare the results with DFT reference calculations from ref. 95. Finally, we use an AIMD trajectory of the CrCl_3 /sapphire interface from ref. 95 to benchmark the different MLIPs through energy and force error metrics.

The bulk validation results are summarized in Table 3. Most MLIPs reproduce the DFT-relaxed hexagonal cell within $\sim 0.4\%$ for a and $\sim 0.6\%$ for c . Interestingly, PET-OMAT-V1 M could not optimize the structure to have forces below the $0.01 \text{ eV/\AA}$ threshold utilized here. The cohesive energies predicted by the MACE and UMA models are clustered near the reported computational and experimental estimates of $6.43 \text{ eV/atom}^{101}$ and $6.38 \text{ eV/atom}^{102}$ respectively. Cohesive energy is an interesting property, as it also checks an MLIP’s ability to predict the isolated atomic energies correctly. Next, we evaluate the performance for the prediction of bulk modulus. The EOS-derived bulk moduli reported in literature $\approx 236 \text{ GPa}^{103,104}$ are softer than the experimental 252.3 GPa value reported for $\alpha\text{-Al}_2\text{O}_3$.¹⁰⁵ Most of the benchmarked MLIPs predict values ($\approx 230 \text{ GPa}$) in good agreement with the reference DFT values, except for the MACE MatPES PBE model. For band gap, the PET-MAD-DOS model yields a value of 5.92 eV , which is close to recent PBE-level values near 6 eV , while remaining below the experimental optical gap of $\sim 8.8 \text{ eV}$, as expected for semilocal electronic-structure references.¹⁰⁶

For isolated CrCl_3 , the molecular spin scan performed using MACE POLAR 1 M, UMA OMOL s1p1, and MACE OMOL XL 1024 placed the minimum at three unpaired electrons ($S = 3/2$, quartet)

Table 3: Bulk α -Al₂O₃ properties from the end-to-end workflow. Lattice parameters are for the hexagonal cell.

Model/reference	a (Å)	c (Å)	E_{coh} (eV/atom)	B_0 (GPa)	E_g (eV)
DFT, this work	4.79	13.06	—	—	—
Literature/experiment	4.76 ¹⁰⁷	12.99 ¹⁰⁷	6.38, ¹⁰² 6.43 ¹⁰¹	236.2, ¹⁰⁴ 252.3 ¹⁰⁵	6.02 (PBE), 8.8 (exp.) ¹⁰⁶
UMA OMAT s1p1	4.81	13.14	6.45	235.03	—
MACE MatPES PBE	4.79	13.07	6.44	191.80	—
MACE MPA Medium	4.80	13.12	6.46	233.95	—
MACE OMAT Medium	4.79	13.10	6.47	237.29	—
ORBv3 OMAT Conv 20	4.81	13.12	5.42	222.44	—
MatterSim V1 S	4.80	13.11	6.01	233.46	—
PET-OMAT-V1 M	—	—	—	231.97	—
PET-MAD-DOS	—	—	—	—	5.92

for the molecular models used for spin-state determination, consistent with the expected high-spin Cr(III) d^3 configuration predicted by DFT. This provides a simple but useful check that MLIP Studio can supply a physically sensible spin initialization before more expensive electronic-structure calculations.

The surface energies are collected in Table 4. For the Al-terminated Al₂O₃(0001) slab, the MLIP surface energies fall between 1.38 and 1.59 J m⁻², below but reasonably close to the present DFT value of 1.84 J m⁻² and the reported 1.65–2.13 J m⁻² range for single-Al-terminated slabs.¹⁰⁸ Additionally, we benchmark the MLIPs for the rotational potential energy curve of the tilted CrCl₃ dimer on the sapphire surface using the data points from ref. 95. The potential energy curves obtained from MLIPs are compared with the reference DFT results in Fig. 12. Most models reproduce the DFT curve qualitatively, with the MACE MatPES PBE curve coming the closest to the reference. Interestingly, the UMA OMAT s1p1 gives a much larger variation of 5.05 eV span and is therefore omitted from the figure for readability.

Finally, Table 5 compares MLIP predictions against the 500-frame AIMD trajectory of the CrCl₃/sapphire interface obtained from ref. 95. Interestingly, PET OMAT v1 M gives the smallest relative-energy error (MAE = 2.2 meV/atom), while UMA OMAT s1p1 gives the best force agreement (MAE = 51.8 eV/Å). MatterSim, MACE OMAT, and MACE MPA remain close in force accuracy, whereas ORBv3 shows a much larger relative-energy and force error for this trajectory. These results illustrate the main practical value of the workflow: the preferred model depends on

Table 4: Al-terminated $\text{Al}_2\text{O}_3(0001)$ surface energies.

Model/reference	Surface energy (J m^{-2})
DFT, this work	1.840
Literature range	1.65–2.13 ¹⁰⁸
UMA OMAT s1p1	1.468
MACE MatPES PBE	1.377
MACE MPA Medium	1.463
MACE OMAT Medium	1.595
ORBv3 OMAT Conv 20	1.442
MatterSim V1 S	1.481




Figure 12: Rotational potential energy curves for a tilted CrCl_3 dimer on sapphire. Energy is relative to the 0° configuration.

the target observable, and MLIP Studio makes that trade-off visible and serves as a useful benchmarking tool.

Table 5: AIMD parity metrics for 500 CrCl₃/sapphire frames. Relative energies are per atom and referenced to the first frame.

Model	Energy MAE (meV/atom)	Energy RMSE (meV/atom)	Force MAE (meV/Å)	Force RMSE (meV/Å)
PET OMAT M v1	2.2	2.4	71.2	103.2
MACE OMAT Medium	3.5	3.8	64.3	96.7
MACE MPA Medium	3.3	3.4	84.7	126.2
MatterSim V1 S	4.3	4.6	77.9	113.7
MACE MatPES PBE	10.6	10.8	111.2	149.0
UMA OMAT s1p1	13.5	13.6	51.8	95.5
ORBv3 OMAT Conv 20	429.5	536.9	165.6	283.2

Together, these illustrative calculations demonstrate how MLIP Studio can be used not only for individual property predictions, but also for constructing a complete workflow for structure optimization, property evaluation, DFT comparison, and model selection on chemically realistic systems.

6 Conclusion and Outlook

In this work, we presented MLIP Studio, a free web-based platform that provides unified access to over 60 universal machine learning interatomic potentials spanning six major model families, including MACE, FairChem (UMA), ORB, MatterSim, and SevenNet, along with semi-empirical (xTB) and classical (UFF) calculators. The platform supports a comprehensive suite of atomistic simulation tasks, including single-point energy, force, and stress evaluations; atomization and cohesive energy calculations; geometry optimization with multiple optimizers; vibrational mode analysis; equation-of-state fitting; and spin-state determination. Through its batch processing and trajectory analysis capabilities, MLIP Studio enables efficient high-throughput screening of molecular configurations and systematic benchmarking of MLIP predictions against DFT reference data, with automated generation of parity plots and error metrics. The ability to upload custom MACE models further extends the platform’s utility for model development and diagnostics. Additionally, built-in wall-time reporting enables systematic benchmarking of model performance across

hardware and model families.

We demonstrated several practical applications that highlight the platform’s value. MLIP-based geometry pre-optimization was shown to accelerate subsequent DFT calculations by factors of $2\times$ to $33\times$ across molecular and periodic systems, offering a compelling alternative to traditional UFF-based pre-relaxation. The spin-determination feature was validated against experimentally known ground states for various transition-metal-containing molecules, and is a useful tool for obtaining initial guess values for DFT codes. Cross-model comparison of potential energy surfaces for 1,000 water-box configurations revealed reassuring qualitative consistency (including identical identification of the lowest- and highest-energy configurations) across eight independently developed MLIPs. Comprehensive CPU and GPU performance benchmarks provided practical guidance for model selection, revealing that inference speed can vary by over an order of magnitude across models and that relative performance rankings shift between CPU and GPU execution. We also showed how the bulk evaluation of ensemble configurations can be used to extract the lowest energy structure, as well as the identification of high-error frames when benchmarking against DFT data, enabling better data curation for finetuning and training. Finally, the $\text{CrCl}_3/\text{sapphire}$ case study demonstrated the use of MLIP Studio as a complete workflow platform for a chemically realistic system, spanning bulk substrate validation, molecular spin-state determination, surface and interface energetics, and AIMD trajectory benchmarking. The results show that model performance is property-dependent: for example, PET-OMAT-v1 M gives the lowest relative-energy error, whereas UMA OMAT s1p1 provides the best force agreement for the $\text{CrCl}_3/\text{sapphire}$ AIMD trajectory, highlighting the importance of task-specific MLIP selection.

Looking forward, several directions for future development are envisioned. First, the integration of molecular dynamics capabilities, including *NVT*, *NPT*, and *NVE* ensembles, would enable the study of finite-temperature properties, diffusion, and phase transitions directly within the platform. Second, the addition of NEB calculations for transition-state searches and minimum-energy path determination would broaden the platform’s applicability to reaction mechanism studies and catalysis. Third, global optimization methods for nanocluster structure prediction represent another

natural extension. Fourth, as new MLIP architectures and models continue to emerge at a rapid pace, the modular plugin architecture of MLIP Studio is designed to accommodate their seamless integration with minimal development effort. Fifth, the development of automated workflows that chain multiple tasks, for example, geometry optimization followed by vibrational analysis to confirm the nature of a stationary point, or equation of state calculations across multiple models for cross-validation, would further enhance productivity. Finally, community-driven contributions of new tasks, post-processing tools, and model integrations are actively encouraged through the official repository.⁴⁵

By lowering the barriers to deploying, comparing, and benchmarking foundation MLIPs, MLIP Studio aims to accelerate the adoption of AI and ML methods in everyday computational chemistry and materials science workflows, and to serve as a shared infrastructure for reproducible and transparent evaluation of the rapidly growing ecosystem of foundation models for atomistic simulation.

Data Availability Statement

All data generated or analysed during this study, including the input structures and the numerical data underlying all results, figures, and tables, are provided in the Supporting Information accompanying this article. The source code for MLIP Studio is publicly available at <https://github.com/mlipstudio/MLIP-Studio>.

Acknowledgement

The authors gratefully acknowledge financial support from the Anusandhan National Research Foundation via grant SPR/2023/000397. The authors acknowledge the Supercomputer Education and Research Centre (SERC) at the Indian Institute of Science for computational facilities.

Use of Large Language Models (LLMs): OpenAI’s ChatGPT (free version) and Anthropic’s Claude (free version) were used for finding synonyms, exploring alternative framing of sentences,

and grammatical refinement. LLMs were not used for factual information or analysis/interpretation of any result. Any content generated by LLMs was reviewed and edited carefully.

Supporting Information Available

Structure files and the output data for the various calculations performed in the article.

Conflict of Interest Statement

The authors declare no conflict of interest.

References

- (1) Jacobs, R.; Morgan, D.; Attarian, S.; Meng, J.; Shen, C.; Wu, Z.; Xie, C. Y.; Yang, J. H.; Artrith, N.; Blaiszik, B.; others A practical guide to machine learning interatomic potentials—Status and future. *Current Opinion in Solid State and Materials Science* **2025**, *35*, 101214.
- (2) Wan, K.; He, J.; Shi, X. Construction of high accuracy machine learning interatomic potential for surface/interface of nanomaterials—a review. *Advanced Materials* **2024**, *36*, 2305758.
- (3) Wang, G.; Wang, C.; Zhang, X.; Li, Z.; Zhou, J.; Sun, Z. Machine learning interatomic potential: Bridge the gap between small-scale models and realistic device-scale simulations. *Iscience* **2024**, *27*.
- (4) Kohn, W.; Sham, L. J. Self-Consistent Equations Including Exchange and Correlation Effects. *Phys. Rev.* **1965**, *140*, A1133.
- (5) Verma, P.; Truhlar, D. G. Status and challenges of density functional theory. *Trends in Chemistry* **2020**, *2*, 302–318.

- (6) Bursch, M.; Mewes, J.-M.; Hansen, A.; Grimme, S. Best-practice DFT protocols for basic molecular computational chemistry. *Angewandte Chemie* **2022**, *134*, e202205735.
- (7) Nørskov, J. K.; Abild-Pedersen, F.; Studt, F.; Bligaard, T. Density functional theory in surface chemistry and catalysis. *Proceedings of the National Academy of Sciences* **2011**, *108*, 937–943.
- (8) Teale, A. M.; Helgaker, T.; Savin, A.; Adamo, C.; Aradi, B.; Arbuznikov, A. V.; Ayers, P. W.; Baerends, E. J.; Barone, V.; Calaminici, P.; others DFT exchange: sharing perspectives on the workhorse of quantum chemistry and materials science. *Physical chemistry chemical physics* **2022**, *24*, 28700–28781.
- (9) Kresse, G.; Furthmüller, J. Efficiency of ab-initio total energy calculations for metals and semiconductors using a plane-wave basis set. *Comput. Mater. Sci.* **1996**, *6*, 15.
- (10) Staroverov, V. N.; Scuseria, G. E.; Tao, J.; Perdew, J. P. Tests of a Ladder of Density Functionals for Bulk Solids and Surfaces. *Phys. Rev. B* **2004**, *69*, 075102.
- (11) Paier, J.; Marsman, M.; Hummer, K.; Kresse, G.; Gerber, I. C.; Ángyán, J. G. Screened hybrid density functionals applied to solids. *J. Chem. Phys.* **2006**, *124*, 154709.
- (12) Haas, P.; Tran, F.; Blaha, P. Calculation of the Lattice Constant of Solids with Semilocal Functionals. *Phys. Rev. B* **2009**, *79*, 085104.
- (13) Burke, K. Perspective on Density Functional Theory. *J. Chem. Phys.* **2012**, *136*, 150901.
- (14) Yuan, E. C.-Y.; Liu, Y.; Chen, J.; Zhong, P.; Raja, S.; Kreiman, T.; Vargas, S.; Xu, W.; Head-Gordon, M.; Yang, C.; others Foundation models for atomistic simulation of chemistry and materials. *Nature Reviews Chemistry* **2026**, 1–19.
- (15) Rhodes, B.; Vandenhaute, S.; Šimkus, V.; Gin, J.; Godwin, J.; Duignan, T.; Neumann, M. Orb-v3: atomistic simulation at scale. *arXiv preprint arXiv:2504.06231* **2025**,

- (16) Neumann, M.; Gin, J.; Rhodes, B.; Bennett, S.; Li, Z.; Choubisa, H.; Hussey, A.; Godwin, J. Orb: A fast, scalable neural network potential. *arXiv preprint arXiv:2410.22570* **2024**,
- (17) Wood, B. M.; Dzamba, M.; Fu, X.; Gao, M.; Shuaibi, M.; Barroso-Luque, L.; Abdelmaq-soud, K.; Gharakhanyan, V.; Kitchin, J. R.; Levine, D. S.; others Uma: A family of universal models for atoms. *arXiv preprint arXiv:2506.23971* **2025**,
- (18) Yang, H.; Hu, C.; Zhou, Y.; Liu, X.; Shi, Y.; Li, J.; Li, G.; Chen, Z.; Chen, S.; Zeni, C.; others Mattersim: A deep learning atomistic model across elements, temperatures and pressures. *arXiv preprint arXiv:2405.04967* **2024**,
- (19) Fu, X.; Wood, B. M.; Barroso-Luque, L.; Levine, D. S.; Gao, M.; Dzamba, M.; Zitnick, C. L. Learning smooth and expressive interatomic potentials for physical property prediction. *arXiv preprint arXiv:2502.12147* **2025**,
- (20) Batatia, I.; Kovacs, D. P.; Simm, G.; Ortner, C.; Csányi, G. MACE: Higher order equivari-ant message passing neural networks for fast and accurate force fields. *Advances in neural information processing systems* **2022**, 35, 11423–11436.
- (21) Park, Y.; Kim, J.; Hwang, S.; Han, S. Scalable parallel algorithm for graph neural network interatomic potentials in molecular dynamics simulations. *Journal of chemical theory and computation* **2024**, 20, 4857–4868.
- (22) Batzner, S.; Musaelian, A.; Sun, L.; Geiger, M.; Mailoa, J. P.; Kornbluth, M.; Molinari, N.; Smidt, T. E.; Kozinsky, B. E (3)-equivariant graph neural networks for data-efficient and accurate interatomic potentials. *Nature communications* **2022**, 13, 2453.
- (23) Steffen, C.; Thomas, K.; Huniar, U.; Hellweg, A.; Rubner, O.; Schroer, A. TmoleX—a graphical user interface for TURBOMOLE. *Journal of computational chemistry* **2010**, 31, 2967–2970.
- (24) Dennington, R.; Keith, T. A.; Millam, J. M. GaussView, Version 6. 2016.

- (25) BURAI-team BURAI: A GUI System for Quantum ESPRESSO. <https://github.com/BURAI-team/burai>, 2016–2018; Accessed: 2026-04-19.
- (26) Smidstrup, S.; Markussen, T.; Vancraeyveld, P.; Wellendorff, J.; Schneider, J.; Gunst, T.; Verstichel, B.; Stradi, D.; Khomyakov, P. A.; Vej-Hansen, U. G.; others QuantumATK: an integrated platform of electronic and atomic-scale modelling tools. *Journal of Physics: Condensed Matter* **2020**, *32*, 015901.
- (27) Hanwell, M. D.; Curtis, D. E.; Lonie, D. C.; Vandermeersch, T.; Zurek, E.; Hutchison, G. R. Avogadro: an advanced semantic chemical editor, visualization, and analysis platform. *Journal of cheminformatics* **2012**, *4*, 17.
- (28) Lehtola, S.; Karttunen, A. J. Free and open source software for computational chemistry education. *Wiley Interdisciplinary Reviews: Computational Molecular Science* **2022**, *12*, e1610.
- (29) Jain, A.; Montoya, J.; Dwaraknath, S.; Zimmermann, N. E.; Dagdelen, J.; Horton, M.; Huck, P.; Winston, D.; Cholia, S.; Ong, S. P.; others The materials project: Accelerating materials design through theory-driven data and tools. In *Handbook of Materials Modeling: Methods: Theory and Modeling*; Springer, 2020; pp 1751–1784.
- (30) Horton, M. K.; Huck, P.; Yang, R. X.; Munro, J. M.; Dwaraknath, S.; Ganose, A. M.; Kingsbury, R. S.; Wen, M.; Shen, J. X.; Mathis, T. S.; others Accelerated data-driven materials science with the Materials Project. *Nature Materials* **2025**, *24*, 1522–1532.
- (31) Chiang, Y.; Kreiman, T.; Weaver, E.; Amin, I.; Kuner, M.; Zhang, C.; Kaplan, A.; Chrzan, D.; Blau, S. M.; Krishnapriyan, A. S.; others MLIP arena: advancing fairness and transparency in machine learning interatomic potentials through an open and accessible benchmark platform. AI for Accelerated Materials Design-ICLR 2025. 2025.
- (32) Riebesell, J.; Goodall, R. E.; Benner, P.; Chiang, Y.; Deng, B.; Lee, A. A.; Jain, A.; Pers-

- son, K. A. Matbench Discovery—A framework to evaluate machine learning crystal stability predictions. *arXiv preprint arXiv:2308.14920* **2023**,
- (33) Khorasani, M.; Abdou, M.; Fernández, J. H. Web application development with streamlit. *Software Development* **2022**, 498, 507.
- (34) Jain, A.; Ong, S. P.; Hautier, G.; Chen, W.; Richards, W. D.; Dacek, S.; Cholia, S.; Gunter, D.; Skinner, D.; Ceder, G.; others Commentary: The Materials Project: A materials genome approach to accelerating materials innovation. *APL materials* **2013**, 1.
- (35) Gjerding, M. N.; Taghizadeh, A.; Rasmussen, A.; Ali, S.; Bertoldo, F.; Deilmann, T.; Knøsgaard, N. R.; Kruse, M.; Larsen, A. H.; Manti, S.; others Recent progress of the computational 2D materials database (C2DB). *2D Materials* **2021**, 8, 044002.
- (36) Talirz, L.; Kumbhar, S.; Passaro, E.; Yakutovich, A. V.; Granata, V.; Gargiulo, F.; Borelli, M.; Uhrin, M.; Huber, S. P.; Zoupanos, S.; others Materials Cloud, a platform for open computational science. *Scientific data* **2020**, 7, 299.
- (37) Rowan: ML-Powered Molecular Design and Simulation. <https://rowansci.com/>, 2026; Official website of the Rowan Scientific computational chemistry platform.
- (38) Quantistry. <https://www.quantistry.ai/>, 2026; Official website of Quantistry Materials Intelligence & Chemical Simulation Platform.
- (39) Matlantis. <https://matlantis.com/>, 2026; Official website of the Matlantis AI-based atomistic simulation platform.
- (40) QpiAI. <https://qpiai.tech/>, 2026; Official website of QpiAI, a quantum computing and AI technology company.
- (41) Pozdnyakov, S.; Ceriotti, M. Smooth, exact rotational symmetrization for deep learning on point clouds. *Advances in Neural Information Processing Systems* **2023**, 36, 79469–79501.

- (42) Mazitov, A.; Bigi, F.; Kellner, M.; Pegolo, P.; Tisi, D.; Fraux, G.; Pozdnyakov, S.; Loche, P.; Ceriotti, M. PET-MAD as a lightweight universal interatomic potential for advanced materials modeling. *Nature Communications* **2025**, *16*, 10653.
- (43) Hjorth Larsen, A.; Jørgen Mortensen, J.; Blomqvist, J.; Castelli, I. E.; Christensen, R.; Duřak, M.; Friis, J.; Groves, M. N.; Hammer, B.; Hargus, C.; others The atomic simulation environment—a Python library for working with atoms. *Journal of Physics: Condensed Matter* **2017**, *29*, 273002.
- (44) Seshadri, K.; Liu, P.; Koes, D. R. The 3Dmol.js Learning Environment: A Classroom Response System for 3D Chemical Structures. *Journal of Chemical Education* **2020**, *97*, 3872–3876.
- (45) Sharma, M.; Punnathanam, S.; Rajan, A. G. MLIP Studio: An Open Platform for Interactive Benchmarking and Simulations Using Machine Learning Interatomic Potentials. GitHub repository, 2026; <https://github.com/mlipstudio/MLIP-Studio>, Accessed: 2026-05-19.
- (46) Geiger, M.; Smidt, T. e3nn: Euclidean neural networks. *arXiv preprint arXiv:2207.09453* **2022**,
- (47) Malosso, C.; Bigi, F.; Pegolo, P.; Abbott, J. W.; Loche, P.; Rossi, M.; Ceriotti, M.; Mazitov, A. High-quality, high-information dataset for universal atomistic machine learning. *arXiv preprint arXiv:2603.02089* **2026**,
- (48) Mazitov, A.; Chorna, S.; Fraux, G.; Bercx, M.; Pizzi, G.; De, S.; Ceriotti, M. Massive Atomic Diversity: a compact universal dataset for atomistic machine learning. *Scientific data* **2025**, *12*, 1857.
- (49) How, W. B.; Febrer, P.; Chong, S.; Mazitov, A.; Bigi, F.; Kellner, M.; Pozdnyakov, S.; Ceriotti, M. A universal machine learning model for the electronic density of states. *Digital Discovery* **2026**,

- (50) Barroso-Luque, L.; Shuaibi, M.; Fu, X.; Wood, B. M.; Dzamba, M.; Gao, M.; Rizvi, A.; Zitnick, C. L.; Ulissi, Z. W. Open materials 2024 (omat24) inorganic materials dataset and models. *arXiv preprint arXiv:2410.12771* **2024**,
- (51) Levine, D. S.; Shuaibi, M.; Spotte-Smith, E. W. C.; Taylor, M. G.; Hasyim, M. R.; Michel, K.; Batatia, I.; Csányi, G.; Dzamba, M.; Eastman, P.; others The open molecules 2025 (omol25) dataset, evaluations, and models. *arXiv preprint arXiv:2505.08762* **2025**,
- (52) Kaplan, A. D.; Liu, R.; Qi, J.; Ko, T. W.; Deng, B.; Riebesell, J.; Ceder, G.; Persson, K. A.; Ong, S. P. A foundational potential energy surface dataset for materials. *arXiv preprint arXiv:2503.04070* **2025**,
- (53) Duval, A.; Mathis, S. V.; Joshi, C. K.; Schmidt, V.; Miret, S.; Malliaros, F. D.; Cohen, T.; Lio, P.; Bengio, Y.; Bronstein, M. A hitchhiker’s guide to geometric gnns for 3d atomic systems. *arXiv preprint arXiv:2312.07511* **2023**,
- (54) Drautz, R. Atomic cluster expansion for accurate and transferable interatomic potentials. *Physical Review B* **2019**, *99*, 014104.
- (55) Eastman, P.; Behara, P. K.; Dotson, D. L.; Galvelis, R.; Herr, J. E.; Horton, J. T.; Mao, Y.; Chodera, J. D.; Pritchard, B. P.; Wang, Y.; others Spice, a dataset of drug-like molecules and peptides for training machine learning potentials. *Scientific Data* **2023**, *10*, 11.
- (56) Batatia, I.; Baldwin, W. J.; Kuryla, D.; Hart, J.; Kasoar, E.; Elena, A. M.; Moore, H.; Gawkowski, M. J.; Shi, B. X.; Kapil, V.; others MACE-POLAR-1: A Polarisable Electrostatic Foundation Model for Molecular Chemistry. *arXiv preprint arXiv:2602.19411* **2026**,
- (57) Kovács, D. P.; Moore, J. H.; Browning, N. J.; Batatia, I.; Horton, J. T.; Pu, Y.; Kapil, V.; Witt, W. C.; Magdau, I.-B.; Cole, D. J.; others Mace-off: Short-range transferable machine learning force fields for organic molecules. *Journal of the American Chemical Society* **2025**, *147*, 17598–17611.

- (58) Batatia, I.; Benner, P.; Chiang, Y.; Elena, A. M.; Kovács, D. P.; Riebesell, J.; Advincula, X. R.; Asta, M.; Avaylon, M.; Baldwin, W. J.; others A foundation model for atomistic materials chemistry. *The Journal of chemical physics* **2025**, 163.
- (59) Kim, J.; Kim, J.; Kim, J.; Lee, J.; Park, Y.; Kang, Y.; Han, S. Data-efficient multifidelity training for high-fidelity machine learning interatomic potentials. *Journal of the American Chemical Society* **2024**, 147, 1042–1054.
- (60) Kim, J.; You, J.; Park, Y.; Lim, Y.; Kang, Y.; Kim, J.; Jeon, H.; Ju, S.; Hong, D.; Lee, S. Y.; others Optimizing cross-domain transfer for universal machine learning interatomic potentials. *Nature Communications* **2026**,
- (61) Kim, S.; Thiessen, P. A.; Bolton, E. E.; Chen, J.; Fu, G.; Gindulyte, A.; Han, L.; He, J.; He, S.; Shoemaker, B. A.; others PubChem substance and compound databases. *Nucleic acids research* **2016**, 44, D1202–D1213.
- (62) Kim, S.; Chen, J.; Cheng, T.; Gindulyte, A.; He, J.; He, S.; Li, Q.; Shoemaker, B. A.; Thiessen, P. A.; Yu, B.; others PubChem 2019 update: improved access to chemical data. *Nucleic acids research* **2019**, 47, D1102–D1109.
- (63) Kim, S.; Chen, J.; Cheng, T.; Gindulyte, A.; He, J.; He, S.; Li, Q.; Shoemaker, B. A.; Thiessen, P. A.; Yu, B.; others PubChem 2023 update. *Nucleic acids research* **2023**, 51, D1373–D1380.
- (64) Kim, S.; Chen, J.; Cheng, T.; Gindulyte, A.; He, J.; He, S.; Li, Q.; Shoemaker, B. A.; Thiessen, P. A.; Yu, B.; others PubChem 2025 update. *Nucleic acids research* **2025**, 53, D1516–D1525.
- (65) Bannwarth, C.; Ehlert, S.; Grimme, S. GFN2-xTB—An accurate and broadly parametrized self-consistent tight-binding quantum chemical method with multipole electrostatics and density-dependent dispersion contributions. *Journal of chemical theory and computation* **2019**, 15, 1652–1671.

- (66) Rappé, A. K.; Casewit, C. J.; Colwell, K. S.; Goddard III, W. A.; Skiff, W. M. UFF, a full periodic table force field for molecular mechanics and molecular dynamics simulations. *Journal of the American chemical society* **1992**, *114*, 10024–10035.
- (67) Takamoto, S.; Shinagawa, C.; Motoki, D.; Nakago, K.; Li, W.; Kurata, I.; Watanabe, T.; Yayama, Y.; Iriguchi, H.; Asano, Y.; Onodera, T.; Ishii, T.; Kudo, T.; Ono, H.; Sawada, R.; Ishitani, R.; Ong, M.; Yamaguchi, T.; Kataoka, T.; Hayashi, A.; Ibuka, T. PFP: Universal Neural Network Potential for Material Discovery. 2021.
- (68) Grimme, S.; Ehrlich, S.; Goerigk, L. Effect of the damping function in dispersion corrected density functional theory. *Journal of computational chemistry* **2011**, *32*, 1456–1465.
- (69) Nocedal, J. Updating quasi-Newton matrices with limited storage. *Mathematics of computation* **1980**, *35*, 773–782.
- (70) Broyden, C. G. The convergence of a class of double-rank minimization algorithms 1. general considerations. *IMA Journal of Applied Mathematics* **1970**, *6*, 76–90.
- (71) Fletcher, R. A new approach to variable metric algorithms. *The computer journal* **1970**, *13*, 317–322.
- (72) Goldfarb, D. A family of variable-metric methods derived by variational means. *Mathematics of computation* **1970**, *24*, 23–26.
- (73) Shanno, D. F. Conditioning of quasi-Newton methods for function minimization. *Mathematics of computation* **1970**, *24*, 647–656.
- (74) Nocedal, J.; Wright, S. J. *Numerical optimization*; Springer, 2006.
- (75) Bitzek, E.; Koskinen, P.; Gähler, F.; Moseler, M.; Gumbusch, P. Structural relaxation made simple. *Physical review letters* **2006**, *97*, 170201.
- (76) Garijo del Río, E.; Mortensen, J. J.; Jacobsen, K. W. Local Bayesian optimizer for atomic structures. *Physical Review B* **2019**, *100*, 104103.

- (77) Meunier, M.; Robertson, S. Materials studio 20th anniversary. 2021.
- (78) Martínez, L.; Andrade, R.; Birgin, E. G.; Martínez, J. M. PACKMOL: A package for building initial configurations for molecular dynamics simulations. *Journal of computational chemistry* **2009**, *30*, 2157–2164.
- (79) Sharma, M.; Mishra, D. CrysX: crystallographic tools for the Android platform. *Journal of Applied Crystallography* **2019**, *52*, 1449–1454.
- (80) Mardirossian, N.; Head-Gordon, M. ω B97M-V: A combinatorially optimized, range-separated hybrid, meta-GGA density functional with VV10 nonlocal correlation. *The Journal of chemical physics* **2016**, *144*.
- (81) Franzke, Y. J.; Holzer, C.; Andersen, J. H.; Begusic, T.; Bruder, F.; Coriani, S.; Della Sala, F.; Fabiano, E.; Fedotov, D. A.; Fürst, S.; others TURBOMOLE: Today and tomorrow. *Journal of chemical theory and computation* **2023**, *19*, 6859–6890.
- (82) Sharma, M.; Franzke, Y. J.; Holzer, C.; Pauly, F.; Sierka, M. Density functional theory for molecular and periodic systems in TURBOMOLE: Theory, implementation, and applications. *The Journal of Physical Chemistry A* **2025**, *129*, 9062–9083.
- (83) Weigend, F.; Ahlrichs, R. Balanced basis sets of split valence, triple zeta valence and quadruple zeta valence quality for H to Rn: Design and assessment of accuracy. *Physical Chemistry Chemical Physics* **2005**, *7*, 3297–3305.
- (84) Perdew, J. P.; Burke, K.; Ernzerhof, M. Generalized Gradient Approximation Made Simple. *Phys. Rev. Lett.* **1996**, *77*, 3865–3868.
- (85) Perdew, J. P.; Burke, K.; Ernzerhof, M. Generalized Gradient Approximation Made Simple [Phys. Rev. Lett. 77, 3865 (1996)]. *Phys. Rev. Lett.* **1997**, *78*, 1396–1396.
- (86) Hafner, J.; Kresse, G. The vienna ab-initio simulation program VASP: An efficient and

- versatile tool for studying the structural, dynamic, and electronic properties of materials. In *Properties of Complex Inorganic Solids*; Springer, 1997; pp 69–82.
- (87) Hafner, J. Ab-initio simulations of materials using VASP: Density-functional theory and beyond. *Journal of computational chemistry* **2008**, 29, 2044–2078.
- (88) Blöchl, P. E. Projector augmented-wave method. *Phys. Rev. B* **1994**, 50, 17953–17979.
- (89) Kresse, G.; Joubert, D. From ultrasoft pseudopotentials to the projector augmented-wave method. *Phys. Rev. B* **1999**, 59, 1758–1775.
- (90) Howard, J. C.; Gray, J. L.; Hardwick, A. J.; Nguyen, L. T.; Tschumper, G. S. Getting down to the fundamentals of hydrogen bonding: Anharmonic vibrational frequencies of (HF) 2 and (H2O) 2 from ab initio electronic structure computations. *Journal of Chemical Theory and Computation* **2014**, 10, 5426–5435.
- (91) Csonka, G. I.; Ruzsinszky, A.; Perdew, J. P. Estimation, computation, and experimental correction of molecular zero-point vibrational energies. *The Journal of Physical Chemistry A* **2005**, 109, 6779–6789.
- (92) Birch, F. Finite elastic strain of cubic crystals. *Physical review* **1947**, 71, 809.
- (93) Murnaghan, F. D. The compressibility of media under extreme pressures. *Proceedings of the National Academy of Sciences* **1944**, 30, 244–247.
- (94) Vinet, P.; Ferrante, J.; Smith, J.; Rose, J. A universal equation of state for solids. *Journal of Physics C: Solid State Physics* **1986**, 19, L467–L473.
- (95) Kumar, V.; Jangid, A.; Sharma, M.; Verma, M.; Pasyanthi, J.; Kumar, K. S.; Sharma, P.; Chiglintsev, E. O.; Panin, M. I.; Punnathanam, S. N.; Chernov, A. I.; Govind Rajan, A.; Singh, A. Tailored Vapor Deposition Unlocks Large-Grain, Wafer-Scale Epitaxial Growth of 2D Magnetic CrCl₃. *Advanced Materials* n/a, e14405.

- (96) Ramakrishnan, R.; Dral, P. O.; Rupp, M.; Von Lilienfeld, O. A. Quantum chemistry structures and properties of 134 kilo molecules. *Scientific data* **2014**, *1*, 1–7.
- (97) Schütt, K. T.; Sauceda, H. E.; Kindermans, P.-J.; Tkatchenko, A.; Müller, K.-R. Schnet—a deep learning architecture for molecules and materials. *The Journal of chemical physics* **2018**, *148*.
- (98) Maron, H.; Ben-Hamu, H.; Serviansky, H.; Lipman, Y. Provably powerful graph networks. *Advances in neural information processing systems* **2019**, *32*.
- (99) Anderson, B.; Hy, T. S.; Kondor, R. Cormorant: Covariant molecular neural networks. *Advances in neural information processing systems* **2019**, *32*.
- (100) Lu, C.; Liu, Q.; Wang, C.; Huang, Z.; Lin, P.; He, L. Molecular property prediction: A multilevel quantum interactions modeling perspective. Proceedings of the AAAI conference on artificial intelligence. 2019; pp 1052–1060.
- (101) Alavi, S.; Sorescu, D. C.; Thompson, D. L. Adsorption of HCl on single-crystal α -Al₂O₃ (0001) surface: a DFT study. *The Journal of Physical Chemistry B* **2003**, *107*, 186–195.
- (102) Nebergall, W. H.; Holtzclaw, H. F.; Robinson, R. W. General chemistry. (*No Title*) **1968**,
- (103) Zhang, Z.; Fu, Z.; Zhang, R.; Legut, D.; Guo, H. Anomalous mechanical strengths and shear deformation paths of Al₂O₃ polymorphs with high ionicity. *RSC advances* **2016**, *6*, 12885–12892.
- (104) Golesorkhtabar, R.; Pavone, P.; Spitaler, J.; Puschnig, P.; Draxl, C. ElaStic: A tool for calculating second-order elastic constants from first principles. *Computer Physics Communications* **2013**, *184*, 1861–1873.
- (105) Smithells, C. J. Metals reference book. V. 3. 1967, 462 P. PLENUM PRESS, 227 W. 17 TH ST., NEW YORK, N. Y. \$21. 00 **1967**,

- (106) Acharya, D.; Perilli, D.; Di Valentin, C. Optimizing the Graphene/ α -Al₂O₃ (0001) Interface through Minimization of Interfacial Stress for Improved Electronic Applications. *ACS Applied Nano Materials* **2025**, 8, 22626–22636.
- (107) Lee, W.; Lagerlof, K. Structural and electron diffraction data for sapphire (α -Al₂O₃). *Journal of electron microscopy technique* **1985**, 2, 247–258.
- (108) Ahmad, A.; Martsinovich, N. Modelling the strength of mineral–organic binding: organic molecules on the α -Al₂O₃ (0001) surface. *RSC advances* **2022**, 12, 27604–27615.

TOC Graphic

