

DeePMD-kit v3: A Multiple-Backend Framework for Machine Learning Potentials

Jinze Zeng,^{*,†,‡,¶} Duo Zhang,^{§,||,⊥} Anyang Peng,[§] Xiangyu Zhang,^{#,⊗} Sensen He,[△] Yan Wang,^{#,⊗} Xinjian Liu,^{||} Hangrui Bi,[∇] Yifan Li,^{††} Chun Cai,[§] Chengqian Zhang,[⊥] Yiming Du,^{#,‡‡} Jia-Xin Zhu,^{¶¶} Pinghui Mo,^{§§} Zhengtao Huang,^{|||} Qiyu Zeng,^{⊥⊥,‡‡} Shaochen Shi,^{⊗⊗} Xuejian Qin,^{△△,∇∇∇} Zhaoxi Yu,^{††} Chenxing Luo,^{‡‡,¶¶} Ye Ding,^{||} Yun-Pei Liu,^{§§§} Ruosong Shi,^{|||} Zhenyu Wang,^{⊥⊥,‡‡} Sigbjørn Løland Bore,^{⊗⊗⊗} Junhan Chang,^{||,△△△} Zhe Deng,^{△△△} Zhaohan Ding,^{||} Siyuan Han,^{∇∇∇} Wanrun Jiang,[§] Guolin Ke,^{||} Zhaoqing Liu,^{△△△} Denghui Lu,^{††,††} Koki Muraoka,^{‡‡,‡‡} Hananeh Oliaei,^{¶¶¶¶} Anurag Kumar Singh,^{§§§§} Haohui Que,^{|||} Weihong Xu,^{§§§} Zhangmancang Xu,^{⊥⊥,⊥⊥} Yong-Bin Zhuang,^{#,‡‡‡} Jiayu Dai,^{⊥⊥,‡‡} Timothy J. Giese,^{⊗⊗⊗⊗} Weile Jia,^{#,} Ben Xu,^{△△△△} Darrin M. York,^{⊗⊗⊗⊗} Linfeng Zhang,^{*,§,||} and Han Wang^{*,∇∇∇∇∇,†††††}

[†]School of Artificial Intelligence and Data Science, University of Science and Technology of China, Hefei 230026, P. R. China

[‡]Suzhou Institute for Advanced Research, University of Science and Technology of China, Suzhou 215123, P. R. China

[¶]Suzhou Big Data & AI Research and Engineering Center, Suzhou 215123, P. R. China

[§]AI for Science Institute, Beijing 100080, P. R. China

^{||}DP Technology, Beijing 100080, P. R. China

[⊥]Academy for Advanced Interdisciplinary Studies, Peking University, Beijing 100871, P. R. China

[#]State Key Lab of Processors, Institute of Computing Technology, Chinese Academy of Sciences, Beijing 100871, P.R. China

[⊗]University of Chinese Academy of Sciences, Beijing 100871, P. R. China

[△]Baidu Inc., Beijing 100085, P. R. China

[∇]Department of Computer Science, University of Toronto, Toronto, Ontario M5S 1A1, Canada

^{††}Department of Chemistry, Princeton University, Princeton, New Jersey 08540, United States

^{‡‡}University of Chinese Academy of Sciences, Beijing 100871, P.R. China

^{¶¶}State Key Laboratory of Physical Chemistry of Solid Surfaces, iChEM, College of Chemistry and Chemical Engineering, Xiamen University, Xiamen 361005, P.R. China

^{§§}College of Integrated Circuits, Hunan University, Changsha, 410082, P.R. China

^{|||}State Key Laboratory of Advanced Technology for Materials Synthesis and Processing, Center for Smart Materials and Device Integration, School of Material Science and Engineering, Wuhan University of Technology, Wuhan, 430070, P.R. China

^{⊥⊥}College of Science, National University of Defense Technology, Changsha 410073, P.R. China

^{##}Hunan Key Laboratory of Extreme Matter and Applications, National University of Defense Technology, Changsha 410073, P.R. China

^{⊗⊗}ByteDance Research, Beijing 100098, P. R. China

^{△△}Ningbo Institute of Materials Technology and Engineering, Chinese Academy of Sciences, Ningbo 315201, P.R. China

^{∇∇∇}College of Materials Science and Opto-Electronic Technology, University of Chinese Academy of Sciences, Beijing 100049, P.R. China

^{†††}Key Laboratory of Theoretical and Computational Photochemistry of Ministry of Education, College of Chemistry, Beijing Normal University, Beijing 100875, P. R. China

^{‡‡‡}Department of Geosciences, Princeton University, Princeton, New Jersey 08544, United States

^{¶¶¶}Department of Applied Physics and Applied Mathematics, Columbia University, New York, New York 10027, United States

^{§§§}Laboratory of AI for Electrochemistry (AI4EC), IKKEM, Xiamen, 361005, Fujian, P. R. China

^{|||}Graduate School of China Academy of Engineering Physics, Beijing 100088, P. R. China

^{⊥⊥⊥}Key Laboratory of Material Simulation Methods and Software of Ministry of Education, College of Physics, Jilin University, Changchun 130012, P.R. China

^{##}International Center of Future Science, Jilin University, Changchun, 130012, P.R. China

^{@@@}Department of Chemistry and Hylleraas Centre for Quantum Molecular Sciences, University of Oslo, 0315 Oslo, Norway

^{△△△}College of Chemistry and Molecular Engineering, Peking University, Beijing 100871, P. R. China

^{∇∇∇∇}New Cornerstone Science Laboratory, State Key Laboratory for Physical Chemistry of Solid Surfaces, Collaborative Innovation Center of Chemistry for Energy Materials, and College of Chemistry and Chemical Engineering, Xiamen University, Xiamen 361005, P. R. China

^{††††}Department of Mechanics and Engineering Science, and HEDPS and CAPT, College of Engineering, Peking University, Beijing 100871, P. R. China

^{‡‡‡‡}Department of Chemical System Engineering, The University of Tokyo, 7-3-1 Hongo, Bunkyo-ku, Tokyo 113-8656, Japan

^{¶¶¶¶}Department of Mechanical Science and Engineering, University of Illinois at Urbana-Champaign, Urbana, IL 61801, United States

^{§§§§}Department of Data Science, Indian Institute of Technology Palakkad, Kerala 678623, India

^{|||}Shanghai Astronomical Observatory, Chinese Academy of Sciences, Shanghai 200030, P.R. China

^{⊥⊥⊥⊥}International School of Materials Science and Engineering, Wuhan University of Technology, Wuhan 430070, P. R. China

^{##}Chaire de Simulation à l'Echelle Atomique (CSEA), Ecole Polytechnique Fédérale de Lausanne (EPFL), Lausanne CH-1015, Switzerland

^{@@@@}Laboratory for Biomolecular Simulation Research, Institute for Quantitative Biomedicine and Department of Chemistry and Chemical Biology, Rutgers University, Piscataway, NJ 08854, USA

^{△△△△}Graduate School of Chinese Academy of Engineering Physics, Beijing 100088, P.R. China

^{∇∇∇∇∇}National Key Laboratory of Computational Physics, Institute of Applied Physics and Computational Mathematics, Fenghao East Road 2, Beijing 100094, P.R. China

^{†††††}HEDPS, CAPT, College of Engineering, Peking University, Beijing 100871, P.R. China

Abstract

In recent years, machine learning potentials (MLPs) have become indispensable tools in physics, chemistry, and materials science, driving the development of software packages for molecular dynamics (MD) simulations and related applications. These packages, typically built on specific machine learning frameworks such as TensorFlow, PyTorch, or JAX, face integration challenges when advanced applications demand communication across different frameworks. The previous TensorFlow-based implementation of DeePMD-kit exemplified these limitations. In this work, we introduce DeePMD-kit version 3, a significant update featuring a multi-backend framework that supports TensorFlow, PyTorch, JAX, and PaddlePaddle backends, and demonstrate the versatility of this architecture through the integration of other MLPs packages and of Differentiable Molecular Force Field. This architecture allows seamless backend switching with minimal modifications, enabling users and developers to integrate DeePMD-kit with other packages using different machine learning frameworks. This innovation facilitates the development of more complex and interoperable workflows, paving the way for broader applications of MLPs in scientific research.

1 Introduction

Over the past decade, machine learning potentials (MLPs) have become increasingly influential in the fields of physics, chemistry, molecular biology, and materials science.^{1–6} This has led to the development of software packages specifically designed for training and employing MLPs in molecular dynamics (MD) and free energy simulations, and other applications requiring accurate potential energy and force calculations.^{7–21} These packages are typically built upon specific machine learning frameworks such as TensorFlow,²² PyTorch,²³ or JAX.²⁴ While a single machine learning framework often meets most requirements, increasingly complex applications necessitate interoperability between packages utilizing different machine learning frameworks. This poses a significant challenge for developers when

these packages rely on disparate frameworks. Additionally, each framework may have unique advantages, such as optimized performance for specific applications or better compatibility with certain hardware. To address these limitations, developers of several MLP packages have created secondary versions based on alternative frameworks. However, introducing new frameworks typically create separate, standalone packages rather than refactoring existing ones, which significantly increases the complexity and maintenance costs for both the packages and their integration with molecular dynamics software. For instance, MACE²⁵ offers both PyTorch-based and JAX-based implementations, but they exist as separate packages, and the integration of LAMMPS with the MACE JAX implementation has not been realized.

The earlier version of the DeePMD-kit package^{7,26} was built on TensorFlow, which posed challenges for integration with other packages utilizing different deep-learning frameworks. To address this limitation, Gao *et al.* developed a JAX-based package²⁷ to enable seamless integration of deep potential models with JAX-MD²⁸ for an end-to-end GPU-accelerated workflow. Similarly, Zhang *et al.* introduced a PyTorch-based package alongside the development of the DPA-2 model²⁹ for better distributed training performance. However, maintaining consistent user interfaces (UIs) and application programming interfaces (APIs) for tasks such as training, inference, and molecular dynamics simulations across these different frameworks proved to be both inconvenient and inefficient.

In this work, we introduce a major new release of the DeePMD-kit package (v3), featuring a multi-backend framework that integrates the existing TensorFlow²² backend with new backends, including PyTorch,²³ JAX,²⁴ and PaddlePaddle.³⁰ These backends are designed to be interchangeable, allowing users and developers to switch between them with minimal modifications. The new framework reuses the well-established interfaces developed in previous versions (see Fig. 1), maintaining seamless integration with other software, such as LAMMPS,³¹ i-PI,³² AMBER,^{33–35} CP2K,³⁶ OpenMM,^{37,38} GROMACS,³⁹ ASE,⁴⁰ and ABACUS.⁴¹ The inclusion of these new backends enables DeePMD-kit to seamlessly interact with other packages that utilize PyTorch, JAX, or PaddlePaddle, thereby expanding its

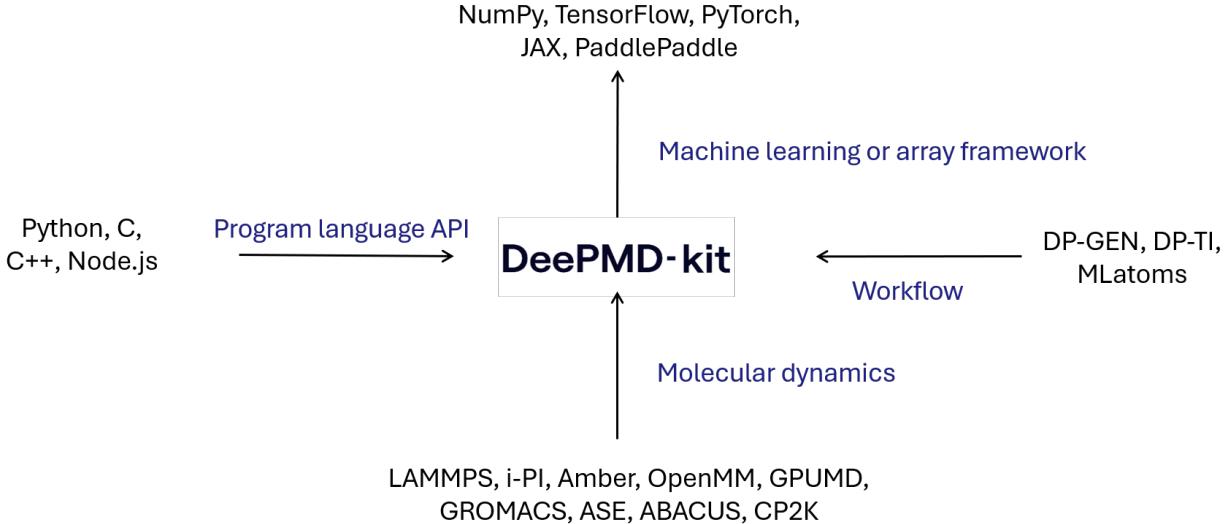


Figure 1: The DeePMD-kit software ecosystem. The arrows indicate dependency flow. Software packages shown in the figure include: (1) DeePMD-kit; (2) Machine learning and array frameworks: NumPy,⁴² TensorFlow,²² PyTorch,²³ JAX,²⁴ and PaddlePaddle;³⁰ (3) Molecular dynamics packages: LAMMPS,³¹ i-PI,³² Amber,³³ OpenMM,³⁷ CP2K,³⁶ GROMACS,³⁹ ASE,⁴⁰ ABACUS,⁴¹ and GPU-MD;⁴³ (4) Workflow packages: DP-GEN,⁴⁴ MLatom,⁴⁵ and DP-TI;⁴⁶ (5) Program language API: Python, C, C++, and Node.js.

compatibility and functionality. Additionally, this update enables DeePMD-kit to leverage unique features and optimizations available in different machine learning frameworks, further enhancing its flexibility and performance.

2 Software Description

2.1 A Multiple-backend Framework

The multiple-backend framework added in DeePMD-kit v3 is aimed to support multiple deep-learning frameworks in a pluggable way while providing a unified interface for users and developers. When introducing the multi-backend framework, no breaking changes were made to the existing Python and C/C++ APIs in DeePMD-kit v2. As a result, the existing interfaces implemented in various molecular dynamics packages can be reused without requiring any modifications.

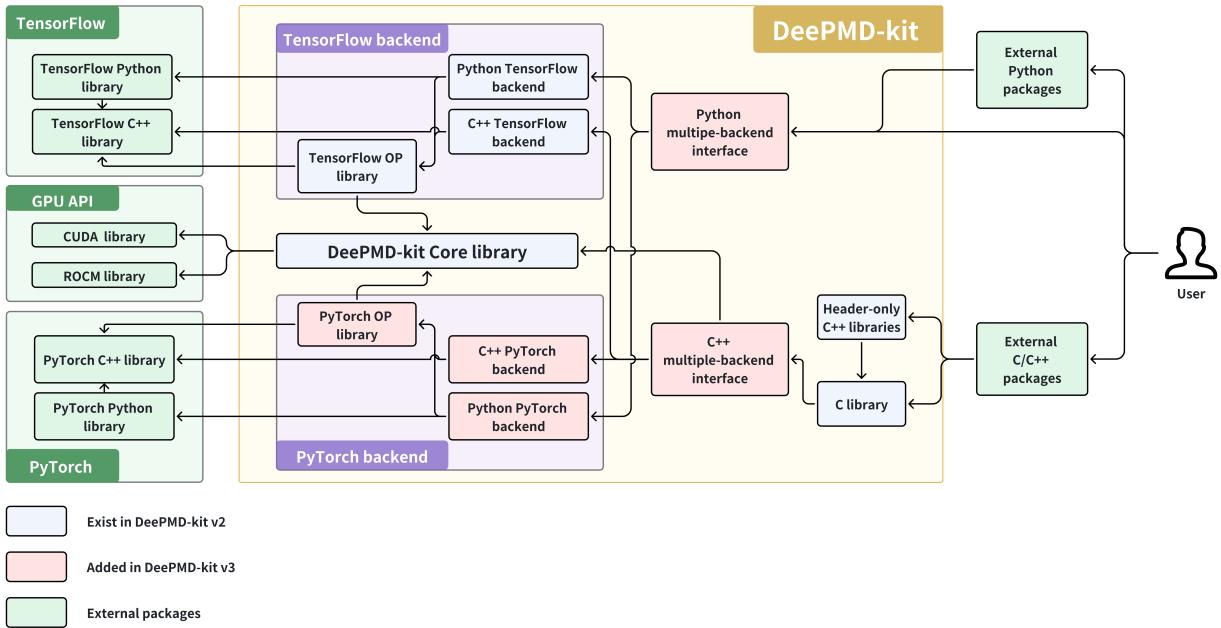


Figure 2: The infrastructure of the TensorFlow and PyTorch backends in the DeePMD-kit software. The arrows indicate dependency flow. Modules in blue existed in DeePMD-kit v2,²⁶ modules in red are newly added in DeePMD-kit v3, and modules in green are machine learning frameworks.

The infrastructure of the TensorFlow and PyTorch backends is illustrated in Fig. 2; other backends are organized in a similar manner. For both training and inference tasks driven by different models, users interact with a unified set of interfaces, regardless of the backend being used. The backend-specific implementation is then invoked to perform the actual computations, leveraging external machine learning frameworks that are highly optimized for performance. In some cases, these backends also share common implementations through customized operators. For example, the model compression functionality⁴⁷ is implemented in the DeePMD-kit core library, utilizing DeePMD-kit CUDA or ROCm libraries for efficient computation.

Users are expected to interact with each backend in a uniform manner, without needing to understand the backend-specific details. During model training and saving via the Python interface, users can simply specify the desired backend. All other training parameters are designed to be backend-agnostic, ensuring a consistent user experience. Once training is

complete, the model is saved in a backend-specific format. When the saved model is loaded for inference through the Python or C++ API during molecular dynamics simulations, the backend is automatically detected based on the model filename, and the appropriate backend module is used. Furthermore, models created with one backend can be easily converted to another backend, providing flexibility and interoperability for various workflows. An example user script is shown as follows:

```
# Use the PyTorch backend to train and freeze the model
dp --pt train input.json
dp --pt freeze -o model.pth
# Test the model
dp test -m model.pth -s dataset
# Convert the model to that in the JAX backend
dp convert-backend model.pth model.savedmodel
```

Developers can incorporate new backends in a modular and pluggable manner, focusing on backend-specific implementations without needing to modify the existing Python or C/C++ APIs or interfaces with external packages. DeePMD-kit provides an implementation in the Array API⁴⁸ (see Section 2.2), simplifying the process of adding new backends based on machine learning frameworks that support the Array API. For developers looking to add new models, the Array API can also be used to implement these models for backends that support it. For backends that do not support the Array API, model implementation can be facilitated with the assistance of large language models (LLMs), as shown in the Supporting Information, streamlining the development process further.

A challenge in the multiple-backend framework is to ensure that the same model driven by different backends produces the same results. To address this issue, we have developed a set of tests that compare the results of the same model driven by different backends. The models can be serialized and deserialized in the Python interface, which is the reason why they can be converted to each other.

2.2 Backends

Five backends are supported in the current version of DeePMD-kit: DP, TensorFlow,²² PyTorch,²³ JAX,²⁴ and PaddlePaddle.³⁰

DP. The DP backend serves as a reference implementation, designed to provide a correct and standardized foundation for model development. It is built using the Array API,⁴⁸ allowing its functionality to be leveraged by other backends (such as JAX) without requiring code duplication. By default, the DP backend uses NumPy,⁴² which does not support gradient computations or GPU acceleration, for computations to minimize dependencies.

TensorFlow. The TensorFlow backend is the original backend of DeePMD-kit. It utilizes the TensorFlow v1 API,²² which employs static computational graphs to optimize performance. These static graphs can be saved into model files and later restored for inference in both Python and C++ interfaces. The customized TensorFlow C++ operators are mainly used to calculate coordinate matrix, force, virial, and embedding network and matrix in compressed models.

PyTorch. The PyTorch backend leverages dynamic computational graphs to provide greater flexibility.²³ It employs TorchScript for model serialization, enabling models to be saved and loaded in both Python and C++ interfaces. The customized PyTorch C++ operators are used for model compression and communication between processors or GPU cards in the graph neural networks of the DPA-2 model.²⁹ PyTorch is friendly to be developed with and thus has a larger user base. PyTorch is widely used in various atomistic packages, including machine learning potential packages,^{12,25,49} force field packages,⁵⁰ semiempirical quantum chemistry packages,⁵¹ and molecular dynamics packages.⁵² The development of the PyTorch backend has made it easier to integrate DeePMD-kit with these existing packages.

JAX. The JAX backend is built on top of the DP backend and the Array API,⁴⁸ using JAX²⁴ as its array library. JAX performs just-in-time (JIT) compilation for enhanced performance. Models created with the JAX backend are saved in a TensorFlow-compatible format using the “jax2tf” converter, allowing them to be loaded via the TensorFlow C++ library.²² JAX has gained popularity in the field of atomistic simulations, supporting machine learning potential packages,⁵⁰ machine learning density functional packages,⁵³ force field packages,^{54,55} and molecular dynamics packages.²⁸

A challenge in the JAX backend is that the JAX JIT compiler requires the input tensors to have static shapes, a constraint not imposed by TensorFlow or PyTorch. However, during molecular dynamics simulations, parameters such as the maximum number of neighbor atoms, the simulation box size, and the number of ghost atoms that do not directly contribute atomic energies are dynamic and not known in advance. To address this limitation, the compiled JIT model is wrapped within a standard TensorFlow model. During each simulation step, the TensorFlow model calculates or post-processes the neighbor list and other tensors with dynamic shapes. The neighbor list is then passed into the compiled JIT model, combining the strengths of both TensorFlow (for handling dynamic inputs) and JAX (for optimized performance through JIT compilation). To handle the dynamic number of ghost atoms, a larger fixed value is initially set, which is adjusted if it becomes insufficient during the simulation.

The choice of integrating JAX with TensorFlow rather than PyTorch is primarily motivated by their shared use of the Accelerated Linear Algebra (XLA) compiler. Initially, JAX was developed based on XLA, which was originally a module within TensorFlow before it became an independent library. Both TensorFlow and JAX inherently utilize XLA for compilation, enabling a straightforward conversion between them via the `jax2tf` function, which compiles JAX functions using XLA and wraps them within a TensorFlow model. In contrast, PyTorch does not include XLA integration by default, making TensorFlow the natural and more efficient choice for creating this hybrid model.

PaddlePaddle. The PaddlePaddle backend, which is based on PaddlePaddle,³⁰ features a similar Python interface to the PyTorch backend, ensuring compatibility and flexibility in model development. PaddlePaddle has introduced dynamic-to-static functionality and PaddlePaddle JIT compiler (CINN) in DeePMD-kit, which allow for dynamic shapes and higher-order differentiation. The dynamic-to-static functionality automatically captures the user’s dynamic graph code and converts it into a static graph. After conversion, the CINN compiler is used to optimize the computational graph, thereby enhancing the efficiency of model training and inference. In our experiments with the DPA-2 model, we achieved approximately a 40% reduction in training time compared to the dynamic graph, effectively improving the model training efficiency. While PaddlePaddle has been utilized in other scientific fields,⁵⁶ DeePMD-kit is the first software package to implement machine learning potentials using PaddlePaddle.

2.3 New Design Principles in version 3

When new backends, including DP, PyTorch, JAX, and PaddlePaddle, were introduced, the design of the DeePMD-kit was thoroughly reconsidered. The following design principles were adopted, differing significantly from those used in the previous TensorFlow backend of DeePMD-kit version 2:

Metaprogramming. To reduce code duplication and enhance maintainability, metaprogramming techniques are employed to generate backend-specific classes from those in the reference DP backend. For instance, in different backends, a neural network is constructed from multiple neural network layers. While the structure of these layers is consistent across backends, the implementation must interact with different machine learning frameworks depending on the backend. By transforming a base neural network class into a backend-specific neural network class using a backend-specific layer class, the implementation of neural networks becomes backend-agnostic. This approach ensures a consistent structure

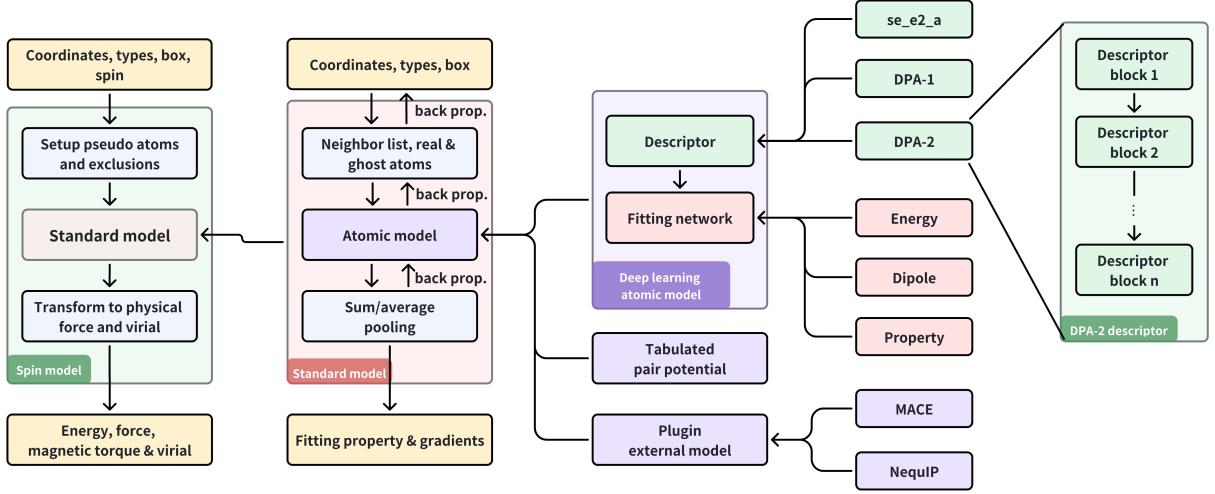


Figure 3: Schematic plot of the model components in DeePMD-kit version 3.

while allowing flexibility for backend-specific details, significantly streamlining development and maintenance.

Atomic model. In version 3, we introduce an innovative component termed the “atomic model” (see Fig. 3). This model is based on the assumption that the physical quantity to be learned, represented by \mathbf{y} , can be decomposed into atomic contributions, expressed as:

$$\mathbf{y} = \sum_i \mathbf{y}_i \quad (1)$$

This approach simplifies the implementation process for model developers, who only need to define and implement the atomic contributions \mathbf{y}_i . The DeePMD-kit subsequently performs the summation of these atomic contributions and computes the generalized force \mathbf{F}_i and virial $\Xi_{\alpha\beta}$, which are defined as the partial derivatives of the physical quantity with respect to atomic coordinates and cell vectors, respectively:

$$\mathbf{F}_i = -\frac{\partial \mathbf{y}}{\partial \mathbf{r}_i}, \quad \Xi_{\alpha\beta} = -\sum_{\gamma} \frac{\partial \mathbf{y}}{\partial h_{\gamma\alpha}} h_{\gamma\beta}. \quad (2)$$

In Equation (2), \mathbf{r}_i represents the Cartesian coordinates of atom i , while $h_{\alpha\beta}$ denotes the β component of the α -th cell vector. It is important to note that the atomic model concept is introduced to facilitate the implementation of physical quantities that can be decomposed into atomic contributions. For cases where atomic decomposition is not suitable, DeePMD-kit supports the implementation of a standard model without relying on atomic decomposition.

Computation of the neighbor list and the coordinate matrix. In the TensorFlow backend, the neighbor list, the coordinate matrix, and the gradient of the coordinate matrix with respect to the coordinates were computed using custom TensorFlow operators, as these operations are not standard in TensorFlow.⁷ This design posed challenges for calculating and training the Hessian, the second gradient of the output with respect to atomic coordinates, $\frac{\partial \mathbf{y}}{\partial \mathbf{x}_i \partial \mathbf{x}_j}$. In the new backends, these computations are performed using standard operators provided by the respective machine learning frameworks. This enables straightforward inference and training of the Hessian using the automatic differentiation capabilities inherent to these frameworks. The neighbor list can be configured to exclude specific pairs of atomic types, thereby implementing range corrections for MD simulations using MLP/MM or QM/MM+ Δ MLP.⁵⁷

Implementation of communications for graph neural network models. The new backends now support graph neural network (GNN) models, such as the DPA-2 model²⁹ and the in-development DPA-3 model. A key challenge in implementing GNN models is managing communication between nodes or GPU cards during molecular dynamics simulations. In a GNN model, atomic features are updated based on atom and edge features within a cutoff radius, which requires access to the neighbor list of surrounding atoms. However, some neighbor atoms may be ghost atoms that do not directly contribute atomic energies, and their neighbor lists are typically unavailable in most molecular dynamics packages.³¹ Some GNN-based MLP packages address this issue by generating ghost atoms within an extended cutoff radius ($r_c \times L$, where r_c is the cutoff radius for each layer and L is the number of layers)

and rebuilding the neighbor lists for the ghost atoms. This approach, however, significantly increases computational cost. Other packages have suggested running simulations on a single GPU card to avoid this issue. To overcome these limitations, the new backends implement a customized C++ operator that uses the message-passing interface (MPI)⁵⁸ to exchange atom and edge features between processors or GPU cards. This solution enables efficient simulations across multiple nodes or GPU cards in high-performance computing environments.

Descriptor block. The complexity of descriptors has increased significantly with the advancement of deep potential models. The DPA-2 model,²⁹ for instance, primarily consists of a representation-initializer block, representation-transformer layers, and a three-body embedding block. The representation-initializer block is derived from the DPA-1 descriptor.⁵⁹ To address the growing complexity and support future development, we introduce a novel component called the “descriptor block” (see Fig. 3), which serves as a modular foundation for constructing complex descriptors. The descriptor block is designed to be flexible and extensible, allowing seamless integration of new descriptors into the framework.

Refactor of the implementation of the DeepSPIN model. In the TensorFlow backend, the DeepSPIN model⁶⁰ was initially integrated into the descriptors and fitting networks in an ad hoc manner, which hindered maintainability and limited the adoption of new features such as the DPA-2 descriptor. To address these limitations, the DeepSPIN model has been refactored in the new backends into a modular component that operates on top of any standard potential energy model.

The DeepSPIN model aims to represent the potential energy of a system as a function of atomic coordinates, spin, and the simulation cell. To achieve this, the model introduces pseudo atoms, which are associated with physical atoms through the relation:

$$\mathbf{r}_{ip} = \mathbf{r}_i + \lambda_{\eta_i} \mathbf{S}_i, \quad i \in \mathcal{A}, \quad (3)$$

where \mathcal{A} denotes the set of physical atoms, \mathbf{S}_i is the spin of atom i , λ_{η_i} is a user-defined, element-specific scaling factor (with η_i representing the element type of atom i), and \mathbf{r}_i is the Cartesian coordinate of physical atom i . Notably, pseudo atoms do not represent physical entities but serve as computational tools to encode the spin states of their associated physical atoms. To ensure type consistency, pseudo atoms inherit mapped element types from their host atoms (e.g., `Fe_spin` from `Fe`), denoted by $\eta_{i^p}^{\text{spin}}$.

The model processes an augmented atomic input, which includes both the augmented coordinates $\{\mathbf{r}_i, \mathbf{r}_{i^p} \mid i \in \mathcal{A}\}$ and the augmented types $\{\eta_i, \eta_{i^p}^{\text{spin}} \mid i \in \mathcal{A}\}$. These inputs are fed into a standard potential energy model to generate augmented outputs, including the augmented atomic energies $\{\mathcal{E}_i, \mathcal{E}_{i^p} \mid i \in \mathcal{A}\}$ and the augmented atomic forces $\{\hat{\mathbf{F}}_i, \hat{\mathbf{F}}_{i^p} \mid i \in \mathcal{A}\}$. Post-processing steps are then applied to convert the augmented predictions into physical quantities. The total energy is computed as the sum of energy contributions from all physical atoms:

$$E = \sum_{i \in \mathcal{A}} E_i. \quad (4)$$

The atomic forces \mathbf{F}_i and magnetic torques $\boldsymbol{\omega}_i$ for atom i are derived, respectively, as:

$$\mathbf{F}_i = -\frac{\partial E}{\partial \mathbf{r}_i} = \hat{\mathbf{F}}_i + \hat{\mathbf{F}}_{i^p}, \quad \boldsymbol{\omega}_i = -\frac{\partial E}{\partial \mathbf{S}_i} = \lambda_{\eta_i} \hat{\mathbf{F}}_{i^p}. \quad (5)$$

The system virial $\mathbf{\Xi}$ is calculated by summing the outer products of the forces and positions of all physical atoms:⁶¹

$$\mathbf{\Xi} = \sum_{i \in \mathcal{A}} \mathbf{F}_i \otimes \mathbf{r}_i. \quad (6)$$

This refactored design, as illustrated by Fig. 3, enhances code clarity and maintainability while maintaining compatibility with standard potential energy model architectures.

Removal of the local frame descriptor. In the previous publication,²⁶ we have indicated that the local frame descriptor⁶² is not continuous at the cutoff radius and the exchanging of the order of two nearest neighbors. Thus, the local frame descriptor is removed in the new backends.

3 Selected Extensions

DeePMD-kit v3 introduces an extensive plugin system designed to seamlessly integrate external models into its various backends. In this section, we demonstrate the versatility of the multi-backend framework by showcasing the integration of other MLPs packages and of Differentiable Molecular Force Field.

3.1 DeePMD-GNN

Most machine learning potential (MLP) packages, including the previous version of DeePMD-kit, are typically restricted to a single type of neural network potential developed by the same team of developers. This limitation imposes additional burdens on users, who must learn new packages, and on developers, who must maintain multiple packages, while also makes it impossible to benchmark models using the same loss function, learning rate, and data sets. In the earlier version of DeePMD-kit, integrating other MLP packages was challenging because most were not based on TensorFlow. This limitation was addressed with the development of the multi-backend framework and the PyTorch backend. Consequently, Zeng *et al.* introduced the DeePMD-GNN package,⁶³ which integrates the MACE²⁵ and NequIP⁴⁹ models into the PyTorch backend. This advancement marked the first time DeePMD-kit supported highly mature external potential models, significantly enhancing its versatility. It is worth noting that the plugin solely extends the trainable model module (see Fig. 3) and reuses the existing loss function and learning rate modules in the DeePMD-kit, allowing different models to be compared easily under the same conditions.

3.2 DMFF plugin

Long-range interactions are critical in many systems but are challenging to incorporate into an MLP model.⁶⁴ Zhu *et al.*⁶⁵ developed a PyTorch version of the Differentiable Molecular Force Field (DMFF) package⁵⁴ that implements Ewald summation⁶⁶ and charge equilibration (QE_q)⁶⁷ methods for handling long-range interactions. These interactions have been integrated into the PyTorch backend of DeePMD-kit, enabling a hybrid approach where long-range interactions are described by traditional force fields, while short- and mid-range interactions are captured by MLPs.

4 Benchmark

A key application of the multi-backend framework is its ability to run molecular dynamics simulations using the most efficient backend. To illustrate this, we present benchmark calculations using the following models: DPA-1 without attention layers ($L=0$) and its compressed version,⁴⁷ DPA-1 with two attention layers ($L=2$),⁵⁹ and DPA-2 (medium).²⁹ These benchmarks were performed in both single- and double-precision modes, employing the TensorFlow, PyTorch, and JAX backends on a single 80 GB NVIDIA H100 GPU card, 40 GB NVIDIA A800 GPU card, and 24 GB NVIDIA 4090 GPU card. The simulations involved a water system of varying atom counts, with each calculation repeated 500 times to obtain an average computational speed. LAMMPS³¹ interfaced with DeePMD-kit v3 is used to perform simulations. Here we note the limitations in the current package when this article was written: only the DPA-1 ($L=0$) model supports model compression, the TensorFlow backend does not support the DPA-2 model, the JAX backend does not support model compression, and the PaddlePaddle backend is still under development.

All of the models use a fitting network that consists of 3 hidden layers with 240, 240, and 240 neurons. The coordinate matrix is encoded from the local environment within a 6 Å cutoff radius and 1 Å of smoothing; The details of DPA-1 models are as below: the

embedding network consists of 3 hidden layers with 25, 50, and 100 neurons; the embedding submatrix size is 12. Specifically, in the DPA-1 ($L=2$) model, the length of hidden vectors in each attention layer is 128. The details of the DPA-2 (medium) model are as below: the representation initializer layer is encoded based on the local environment within a 6 Å cutoff radius and smoothed over 1 Å; the model uses 6 representation transformer (Reperformer) layers, with each layer calculated using a 4 Å cutoff and 1 Å of smoothing; three-body embedding is incorporated within a 4 Å cutoff; the embedding network consists of three hidden layers with 25, 50, and 100 neurons, and the embedding submatrix has a size of 12; the fitting network also has three hidden layers, each containing 240 neurons; the dimensions for the invariant single-atom and pair-atom representations are set to 120 and 32, respectively. Additionally, the localized single-atom representation update mechanism excludes the self-attention layer, while the Reperformer pair-atom representation is updated using a gated self-attention layer.

The performance on the H100, A800, and 4090 cards is summarized in Table S1, S2, and S3, respectively, and illustrated in Fig. 4, S1, and S2, respectively. Across TensorFlow, PyTorch, and JAX, the three GPU cards exhibit similar relative performance. The results indicate that for the DPA-1 model without attention layers, the compressed TensorFlow model is the fastest. Although model compression improves performance by a similar amount on both the TensorFlow and PyTorch backends, other components of the model cause the final performance to differ. For more computationally demanding models, such as the DPA-1 model with attention layers and the DPA-2 model, the JAX backend generally achieves the highest performance. However, two exceptions were observed: in the DPA-2 FP32 case with 12,288 atoms on the H100 card, where PyTorch outperformed JAX; in the DPA-1 ($L=2$) FP64 case on the 4090 card, where TensorFlow outperformed JAX. Additionally, different backends exhibit varying GPU memory usage, which can influence the choice of backend depending on the computational resources available. These findings demonstrate that no single backend consistently outperforms others across all models. Therefore, the

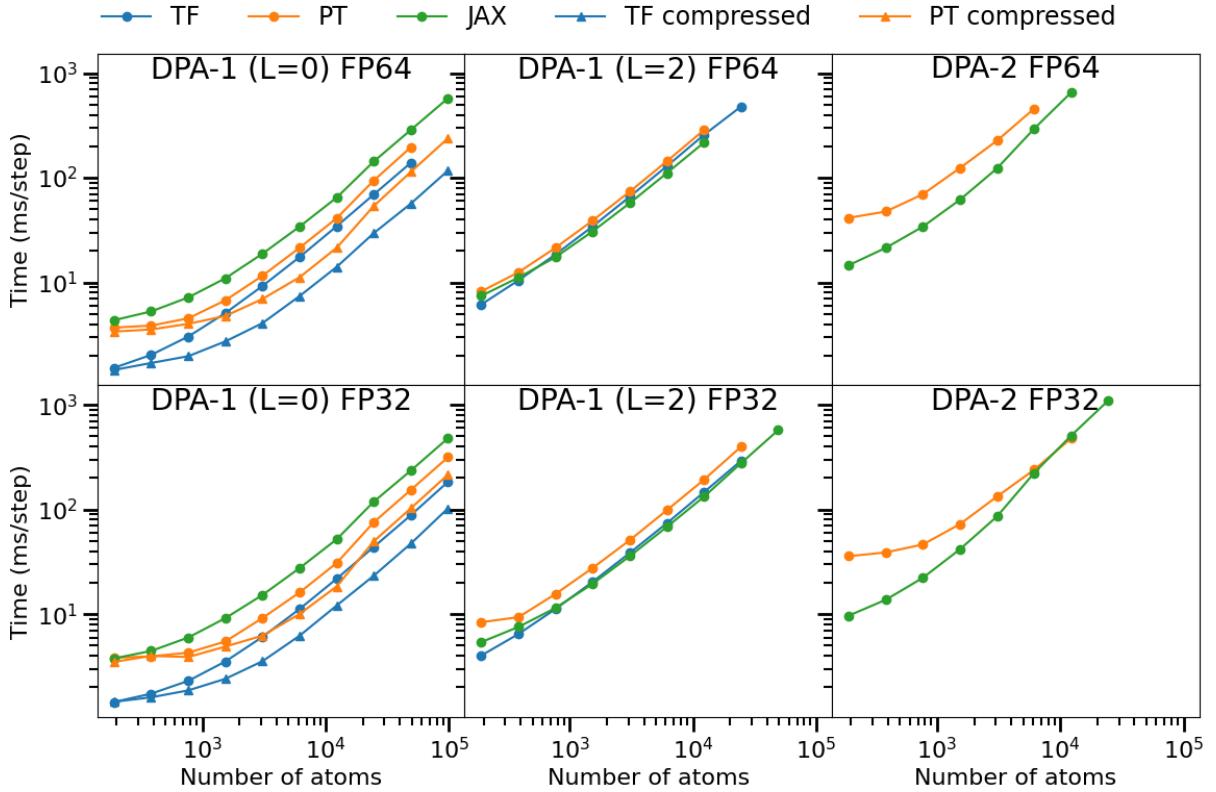


Figure 4: The molecular dynamics performance (ms/step) of DPA-1 ($L=0$), DPA-1 ($L=2$), and DPA-2 models in double (FP64) and single (FP32) precisions in TensorFlow (TF), PyTorch (PT), and JAX backends on a single 80 GB NVIDIA H100 GPU card. The water system with different numbers of atoms is used for simulations.

multi-backend framework is highly valuable, enabling users to identify and utilize the most efficient backend for their specific needs.

5 Conclusions

In this work, we introduced DeePMD-kit v3, a significant enhancement to the widely used machine learning potential package. This new version addresses critical limitations of its predecessor by incorporating a multi-backend framework supporting TensorFlow, PyTorch, JAX, and PaddlePaddle. By enabling seamless backend switching with minimal changes, DeePMD-kit v3 provides a flexible and interoperable platform for researchers and developers. This advancement not only facilitates the integration of DeePMD-kit with diverse software

ecosystems but also empowers the development of more complex and efficient workflows for molecular dynamics simulations and other applications requiring machine learning potentials. The best performance for molecular dynamics can be archeived by selecting the most efficient backend for a specific model and computational resource. We anticipate that DeePMD-kit v3 will significantly expand the accessibility and applicability of MLPs in physics, chemistry, and materials science, fostering innovation and collaboration in these fields.

Data Availability

Source code for the project can be found at <https://github.com/deepmodeling/deepmd-kit>. Source code for the benchmark can be found at <https://github.com/njzjz/benchmark-dpv3>.

Acknowledgement

The authors thank Dr. Rocco Meli for his code contributions. ChatGPT 4o was used to edit the English (prompt: polish the following paragraph in a research article). J. Z. acknowledges 2023 Chinese Government Award for Outstanding Self-financed Students Aboard presented by the China Scholarship Council. S. H. acknowledges help from Yanjun Ma, Tiezhu Gao, Xiaoguang Hu. J.-X. Z. gratefully acknowledges Xiamen University and iChEM for a Ph.D. studentship. A. K. S. acknowledges help from Dr. Manish Modani (Principal Solution Architect at NVIDIA), Dr. Unnikrishnan C (Assistant Professor at IIT, Palakkad), and C-DAC Centre, Patna, India. The work of Q. Z. and J. D. was supported by the the fundation the Science and Technology Innovation Program of Hunan Province under Grant No. 2021RC4026. The work of S. L. B. was supported by the Research Council of Norway through the Centre of Excellence Hylleraas Centre for Quantum Molecular Sciences (Grant 262695) and the Young Researcher Talent (Grant 344993), as well as by the EuroHPC Joint Undertaking (Grant EHPC-REG-2023R02-088). The work of W. J. was supported by the Natural Science foundation of China (No. 92270206). The work of T. J. G. and D. M. Y. was

supported by the National Institutes of Health (No. GM107485 to D.M.Y.) and the National Science Foundation (CSSI Frameworks Grant No. 2209718 to D.M.Y.). The work of H. W. was supported by the National Key R&D Program of China (Grant No. 2022YFA1004300) and the National Natural Science Foundation of China (Grant No. 12122103). Computational resources are provided by: ByteDance Volcano Engine Cloud; Alibaba Cloud PAI; IKKEM Intelligent Computing Center, for the work of Y.-P. L.

Supporting Information Available

The molecular dynamics performance (ms/step) of DPA-1 ($L=0$), DPA-1 ($L=2$), and DPA-2 models in double (FP64) and single (FP32) precisions in TensorFlow (TF), PyTorch (PT), and JAX backends on a single 40 GB NVIDIA H100, A800, and 4090 GPU card. Example of Using Large Language Models for Code Conversion between Backends.

References

- (1) Behler, J. Perspective: Machine learning potentials for atomistic simulations. *J. Chem. Phys.* **145**, 170901.
- (2) Butler, K. T.; Davies, D. W.; Cartwright, H.; Isayev, O.; Walsh, A. Machine learning for molecular and materials science. *Nature* **559**, 547–555.
- (3) Noé, F.; Tkatchenko, A.; Müller, K.-R.; Clementi, C. Machine Learning for Molecular Simulation. *Annu. Rev. Phys. Chem.* **71**, 361–390.
- (4) Pinheiro Jr, M.; Ge, F.; Ferré, N.; Dral, P. O.; Barbatti, M. Choosing the right molecular machine learning potential. *Chem. Sci.* **12**, 14396–14413.
- (5) Manzhos, S.; Carrington Jr, T. Neural Network Potential Energy Surfaces for Small Molecules and Reactions. *Chem. Rev.* **121**, 10187–10217.

(6) Zeng, J.; Cao, L.; Zhu, T. Neural network potentials. In *Quantum Chemistry in the Age of Machine Learning*; Dral, P. O., Ed.; Elsevier, Chapter 12, pp 279–294.

(7) Wang, H.; Zhang, L.; Han, J.; E, W. DeePMD-kit: A deep learning package for many-body potential energy representation and molecular dynamics. *Comput. Phys. Commun.* **228**, 178–184.

(8) Schütt, K. T.; Kessel, P.; Gastegger, M.; Nicoli, K. A.; Tkatchenko, A.; Müller, K.-R. SchNetPack: A Deep Learning Toolbox For Atomistic Systems. *J. Chem. Theory Comput.* **15**, 448–455.

(9) Chmiela, S.; Sauceda, H. E.; Poltavsky, I.; Müller, K.-R.; Tkatchenko, A. sGDML: Constructing accurate and data efficient molecular force fields using machine learning. *Computer Physics Communications* **240**, 38–45.

(10) Unke, O.; Meuwly, M. PhysNet: A Neural Network for Predicting Energies, Forces, Dipole Moments, and Partial Charges. *J. Chem. Theory Comput.* **15**, 3678–3693.

(11) Lee, K.; Yoo, D.; Jeong, W.; Han, S. SIMPLE-NN: An efficient package for training and executing neural- network interatomic potentials. *Computer Physics Communications* **242**, 95–103.

(12) Gao, X.; Ramezanghorbani, F.; Isayev, O.; Smith, J. S.; Roitberg, A. E. TorchANI: A free and open source PyTorch-based deep learning implementation of the ANI neural network potentials. *Journal of chemical information and modeling* **60**, 3408–3415.

(13) Dral, P. O.; Ge, F.; Xue, B.-X.; Hou, Y.-F.; Pinheiro Jr, M.; Huang, J.; Barbatti, M. MLatom 2: An Integrative Platform for Atomistic Machine Learning. *Top. Curr. Chem. (Cham)* **379**, 27.

(14) Singraber, A.; Behler, J.; Dellago, C. Library-Based LAMMPS Implementation of High-Dimensional Neural Network Potentials. *J. Chem. Theory Comput.* **15**, 1827–1840.

(15) Zhang, Y.; Xia, J.; Jiang, B. REANN: A PyTorch-based end-to-end multi-functional deep neural network package for molecular, reactive, and periodic systems. *J. Chem. Phys.* **156**, 114801.

(16) Schütt, K. T.; Hessmann, S. S. P.; Gebauer, N. W. A.; Lederer, J.; Gastegger, M. SchNetPack 2.0: A neural network toolbox for atomistic machine learning. *J. Chem. Phys.* **158**, 144801.

(17) Fan, Z.; Wang, Y.; Ying, P.; Song, K.; Wang, J.; Wang, Y.; Zeng, Z.; Xu, K.; Lindgren, E.; Rahm, J. M.; Gabourie, A. J.; Liu, J.; Dong, H.; Wu, J.; Chen, Y.; Zhong, Z.; Sun, J.; Erhart, P.; Su, Y.; Ala-Nissila, T. GPUMD: A package for constructing accurate machine-learned potentials and performing highly efficient atomistic simulations. *J. Chem. Phys.* **157**, 114801.

(18) Novikov, I. S.; Gubaev, K.; Podryabinkin, E. V.; Shapeev, A. V. The MLIP package: moment tensor potentials with MPI and active learning. *Mach. Learn.: Sci. Technol.* **2**, 025002.

(19) Yanxon, H.; Zagaceta, D.; Tang, B.; Matteson, D. S.; Zhu, Q. PyXtal_FF: a python library for automated force field generation. *Mach., Learn.,: Sci., Technol.* **2021**, *2*, 027001.

(20) Zeng, J.; Tao, Y.; Giese, T. J.; York, D. M. QD π : A Quantum Deep Potential Interaction Model for Drug Discovery. *J. Chem. Theory Comput.* **19**, 1261–1275.

(21) Zeng, J.; Tao, Y.; Giese, T. J.; York, D. M. Modern semiempirical electronic structure methods and machine learning potentials for drug discovery: Conformers, tautomers, and protonation states. *J. Chem. Phys.* **158**, 124110.

(22) Abadi, M.; Agarwal, A.; Barham, P.; Brevdo, E.; Chen, Z.; Citro, C.; Corrado, G. S.; Davis, A.; Dean, J.; Devin, M.; Ghemawat, S.; Goodfellow, I.; Harp, A.; Irving, G.;

Isard, M.; Jia, Y.; Jozefowicz, R.; Kaiser, L.; Kudlur, M.; Levenberg, J.; Mané, D.; Monga, R.; Moore, S.; Murray, D.; Olah, C.; Schuster, M.; Shlens, J.; Steiner, B.; Sutskever, I.; Talwar, K.; Tucker, P.; Vanhoucke, V.; Vasudevan, V.; Viégas, F.; Vinyals, O.; Warden, P.; Wattenberg, M.; Wicke, M.; Yu, Y.; Zheng, X. TensorFlow: Large-Scale Machine Learning on Heterogeneous Systems. <https://www.tensorflow.org/>, Software available from tensorflow.org.

(23) Paszke, A.; Gross, S.; Massa, F.; Lerer, A.; Bradbury, J.; Chanan, G.; Killeen, T.; Lin, Z.; Gimelshein, N.; Antiga, L.; Desmaison, A.; Köpf, A.; Yang, E.; DeVito, Z.; Raison, M.; Tejani, A.; Chilamkurthy, S.; Steiner, B.; Fang, L.; Bai, J.; Chintala, S. PyTorch: An Imperative Style, High-Performance Deep Learning Library. *arXiv* **2019**, 1912.01703.

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

(25) Batatia, I.; Kovács, D. P.; Simm, G. N. C.; Ortner, C.; Csányi, G. MACE: higher order equivariant message passing neural networks for fast and accurate force fields. Proceedings of the 36th International Conference on Neural Information Processing Systems. Red Hook, NY, USA, 2024.

(26) Zeng, J.; Zhang, D.; Lu, D.; Mo, P.; Li, Z.; Chen, Y.; Rynik, M.; Huang, L.; Li, Z.; Shi, S.; Wang, Y.; Ye, H.; Tuo, P.; Yang, J.; Ding, Y.; Li, Y.; Tisi, D.; Zeng, Q.; Bao, H.; Xia, Y.; Huang, J.; Muraoka, K.; Wang, Y.; Chang, J.; Yuan, F.; Bore, S. L.; Cai, C.; Lin, Y.; Wang, B.; Xu, J.; Zhu, J.-X.; Luo, C.; Zhang, Y.; Goodall, R. E. A.; Liang, W.; Singh, A. K.; Yao, S.; Zhang, J.; Wentzcovitch, R.; Han, J.; Liu, J.; Jia, W.; York, D. M.; E, W.; Car, R.; Zhang, L.; Wang, H. DeePMD-kit v2: A software package for deep potential models. *J. Chem. Phys.* **2023**, 159, 054801.

(27) Gao, R.; Li, Y.; Car, R. Enhanced deep potential model for fast and accurate molecular dynamics: application to the hydrated electron. *Phys. Chem. Chem. Phys.* **2024**, *26*, 23080–23088.

(28) Schoenholz, S. S.; Cubuk, E. D. JAX, M.D. A framework for differentiable physics*. *J. Stat. Mech.* **2021**, *2021*, 124016.

(29) Zhang, D.; Liu, X.; Zhang, X.; Zhang, C.; Cai, C.; Bi, H.; Du, Y.; Qin, X.; Peng, A.; Huang, J.; Li, B.; Shan, Y.; Zeng, J.; Zhang, Y.; Liu, S.; Li, Y.; Chang, J.; Wang, X.; Zhou, S.; Liu, J.; Luo, X.; Wang, Z.; Jiang, W.; Wu, J.; Yang, Y.; Yang, J.; Yang, M.; Gong, F.-Q.; Zhang, L.; Shi, M.; Dai, F.-Z.; York, D. M.; Liu, S.; Zhu, T.; Zhong, Z.; Lv, J.; Cheng, J.; Jia, W.; Chen, M.; Ke, G.; E, W.; Zhang, L.; Wang, H. DPA-2: a large atomic model as a multi-task learner. *npj Comput. Mater.* **2024**, *10*, 293.

(30) Ma, Y.; Yu, D.; Wu, T.; Wang, H. PaddlePaddle: An Open-Source Deep Learning Platform from Industrial Practice. *Frontiers of Data and Computing* **2019**, *1*, 105–115.

(31) Thompson, A. P.; Aktulga, H. M.; Berger, R.; Bolintineanu, D. S.; Brown, W. M.; Crozier, P. S.; in 't Veld, P. J.; Kohlmeyer, A.; Moore, S. G.; Nguyen, T. D.; Shan, R.; Stevens, M. J.; Tranchida, J.; Trott, C.; Plimpton, S. J. LAMMPS - a flexible simulation tool for particle-based materials modeling at the atomic, meso, and continuum scales. *Computer Physics Communications* **271**, 108171.

(32) Litman, Y.; Kapil, V.; Feldman, Y. M. Y.; Tisi, D.; Begušić, T.; Fidanyan, K.; Fraux, G.; Higer, J.; Kellner, M.; Li, T. E.; Pós, E. S.; Stocco, E.; Trenins, G.; Hirshberg, B.; Rossi, M.; Ceriotti, M. i-PI 3.0: A flexible and efficient framework for advanced atomistic simulations. *J. Chem. Phys.* **2024**, *161*, 062504.

(33) Case, D. A.; Aktulga, H. M.; Belfon, K.; Cerutti, D. S.; Cisneros, G. A.; Cruzeiro, V. W. D.; Forouzesh, N.; Giese, T. J.; Götz, A. W.; Gohlke, H.; Izadi, S.; Kasavajhala, K.; Kaymak, M. C.; King, E.; Kurtzman, T.; Lee, T.-S.; Li, P.; Liu, J.; Luchko, T.; Luo, R.;

Manathunga, M.; Machado, M. R.; Nguyen, H. M.; O’Hearn, K. A.; Onufriev, A. V.; Pan, F.; Pantano, S.; Qi, R.; Rahnamoun, A.; Risheh, A.; Schott-Verdugo, S.; Shajian, A.; Swails, J.; Wang, J.; Wei, H.; Wu, X.; Wu, Y.; Zhang, S.; Zhao, S.; Zhu, Q.; Cheatham 3rd, T. E.; Roe, D. R.; Roitberg, A.; Simmerling, C.; York, D. M.; Nagan, M. C.; Merz Jr, K. M. AmberTools. *J. Chem. Inf. Model.* **2023**, *63*, 6183–6191.

(34) Tao, Y.; Giese, T. J.; Ekesan, S.; Zeng, J.; Aradi, B.; Hourahine, B.; Aktulga, H. M.; Götz, A. W.; Merz Jr, K. M.; York, D. M. Amber free energy tools: Interoperable software for free energy simulations using generalized quantum mechanical/molecular mechanical and machine learning potentials. *J. Chem. Phys.* **2024**, *160*, 224104.

(35) Giese, T. J.; Zeng, J.; Lerew, L.; McCarthy, E.; Tao, Y.; Ekesan, S.; York, D. M. Software Infrastructure for Next-Generation QM/MM-ΔMLP Force Fields. *J. Phys. Chem., B* **2024**, *128*, 6257–6271.

(36) Kühne, T. D.; Iannuzzi, M.; Del Ben, M.; Rybkin, V. V.; Seewald, P.; Stein, F.; Laino, T.; Khaliullin, R. Z.; Schütt, O.; Schiffmann, F.; Golze, D.; Wilhelm, J.; Chulkov, S.; Bani-Hashemian, M. H.; Weber, V.; Borštník, U.; Taillefumier, M.; Jakobovits, A. S.; Lazzaro, A.; Pabst, H.; Müller, T.; Schade, R.; Guidon, M.; Andermatt, S.; Holmberg, N.; Schenter, G. K.; Hehn, A.; Bussy, A.; Belleflamme, F.; Tabacchi, G.; Glöß, A.; Lass, M.; Bethune, I.; Mundy, C. J.; Plessl, C.; Watkins, M.; VandeVondele, J.; Krack, M.; Hutter, J. CP2K: An electronic structure and molecular dynamics software package - Quickstep: Efficient and accurate electronic structure calculations. *J. Chem. Phys.* **2020**, *152*, 194103.

(37) Eastman, P.; Galvelis, R.; Peláez, R. P.; Abreu, C. R. A.; Farr, S. E.; Gallicchio, E.; Gorenko, A.; Henry, M. M.; Hu, F.; Huang, J.; Krämer, A.; Michel, J.; Mitchell, J. A.; Pande, V. S.; Rodrigues, J. P.; Rodriguez-Guerra, J.; Simmonett, A. C.; Singh, S.; Swails, J.; Turner, P.; Wang, Y.; Zhang, I.; Chodera, J. D.; De Fabritiis, G.; Mark-

land, T. E. OpenMM 8: Molecular Dynamics Simulation with Machine Learning Potentials. *J. Phys. Chem., B* **2024**, *128*, 109–116.

(38) Ding, Y.; Huang, J. Implementation and Validation of an OpenMM Plugin for the Deep Potential Representation of Potential Energy. *Int. J. Mol. Sci.* **2024**, *25*, 1448.

(39) Abraham, M. J.; Murtola, T.; Schulz, R.; Páll, S.; Smith, J. C.; Hess, B.; Lindahl, E. GROMACS: High performance molecular simulations through multi-level parallelism from laptops to supercomputers. *SoftwareX* **19–25**.

(40) Hjorth Larsen, A.; Jørgen Mortensen, J.; Blomqvist, J.; Castelli, I. E.; Christensen, R.; Dułak, M.; Friis, J.; Groves, M. N.; Hammer, B.; Hargus, C.; Hermes, E. D.; Jennings, P. C.; Bjerre Jensen, P.; Kermode, J.; Kitchin, J. R.; Leonhard Kolsbjerg, E.; Kubal, J.; Kaasbjerg, K.; Lysgaard, S.; Bergmann Maronsson, J.; Maxson, T.; Olsen, T.; Pastewka, L.; Peterson, A.; Rostgaard, C.; Schiøtz, J.; Schütt, O.; Strange, M.; Thygesen, K. S.; Vegge, T.; Vilhelmsen, L.; Walter, M.; Zeng, Z.; Jacobsen, K. W. The atomic simulation environment—a Python library for working with atoms. *J. Phys. Condens. Matter* **29**, 273002.

(41) Li, P.; Liu, X.; Chen, M.; Lin, P.; Ren, X.; Lin, L.; Yang, C.; He, L. Large-scale ab initio simulations based on systematically improvable atomic basis. *Computational Materials Science* **112**, 503–517.

(42) Harris, C. R.; Millman, K. J.; van der Walt, S. J.; Gommers, R.; Virtanen, P.; Cournapeau, D.; Wieser, E.; Taylor, J.; Berg, S.; Smith, N. J.; Kern, R.; Picus, M.; Hoyer, S.; van Kerkwijk, M. H.; Brett, M.; Haldane, A.; Del Río, J. F.; Wiebe, M.; Peterson, P.; Gérard-Marchant, P.; Sheppard, K.; Reddy, T.; Weckesser, W.; Abbasi, H.; Gohlke, C.; Oliphant, T. E. Array programming with NumPy. *Nature* **2020**, *585*, 357–362.

(43) Fan, Z.; Chen, W.; Vierimaa, V.; Harju, A. Efficient molecular dynamics simulations

with many-body potentials on graphics processing units. *Comput. Phys. Commun.* **2017**, *218*, 10–16.

(44) Zhang, Y.; Wang, H.; Chen, W.; Zeng, J.; Zhang, L.; Han, W.; E, W. DP-GEN: A concurrent learning platform for the generation of reliable deep learning based potential energy models. *Comput. Phys. Commun.* **253**, 107206.

(45) Dral, P. O.; Ge, F.; Hou, Y.-F.; Zheng, P.; Chen, Y.; Barbatti, M.; Isayev, O.; Wang, C.; Xue, B.-X.; Pinheiro Jr, M.; Su, Y.; Dai, Y.; Chen, Y.; Zhang, L.; Zhang, S.; Ullah, A.; Zhang, Q.; Ou, Y. MLatom 3: A Platform for Machine Learning-Enhanced Computational Chemistry Simulations and Workflows. *J. Chem. Theory Comput.* **2024**, *20*, 1193–1213.

(46) Fengbo Yuan, J. Z. L. Z., Yifan Li; Wang, H. DPTI: A Python Package to Automate Thermodynamic Integration for Free Energy Calculations. Manuscript in preparation.

(47) Lu, D.; Jiang, W.; Chen, Y.; Zhang, L.; Jia, W.; Wang, H.; Chen, M. DP Compress: A Model Compression Scheme for Generating Efficient Deep Potential Models. *J. Chem. Theory Comput.* **18**, 5559–5567.

(48) Meurer, A.; Reines, A.; Gommers, R.; Fang, Y.-L.; Kirkham, J.; Barber, M.; Hoyer, S.; Müller, A.; Zha, S.; Shanabrook, S.; Gacha, S.; Lezcano-Casado, M.; Fan, T.; Reddy, T.; Passos, A.; Kwon, H.; Oliphant, T.; Standards, C. Python Array API Standard: Toward Array Interoperability in the Scientific Python Ecosystem. 2023; pp 8–17.

(49) 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. *Nat. Commun.* **13**, 2453.

(50) Orlando, G.; Serrano, L.; Schymkowitz, J.; Rousseau, F. Integrating physics in deep learning algorithms: a force field as a PyTorch module. *Bioinform. (Oxford Engl.)* **2024**, *40*, btae160.

(51) Friede, M.; Hölzer, C.; Ehlert, S.; Grimme, S. dxtb-An efficient and fully differentiable framework for extended tight-binding. *J. Chem. Phys.* **2024**, *161*, 062501.

(52) Doerr, S.; Majewski, M.; Pérez, A.; Krämer, A.; Clementi, C.; Noe, F.; Giorgino, T.; De Fabritiis, G. TorchMD: A Deep Learning Framework for Molecular Simulations. *J. Chem. Theory Comput.* **2021**, *17*, 2355–2363.

(53) Chen, Y.; Zhang, L.; Wang, H.; E, W. DeePKS-kit: A package for developing machine learning-based chemically accurate energy and density functional models. *Computer Physics Communications* **282**, 108520.

(54) Wang, X.; Li, J.; Yang, L.; Chen, F.; Wang, Y.; Chang, J.; Chen, J.; Feng, W.; Zhang, L.; Yu, K. DMFF: An Open-Source Automatic Differentiable Platform for Molecular Force Field Development and Molecular Dynamics Simulation. *J. Chem. Theory Comput.* **2023**, *19*, 5897–5909.

(55) Kaymak, M. C.; Rahnamoun, A.; O’Hearn, K. A.; van Duin, A. C. T.; Merz Jr, K. M.; Aktulga, H. M. JAX-ReaxFF: A Gradient-Based Framework for Fast Optimization of Reactive Force Fields. *J. Chem. Theory Comput.* **2022**, *18*, 5181–5194.

(56) Xue, Y.; Liu, Z.; Fang, X.; Wang, F. Multimodal Pre-Training Model for Sequence-based Prediction of Protein-Protein Interaction. Proceedings of the 16th Machine Learning in Computational Biology meeting. 2022; pp 34–46.

(57) Zeng, J.; Giese, T. J.; Ekesan, c.; York, D. M. Development of Range-Corrected Deep Learning Potentials for Fast, Accurate Quantum Mechanical/molecular Mechanical Simulations of Chemical Reactions in Solution. *Journal of Chemical Theory and Computation* **17**, 6993–7009.

(58) Message Passing Interface Forum, MPI: A Message-Passing Interface Standard Version 4.1. 2023.

(59) Zhang, D.; Bi, H.; Dai, F.-Z.; Jiang, W.; Liu, X.; Zhang, L.; Wang, H. Pretraining of attention-based deep learning potential model for molecular simulation. *npj Comput. Mater.* **2024**, *10*, 94.

(60) Yang, T.; Cai, Z.; Huang, Z.; Tang, W.; Shi, R.; Godfrey, A.; Liu, H.; Lin, Y.; Nan, C.-W.; Ye, M.; Zhang, L.; Wang, K.; Wang, H.; Xu, B. Deep learning illuminates spin and lattice interaction in magnetic materials. *Phys. Rev. B* **2024**, *110*, 64427.

(61) Tuckerman, M. *Statistical Mechanics: Theory and Molecular Simulation*; OUP Oxford, 2010.

(62) Zhang, L.; Han, J.; Wang, H.; Car, R.; E, W. Deep potential molecular dynamics: a scalable model with the accuracy of quantum mechanics. *Phys. Rev. Lett.* **120**, 143001.

(63) Zeng, J.; Giese, T. J.; Zhang, D.; Wang, H.; York, D. M. DeePMD-GNN: A DeePMD-kit Plugin for External Graph Neural Network Potentials. *J. Chem. Inf. Model.* **2025**, *65*, 3154–3160.

(64) Anstine, D. M.; Isayev, O. Machine Learning Interatomic Potentials and Long-Range Physics. *J. Phys. Chem., A* **2023**, *127*, 2417–2431.

(65) Zhu, J.-X.; Cheng, Z.; Yu, K. DMFF plugin for DeePMD-kit (torch backend). 2025; <https://github.com/chiahsinchu/dp\%5Fdmff>.

(66) Darden, T.; York, D.; Pedersen, L. Particle mesh Ewald: An $N \log(N)$ method for Ewald sums in large systems. *J. Chem. Phys.* **98**, 10089–10092.

(67) Rappe, A. K.; Goddard, W. A. Charge equilibration for molecular dynamics simulations. *J. Phys. Chem.* **1991**, *95*, 3358–3363.

TOC Graphic

