



HAL
open science

FeNNol: an Efficient and Flexible Library for Building Force-field-enhanced Neural Network Potentials

Thomas Plé, Olivier Adjoua, Louis Lagardère, Jean-Philip Piquemal

► To cite this version:

Thomas Plé, Olivier Adjoua, Louis Lagardère, Jean-Philip Piquemal. FeNNol: an Efficient and Flexible Library for Building Force-field-enhanced Neural Network Potentials. *The Journal of Chemical Physics*, 2024, 161 (4), pp.042502. 10.1063/5.0217688 . hal-04567115

HAL Id: hal-04567115

<https://hal.science/hal-04567115v1>

Submitted on 25 Jul 2024

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



Distributed under a Creative Commons Attribution 4.0 International License

FeNNol: an Efficient and Flexible Library for Building Force-field-enhanced Neural Network Potentials

Thomas Plé,^{1, a)} Olivier Adjoua,¹ Louis Lagardère,¹ and Jean-Philip Piquemal^{1, b)}
Sorbonne Université, LCT, UMR 7616 CNRS, 75005, Paris, France

(Dated: 2 July 2024)

Neural network interatomic potentials (NNPs) have recently proven to be powerful tools to accurately model complex molecular systems while bypassing the high numerical cost of *ab-initio* molecular dynamics simulations. In recent years, numerous advances in model architectures as well as the development of hybrid models combining machine-learning (ML) with more traditional, physically-motivated, force-field interactions have considerably increased the design space of ML potentials. In this paper, we present FeNNol, a new library for building, training and running force-field-enhanced neural network potentials. It provides a flexible and modular system for building hybrid models, allowing to easily combine state-of-the-art embeddings with ML-parameterized physical interaction terms without the need for explicit programming. Furthermore, FeNNol leverages the automatic differentiation and just-in-time compilation features of the Jax Python library to enable fast evaluation of NNPs, shrinking the performance gap between ML potentials and standard force-fields. This is demonstrated with the popular ANI-2x model reaching simulation speeds nearly on par with the AMOEBA polarizable force-field on commodity GPUs (GPU=Graphics processing unit). We hope that FeNNol will facilitate the development and application of new hybrid NNP architectures for a wide range of molecular simulation problems.

I. INTRODUCTION

The field of large-scale molecular dynamics (MD), has traditionally relied on the use of force fields (FFs) to describe interatomic interactions. Force fields can leverage different functional forms, from simple classical fixed charge models such as AMBER¹, CHARMM² or OPLS³ to the more *ab-initio*-inspired ones such as SIBFA^{4,5}, GEM^{6,7} or MB-pol^{8,9}, with a spectrum of intermediate ones such as AMOEBA^{10,11}, AMOEBA+^{12,13} or ARROW¹⁴. Each form strikes a unique balance between representational power and computational efficiency and the choice between force fields depends on the complexity of the system under study,^{15–17} as well as the availability of model parameters. Force-fields have reached a high maturity over the years, with automatic parameterization tools^{18–21} and efficient implementations^{22–27}, enabling their use for larger and more complex systems. Despite continuous refinements, traditional FFs have inherent limitations due to their fixed functional form which often relies on *ad-hoc* atom types, and to the complexity of the parameterization task.

In recent years, a new breed of potentials has emerged, powered by machine learning (ML) techniques. These ML potentials, including for example neural network interatomic potentials (NNPs), leverage flexible functional forms and large amounts of *ab initio* data to reach very high accuracy on broad classes of systems. This makes them a promising alternative to traditional FFs, especially for modeling exotic systems or chem-

ical reactivity. Stemming from foundational architectures such as the high-dimensional Behler-Parinello Neural Network²⁸ (BPNN) or the Gaussian Approximation Potential²⁹ (GAP), a large variety of model architectures have been developed, using various computational frameworks such as kernel models^{29–31}, NNPs with fixed-form symmetry functions^{28,32,33}, message-passing neural networks^{34–38}, E(3)-equivariant tensor computations^{39–44} or multi-scale representations^{45,46}. The development of hybrid models, which combine ML with more traditional, physically-motivated force-field interactions, has further expanded the design space of ML potentials^{9,41,47–57}. Comprehensive reviews of ML potential architectures and the chronology of their development can be found in refs. 58 and 59.

The advent of these new models and architectures has however brought its own set of challenges. Prominently, numerical implementations of these architectures are currently scattered in many packages provided by the authors of the methods. This makes a direct comparison of the intrinsic computational complexity of different models a difficult task. A few existing libraries provide a wider scope and more modular implementations such as SchnetPack⁶⁰, MLAtom⁶¹ or DeepMD-kit⁶² but they are still limited in the range of models offered to the user. Furthermore, designing models for non-standard workflows usually requires custom implementations, thus imposing barriers at entry for non-expert users. Examples of such applications include, for example, fitting distributed multipoles^{63–65} or a decomposition of the energy, or feeding force-field terms with ML-generated parameters, as was done in our previously proposed FeNNix model⁵⁷.

To address these challenges, we present FeNNol, a new Python library designed for building, training, and running atomistic machine-learning models, with a partic-

^{a)}Corresponding author: thomas.ple@sorbonne-universite.fr

^{b)}Corresponding author: jean-philip.piquemal@sorbonne-universite.fr

This is the author's peer reviewed, accepted manuscript. However, the online version of record will be different from this version once it has been copyedited and typeset.

PLEASE CITE THIS ARTICLE AS DOI: 10.1063/5.0217688

ular focus on physics-enhanced neural network potentials. FeNNol provides a flexible and modular system that allows users to easily build custom models, allowing for example the combination of state-of-the-art atomic embeddings with ML-parameterized physical interaction terms, without the need for explicit programming. FeNNol leverages the Jax framework^{66,67} and its just-in-time compilation capabilities, and provides a collection of efficient and configurable modules that can be composed to form complex models.

The paper is structured as follows: in section II, we describe FeNNol's software architecture, detailing its design principles and the underlying technologies it leverages. We focus in particular on FeNNol "modules" which are at the core of model building. In section III, we introduce a novel multi-paradigm atomic embedding which serves as a demonstration of the highly configurable nature of FeNNol modules. In section IV, we describe FeNNol's training system that enables users to define complex models and train them on generic tasks. In section V, we present the different ways to use FeNNol models in molecular dynamics simulation. We then demonstrate the performance of FeNNol's models and native MD engine by showing that our implementation of the popular ANI-2x model⁶⁸ reaches simulation speeds close to the optimized GPU-accelerated (GPU=Graphics Processing Unit) Tinker-HP implementation of the AMOEBA force-field²⁷. We also showcase our interface with Deep-HP⁶⁹ that enables efficient NNP simulations of systems reaching half a million atoms. Finally, in section VI, we provide concrete examples of model training with FeNNol, demonstrating how one can easily build non-standard models and how state-of-the-art results can be achieved within a few hours of training on a commodity GPU.

II. SOFTWARE ARCHITECTURE

FeNNol is built using the Jax library⁶⁶ that enables developers to easily write GPU-accelerated and automatically differentiable Python functions. More specifically, it relies on the Flax ML framework⁷⁰ to build parameterized models (such as neural networks) in a compact, highly configurable and modular way. Model building in FeNNol is organized around Flax Modules that each implement a specific functionality (for example computing an atomistic embedding, aggregating pairwise information, evaluating a neural network, etc...). By convention, modules take as input a dictionary containing information on the molecular system (coordinates, species, batch indices, etc...) and all information provided by previous modules. Each module thus uses this information to perform its computation and returns a new enriched dictionary that is passed to subsequent modules. In the following, we shall refer to this dictionary as the *system state* (as opposed to the *model state* that contains the different parameters and mutable state variables of the model).

An important feature of Jax is its ability to just-in-time (JIT) compile and optimize code (using the XLA compiler), thus merging modules and removing unused code at execution. This allows for very configurable modules while retaining efficiency at runtime. The principal drawback of this approach (compared to eager execution in Pytorch for example) is that all array sizes must be known at compilation time. The main restriction that this imposes is for the construction of neighbour lists (the list of pairs of atoms with distance under a certain cutoff) which size may be unknown in advance. The construction of neighbour lists thus cannot be fully compiled and we resolve this issue by using a preprocessing pipeline that is not fully JIT compiled nor differentiated (details in section II A).

A global description of the software architecture and data flows is provided in figure 1. In the following subsections, we describe the different modules available to build models, starting with the preprocessing modules that handle operations on neighbour lists and following with the computation modules that handle the bulk of models' logic and are divided in a few categories: atomistic embeddings, chemical and radial encodings, physics modules, neural networks and operation modules. We will then describe how one can extend FeNNol by registering custom modules. Finally, we provide an example of how to build a full model with the FENNIX constructor which is the main entry point for model building.

A. Preprocessing Modules

As explained earlier, a preprocessing pipeline that is not compiled is necessary to build neighbour lists without assuming a maximum size in advance. The preprocessing modules thus handle operations on neighbour lists – or more generally molecular graphs – such as building, filtering, padding or extending.

The main preprocessing module is the `GraphGenerator` which constructs a molecular graph given a cutoff distance, a set of atom coordinates, and optionally cell vectors for periodic boundary conditions. To avoid padding the array of coordinates when a batch of different molecular structures is provided, we chose to handle a concatenated array of all the atoms, accompanied by an array keeping track of the batch index of each atom. A molecular graph consists *a minima* of two arrays of atom indices: the source and destination of each edge in the molecular graph. An edge is assigned between two atoms when they have the same batch index and the distance between them is lower than the cutoff distance. Note that in the current implementations, all graphs are directed (meaning that each pair of atoms appears twice in the neighbour list, with source and destination atoms reversed). If cell vectors are provided, edges are assigned for each pair

FeNNol Architecture and workflow

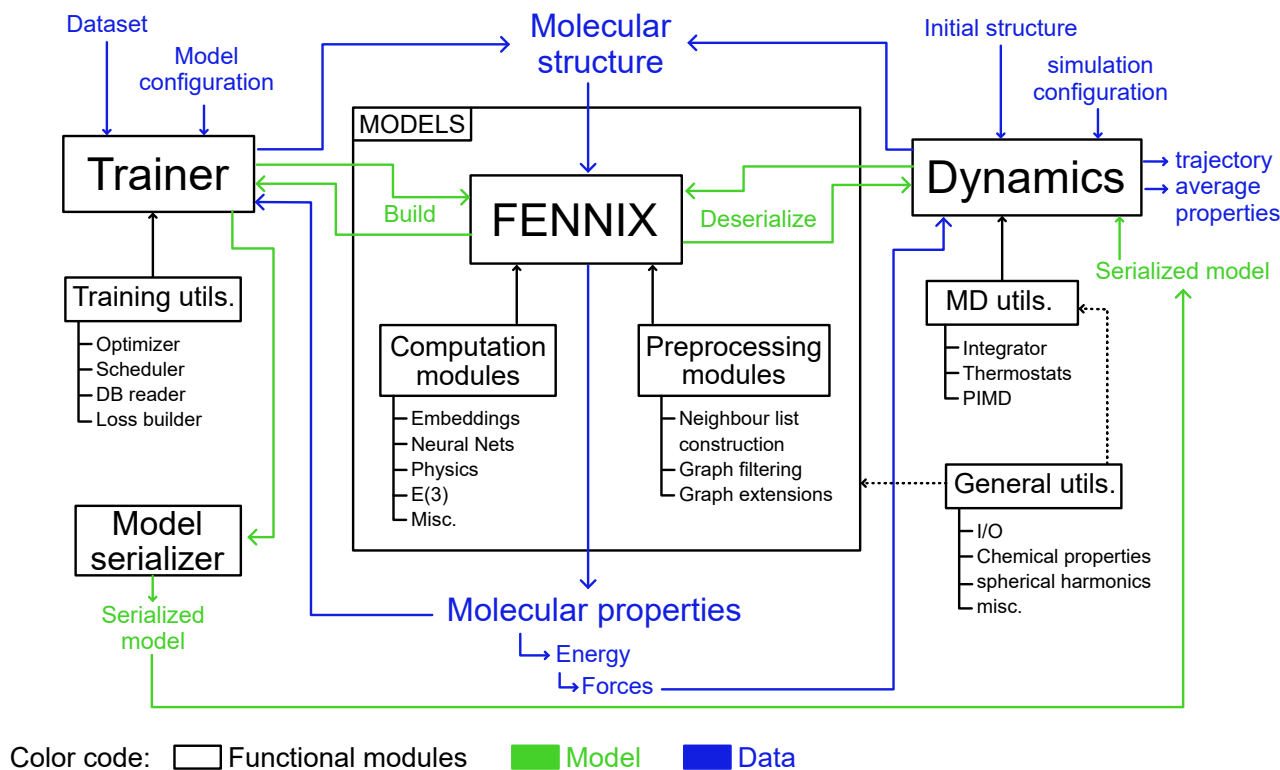


FIG. 1. Architecture and data flow of the FeNNol library. Black blocks represent functional modules, green arrows represent model operations and blue arrows represent data flows.

of images within the cutoff distance (with at least one of the atom in the main periodic cell) or according to the minimum image convention if required by the user. When periodic boundary conditions (PBCs) are applied, an array containing the shift (in units of cell vectors) to apply to the vector connecting the source and destination atoms is added to the graph. In the current implementation, neighbour lists are obtained by first computing all distances between pairs of atoms (that have the same batch index) and then filtering to the cutoff distance. This implementation is efficient for small number of atoms and, with GPU acceleration (and using a neighbour list skin in MD simulations to compute all distances less often), we found that systems up to 20,000 atoms do not suffer too much from the quadratic scaling. Future releases will use cell lists to improve the scaling of neighbour lists construction. Finally, when the model uses long-range interactions that require Ewald summations⁷¹, the user can configure the `GraphGenerator` module to construct vectors from the reciprocal lattice and filter them according to a magnitude cutoff.

In some cases, it can be useful to filter the molecular graph according to a shorter cutoff distance. This would for example happen when long-range physical interaction (for example dispersion interactions) are cou-

pled to a shorter-range ML embedding. In that case, the `GraphGenerator` module would construct the longer-range graph and the `GraphFilter` module can be used to select only the pairs within the shorter cutoff. This outputs a separate graph (which is marked as a subgraph of the original longer-range graph) with the filtered source and destination arrays and a new array containing a mapping of edges in the new graph to indices in the original graph. This map can be used by subsequent modules to transfer edge properties between the two graphs.

Finally, we provide the `GraphAngularExtension` module that builds a list of pairs of edges with a common central atom. This is used when explicit three-body terms are computed by the model, for example in Behler-Parrinello²⁸ or ANI³³ symmetry functions. This module does not produce a new graph but simply enriches the input graph.

In order to minimize re-compilation of the main modules, the size of neighbour lists must not change often. Thus, each array which size is the number of edges in the graph is padded to the largest neighbour list size encountered up to the current execution. A state variable thus keeps track of this maximum size which is updated to the current list size if it exceeds it (with an optional multiplicative factor to avoid too frequent re-

compilations). In practice, for example when performing molecular dynamics, this implies multiple re-compilation in the beginning of the simulation. When the maximum size of the neighbour list is found, the main modules do not need to re-compile and the simulation achieves its full efficiency. This can be achieved faster by increasing the multiplicative factor for the padding. The additional edges from the padding however impose an overhead on the computation and a balance between fast convergence of the list size and minimal padding must be found for the specific model and simulated system.

Finally, since, the preprocessing modules are not differentiated, the distances computed in these modules cannot be used to compute forces (or coordinates-dependent derivatives in general) as they do not explicitly depend on the coordinates array passed to the main modules. Thus, for each preprocessing module, a matching *processing* module is automatically added to the list of main modules when constructing a full model (see section IID). These processing modules compute edge vectors and distances with the provided coordinates and neighbour lists, as well as angles between pairs of edges if necessary. Furthermore, they compute a switching function that depends on the distance and goes smoothly to zero as it approaches the cutoff distance. This switching function is used by subsequent modules to avoid discontinuities in molecular properties when atoms enter or leave an atomic neighbourhood. The switching function can be configured by the user with optional trainable parameters. Functions available in the current implementation are the cosine function (as used in the ANI³³ models for example), the polynomial envelope from ref. 72 and the exponential function from SpookyNet⁴¹.

B. Computation Modules

In this section, we describe the main modules responsible for the bulk of the computation in a model. A model is constituted by a chain of such modules that perform operations on the molecular graphs and system state such as computing an atomic environment embedding (or descriptor) or processing it with a multi-layer perceptron to compute the energy. This chain of modules is encapsulated in a top-level Flax module which is JIT-compiled and acts as a pure function that can be automatically differentiated either with respect to model parameters (for training) or to inputs (for example to compute forces). In the following, we list the different categories of modules and provide some examples.

a. Atomic Embeddings

Atomic embedding modules process the molecular graph to compute descriptors of the atoms' environments that can later be used by other modules (such as neural networks) to compute properties of the system (for

example the energy). Embeddings are usually of atom-wise or pair-wise nature and can be either local (that only process directly connected atoms), message-passing (that capture longer-range interactions via iterative exchanges of information with neighbouring atoms) or completely non-local (for example global embeddings such as Coulomb matrices⁷³ or non-local atom-wise like in the LODE framework^{45,46} or the attention mechanism in SpookyNet⁴¹). FeNNol provides various embeddings from the literature: local models such as ANI AEVs³³, Allegro⁴⁴, EEACSF⁷⁴ and DeepPot-SE⁷⁵ and message-passing models such as AIMNet³⁸, HIP-NN³⁷, PaiNN⁷⁶ and NewtonNet⁷⁷. When necessary, E(3)-equivariant operations (such as tensor products) are performed using the `e3nn-jax`⁷⁸ library. In some simple cases, we implemented equivariant operations from scratch to increase performance (for example, the filtered tensor product proposed in ref. 79). We also implemented a new multi-paradigm embedding denoted as CRATE that will be described in section III.

b. Chemical and Radial Encodings

Basic building blocks of most atomic embeddings are chemical and radial encodings that capture information about the chemical species of each atom and the distance between pairs of atoms.

The chemical encoding maps each chemical species to a multidimensional vector that can be either fixed or learned during training and that will be further processed by the atomic embedding to combine it with geometric information. Different chemical encodings will provide different prior knowledge to the model about the relations between chemical species. For example, a simple one-hot encoding makes no particular assumptions while the positional encoding of ref. 57 or the electronic structure encoding of ref. 41 assume similarities for atoms close together in the periodic table. FeNNol provides a few different chemical encodings: the one-hot encoding, an encoding based on electronic occupancy as introduced in TeaNet⁸⁰, the electronic structure encoding from 41, the positional encoding from ref. 57, a 4-dimensional encoding based on the Stowe-Janet-Scerri periodic table. The encoding can also be initialized as a random vector for each species. These encodings can be combined together via concatenation and optionally be optimized during training.

The radial encoding is a projection of the distance between two atoms on a set of smooth basis functions. FeNNol implements multiple families of basis functions: the simple Gaussian basis used for example in the ANI models³³, the Bessel basis introduced in ref. 72, the inverse-distance Gaussian used in HIP-NN³⁷, a Fourier expansion and the Bernstein polynomials used in ref. 41. As the nature and number of basis functions directly influence the spatial regularity of the model, it can drastically impact the fit of target properties so that comparison between different radial basis can be useful when designing a model.

c. Physics Modules

In order to build hybrid models that incorporate long-range effects such as FENNIX-OP1⁵⁷, FeNNol offers multiple physically-motivated interaction modules. It provides for example a Coulomb module that computes electrostatic interactions between distributed charges with different short-range damping schemes available. In this module, long-range interactions in periodic systems are handled using Ewald summations (later versions of FeNNol will include an implementation of Smooth Particle Mesh Ewald⁸¹ for more efficient interactions). Other available physics modules include the recently proposed QDO dispersion and exchange interactions⁸², charge equilibration (as defined in DFT-D4⁸³) and the ZBL repulsive pair-potential⁸⁴. Most of the physics modules can be parameterized using dynamically computed properties (for example obtained via a neural network) as is done for example with atomic volumes parameterizing electrostatics and dispersion interaction in ref. 57.

d. Neural Networks

The main use of neural networks in atomistic models is to process embeddings and produce outputs (for example atomic energies, charges, etc...). The neural network weights can then be fitted against relevant targets in a dataset or used to parameterize physical interactions. FeNNol provides multiple neural network architectures: a simple multi-layer perceptron, a residual neural network and various chemical-species-dependent neural networks. For the latter, different levels of species specialization are provided: full species-dependent network parameters, shared parameters with species-dependent activation scales and bias or low-rank chemical adaptation (ZLoRaNet) inspired from the LoRa methodology of fine-tuning⁸⁵.

e. Operation Modules

Finally, we implemented various miscellaneous modules that perform simple operations on the system state such as reshaping an entry, adding multiple entries together, sum bond-wise properties to atomic centers or apply various shifts, scaling and activation functions. These can be used to construct more complex models from simple building blocks directly in the model input instead of coding explicit custom modules. A basic example would be to use the `Concatenate` module to add charge information to a short-range embedding as is done in 4G-HDNNP⁴⁹. This flexibility aims at letting the user quickly experiment and iterate with the model architecture.

C. Extensions with custom modules

As mentioned in the beginning of section II, a FeNNol module is simply a Flax module (i.e. a class inheriting from `flax.linen.Module`) which `__call__` method takes the system state (the dictionary containing all previously computed properties of the system) as input and returns a new processed system state as output. As an example, the following code shows the `Activation` module :

```
class Activation(flax.linen.Module):
    key: str
    activation: Union[Callable, str] = nn.silu
    scale_out: float = 1.0
    shift_out: float = 0.0
    output_key: Optional[str] = None

    ### FeNNol ID for automatic registration
    FID: str = "ACTIVATION"

    @flax.linen.compact
    def __call__(self, inputs: Dict[str, Any]):

        ## get the entry to which we want
        ## to apply the activation function
        x = inputs[self.key]

        ## get activation function
        activation = (
            activation_from_str(self.activation)
            if isinstance(self.activation, str)
            else self.activation
        )

        ## apply activation and
        ## linear transformation
        output = (
            self.scale_out * activation(x)
            + self.shift_out
        )

        ## determine output key
        if self.output_key is None:
            output_key = self.key
        else:
            output_key = self.output_key

        ## return new system state
        return {**inputs, output_key: output}
```

This module applies an activation function followed by a linear transformation to an entry in the system state. As flax modules are frozen dataclasses, all their attributes are specified and explicitly typed in the class definition, the class constructor is automatically built and attributes are immutable after initialization (see

the Flax documentation for more details). Note that for a flax module to behave properly, the `__call__` method must be a pure function. This is why we do not explicitly modify the `inputs` dictionary but rather return a shallow copy with additional keys. A key advantage of the jax/flax framework is that all non-jax operations are considered static by the jit compiler and thus executed only once at compile-time. This allows modules to be extremely configurable without imposing computational overheads. For example, the `activation_from_str` function in the previous code example parses the `self.activation` attribute to build an activation function from a character string. By construction, the attributes of flax modules are immutable so that `self.activation` is guaranteed to never change and it would thus be wasteful to parse it again at each call of the model. Thanks to the jit system, this parsing is only done once at compile-time.

When importing the FeNNol library, all the flax modules in the source code with the FID (Fennol Identifier) attribute are automatically registered as FeNNol modules and available when building a model. FeNNol also automatically registers modules from the Python files found in repositories listed in the environment variable `FENNOL_MODULES_PATH`. The user can also manually register a module after initialization by providing their custom module and an FID to the `register_fennix_module` function before building or loading a model.

D. Building a FeNNix model

In FeNNol, a model is constructed by calling the `FENNIX` constructor which most important arguments are two ordered dictionaries: one that lists the preprocessing modules and their properties and another one for the computation modules with their properties:

```
from fennol import FENNIX
import jax
rng_key = jax.random.PRNGKey(123456)

model = FENNIX(
    cutoff=5.,
    modules={...},
    preprocessing={...},
    rng_key=rng_key,
    energy_terms=["energy"],
)
```

The model automatically builds a primary graph with a cutoff radius specified by the `cutoff` argument. This primary graph can optionally be customized using the preprocessing parameters. The argument `rng_key` is a key for the random number generator that is used to initialize the model parameters. The `energy_terms` argument defines which keys in the final system state will

be used as energy components and summed to form the total energy. If only one graph is required (i.e. no filtering and no angle information needed), the preprocessing argument can be omitted. The `modules` dictionary has the following structure:

```
modules = {
    "module1":{
        "FID": MODULE_1_FID
        "param1": ...
        "param2": ...
    },
    "module2":{
        ...
    },
    ...
}
```

The key for each module will be used as the `name` attribute of the underlying flax module and may be used to setup output keys. The "FID" key is the identifier of the module. Finally, the remaining entries set values for the module parameters (i.e. the dataclass attributes). A list of registered modules can be obtained by calling the `available_fennix_modules` function.

After manipulations (for example training), a model can be serialized and written to a file using the `FENNIX.save` method as

```
model.save("model_name.fnx")
```

By convention, we use the `.fnx` extension to identify model files. This file contains all the information needed to reconstruct the model. This can be achieved with the `FENNIX.load` class method as

```
model = FENNIX.load("model_name.fnx")
```

A flow diagram of the execution of a ANI-like model built with FeNNol is provided in figure 2. The input is built from chemical species (atomic numbers) and coordinates. In case where multiple structures are provided (batching), all the atoms are concatenated in the `species` and `coordinates` arrays and an additional `natoms` array is provided containing the number of atoms of each structure in the batch and is used to compute the batch index of each atom. The first modules on the left are the preprocessing modules that build the neighbour list (with a cutoff of 5 Å), filter it to a new graph with a cutoff of 3.5 Å and finally extend this graph with indices of pairs of neighbours forming angles that will be used to construct the environment vector in the `ANI_AEV` module. The blue arrows indicate graph processing modules that are automatically added by the corresponding preprocessing modules. As explained earlier, these processing modules compute distances and angles using provided edge indices that will be automatically differentiated (green box in the diagram) to obtain atomic forces. The next two modules are the user-defined main modules (here `ANI_AEV` that computes the atomic environment vectors of the ANI model³³ and `ChemicalNet` that

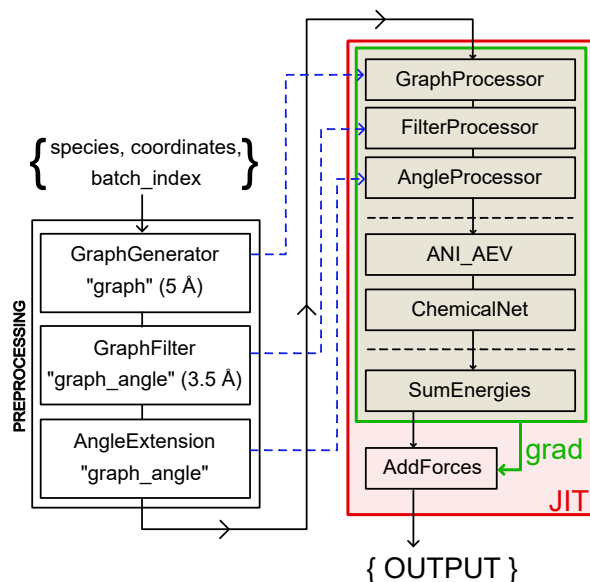


FIG. 2. Flow diagram of the execution of a ANI-like model built with FeNNol

applies a species-specialized neural network to the AEVs to obtain atomic energies). The last automatically differentiated block simply sums the energy contributions of each atom to the corresponding structure's total energy. The final module adds atomic forces obtained from the `jax.grad` transform to the output.

III. THE "CRATE" MULTI-PARADIGM EMBEDDING

In order to exploit the highly configurable nature of flax modules, we developed the CRATE (Configurable Resources from ATomic Environment) atomic embedding. The idea behind CRATE is to combine chemical and geometric information from different sources: chemical information from species embeddings, radial and angular information from radial basis and angle embeddings, tensorial information from E(3)-equivariant operations and long-range information using the LODÉ methodology. The user can specify which resources they need, depending on the nature of the system, their fitting targets (for example fitting distributed multipoles would require tensorial information while energies might only require radial and angular information) and their computational budget. These resources can also be iteratively combined via interaction layers using message-passing or, if locality is important, using simulated message-passing with a local key-query mechanism that we detail in Appendix A or with directed edge embeddings. Finally, we designed the CRATE components to minimize expensive operations on graph edges or triplets. For examples, neural networks applied to high-dimensional vectors or expensive E(3) tensor products are only performed on atomic centers. Edge or triplet

operations then mostly consist in element-wise addition, multiplication or function activation, as well as (generalized) outer products. A detailed description of the model architecture and the composition of the different resources can be found in Appendix A.

By selecting specific resources and interaction layers, the CRATE embedding can be made to operate similarly to well-established architectures. For example, selecting radial and angular information with only one interaction layer, the model behaves like the ANI³³ or BPNN²⁸ architectures. Keeping the same resources and adding multiple message-passing layers leads to an architecture that resembles AIMNet³⁸. Finally, if we select E(3)-equivariant resources, refine them with local atom-wise tensor products and combine them using message-passing interactions, we obtain a model close to the recently proposed MACE⁸⁶. Each interaction layer can be configured separately so that many combinations are possible and one can tailor the architecture for their data and computational efficiency requirements.

IV. MODEL TRAINING

FeNNol provides a basic training system that should handle most simple model training tasks. This training system can be invoked with the command `fennol_train` that is installed with the FeNNol Python package. The input to the training program is a structured file (either yaml, json or toml) that describes both the training parameters and the model structure. The latter is the same as presented in section IID. The training parameters can be subdivided in three categories: general parameters, optimizer and schedule parameters and loss function definition.

General Parameters General parameters include the path to the dataset, the batch size, number of batches per epoch, maximum number of epochs, the names of the output directory and log file, seed for the random number generator, etc...

Optimizer and schedule parameters Model optimization is performed via stochastic gradient descent using optimizers provided by the Optax library⁶⁷. The default optimizer is the Adabelief⁸⁷ which we found to perform slightly better than the more common Adam optimizer. The optimizer can be chosen and configured by the user by providing the name of the Optax optimizer and its parameters, as described in the Optax documentation. Decoupled weight decay⁸⁸ can be applied (independently of the chosen optimizer) to all or a subset of user-selected parameters. Additionally, parameter gradients can be clipped using the adaptive gradient clipping methodology⁸⁹. Finally, the user can define a subset of modules and/or parameters to be frozen, meaning that their parameters will not be updated (this is useful for example in a fine-tuning or transfer-learning stage).

The learning-rate scheduler can also be configured by the user with a choice between a constant learning-rate, a cosine-one-cycle schedule⁹⁰ or an adaptive schedule similar to Pytorch's `reduce_lr_on_plateau`.

Loss definition The full loss function for the model optimization is defined in the training configuration file. This definition takes the form of a dictionary of loss function components with parameters such as the reference key in the dataset, the key to use in the model output, the weight of the loss component and the type of loss function to use. The following listing (in yaml format) provides an example for a mean square error loss function on both energy and forces:

```
loss:
  e:
    key: total_energy
    ref: formation_energy
    weight: 1.e-3
    type: mse
    per_atom: True
  f:
    key: forces
    ref: forces
    weight: 1.
    type: mse
```

Multiple loss types are provided such as means squared error, mean absolute error, log hyperbolic cosine, negative log-likelihood or CRPS for ensemble training⁹¹ or evidential losses^{92,93}. The loss type can be set for all components at once using the `default_loss_type` parameter. The `per_atom` directive only applies to system-wise properties and divides the loss by the number of atoms. This can be useful when training on systems of widely varying sizes.

Training stages When training a model on a multi-component loss function, it may be useful to define a multi-stage training. As a concrete example, we can mention the training of the FENNIX-OP1 model⁵⁷ which required to first train on short-range energies and partial charges, and then training again with long-range energy contributions activated. Another application is transfer learning where we would for example pre-train a model on a large DFT dataset and then transfer the model to CCSD(T) accuracy on a smaller dataset. In order to facilitate the reproducibility of such multi-stage training, it can be fully defined in a single input file to the FeNNol training system. The following listing shows an example of a possible transfer-learning task:

```
training:
  batch_size: 128
  default_loss_type: mse
  lr: 1.e-3

stages:
```

```
train_dft:
  dspath: dft_dataset.pkl
  schedule_type: reduce_on_plateau
  end_event: [lr, 1.e-6]
  max_epochs: 100000
  loss:
    e:
      key: total_energy
      ref: formation_energy

train_ccsdt:
  dspath: ccsdt_dataset.pkl
  schedule_type: cosine_onecycle
  end_event: none
  max_epochs: 100
  frozen: [embedding]
```

In this example, the first stage `train_dft` trains the model on formation energies from the (hypothetical) `dft_dataset.pkl` dataset. This stage lasts a maximum of 100000 epochs or ends when the learning rate drops under 10^{-6} . The second stage trains the model with a fixed schedule on the `ccsdt_dataset.pkl`. Training parameters between stages are cumulative so that the loss did not need to be redefined in the second stage (provided that the keys in both datasets are the same). Finally, the second stage freezes the parameters of the model's embedding (defined by the module name `embedding`) to limit the number of parameters to update for the transfer-learning task.

Structure of accepted datasets In its current version, FeNNol only accepts a simple but specific structure for datasets. The root of the dataset should be a dictionary with at least the two entries `training` and `validation` that are both lists containing the molecular configurations (one per element in the list) that will be used to train and validate the model respectively. Each molecular configuration is itself a dictionary which should contain the entries `species` (an array of atomic numbers of the atoms in the configuration) and `coordinates` (the array of 3D coordinates of the atoms), as well as all the keys used as reference in the definition of the training loss. The dataset must be formatted by the user and can be provided in multiple formats. The simplest is to use the generic `pickle` format to save the dataset from a python script. This format will be completely loaded in memory before training and loading batches is thus very fast. This format is best suited for small to moderately-large datasets as it must fit in memory and initial loading times can be very long for large datasets. In most applications, it should however be the simplest format and we found it practical even for quite large datasets (for example the ANI-2x dataset⁶⁸ with almost ten million entries). For larger datasets, one can provide a hdf5 file or a sqlite database that can be queried without loading it in memory completely. These formats may however have slower access times.

V. RUNNING MOLECULAR DYNAMICS SIMULATIONS

FeNNol offers multiple ways of running molecular dynamics (MD) simulations with FeNNix models. One can either use custom python scripts, the provided calculator for the Atomic Simulation Environment⁹⁴ (ASE) package, the Tinker-HP MD engine^{26,27} via the Deep-HP interface⁶⁹ or FeNNol's native MD engine.

A. Custom scripts

FeNNix models define easy-to-use shorthand functions for computing energies, forces and stress tensor that can be used to build a custom MD engine in python. These can be accessed via the class methods: `FENNIX.total_energy`, `FENNIX.energy_and_forces` and `FENNIX.energy_and_forces_and_virial`. They all accept keyword arguments which constitute the initial system state (so at least atomic species and coordinates) and return the total energy, forces and virial along with the final system state. Under the hood, these methods handle padding (and unpadding) of arrays to prevent too frequent recompilations of the model. Furthermore, these routines can be configured with the method `FENNIX.set_energy_terms` to sum different keys of the final system state that are interpreted as energy components. For example, one could call

```
model.set_energy_terms(["nn_energy",
                       "coulomb_energy"])
```

to define the total energy as the sum of a neural network potential (with `nn_energy` the key referring to the output of the neural network) and electrostatic interactions via a Coulomb term (stored in the `coulomb_energy` key).

B. ASE calculator

FeNNol provides a `FENNIXCalculator` class that can be used with the Atomic Simulation Environment⁹⁴ (ASE) package. This enables the use of all the advanced functionalities provided by ASE such as geometry optimization, phonon calculations and various MD algorithms in a very user-friendly manner. For example, a geometry optimization followed by a short MD simulation with a FeNNix model can be obtained with the following simple code:

```
import ase
from fennol.ase import FENNIXCalculator

# initialization
atoms = ase.io.read("aspirin.xyz")
atoms.calc = FENNIXCalculator(
    "model_file.fnx"
)
```

```
# geometry optimization
opt=ase.optimize.BFGS(atoms)
opt.run(fmax=0.01)
print(atoms.get_potential_energy())

# MD simulation
dyn = ase.md.verlet.VelocityVerlet(
    atoms, 0.5 * ase.units.fs)
dyn.run(1000)
```

Since ASE is able to communicate with the i-PI software⁹⁵, simulations including nuclear quantum effects (NQEs) are directly available with this calculator.

C. Deep-HP interface

We modified the Deep-HP interface introduced in ref. 69 to enable the use of FeNNix models as energy components in Tinker-HP. This interface has multiple benefits. First, it allows the use of the different integrators, thermostats, barostats and accelerated sampling methods provided by Tinker-HP^{26,27}. Efficient methods for simulating NQEs are also readily available with the Quantum-HP module⁹⁶. Second, it enables the parallelization of FeNNix models over multiple GPUs and multiple nodes, opening up simulations of very large systems with ML models. Note that currently, multi-GPU parallelization is only safe with purely local models such as ANI or Allegro. Finally, FeNNix models can be combined with traditional force fields, either via an ML/MM embedding^{69,97,98}, or as energy components in a hybrid model^{99,100}.

As in ref. 69, we directly convert GPU pointers allocated from Tinker-HP (for example coordinates or neighbour lists) to Jax arrays via the DLPack protocol in order to minimize data movements and CPU/GPU communications. Note that the Tinker-HP/FeNNol interface is currently in early stages of development and will be refined and optimized in future works.

D. Native MD engine

FeNNol provides an efficient native engine to perform molecular dynamics simulations, which can be invoked with the `fennol_md` command. The engine is written purely in Jax, is completely GPU-resident and leverages Jax's asynchronous dispatch capabilities for efficient dynamics even for small systems. Multiple thermostats are available (Langevin, Bussi-Parinello or Nose-Hoover) and the integration scheme used is based on symplectic time-splitting methods and is adapted depending on the thermostat (for example it reverts to the velocity Verlet integrator for NVE MD and to the BAOAB integrator¹⁰¹ for Langevin MD). Simulations can be performed either in vacuum or with periodic boundary conditions (with

interactions with all images or minimum image convention). A current limitation of the engine, however, is that constant pressure simulations are not yet implemented. Furthermore, multi-GPU parallelization is not yet available.

For simulations including NQEs, we implemented two state-of-the-art methods: the adaptive Quantum Thermal Bath¹⁰² (adQTB) and Ring-Polymer Molecular Dynamics¹⁰³ (RPMD). The adQTB method includes approximate NQEs in a classical-like simulations via a colored-noise Langevin thermostat. Our implementation incurs only a small overhead compared to classical Langevin MD (between 5% and 20% depending on system size, in agreement with previous studies¹⁰⁴). On the other hand, RPMD uses replicas of the original system to achieve asymptotically exact (in the limit of infinite number of replicas) static properties of the quantum nuclei. This method is thus much more compute-intensive as it requires simulating a much larger system than in classical or adQTB MD (typical numbers of replicas are between 8 and 100 depending on the system and the simulated conditions). In our implementation, the potential energy of each of the replicas is efficiently obtained using FeNNol's batching capabilities.

E. Performance benchmarks

In the remaining of this section, we provide a brief study of simulation accuracy with respect to Jax matrix multiplication level of precision (via energy conservation tests) and performance benchmarks for our implementation of the ANI-2x⁶⁸ model. The model is implemented only with standard FeNNol modules (*i.e.* without custom CUDA kernels). Parameters for the neural network are obtained from the TorchANI¹⁰⁵ implementation and converted to Flax variables. We only use the first from the ensemble of eight neural networks that compose the ANI-2x model. All our tests were performed with Jax 0.4.25 with jaxlib 0.4.23 and CUDA 11.8. We used either a NVIDIA RTX 3090 or A100 GPU with respectively 24GB and 40GB of memory.

For our accuracy tests, we compare single (`float32`) and double precision (`float64`) representations. For single precision, we also compare the various matrix multiplications precisions provided by Jax that are denoted as "default" (which uses `bfloat16` numbers), "high" (which uses `tensorfloat32`) and "highest" (which uses full `float32`). Figure 3 shows the total energy of a periodic box of 216 water molecules over a 1 ns NVE simulation with a timestep of 0.1 fs (this small timestep was chosen to minimize integration error and focus on precision errors) with the different precision configurations. All the simulations start from the same off-equilibrium configuration and velocities are initialized to zero (this results in an average temperature of 315K

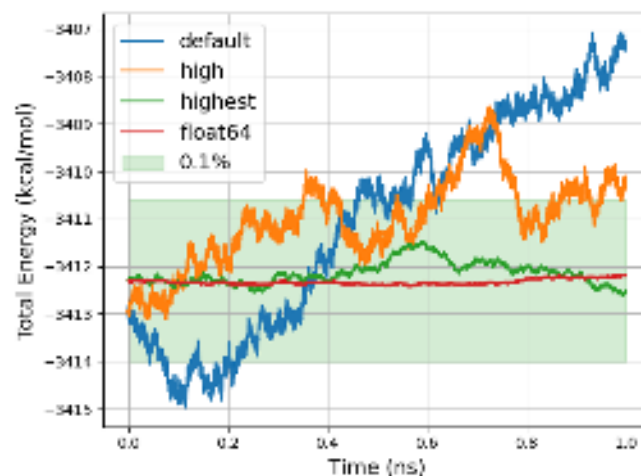


FIG. 3. Total energy of a box of 216 water molecules as a function of time for different simulation precisions. The float64 curve corresponds to full double precision. The other curves correspond to single precision with different Jax settings for matrix multiplication precision (default=`bfloat16`, high=`tensorfloat32`, highest=`float32`). All simulations use a timestep of 0.1fs and the ANI-2x model. The height of the green box represents 0.1% of the initial total energy.

with a standard deviation of 12K). It is clear from the figure, that for this simulation setup, both "default" and "high" configurations are not precise enough to obtain good energy conservation. Furthermore, we see a non-negligible shift in the initial energy compared to `float64`. The "highest" `float32` configuration provides acceptable energy conservation (fluctuations below 0.1% of the initial energy, as indicated by the green box in the figure) and no initial energy shift compared to the full `float64` simulation. From a performance perspective, differences between the three single precision configurations are small, while double precision is significantly more expensive on the RTX 3090 GPU (about a factor of 4). Note that this difference between single and double precision is less important with HPC GPUs. For example using double precision on a A100 GPU only adds a 15% overhead. Given these tests, we set the default precision for FeNNol's MD engine to "highest" `float32` and warn users of Jax that the default precision settings might not be sufficient in the context of MD simulations.

We now test the performance of our implementation on water boxes of varying sizes (from 648 atoms to 500K atoms) and on the DHFR protein¹⁰⁶ explicitly solvated in water (around 20K atoms). We compare our implementation (with ASE, Deep-HP or our native MD engine) to the base TorchANI implementation (with ASE or Deep-HP as MD engine). We also compare the performance of ANI-2x to the AMOEBA polarizable force field¹¹ simulated using Tinker-HP. Table I summarizes the differ-

	FeNNol	FeNNol+skin	FeNNol+ASE	FeNNol+Deep-HP	TorchANI	TorchANI+Deep-HP
model implementation	Jax	Jax	Jax	Jax	Pytorch	Pytorch
neighbour list	$\mathcal{O}(N^2)$	$\mathcal{O}(N^2)/N_{\text{update}}$	$\mathcal{O}(N^2)$	$\mathcal{O}(N)$	$\mathcal{O}(N^2)$	$\mathcal{O}(N)$
MD Engine	FeNNol(Jax)	FeNNol(Jax)	ASE(Numpy)	Tinker-HP(CUDA)	ASE(Numpy)	Tinker-HP(CUDA)

TABLE I. Summary of the different implementations that are compared in the performance benchmark of section V E. The "neighbour list" line shows the complexity of the neighbour list construction. The column headers refer to the legend in Figure 4.

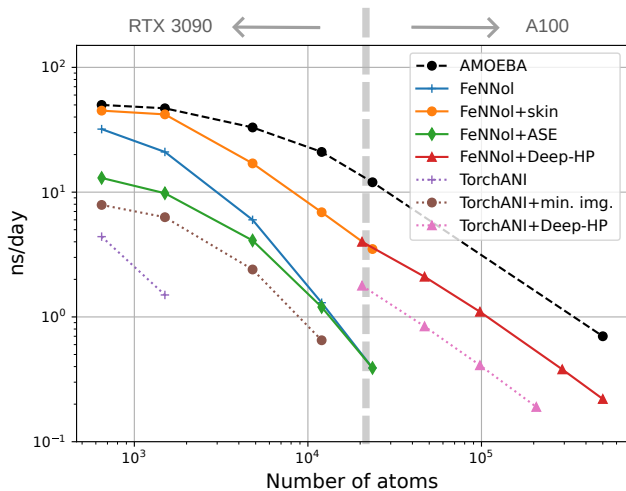


FIG. 4. Performance (in ns simulated per day) for various system sizes. The black dashed line corresponds to the AMOEBA force field simulated using Tinker-HP. The solid lines correspond to ANI-2x with FeNNol's implementation and different MD engines. The dotted lines correspond to ANI-2x with TorchANI's implementation (min. img.=minimum image convention). All simulations use float32 precision, a 0.5fs timestep and a Langevin thermostat set at 300K. Simulations with TorchANI and FeNNol's native MD engine (under 20K atoms) were performed on a RTX 3090 GPU. Simulations with Deep-HP (above 20K atoms) were performed on a A100 GPU.

ences between the implementations that we compare in this section. All simulations use float32 precision, a 0.5 fs timestep, a Langevin thermostat set at 300K and periodic boundary conditions. Both FeNNol and TorchANI are set to only use the first out of the eight neural networks that compose the ANI-2x model. Simulations with TorchANI and FeNNol's native MD engine were performed on a RTX 3090 GPU with 24GB of memory. Simulations with FeNNol+Deep-HP were performed on a A100 GPU with 40GB of memory.

Figure 4 shows the simulation performance (in ns simulated per day) for the different setups. In order to obtain fairer comparisons, we implemented the minimum image convention in TorchANI, which already provides a significant increase in performance for these systems. Note that we could not simulate the larger systems with TorchANI due to allocations exceeding our GPU's memory. If we first compare TorchANI to FeNNol both running with the ASE MD engine, we see that FeNNol gains

almost a factor of two in performance which is mainly due to the efficient JIT compilation of Jax (note that a more direct comparison could be achieved by using the recently introduced JIT compilation tools in Pytorch 2 but that would require a complete rewrite of TorchANI with fixed array shapes). Since ASE integrates the equations of motion on the CPU, using this engine requires frequent data transfers between CPU and GPU. Using FeNNol's native MD engine, we alleviate this issue with a fully GPU-resident and asynchronous simulation, which considerably increase performance for smaller systems (nearly a factor of three for the smallest system). For larger systems, the main bottleneck becomes the construction of neighbour lists, which in the current implementation, scales quadratically with system size (compared to the linear scaling of the ANI-2x model itself). The impact of this quadratic scaling can be pushed back to larger system sizes by using a neighbour list "skin" and to reconstruct the full neighbour list only once every few steps. If we add a buffer of 2 Å and, following the rule used in Tinker-HP, reconstruct the neighbour list once every 40 fs (80 steps), we obtain a significant improvement, reaching performance close to the AMOEBA force field on the smaller systems. This performance increase is due to both a gain in neighbour list construction and a better leveraging of Jax's asynchronous dispatch when using the additional buffer. Note that this comparison only takes into account the raw performance of the models with a fixed integration scheme. Typical MD runs with force fields however can usually afford larger timesteps and use advanced integration schemes (for example multi-timestepping¹⁰⁷) which are not compatible with current ML potentials and will require architectural evolutions.

For systems of more than 20K atoms, the quadratic scaling of neighbour lists construction in FeNNol's native MD engine takes a significant toll on performance. We thus switched to Deep-HP – which leverages the efficient cell lists implemented in Tinker-HP – for systems between 20K and 500K atoms. For these large systems, we see similar scaling trends with FeNNol+Deep-HP compared to AMOEBA with about a factor of 3 in raw performance. Note that the FeNNol+Deep-HP interface is in very early development, so that we expect this gap to shrink with further optimizations and a tighter integration between both codes. In particular, we expect that performing all the model preprocessing (which does not require automatic differentiation) inside Tinker-HP will provide a boost in performance and significant mem-

ory savings. When comparing to the raw performance of TorchANI+Deep-HP, we see again a factor ~ 2.5 with respect to FeNNol.

The results reported here were measured with a RTX 3090 GPU for smaller systems and a A100 HPC GPU for larger systems. We also tested the A100 GPU on the smaller systems and did not see significant differences with respect to the RTX GPU. For both GPUs, utilization is constantly close to 100% for all systems except the smallest one. For the A100 GPU, power usage reaches 70% of the cap value (280W / 400W) for the larger systems. For the RTX 3090, power usage stays close to the cap value (350W) for all systems except the smallest one, showing nearly maximal GPU occupancy. We can thus expect performance gains with the newer generation RTX 4090 that are close to theoretical values (i.e. at least a factor of two).

This brief study shows that neural network potentials can indeed reach performance close to force fields (at least the more expensive and accurate polarizable ones) and that the efficiency of the model implementation and the MD engine are both critical.

VI. USAGE EXAMPLES

In this section, we show three examples of models that can be trained with the FeNNol library and display training performance and timings on standard datasets. The first model is a message-passing CRATE model trained on the QM7-X dataset¹⁰⁸. The second model is an equivariant force-field-enhanced model similar to the FENNIX-OP1 model from ref. 57 which is trained on the ANI-1x dataset¹⁰⁹. Finally, we show how FeNNol can be used to model other properties than the potential energy with the task of learning distributed multipoles from a subset of the SPICE dataset¹¹⁰. Note that the goal of this section is not to release new models but to give the reader an idea of the diversity of models that can be built with FeNNol, as well as the time and computer resources required for training.

A. Training a CRATE model on the QM7-X dataset

For our first example, we train a CRATE model to reproduce formation energies and forces from the QM7-X dataset¹⁰⁸. The QM7-X dataset is composed of 4.2 million out-of-equilibrium configurations of small molecules (up to 7 heavy atoms with species H, C, N, O, S, Cl) with energies and forces computed at the PBE0+MDB level of theory. We trained two CRATE models with radial and angular resources, two or three layers of message-passing interactions with a cutoff radius of 5.3 Å and short-range trainable ZBL repulsion⁸⁴. Since this dataset only comprises isolated molecules, we did not include explicit long-range force-field terms in the model. The model was trained using batches of 128 configurations and a mean

square error loss with a weight ratio of 1:100 between total energies and forces. We refer the reader to the archive file provided in a Zenodo repository (see Data Availability section IX) for the complete input file defining the model, training parameters and loss function. We used the Adabelief⁸⁷ optimizer with a fixed learning rate schedule starting from 2.e-4 and decreasing to 1.e-6 over the course of 10^6 random batches. We reserved a random 20% of the dataset for validation. Training completed in 3.7 hours and 4 hours for the two- and three-layers model respectively, on a single NVIDIA RTX 3090 GPU. The two-layers model achieved a mean absolute error (MAE) on the validation set of 0.70 kcal/mol/Å for forces and 0.59 kcal/mol on energies. The three-layer model achieved a slightly better MAE of 0.49 kcal/mol/Å on forces and 0.39 kcal/mol on energies. These results are in line with those previously reported in ref. 41 for the PaiNN and SpookyNet models and show that state-of-the-art results can be achieved in a few hours of training on a single commodity GPU.

B. Force-field-enhanced equivariant model

For this second example, we trained a model similar to the FENNIX-OP1 model in ref.⁵⁷. We use an Allegro embedding⁴⁴ with 256 scalar features, 16 equivariant channels with maximum degree $l_{\max} = 2$ and three local interaction layers with a cutoff of 5.3Å. For this model, we included four energy components: local pairwise energies obtained via a neural network, a short-range ZBL repulsion energy, Coulomb interactions using fluctuating atom-centered charges and dispersion interactions via the two-body Optimized Quantum Drude Oscillator model⁸². The Coulomb and dispersion interactions are parameterized via charges and atomic volumes learned from MB-ISA¹¹¹ density partitioning provided in the dataset. The sum of partial charges is constrained to the total charge of the system using a neural charge equilibration inspired from ref. 112. We train the model on the ANI-1x dataset¹⁰⁹ using the same training setup as in section VIA. The weight ratio on the different loss components is 1:100:5:5 for total energies, force components, charges and volume ratios respectively. Training completed in 18 hours with a MAE of 1.10 kcal/mol/Å on forces, 0.85 kcal/mol on total energies, 0.0060 e on charges and 0.0059 on volume ratios.

For comparison, we trained a CRATE model with radial and equivariant resources and two message-passing layers with a cutoff of 5.3Å (yielding a receptive field of 10.6Å). Similarly to the Allegro model, we used 16 equivariant channels with $l_{\max} = 2$. At each layer, we perform two atom-centered channel-wise tensor products. Training completed in 5 hours which is comparable to the models in the previous section and much faster than for the Allegro model. This difference is mainly due to the pairwise nature of the Allegro model which is more computationally demanding than an atom-centered model.

The CRATE model achieved a MAE of 1.14 kcal/mol/Å on forces, 0.81 kcal/mol on total energies, 0.0050 e on charges and 0.0043 on volume ratios, showing similar performance to the Allegro model.

C. Learning distributed multipoles

For the last example, we trained an equivariant model to reproduce distributed MB-ISA multipoles (up to quadrupoles) from the Pubchem subset of the SPICE dataset¹¹⁰. We selected a CRATE embedding with radial, angular and equivariant resources with three interaction layers of cutoff 3.5Å. We used 32 equivariant channels with maximum degree $l_{\max} = 2$ and perform one tensor product per layer. We obtain the atomic dipoles and quadrupoles from a linear combination of the equivariant channels of the corresponding irreducible representation. To obtain total-charge-constrained partial charges, we use the scalar embedding produced by CRATE to define a weight on each atom that is then used to distribute all the valence electrons in the system (including eventual ionic charges). This constrains the total number of electrons and prevents the atoms from unphysically depleting more than their number of valence electrons with too large positive charges. We kept 10% of the 1.4 million configurations as a validation set. Furthermore, we pre-processed the quadrupoles in the dataset to make them traceless. We used a similar training setup as in the previous two examples and applied the same weight to all the multipole components in the loss function. Training completed in 14 hours and the model achieved a MAE of 0.004 e on charges, 0.0017 $e\text{Å}$ on dipole components and 0.0013 $e\text{Å}^2$ on quadrupole components on the validation set. For reference, the MAE for a trivial model that predicts only neutral spherical atoms is 0.27 e for charges, 0.019 $e\text{Å}$ for dipoles and 0.0087 $e\text{Å}^2$ for quadrupoles.

VII. CONCLUSION

We introduced FeNNol, a new library for building, training and running molecular dynamics simulations with machine-learning potentials, with a particular focus on force-field-enhanced neural network potentials. It leverages the Jax library⁶⁶ and the Flax⁷⁰ machine-learning framework to construct parameterized functions that efficiently run on accelerators and can be automatically differentiated. FeNNol enables the user to build complex models in a declarative way, without the need for explicit programming. The library provides an extensive collection of modules that cover atomic neighborhood descriptors (or embeddings), various architectures of neural networks, encoding schemes for spatial and chemical information, operation modules and physics modules (for example force-field interaction terms). These modules can be composed to build models that are tailored to the user's requirements. More advanced users can also

register custom modules that can be made very efficient without resorting to low-level implementation thanks to Jax's JIT compilation capabilities.

In order to showcase the highly configurable nature of FeNNol modules, we developed the CRATE atomic embedding that can be set up to exploit information from various chemical and geometric resources (radial basis, angular information, E(3) equivariant tensors, etc...) and combine them via either local or message-passing interaction layers. This embedding can thus be adjusted by the user depending on the complexity of the learning task and the computational cost requirements.

FeNNol provides a complete training system that allows defining multi-objective and multi-stage training in an easy and reproducible way. We showed with a few examples that it can be used to train models for different tasks (potential energy surfaces, atom-in-molecule properties, distributed multipoles) with state-of-the-art results in a few hours on a single commodity GPU.

Finally, we provide three main approaches for running MD simulations with ML models: an ASE calculator, an interface to Tinker-HP via Deep-HP⁶⁹ and an efficient native MD engine. The ASE calculator allows to transparently use all of the functionalities and benefit from the flexibility of the Atomic Simulation Environment⁹⁴. Similarly, the Deep-HP interface enables the use of all of Tinker-HP's advanced functionalities and its massive parallelization capabilities. Furthermore it allows for mixed ML/MM simulations as shown in ref. 69. We showed that our early version of the interface enables simulations of systems up to 500K atoms on a single A100 GPU with performance and scaling comparable to the AMOEBA force fields. Finally, FeNNol's native MD engine allows to perform very efficient MD simulations of intermediate size systems (more than 20K atoms) directly within the FeNNol package by leveraging the full capabilities of Jax's JIT compilation and asynchronous dispatch. We showed that our implementation of the ANI-2x model can reach simulation speeds nearly on par with an optimized version of the AMOEBA polarizable force field.

FeNNol thus constitutes a very modular framework for modeling complex chemical systems with physics-enhanced machine-learning models, allowing to easily experiment with different architectures and to iterate quickly, while remaining highly efficient. We hope that this library will accelerate the development and application of ML models and facilitate their access to a wider audience.

VIII. CODE AVAILABILITY

The FeNNol library is available on Github <https://github.com/thomasple/FeNNol> under the LGPL3 licence. Example input files are provided both for model training and molecular dynamics simulations.

IX. DATA AVAILABILITY

The input and output files for the three examples of section VI are provided in an archive file stored in the Zenodo repository at <https://zenodo.org/records/11146435>. Datasets used in these examples are publicly available and can be freely obtained from the relevant repositories. The preprocessed dataset files (formatted and with the specific data split used here) can be provided upon request.

X. ACKNOWLEDGEMENT

This work was made possible thanks to funding from the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation program (grant agreement No 810367), project EMC2. Computations have been performed at GENCI (IDRIS, Orsay, France and TGCC, Bruyères le Chatel) on grant no A0130712052.

Appendix A: Detailed description of the CRATE embedding

CRATE defines an embedding vector for each atom in the system that is iteratively refined over L layers by combining chemical and geometric "resources" from its local environment. The embedding x_i for atom i is initialized with an encoding of its chemical species e_{Z_i} (which can be any of the species encodings provided in FeNNol).

At each interaction layer l , the embedding is projected to two lower-dimensional spaces r_i, s_i via a learnable affine transform:

$$r_i = W_r^{(l)} x_i + b_r^{(l)} \quad (\text{A1})$$

$$s_i = W_s^{(l)} x_i + b_s^{(l)} \quad (\text{A2})$$

The r_i vector is used to retain information from the previous layer while the s_i vector will be combined with local resources. For each neighbour j of atom i , we then form a neighbour embedding s_{ij} . When using message passing, this neighbour embedding is simply s_j . Alternatively, when a fully local model is required, we need to form a neighbour embedding that accounts both for the current embedding of atom i and available local information on atom j , i.e. its chemical species. For efficiency's sake, we use a key-query mechanism similar to the one used in attention:

$$Q_i = W_Q^{(l)} s_i \quad ; \quad K_j = W_K^{(l)} e_{Z_j} \quad (\text{A3})$$

$$s_{ij} = Q_i^T K_j / \sqrt{d} \quad (\text{A4})$$

with $W_Q^{(l)} \in \mathbb{R}^{d \times \dim(s)}$ and $W_K^{(l)} \in \mathbb{R}^{d \times \dim(e_z)}$ two trainable matrices (with d an hyperparameter intermediate dimension).

The local geometric resources are then computed and concatenated with r_i to form a vector R_i containing all the resources from this layer.

$$R_i = r_i \parallel \left(\parallel_{g \in \text{res}(l)} g(x_i, \{s_{ij}\}) \right) \quad (\text{A5})$$

These resources are then mixed via a multi-layer perceptron $MLP^{(l)}$ to form the embedding update:

$$x_i^{(l+1)} = \sigma(F^{(l)}) \odot x_i^{(l)} + u(MLP^{(l)}(R_i)) \quad (\text{A6})$$

where u is an optional activation function, $F^{(l)} \in \mathbb{R}^{\dim(x)}$ and $\sigma(F^{(l)})$ is a trainable "forget gate" (with σ the sigmoid function and \odot the element-wise multiplication). The current implementation has four available resources: radial, angular, equivariant and long-range.

Radial resources radial resources are simply obtained by an outer product of the neighbour embedding with a projection of the distance into a radial basis $B(r_{ij})$. These pairwise chemical-radial information are then multiplied by a switching function $f(r_{ij})$ to ensure continuity when atoms enter or leave the neighborhood and then summed over the neighbours:

$$g_i^{\text{rad}} = \sum_j s_{ij} \otimes B(r_{ij}) f(r_{ij}) \quad (\text{A7})$$

The radial basis and the switching function can be chosen by the user from FeNNol's provided functional forms.

Angular resources angular resources are obtained by combining information from triplets of atoms. Importantly, the user can choose the cutoff defining the neighborhood for triplets of atoms to be smaller than for the radial resources. This is significant in the condensed phase as the number of triplets of atoms quickly grows with the cutoff radius. To build angular resources, we start by forming a reduced chemical-radial basis D_{ij} and D_{ik} for the two edges:

$$[D_{ij}]_c = \tilde{f}(r_{ij}) \sum_{ab} [\tilde{B}(r_{ij})]_a [s_{ij}]_b [W_{\text{ang}}^{(l)}]_{abc} \quad (\text{A8})$$

with $W_{\text{ang}}^{(l)} \in \mathbb{R}^{\dim(\tilde{B}) \times \dim(s) \times d_a}$ a trainable tensor and d_a a hyperparameter of the model. \tilde{f} is the switching function adapted to the cutoff defined for triplets and \tilde{B} is a radial basis for the triplet graph (that can be identical to B or not). We then project the angle between the two edges into an angular basis Θ_{ijk} (this basis can be either a small fourier expansion $[\Theta_{ijk}]_n = \cos(n\theta_{ijk})$ or the angular symmetry function used in the ANI models³³ $[\Theta_{ijk}]_n = (0.5 + 0.5 \cos(\theta_{ijk} - \theta_n))^5$) and combine the information as:

$$g_i^{\text{ang}} = \sum_{\{jk\}} \Theta_{ijk} \otimes (D_{ij} \odot D_{ik}) \quad (\text{A9})$$

equivariant resources equivariant resources are obtained by forming an equivariant neighborhood basis from a chemical-radial basis and spherical harmonics and iteratively refining an equivariant embedding via multiple atom-centered tensor products, similarly to what is proposed in the MACE model⁸⁶. The equivariant neighborhood basis $\hat{\rho}_i$ (we denote equivariant tensors with hat notation) is initialized to zero and updated at each layer with:

$$b_{ij}^{(l)} = W_e^{(l)}(s_{ij} \otimes B(r_{ij})f(r_{ij})) \quad (\text{A10})$$

$$\hat{V}_{ij}^{(l)} = \begin{cases} \hat{V}_j^{(l-1)} & \text{if } l > 1 \text{ and message-passing,} \\ \text{else } \hat{Y}_{ij} & \end{cases} \quad (\text{A11})$$

$$\hat{\rho}_i^{(l)} = \hat{\rho}_i^{(l-1)} + \sum_j b_{ij}^{(l)} \otimes \hat{V}_{ij}^{(l)} \quad (\text{A12})$$

with \hat{V}_j the current equivariant embedding of atom j and \hat{Y}_{ij} the tensor representing the coefficients of the direction between atom i and j in the real spherical harmonics basis up to a user-defined maximum degree l_{\max} . The trainable matrix $W_e^{(l)}$ converts the radial-chemical basis to a user-selected N_{channels} equivariant channels. This number of channels is typically kept small to limit the high computational cost of equivariant tensor products. For the first interaction layer, we initialize the equivariant embedding \hat{V}_i as a linear combination of the irreps of each channel: $[\hat{V}_i^{(0)}]_{n\lambda m} = \sum_{n'} [W_v^{(0)}]_{nn'\lambda} [Y_{ij}]_{n'\lambda m}$.

After forming the neighborhood basis, we perform a series of N_{TP} tensor products to update V_i :

$$[\hat{L}_i]_{n\lambda m}^{(l,t)} = \sum_{p=(\lambda_1, \lambda_2 \rightarrow \lambda)} [W_{\text{path}}^{(l,t)}]_p \left(\sum_{m_1, m_2} C_{m_1 m_2 m}^{\lambda_1 \lambda_2 \lambda} [\hat{\rho}_i]_{n\lambda_1 m_1} [\hat{V}_i^{(l,t-1)}]_{n\lambda_2 m_2} \right) \quad (\text{A13})$$

$$[\hat{V}_i^{(l,t)}]_{n\lambda m} = \sum_{n'} [W_v^{(l,t)}]_{nn'} [\hat{V}_i^{(l,t-1)} + \hat{L}_i^{(l,t)}]_{n'\lambda m} \quad (\text{A14})$$

where $C_{m_1 m_2 m}^{\lambda_1 \lambda_2 \lambda}$ are Clebsch–Gordan coefficients for the basis of real spherical harmonics, the index p enumerates tensor product paths that produce a certain irrep of degree λ (the user can choose to include either all the paths that produce a certain λ or only the ones which conserve the parity of the output irrep). The trainable parameters $W_{\text{path}}^{(l,t)}$ combine the different tensor product paths to conserve the number of irreps. This is similar to the scheme proposed in ref. 79 and is key to prevent the explosion of the number of paths after iterating tensor products. Equation (A14) corresponds to an update of V_i followed by channel mixing. We initialize as $\hat{V}_i^{(l,0)} = \hat{V}_i^{(l-1)}$ and keep the last iteration as the current equivariant embedding $\hat{V}_i^{(l)} = \hat{V}_i^{(l, N_{\text{TP}})}$.

Finally, equivariant resources at layer l are obtained by concatenating the scalar irreps from the different $\hat{L}_i^{(l,t)}$ as:

$$g_i^{\text{E3}} = \parallel_{t,n} [\hat{L}_i^{(l,t)}]_{n00} \quad (\text{A15})$$

Long-range resources long-range resources on atom i are obtained as the electrostatic potential generated by a set of fictitious point charges placed on all the other atoms. We start by obtaining N_q sets of point charges and short-range damping coefficients as:

$$[q_i]_n, [\tilde{\gamma}_i]_n = \left[\text{MLP}^{(q)}(e_{Z_i}) \right]_n \quad (\text{A16})$$

$$\gamma_i = \text{softplus}(\tilde{\gamma}_i) \quad (\text{A17})$$

where $\text{MLP}^{(q)}$ is a multi-layer perceptron outputting $2N_q$ numbers. For each atom, we also have N_q trainable effective radii a_i that are initialized to the atom's Van der Waals radius and we obtain a pairwise damping as $\gamma_{ij} = (a_i^2 + a_j^2)^{-1/2}$. The long-range resources are then obtained as:

$$[g_i^{lr}]_n = \sum_j [q_j]_n \frac{\text{erf}([\gamma_i]_n \gamma_{ij} r_{ij})}{r_{ij}} \quad (\text{A18})$$

This scheme is inspired from the LODE framework^{45,46} but we simplified it to keep only invariant quantities instead of full equivariant tensors. In the current implementation, long-range resources are only computed at the first iteration and fed back unchanged into the embedding update of equation (A5) at each layer. This avoids multiple iterations of costly long-range interactions. For now, the long-range embedding is not compatible with periodic boundary conditions.

¹J. Wang, R. M. Wolf, J. W. Caldwell, P. A. Kollman, and D. A. Case, *Journal of computational chemistry* **25**, 1157 (2004).

²K. Vanommeslaeghe, E. Hatcher, C. Acharya, S. Kundu, S. Zhong, J. Shim, E. Darian, O. Guvench, P. Lopes, I. Vorobyov, *et al.*, *Journal of computational chemistry* **31**, 671 (2010).

³W. L. Jorgensen, D. S. Maxwell, and J. Tirado-Rives, *Journal of the American Chemical Society* **118**, 11225 (1996), <https://doi.org/10.1021/ja9621760>.

⁴N. Gresh, G. A. Cisneros, T. A. Darden, and J.-P. Piquemal, *Journal of Chemical Theory and Computation* **3**, 1960 (2007).

This is the author's peer reviewed, accepted manuscript. However, the online version of record will be different from this version once it has been copyedited and typeset.

PLEASE CITE THIS ARTICLE AS DOI: 10.1063/5.0217688

- ⁵S. Naseem-Khan, L. Lagardère, C. Narth, G. A. Cisneros, P. Ren, N. Gresh, and J.-P. Piquemal, arXiv preprint arXiv:2201.00804 (2022).
- ⁶J.-P. Piquemal, G. A. Cisneros, P. Reinhardt, N. Gresh, and T. A. Darden, *The Journal of chemical physics* **124** (2006).
- ⁷R. E. Duke, O. N. Starovoytov, J.-P. Piquemal, and G. A. Cisneros, *Journal of Chemical Theory and Computation* **10**, 1361 (2014).
- ⁸S. K. Reddy, S. C. Straight, P. Bajaj, C. Huy Pham, M. Riera, D. R. Moberg, M. A. Morales, C. Knight, A. W. Götz, and F. Paesani, *The Journal of Chemical Physics* **145**, 194504 (2016).
- ⁹X. Zhu, M. Riera, E. F. Bull-Vulpe, and F. Paesani, *Journal of Chemical Theory and Computation* (2023).
- ¹⁰P. Ren and J. W. Ponder, *The Journal of Physical Chemistry B* **107**, 5933 (2003), <https://doi.org/10.1021/jp027815+>.
- ¹¹J. W. Ponder, C. Wu, P. Ren, V. S. Pande, J. D. Chodera, M. J. Schnieders, I. Haque, D. L. Mobley, D. S. Lambrecht, R. A. DiStasio Jr, *et al.*, *The Journal of Physical Chemistry B* **114**, 2549 (2010).
- ¹²C. Liu, J.-P. Piquemal, and P. Ren, *Journal of Chemical Theory and Computation* **15**, 4122 (2019), pMID: 31136175, <https://doi.org/10.1021/acs.jctc.9b00261>.
- ¹³C. Liu, J.-P. Piquemal, and P. Ren, *The Journal of Physical Chemistry Letters* **11**, 419 (2020), pMID: 31865706, <https://doi.org/10.1021/acs.jpcllett.9b03489>.
- ¹⁴G. Nawrocki, I. Leontyev, S. Sakipov, M. Darkhovskiy, I. Kurnikov, L. Pereyaslavets, G. Kamath, E. Voronina, O. Butin, A. Illarionov, *et al.*, *Journal of Chemical Theory and Computation* **18**, 7751 (2022).
- ¹⁵J. Melcer and J.-P. Piquemal, *Frontiers in molecular biosciences* **6**, 143 (2019).
- ¹⁶Y. Shi, P. Ren, M. Schnieders, and J.-P. Piquemal, “Polarizable force fields for biomolecular modeling,” in *Reviews in Computational Chemistry Volume 28* (John Wiley and Sons, Ltd, 2015) Chap. 2, pp. 51–86.
- ¹⁷Z. Jing, C. Liu, S. Y. Cheng, R. Qi, B. D. Walker, J.-P. Piquemal, and P. Ren, *Ann. Rev. Biophys.* **48**, 371 (2019).
- ¹⁸J. Wang, W. Wang, P. A. Kollman, and D. A. Case, *J. Am. Chem. Soc* **222**, 2001 (2001).
- ¹⁹B. Walker, C. Liu, E. Wait, and P. Ren, *Journal of computational chemistry* **43**, 1530 (2022).
- ²⁰Y. Wang, J. Fass, B. Kaminow, J. E. Herr, D. Rufa, I. Zhang, I. Pulido, M. Henry, H. E. B. Macdonald, K. Takaba, *et al.*, *Chemical Science* **13**, 12016 (2022).
- ²¹G. Chen, T. J. Inizan, T. Plé, L. Lagardere, J.-P. Piquemal, and Y. Maday, (2023).
- ²²A. P. Thompson, H. M. Aktulga, R. Berger, D. S. Bolintineanu, W. M. Brown, P. S. Crozier, P. J. In't Veld, A. Kohlmeyer, S. G. Moore, T. D. Nguyen, *et al.*, *Computer Physics Communications* **271**, 108171 (2022).
- ²³D. A. Case, T. A. Darden, T. E. Cheatham, C. L. Simmerling, J. Wang, R. E. Duke, R. Luo, M. Crowley, R. C. Walker, W. Zhang, *et al.*, (2008).
- ²⁴D. Van Der Spoel, E. Lindahl, B. Hess, G. Groenhof, A. E. Mark, and H. J. Berendsen, *Journal of computational chemistry* **26**, 1701 (2005).
- ²⁵J. C. Phillips, R. Braun, W. Wang, J. Gumbart, E. Tajkhorshid, E. Villa, C. Chipot, R. D. Skeel, L. Kale, and K. Schulten, *Journal of computational chemistry* **26**, 1781 (2005).
- ²⁶L. Lagardère, L.-H. Jolly, F. Lipparini, F. Aviat, B. Stamm, Z. F. Jing, M. Harger, H. Torabifard, G. A. Cisneros, M. J. Schnieders, N. Gresh, Y. Maday, P. Y. Ren, J. W. Ponder, and J.-P. Piquemal, *Chemical Science* **9**, 956 (2018).
- ²⁷O. Adjoua, L. Lagardère, L.-H. Jolly, A. Durocher, T. Very, I. Dupays, Z. Wang, T. Jaffrelot Inizan, F. Célerse, P. Ren, J. W. Ponder, and J.-P. Piquemal, *Journal of Chemical Theory and Computation* **17**, 2034 (2021).
- ²⁸J. Behler and M. Parrinello, *Physical Review Letters* **98**, 146401 (2007).
- ²⁹A. P. Bartók, M. C. Payne, R. Kondor, and G. Csányi, *Physical review letters* **104**, 136403 (2010).
- ³⁰S. Chmiela, H. E. Sauceda, K.-R. Müller, and A. Tkatchenko, *Nature Communications* **9**, 1 (2018).
- ³¹F. Bigi, S. N. Pozdnyakov, and M. Ceriotti, arXiv preprint arXiv:2303.04124 (2023).
- ³²K. Shakouri, J. Behler, J. Meyer, and G.-J. Kroes, *The journal of physical chemistry letters* **8**, 2131 (2017).
- ³³J. S. Smith, O. Isayev, and A. E. Roitberg, *Chemical Science* **8**, 3192 (2017).
- ³⁴K. T. Schütt, F. Arbabzadah, S. Chmiela, K. R. Müller, and A. Tkatchenko, *Nature communications* **8**, 13890 (2017).
- ³⁵K. Schütt, P.-J. Kindermans, H. E. Sauceda Felix, S. Chmiela, A. Tkatchenko, and K.-R. Müller, *Advances in neural information processing systems* **30** (2017).
- ³⁶J. Gilmer, S. S. Schoenholz, P. F. Riley, O. Vinyals, and G. E. Dahl, in *International conference on machine learning* (PMLR, 2017) pp. 1263–1272.
- ³⁷N. Lubbers, J. S. Smith, and K. Barros, *The Journal of Chemical Physics* **148** (2018).
- ³⁸R. Zubatyuk, J. S. Smith, J. Leszczynski, and O. Isayev, *Science advances* **5**, eaav6490 (2019).
- ³⁹V. Zaverkin and J. Kästner, *Journal of Chemical Theory and Computation* **16**, 5410 (2020).
- ⁴⁰Z. Qiao, A. S. Christensen, M. Welborn, F. R. Manby, A. Anandkumar, and T. F. Miller III, arXiv preprint arXiv:2105.14655 (2021).
- ⁴¹O. T. Unke, S. Chmiela, M. Gastegger, K. T. Schütt, H. E. Sauceda, and K.-R. Müller, *Nature Communications* **12**, 1 (2021).
- ⁴²J. Gastegger, F. Becker, and S. Günnemann, *Advances in Neural Information Processing Systems* **34**, 6790 (2021).
- ⁴³S. Batzner, A. Musaelian, L. Sun, M. Geiger, J. P. Mailoa, M. Kornbluth, N. Molinari, T. E. Smidt, and B. Kozinsky, *Nature Communications* **13**, 1 (2022).
- ⁴⁴A. Musaelian, S. Batzner, A. Johansson, L. Sun, C. J. Owen, M. Kornbluth, and B. Kozinsky, *Nature Communications* **14**, 579 (2023).
- ⁴⁵A. Grisafi and M. Ceriotti, *The Journal of Chemical Physics* **151**, 204105 (2019).
- ⁴⁶A. Grisafi, J. Nigam, and M. Ceriotti, *Chemical Science* **12**, 2078 (2021).
- ⁴⁷C. M. Handley and P. L. Popelier, *Journal of chemical theory and computation* **5**, 1474 (2009).
- ⁴⁸N. Artrith, T. Morawietz, and J. Behler, *Physical Review B* **83**, 153101 (2011).
- ⁴⁹T. W. Ko, J. A. Finkler, S. Goedecker, and J. Behler, *Nature Communications* **12**, 1 (2021).
- ⁵⁰O. T. Unke and M. Meuwly, *Journal of Chemical Theory and Computation* **15**, 3678 (2019).
- ⁵¹C. Qu, Q. Yu, P. L. Houston, R. Conte, A. Nandi, and J. M. Bowman, *Journal of Chemical Theory and Computation* (2023).
- ⁵²L. Yang, J. Li, F. Chen, and K. Yu, *The Journal of Chemical Physics* **157**, 214108 (2022).
- ⁵³L. Zhang, H. Wang, M. C. Muniz, A. Z. Panagiotopoulos, R. Car, and W. E, *The Journal of Chemical Physics* **156**, 124107 (2022).
- ⁵⁴J. M. Bowman, C. Qu, R. Conte, A. Nandi, P. L. Houston, and Q. Yu, *Journal of Chemical Theory and Computation* **19**, 1 (2022).
- ⁵⁵J. Chen and K. Yu, *Journal of Chemical Theory and Computation* **20**, 253 (2023).
- ⁵⁶Y. Wang, T. J. Inizan, C. Liu, J.-P. Piquemal, and P. Ren, *The Journal of Physical Chemistry B* **128**, 2381 (2024), pMID: 38445577, <https://doi.org/10.1021/acs.jpcc.3c08166>.
- ⁵⁷T. Plé, L. Lagardère, and J.-P. Piquemal, *Chemical Science* **14**, 12554 (2023).
- ⁵⁸M. Pinheiro, F. Ge, N. Ferré, P. O. Dral, and M. Barbatti, *Chemical Science* **12**, 14396 (2021).

This is the author's peer reviewed, accepted manuscript. However, the online version of record will be different from this version once it has been copyedited and typeset.

PLEASE CITE THIS ARTICLE AS DOI: 10.1063/5.0217688

- ⁵⁹E. Kocer, T. W. Ko, and J. Behler, Annual review of physical chemistry **73**, 163 (2022).
- ⁶⁰K. T. Schütt, S. S. Hessmann, N. W. Gebauer, J. Lederer, and M. Gastegger, The Journal of Chemical Physics **158** (2023).
- ⁶¹P. O. Dral, F. Ge, Y.-F. Hou, P. Zheng, Y. Chen, M. Barbatti, O. Isayev, C. Wang, B.-X. Xue, M. Pinheiro Jr, *et al.*, Journal of Chemical Theory and Computation (2024).
- ⁶²J. Zeng, D. Zhang, D. Lu, P. Mo, Z. Li, Y. Chen, M. Rynik, L. Huang, Z. Li, S. Shi, *et al.*, The Journal of Chemical Physics **159** (2023).
- ⁶³T. Bereau, D. Andrienko, and O. A. Von Lilienfeld, Journal of chemical theory and computation **11**, 3225 (2015).
- ⁶⁴Z. L. Glick, A. Koutsoukas, D. L. Cheney, and C. D. Sherrill, The Journal of Chemical Physics **154** (2021).
- ⁶⁵M. Thurlemann, L. Boselt, and S. Riniker, Journal of Chemical Theory and Computation **18**, 1701 (2022).
- ⁶⁶J. Bradbury, R. Frostig, P. Hawkins, M. J. Johnson, C. Leary, D. Maclaurin, G. Neula, A. Paszke, J. VanderPlas, S. Wanderman-Milne, and Q. Zhang, "JAX: composable transformations of Python+NumPy programs," (2018).
- ⁶⁷DeepMind, I. Babuschkin, K. Baumli, A. Bell, S. Bhupatiraju, J. Bruce, P. Buchlovsky, D. Budden, T. Cai, A. Clark, I. Danihelka, A. Dedieu, C. Fantacci, J. Godwin, C. Jones, R. Hemsley, T. Hennigan, M. Hessel, S. Hou, S. Kapturowski, T. Keck, I. Kemaev, M. King, M. Kunesch, L. Martens, H. Merzic, V. Mikulik, T. Norman, G. Papamakarios, J. Quan, R. Ring, F. Ruiz, A. Sanchez, L. Sartran, R. Schneider, E. Sezener, S. Spencer, S. Srinivasan, M. Stanojević, W. Stokowiec, L. Wang, G. Zhou, and F. Viola, "The DeepMind JAX Ecosystem," (2020).
- ⁶⁸C. Devereux, J. S. Smith, K. K. Huddleston, K. Barros, R. Zubatyuk, O. Isayev, and A. E. Roitberg, Journal of Chemical Theory and Computation **16**, 4192 (2020).
- ⁶⁹T. Jaffrelot Inizan, T. Plé, O. Adjoua, P. Ren, H. Gökcän, O. Isayev, L. Lagardère, and J.-P. Piquemal, Chemical Science **14**, 5438 (2023).
- ⁷⁰J. Heek, A. Levskaya, A. Oliver, M. Ritter, B. Rondepierre, A. Steiner, and M. van Zee, "Flax: A neural network library and ecosystem for JAX," (2023).
- ⁷¹P. P. Ewald, Ann. Phys **369**, 1 (1921).
- ⁷²J. Gasteiger, J. Groß, and S. Günnemann, arXiv preprint arXiv:2003.03123 (2020).
- ⁷³M. Rupp, A. Tkatchenko, K.-R. Müller, and O. A. Von Lilienfeld, Physical Review Letters **108**, 058301 (2012).
- ⁷⁴M. Eckhoff and M. Reiher, Journal of Chemical Theory and Computation **19**, 3509 (2023).
- ⁷⁵L. Zhang, J. Han, H. Wang, W. Saidi, R. Car, *et al.*, Advances in neural information processing systems **31** (2018).
- ⁷⁶K. Schütt, O. Unke, and M. Gastegger, in *International Conference on Machine Learning* (PMLR, 2021) pp. 9377–9388.
- ⁷⁷M. Haghighatlari, J. Li, X. Guan, O. Zhang, A. Das, C. J. Stein, F. Heidar-Zadeh, M. Liu, M. Head-Gordon, L. Bertels, *et al.*, Digital Discovery **1**, 333 (2022).
- ⁷⁸M. Geiger and T. Smidt, arXiv preprint arXiv:2207.09453 (2022).
- ⁷⁹B. Kozinsky, A. Musaelian, A. Johansson, and S. Batzner, in *Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis* (2023) pp. 1–12.
- ⁸⁰S. Takamoto, S. Izumi, and J. Li, Computational Materials Science **207**, 111280 (2022).
- ⁸¹U. Essmann, L. Perera, M. L. Berkowitz, T. Darden, H. Lee, and L. G. Pedersen, The Journal of chemical physics **103**, 8577 (1995).
- ⁸²A. Khabibrakhmanov, D. V. Fedorov, and A. Tkatchenko, Journal of Chemical Theory and Computation **19**, 7895 (2023).
- ⁸³E. Caldeweyher, S. Ehlert, A. Hansen, H. Neugebauer, S. Spicher, C. Bannwarth, and S. Grimme, The Journal of Chemical Physics **150** (2019).
- ⁸⁴J. F. Ziegler and J. P. Biersack, in *Treatise on heavy-ion science: volume 6: astrophysics, chemistry, and condensed matter* (Springer, 1985) pp. 93–129.
- ⁸⁵E. J. Hu, Y. Shen, P. Wallis, Z. Allen-Zhu, Y. Li, S. Wang, L. Wang, and W. Chen, arXiv preprint arXiv:2106.09685 (2021).
- ⁸⁶I. Batatia, D. P. Kovacs, G. Simm, C. Ortner, and G. Csányi, Advances in Neural Information Processing Systems **35**, 11423 (2022).
- ⁸⁷J. Zhuang, T. Tang, Y. Ding, S. C. Tatikonda, N. Dvornek, X. Papademetris, and J. Duncan, Advances in neural information processing systems **33**, 18795 (2020).
- ⁸⁸I. Loshchilov and F. Hutter, arXiv preprint arXiv:1711.05101 (2017).
- ⁸⁹A. Brock, S. De, S. L. Smith, and K. Simonyan, in *International Conference on Machine Learning* (PMLR, 2021) pp. 1059–1071.
- ⁹⁰L. N. Smith and N. Topin, in *Artificial intelligence and machine learning for multi-domain operations applications*, Vol. 11006 (SPIE, 2019) pp. 369–386.
- ⁹¹M. Kellner and M. Ceriotti, arXiv preprint arXiv:2402.16621 (2024).
- ⁹²A. Amini, W. Schwarting, A. Soleimany, and D. Rus, Advances in neural information processing systems **33**, 14927 (2020).
- ⁹³N. Meinert and A. Lavin, arXiv preprint arXiv:2104.06135 (2021).
- ⁹⁴A. H. Larsen, J. J. Mortensen, J. Blomqvist, I. E. Castelli, R. Christensen, M. Dulak, J. Friis, M. N. Groves, B. Hammer, C. Hargus, *et al.*, Journal of Physics: Condensed Matter **29**, 273002 (2017).
- ⁹⁵V. Kapil, M. Rossi, O. Marsalek, R. Petraglia, Y. Litman, T. Spura, B. Cheng, A. Cuzzocrea, R. H. Meißner, D. M. Wilkins, *et al.*, Computer Physics Communications **236**, 214 (2019).
- ⁹⁶T. Plé, N. Mauger, O. Adjoua, T. J. Inizan, L. Lagardère, S. Huppert, and J.-P. Piquemal, Journal of Chemical Theory and Computation **19**, 1432 (2023), PMID: 36856658, <https://doi.org/10.1021/acs.jctc.2c01233>.
- ⁹⁷S.-L. J. Lahey and C. N. Rowley, Chemical Science **11**, 2362 (2020).
- ⁹⁸R. Galvelis, A. Varela-Rial, S. Doerr, R. Fino, P. Eastman, T. E. Markland, J. D. Chodera, and G. De Fabritiis, Journal of Chemical Information and Modeling **63**, 5701 (2023).
- ⁹⁹A. Illarionov, S. Sakipov, L. Pereyaslavets, I. V. Kurnikov, G. Kamath, O. Butin, E. Voronina, I. Ivahnenko, I. Leontyev, G. Nawrocki, *et al.*, Journal of the American Chemical Society **145**, 23620 (2023).
- ¹⁰⁰Y. Wang, T. J. Inizan, C. Liu, J.-P. Piquemal, and P. Ren, The Journal of Physical Chemistry B **128**, 2381 (2024).
- ¹⁰¹B. Leimkuhler and C. Matthews, Applied Mathematics Research eXpress **2013**, 34 (2013).
- ¹⁰²E. Mangaud, S. Huppert, T. Plé, P. Depondt, S. Bonella, and F. Finocchi, Journal of Chemical Theory and Computation **15**, 2863 (2019).
- ¹⁰³S. Habershon, D. E. Manolopoulos, T. E. Markland, and T. F. Miller III, Annual Review of Physical Chemistry **64**, 387 (2013).
- ¹⁰⁴N. Mauger, T. Plé, L. Lagardère, S. Bonella, E. Mangaud, J.-P. Piquemal, and S. Huppert, The Journal of Physical Chemistry Letters **12**, 8285 (2021), PMID: 34427440, <https://doi.org/10.1021/acs.jpcltt.1c01722>.
- ¹⁰⁵X. Gao, F. Ramezanghorbani, O. Isayev, J. S. Smith, and A. E. Roitberg, Journal of Chemical Information and Modeling **60**, 3408 (2020).
- ¹⁰⁶R. Salomon-Ferrer, D. A. Case, and R. C. Walker, Wiley Interdisciplinary Reviews: Computational Molecular Science **3**, 198 (2013).
- ¹⁰⁷L. Lagardère, F. Aviat, and J.-P. Piquemal, The journal of physical chemistry letters **10**, 2593 (2019).
- ¹⁰⁸J. Hoja, L. Medrano Sandonas, B. G. Ernst, A. Vazquez-Mayagoitia, R. A. DiStasio Jr, and A. Tkatchenko, Scientific data **8**, 43 (2021).
- ¹⁰⁹J. S. Smith, R. Zubatyuk, B. Nebgen, N. Lubbers, K. Barros, A. E. Roitberg, O. Isayev, and S. Tretiak, Scientific data **7**, 134

(2020).
¹¹⁰P. Eastman, P. K. Behara, D. L. Dotson, R. Galvelis, J. E. Herr, J. T. Horton, Y. Mao, J. D. Chodera, B. P. Pritchard, Y. Wang, *et al.*, *Scientific Data* **10**, 11 (2023).

¹¹¹T. Verstraelen, S. Vandenbrande, F. Heidar-Zadeh, L. Vanduyfhuys, V. Van Speybroeck, M. Waroquier, and P. W. Ayers, *Journal of Chemical Theory and Computation* **12**, 3894 (2016).

¹¹²R. Zubatyuk, J. S. Smith, B. T. Nebgen, S. Tretiak, and O. Isayev, *Nature Communications* **12**, 1 (2021).

This is the author's peer reviewed, accepted manuscript. However, the online version of record will be different from this version once it has been copyedited and typeset.

PLEASE CITE THIS ARTICLE AS DOI: 10.1063/5.0217688